



**Publicly Available Specification (PAS);
O-RAN Fronthaul Control, User and Synchronization Plane
Specification v12.01;
(O-RAN.WG4.CUS.0-R003-v12.01)**

CAUTION

The present document has been submitted to ETSI as a PAS produced by O-RAN Alliance and approved by the ETSI Technical Committee Mobile Standards Group (MSG).

ETSI had been assigned all the relevant copyrights related to the document O-RAN.WG4.CUS.0-R003-v12.01 on an "as is basis". Consequently, to the fullest extent permitted by law, ETSI disclaims all warranties whether express, implied, statutory or otherwise including but not limited to merchantability, non-infringement of any intellectual property rights of third parties. No warranty is given about the accuracy and the completeness of the content of the present document.

Reference

RTS/MSG-001142

Keywords

control, front haul, management plane, PAS

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Important notice

The present document can be downloaded from:
<https://www.etsi.org/standards-search>

The present document may be made available in electronic versions and/or in print. The content of any electronic and/or print versions of the present document shall not be modified without the prior written authorization of ETSI. In case of any existing or perceived difference in contents between such versions and/or in print, the prevailing version of an ETSI deliverable is the one made publicly available in PDF format at www.etsi.org/deliver.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at <https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx>

If you find errors in the present document, please send your comment to one of the following services:
<https://portal.etsi.org/People/CommitteeSupportStaff.aspx>

If you find a security vulnerability in the present document, please report it through our Coordinated Vulnerability Disclosure Program:
<https://www.etsi.org/standards/coordinated-vulnerability-disclosure>

Notice of disclaimer & limitation of liability

The information provided in the present deliverable is directed solely to professionals who have the appropriate degree of experience to understand and interpret its content in accordance with generally accepted engineering or other professional standard and applicable regulations.

No recommendation as to products and services or vendors is made or should be implied.

No representation or warranty is made that this deliverable is technically accurate or sufficient or conforms to any law and/or governmental rule and/or regulation and further, no representation or warranty is made of merchantability or fitness for any particular purpose or against infringement of intellectual property rights.

In no event shall ETSI be held liable for loss of profits or any other incidental or consequential damages.

Any software contained in this deliverable is provided "AS IS" with no warranties, express or implied, including but not limited to, the warranties of merchantability, fitness for a particular purpose and non-infringement of intellectual property rights and ETSI shall not be held liable in any event for any damages whatsoever (including, without limitation, damages for loss of profits, business interruption, loss of information, or any other pecuniary loss) arising out of or related to the use of or inability to use the software.

Copyright Notification

No part may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm except as authorized by written permission of ETSI.

The content of the PDF version shall not be modified without the written authorization of ETSI.
The copyright and the foregoing restriction extend to reproduction in all media.

© ETSI 2024.
All rights reserved.

Contents

Intellectual Property Rights	15
Foreword.....	15
Modal verbs terminology.....	15
1 Scope	16
2 References	16
2.1 Normative references	16
2.2 Informative references.....	18
3 Definition of terms, symbols and abbreviations.....	19
3.1 Terms.....	19
3.2 Symbols.....	21
3.3 Abbreviations	21
3.4 Document conventions	22
3.4.1 Terminologies	22
3.4.2 Fields and bitmasks in messages.....	23
4 Architecture & requirements	24
4.1 Architectural aspects	24
4.1.1 Architectural entities	24
4.2 Functional split.....	24
4.2.1 Introduction.....	24
4.2.2 Split 7-2x (DL)	25
4.2.2.1 Description	25
4.2.2.2 Benefits and justification (informative)	27
4.2.3 Split 7-2x (UL)	28
4.2.3.1 Description	28
4.2.3.2 Benefits and justification (informative)	29
4.3 Data flows	29
4.4 Latency requirements	31
4.4.1 Introduction.....	31
4.4.2 Timing parameter relationships	32
4.4.2.1 General	32
4.4.2.2 O-DU Transmission Window.....	33
4.4.3 U-Plane/ C-Plane timing	34
4.4.4 Computed latency methods.....	39
4.4.4.1 Introduction.....	39
4.4.4.2 Fronthaul timing domain.....	41
4.4.4.3 Defined Transport method	42
4.4.4.4 Measured Transport method.....	43
4.4.5 Latency categories for O-DU with dynamic timing advance.....	46
4.4.6 Latency categories for O-DU with fixed timing advance	48
4.4.7 Non Delay-managed U-Plane traffic	49
4.5 Reception window monitoring	50
4.6 Transmission windows	51
4.6.1 Normal transmission	51
4.6.2 Uniformly distributed transmission	51
4.6.3 Ordered transmission	51
4.6.4 Scheduled transmission.....	52
4.7 O-RU external antenna delay handling	53
4.7.1 Introduction.....	53
4.7.2 Minimal O-DU impact method (defined transport method)	53
4.7.2.1 Description.....	53
4.7.2.2 Minimal O-DU impact method - Example.....	54
4.8 UL frame timing to DL frame	57
4.8.1 Overview	57
4.8.2 Carrier types and default values of n-ta-offset.....	58

5	Transport & Protocol Architecture.....	59
5.1	Transport encapsulation types.....	59
5.1.1	Ethernet encapsulation.....	59
5.1.2	IP/UDP encapsulation.....	59
5.1.3	Transport headers.....	60
5.1.3.1	Introduction.....	60
5.1.3.2	eCPRI transport header.....	60
5.1.3.2.1	Overview.....	60
5.1.3.2.2	ecpriVersion (eCPRI protocol revision).....	61
5.1.3.2.3	ecpriReserved (eCPRI reserved).....	61
5.1.3.2.4	ecpriConcatenation (eCPRI concatenation indicator).....	61
5.1.3.2.5	ecpriMessage (eCPRI message type).....	61
5.1.3.2.6	ecpriPayload (eCPRI payload size).....	61
5.1.3.2.7	ecpriRtcid / ecpriPcid (real time control data / IQ data transfer message series identifier).....	62
5.1.3.2.8	ecpriSeqid (message identifier).....	64
5.1.3.3	IEEE 1914.3 transport header.....	66
5.1.3.3.1	Overview.....	66
5.1.3.3.2	RoEsubType (sub type / message type).....	66
5.1.3.3.3	RoEflowID (flow identifier).....	67
5.1.3.3.4	RoElength (length).....	67
5.1.3.3.5	RoEorderInfo (order information).....	67
5.2	Protocol architecture.....	68
5.2.1	C-Plane.....	68
5.2.2	U-Plane.....	68
5.2.3	S-Plane.....	69
5.3	Quality of service.....	69
5.4	Data flow identification.....	70
5.5	Fragmentation.....	70
5.5.1	Introduction.....	70
5.5.2	Application layer fragmentation.....	70
5.5.3	Radio transport layer (eCPRI or IEEE-1914.3) fragmentation.....	71
5.5.4	Fragmentation considerations.....	71
6	Security.....	73
6.1	General.....	73
7	C-Plane Protocol.....	74
7.1	General.....	74
7.1.1	C-Plane transport.....	74
7.2	Elementary procedures.....	74
7.2.1	Scheduling and beamforming commands transfer procedure.....	74
7.2.2	Symbol numbering and duration.....	75
7.2.3	Mixed numerology and PRACH handling.....	75
7.2.3.1	Overview.....	75
7.2.3.2	Frequency indexing across multiple numerologies.....	75
7.2.3.3	Slot indexing with mixed numerologies.....	78
7.2.3.4	PRACH formats with multiple PRACH Symbols.....	80
7.2.4	DL precoding configuration parameters and indications.....	80
7.2.5	LAA commands transfer procedure.....	82
7.2.5.1	Introduction.....	82
7.2.5.2	LBT procedure overview.....	82
7.2.5.3	Definitions.....	83
7.2.5.4	General guidelines for the LAA-procedure.....	84
7.2.5.4.1	PDSCH Transmission.....	84
7.2.5.4.2	DRS Transmission.....	84
7.2.5.4.3	Congestion window information transmission.....	85
7.2.6	Dynamic Spectrum Sharing (DSS).....	85
7.2.6.1	Overview.....	85
7.2.6.2	Dynamic Spectrum Sharing (DSS) via dedicated endpoints.....	85
7.2.6.3	Dynamic Spectrum Sharing (DSS) via Section Extension 9 for DSS.....	85
7.2.7	Channel information based beamforming.....	86
7.2.8	ACK/NACK Procedure for C-Plane Messages.....	86

7.2.9	Section Type 4: Slot Configuration Control	87
7.2.9.1	Overview	87
7.2.9.2	Section Type 4 commands	88
7.2.9.2.1	TIME_DOMAIN_BEAM_CONFIG (Time-domain beamforming command)	88
7.2.9.2.2	TDD_CONFIG_PATTERN (TDD configuration command)	88
7.2.9.3	ACK/NACK support	88
7.2.9.4	O-RU handling	88
7.2.9.5	Latency requirements	89
7.3	Elements for the C-Plane protocol	89
7.3.1	Section Types	89
7.3.2	Section Extensions	90
7.4	Section Type elements	90
7.4.1	Overview	90
7.4.2	Section Type 0 elements	90
7.4.3	Section Type 1 elements	92
7.4.4	reserved (reserved for future use)	94
7.4.5	Section Type 3 elements	94
7.4.6	Section Type 4 elements	95
7.4.7	Section Type 5 elements	99
7.4.8	Section Type 6 elements	100
7.4.9	Section Type 7 elements	103
7.4.10	Section Type 8 elements	107
7.5	Coding of Section Type IEs	108
7.5.1	Transport layer	108
7.5.2	Application layer, common	108
7.5.2.1	dataDirection (data direction (gNB Tx/Rx))	108
7.5.2.2	payloadVersion (payload version)	108
7.5.2.3	filterIndex (filter index)	108
7.5.2.4	frameId (frame identifier)	109
7.5.2.5	subframeId (subframe identifier)	109
7.5.2.6	slotId (slot identifier)	109
7.5.2.7	startSymbolId (start symbol identifier)	110
7.5.2.8	numberOfsections (number of sections)	110
7.5.2.9	sectionType (Section Type)	110
7.5.2.10	udCompHdr (user data compression header)	110
7.5.2.11	numberOfUEs (number Of UEs)	111
7.5.2.12	timeOffset (time offset)	111
7.5.2.13	frameStructure (frame structure)	111
7.5.2.14	cpLength (cyclic prefix length)	112
7.5.2.15	ciCompHdr (channel information compression header)	113
7.5.2.16	scs (sub-carrier spacing)	114
7.5.2.17	cmdScope (command scope)	114
7.5.2.18	numberOfST4Cmds (Number of Section Type 4 commands)	115
7.5.3	Application layer, sections	115
7.5.3.1	sectionId (section identifier)	115
7.5.3.2	rb (resource block indicator)	116
7.5.3.3	symInc (symbol number increment command)	116
7.5.3.4	startPrbc (starting PRB of data section description)	117
7.5.3.5	reMask (resource element mask)	117
7.5.3.6	numPrbc (number of contiguous PRBs per data section description)	118
7.5.3.7	numSymbol (number of symbols)	118
7.5.3.8	ef (extension flag)	119
7.5.3.9	beamId (beam identifier)	119
7.5.3.10	ueId	119
7.5.3.11	freqOffset (frequency offset)	119
7.5.3.12	regularizationFactor (regularization factor)	120
7.5.3.13	ciIsample, ciQsample (channel information I and Q values)	120
7.5.3.14	laaMsgType (LAA message type)	121
7.5.3.15	laaMsgLen (LAA message length)	121
7.5.3.16	lbtHandle	121
7.5.3.17	lbtDeferFactor (listen-before-talk defer factor)	121
7.5.3.18	lbtBackoffCounter (listen-before-talk backoff counter)	122

7.5.3.19	lbtOffset (listen-before-talk offset)	122
7.5.3.20	MCOT (maximum channel occupancy time)	122
7.5.3.21	lbtMode (LBT mode)	122
7.5.3.22	lbtPdschRes (LBT PDSCH result)	122
7.5.3.23	sfStatus (subframe status)	123
7.5.3.24	lbtDrsRes (LBT DRS result)	123
7.5.3.25	initialPartialSF (Initial partial SF)	123
7.5.3.26	lbtBufErr (LBT buffer error)	123
7.5.3.27	sfnSf (SFN/SF end)	123
7.5.3.28	lbtCWConfig_H (HARQ parameters for congestion window management)	124
7.5.3.29	lbtCWConfig_T (TB parameters for congestion window management)	124
7.5.3.30	lbtTrafficClass (traffic class priority for congestion window management)	124
7.5.3.31	lbtCWR_Rst (notification about packet reception successful or not)	124
7.5.3.32	ciCompParam (channel information compression parameter)	124
7.5.3.33	reserved (reserved for future use)	125
7.5.3.34	numberOfAcks (number of ACKs)	125
7.5.3.35	numberOfNacks (number of NACKs)	125
7.5.3.36	ackId (ACK Id)	125
7.5.3.37	nackId (NACK Id)	125
7.5.3.38	st4CmdType (Section Type 4 command type)	126
7.5.3.39	st4CmdLen (Section Type 4 command length)	126
7.5.3.40	numSlots (number of slots)	126
7.5.3.41	ackNackReqId (ACK/NACK Request ID)	126
7.5.3.42	symbolMask (TIME_DOMAIN_BEAM_CONFIG symbol mask)	127
7.5.3.43	disableTDBFNs (disable sending tdBeamNum)	127
7.5.3.44	tdBeamGrp (time-domain beam group identity)	127
7.5.3.45	bfwCompHdr (beamforming weight compression header)	127
7.5.3.46	disableTDBFWs (disable sending time-domain beamforming weights)	128
7.5.3.47	bfwCompParam (beamforming weight compression parameter)	128
7.5.3.48	bfWI, bfWQ (beamforming weights I and Q values)	128
7.5.3.49	dirPattern (symbol data direction (gNB Tx/Rx) pattern)	128
7.5.3.50	guardPattern (guard pattern bitmask)	128
7.6	Section Extension elements	129
7.6.1	Overview	129
7.6.2	Common parameters	132
7.6.2.1	extType (extension type)	132
7.6.2.2	ef (extension flag)	132
7.6.2.3	extLen (extension length)	132
7.7	Coding of Section Extension IEs	133
7.7.1	SE 1: Beamforming weights	133
7.7.1.1	Overview	133
7.7.1.2	bfwCompHdr (beamforming weight compression header)	133
7.7.1.3	bfwCompParam (beamforming weight compression parameter)	134
7.7.1.4	bfwI (beamforming weight in-phase value)	135
7.7.1.5	bfwQ (beamforming weight quadrature value)	135
7.7.2	SE 2: Beamforming attributes	135
7.7.2.1	Overview	135
7.7.2.2	bfaCompHdr (beamforming attributes compression header)	135
7.7.2.3	bfAzPt (beamforming azimuth pointing parameter)	136
7.7.2.4	bfZePt (beamforming zenith pointing parameter)	137
7.7.2.5	bfAz3dd (beamforming azimuth beamwidth parameter)	137
7.7.2.6	bfZe3dd (beamforming zenith beamwidth parameter)	137
7.7.2.7	bfAzSl (beamforming azimuth sidelobe parameter)	137
7.7.2.8	bfZeSl (beamforming zenith sidelobe parameter)	137
7.7.2.9	zero-padding	138
7.7.3	SE 3: DL precoding parameters	138
7.7.3.1	Overview	138
7.7.3.2	codebookIndex (precoder codebook used for transmission)	139
7.7.3.3	layerID (Layer ID for DL transmission)	139
7.7.3.4	txScheme (transmission scheme)	139
7.7.3.5	numLayers (number of layers used for DL transmission)	140
7.7.3.6	crsReMask (CRS resource element mask)	140

7.7.3.7	crsSymNum (CRS symbol number indication).....	140
7.7.3.8	crsShift (crsShift used for DL transmission).....	140
7.7.3.9	beamIdAP1 (beam id to be used for antenna port 1).....	141
7.7.3.10	beamIdAP2 (beam id to be used for antenna port 2).....	141
7.7.3.11	beamIdAP3 (beam id to be used for antenna port 3).....	141
7.7.4	SE 4: Modulation compression parameters	141
7.7.4.1	Overview	141
7.7.4.2	csf (constellation shift flag).....	142
7.7.4.3	modCompScaler (modulation compression scaler value)	142
7.7.5	SE 5: Modulation compression additional parameters.....	143
7.7.5.1	Overview	143
7.7.5.2	mcScaleReMask (modulation compression power scale RE mask).....	144
7.7.5.3	csf (constellation shift flag).....	144
7.7.5.4	mcScaleOffset (scaling value for modulation compression)	144
7.7.6	SE 6: Non-contiguous PRB allocation in time and frequency domain	145
7.7.6.1	Overview	145
7.7.6.2	rbgSize (resource block group size).....	146
7.7.6.3	rbgMask (resource block group bit mask).....	146
7.7.6.4	symbolMask (symbol bit mask).....	147
7.7.6.5	priority.....	148
7.7.6.6	repetition (repetition flag)	148
7.7.7	SE 7: eAxC mask.....	148
7.7.7.1	Overview	148
7.7.7.2	eAxCmask (eAxC Mask)	148
7.7.8	SE 8: Regularization factor.....	149
7.7.8.1	Overview	149
7.7.8.2	regularizationFactor (regularization factor)	149
7.7.9	SE 9: Dynamic Spectrum Sharing parameters	150
7.7.9.1	Introduction.....	150
7.7.9.2	technology (interface name).....	150
7.7.10	SE 10: Group configuration of multiple ports.....	150
7.7.10.1	Overview	150
7.7.10.2	beamGroupType.....	151
7.7.10.3	numPortc	151
7.7.10.4	Interaction with other Section Extensions.....	152
7.7.11	SE 11: Flexible beamforming weights.....	153
7.7.11.1	Overview	153
7.7.11.2	bfwCompHdr (beamforming weight compression header)	154
7.7.11.3	bfwCompParam for PRB bundle x (beamforming weight compression parameter).....	154
7.7.11.4	numBundPrb (Number of bundled PRBs per beamforming weights).....	154
7.7.11.5	bfwI (beamforming weight in-phase value)	154
7.7.11.6	bfwQ (beamforming weight quadrature value).....	155
7.7.11.7	disableBFWs (disable beamforming weights).....	155
7.7.11.8	RAD (Reset after PRB discontinuity)	155
7.7.11.9	contInd (PRB region continuity flag).....	155
7.7.12	SE 12: Non-contiguous PRB allocation with frequency ranges	156
7.7.12.1	Overview	156
7.7.12.2	priority (priority of section description).....	157
7.7.12.3	symbolMask (symbol bit mask).....	158
7.7.12.4	offStartPrb(r) (offset of PRB range start).....	158
7.7.12.5	numPrb(r) (number of PRBs in PRB range)	158
7.7.13	SE 13: PRB allocation with frequency hopping	158
7.7.13.1	Overview	158
7.7.13.2	nextSymbolId(n) (next hop start symbol)	159
7.7.13.3	nextStartPrbc(n) (next hop startPrbc).....	159
7.7.14	SE 14: Nulling layer information for ueld based beamforming.....	160
7.7.14.1	Overview	160
7.7.14.2	nullLayerInd (null layer indication)	160
7.7.15	SE 15: Mixed numerology information for ueld based beamforming	161
7.7.15.1	Overview	161
7.7.15.2	frameStructure (frame structure).....	161
7.7.15.3	freqOffset (frequency offset).....	161

7.7.15.4	cpLength (cyclic prefix length).....	161
7.7.16	SE 16: Antenna mapping in UE channel information based UL beamforming.....	161
7.7.16.1	Overview.....	161
7.7.16.2	antMask.....	162
7.7.17	SE 17: Indication of user port group.....	162
7.7.17.1	Overview.....	162
7.7.17.2	numUeID.....	162
7.7.18	SE 18: Uplink transmission management.....	163
7.7.18.1	Introduction.....	163
7.7.18.2	transmissionWindowOffset.....	163
7.7.18.3	transmissionWindowSize.....	163
7.7.18.4	toT (type of transmission).....	164
7.7.18.5	Interaction with other Section Extensions.....	164
7.7.19	SE 19: Compact beamforming information for multiple port.....	165
7.7.19.1	Overview.....	165
7.7.19.2	disableBFWs (disable beamforming weights).....	166
7.7.19.3	repetition (repeat port info flag).....	166
7.7.19.4	numPortc.....	167
7.7.19.5	priority (priority of section description).....	167
7.7.19.6	symbolMask (resource symbol bitmask).....	167
7.7.19.7	bfwCompHdr.....	167
7.7.19.8	portReMask (RE bitmask per port).....	167
7.7.19.9	portSymbolMask (symbol bitmask per port).....	167
7.7.19.10	bfwCompParam (beamforming weight compression parameter).....	168
7.7.19.11	beamId (beam identifier for a port).....	168
7.7.19.12	bfwI (beamforming weight in-phase value).....	168
7.7.19.13	bfwQ (beamforming weight quadrature-phase value).....	168
7.7.19.14	Interaction with Other Section Extensions.....	169
7.7.20	SE 20: Puncturing extension.....	170
7.7.20.1	Overview.....	170
7.7.20.2	numPuncPatterns (number of puncturing pattern).....	171
7.7.20.3	symbolMask (puncturing pattern symbol mask).....	171
7.7.20.4	startPuncPrb (starting PRB to which one puncturing pattern applies).....	171
7.7.20.5	numPuncPrb (number of contiguous PRBs to which one puncturing pattern applies).....	171
7.7.20.6	puncReMask (puncturing pattern RE mask).....	171
7.7.20.7	rb (resource block indicator).....	171
7.7.20.8	rbgIncl (rbg included flag).....	172
7.7.20.9	rbgSize (rbg size).....	172
7.7.20.10	rbgMask (rbg bitmask).....	172
7.7.20.11	Interaction with Other Extensions.....	172
7.7.21	SE 21: Variable PRB group size for channel information.....	172
7.7.21.1	Overview.....	172
7.7.21.2	ciPrbGroupSize (channel information PRB group size).....	172
7.7.22	SE 22: ACK/NACK request.....	173
7.7.22.1	Overview.....	173
7.7.22.2	ackNackReqId (ACK/NACK request Id).....	173
7.7.23	SE 23: Arbitrary symbol pattern modulation compression parameters.....	173
7.7.23.1	Overview.....	173
7.7.23.2	numSymPrbPattern (number of symbol and resource block patterns).....	174
7.7.23.3	symMask (symbol mask part of symPrbPattern).....	175
7.7.23.4	prbPattern (resource block pattern part of symPrbPattern).....	175
7.7.23.5	numMcScaleOffset (number of modulation compression scaling value per symPrbPattern).....	175
7.7.23.6	mcScaleReMask (modulation compression power scale RE mask).....	175
7.7.23.7	csf (constellation shift flag).....	175
7.7.23.8	Interaction with Other Section Extensions.....	175
7.8	Aspects of C-Plane operation.....	176
7.8.1	Coupling of C-Plane and U-Plane.....	176
7.8.1.1	Overview.....	176
7.8.1.2	Coupling via sectionId value.....	176
7.8.1.3	Coupling via frequency and time.....	177
7.8.1.4	Coupling via frequency and time with priorities.....	178
7.8.1.5	Coupling via frequency and time with priorities (optimized).....	179

7.8.2	O-RUs per endpoint and per C-Plane message limits	179
7.8.2.1	O-RU per endpoint processing limits	179
7.8.2.1.1	Overview	179
7.8.2.1.2	Limits on Number of PRB ranges	179
7.8.2.2	O-RU C-Plane message limits.....	180
7.9	C-Plane optimizations	180
7.9.1	C-Plane optimization using Section Extension 6	180
7.9.2	C-Plane Optimization using Section Extension 7	180
7.9.3	C-Plane optimization using Section Extension 10	180
7.9.3.1	Overview	180
7.9.4	C-Plane optimization using Section Extension 11	181
7.9.4.1	Overview	181
7.9.4.2	Interaction between Section Extension 11 and Section Extension 6	181
7.9.4.2.1	Overview	181
7.9.4.2.2	Interaction when RAD = 1.....	181
7.9.4.2.3	Interaction when RAD = 0.....	182
7.9.4.3	Interaction between Section Extension 11 and Section Extension 12	183
7.9.4.4	Interaction between Section Extension 11 and Section Extension 13	185
7.9.5	C-Plane optimization using Section Extension 12	189
7.9.6	C-Plane optimization using Section Extension 13	189
7.9.7	Coupling via frequency and time with priorities.....	189
7.9.8	Coupling via frequency and time with priorities (optimized)	191
7.9.9	U-Plane Operation Without C-Plane	192
7.9.10	Modulation compression with Section Extension 10.....	192
7.9.11	Optimization with Section Extension 19	196
7.9.12	Optimizations with Section Extension 20.....	197
7.9.13	Optimizations with ueId indicating non-scheduled PRB range	197
8	U-Plane Protocol	200
8.1	General	200
8.1.1	U-Plane transport	200
8.1.2	U-Plane data compression	200
8.1.3	Digital power scaling	200
8.1.3.1	Definition of IQ power in dBFS.....	200
8.1.3.2	Definition of Gain over Fronthaul Interface.....	201
8.1.3.2.1	DL and UL gain definition commonality	201
8.1.3.2.2	DL gain definition	203
8.1.3.2.3	UL gain definition	204
8.1.3.3	TX power budget for Category A and Category B O-RUs	205
8.2	Elementary procedures	206
8.2.1	IQ data transfer procedure	206
8.2.2	IQ data transfer procedure without C-Plane.....	208
8.3	Elements for the U-Plane protocol	209
8.3.1	General.....	209
8.3.2	DL/UL data	209
8.3.3	UL/DL data coding of information elements	212
8.3.3.0	General	212
8.3.3.1	dataDirection (data direction (gNB Tx/Rx))	212
8.3.3.2	payloadVersion (payload version)	212
8.3.3.3	filterIndex (filter index).....	212
8.3.3.4	frameId (frame identifier)	212
8.3.3.5	subframeId (subframe identifier)	212
8.3.3.6	slotId (slot identifier).....	212
8.3.3.7	symbolId (symbol identifier).....	212
8.3.3.8	sectionId (section identifier).....	212
8.3.3.9	rb (resource block indicator)	212
8.3.3.10	symInc (symbol number increment command).....	212
8.3.3.11	startPrbu (starting PRB of user plane section)	213
8.3.3.12	numPrbu (number of PRBs per user plane section)	213
8.3.3.13	udCompHdr (user data compression header)	214
8.3.3.14	reserved (reserved for future use).....	214
8.3.3.15	udCompParam (user data compression parameter)	215

8.3.3.16	iSample (in-phase sample)	215
8.3.3.17	qSample (quadrature sample).....	215
8.3.3.18	sReSMask (selective RE sending mask)	215
8.3.3.19	udCompLen (PRB field length)	216
8.3.4	DL data precoding	216
8.3.5	Data transfer for special cases.....	217
8.3.5.1	Data message mapping and packetization.....	217
8.3.5.2	Uplink data transfer.....	217
8.3.5.3	PRACH data transfer without C-Plane.....	217
8.3.5.4	SRS data transfer without C-Plane.....	217
8.4	U-Plane optimizations	217
8.4.1	Coupling via frequency and time	217
8.4.2	Data blanking.....	218
8.4.2.1	Overview	218
8.4.2.2	PRB blanking	218
8.4.2.3	Transmission blanking	218
8.5	Aspects of U-Plane operation.....	218
8.5.1	O-RU U-Plane limits	218
8.5.1.1	O-RU U-Plane message limits	218
9	Counters and KPIs.....	220
9.1	Counters	220
10	Specification mandatory and optional capabilities.....	223
10.1	General	223
10.2	Mandatory and optional capabilities.....	224
11	S-Plane Protocol.....	232
11.1	General	232
11.1.1	Overview	232
11.2	Synchronization baseline.....	232
11.2.1	List of reference documents.....	232
11.2.2	Clock model and synchronization topology.....	232
11.2.2.1	Introduction.....	232
11.2.2.2	Topology configuration LLS-C1 and LLS-C2 synchronization.....	233
11.2.2.3	Topology configuration LLS-C3 synchronization	236
11.2.2.4	Topology configuration LLS-C4 synchronization	238
11.2.3	Clock Synchronization.....	238
11.2.4	Profiles.....	239
11.2.4.1	Physical layer frequency signals (PLFS).....	239
11.2.4.2	PTP.....	240
11.2.4.2.1	Full timing support	240
11.2.4.2.2	Partial timing support	240
11.2.5	Synchronization accuracy	241
11.2.5.1	Introduction.....	241
11.2.5.2	Jitter.....	242
11.2.5.3	Wander.....	242
11.2.5.4	Air interface frequency error.....	242
11.2.5.5	Air interface maximum time error	242
11.3	Time and frequency synchronization requirements.....	245
11.3.1	Allowed PTP and PLFS clock types and clock classes.....	245
11.3.1.1	Overview	245
11.3.1.2	O-DU synchronization capability options:.....	246
11.3.1.3	O-RU synchronization capability options:	246
11.3.2	Frequency and time synchronization requirements across fronthaul network elements	247
11.3.2.1	Configurations LLS-C1 and LLS-C2.....	247
11.3.2.2	Configuration LLS-C3	250
11.3.2.3	Configuration LLS-C4	253
11.4	Node synchronization state behavior.....	255
11.4.1	Configurations LLS-C1 and LLS-C2.....	255
11.4.1.1	Overview	255
11.4.1.2	M-Plane disconnected	255
11.4.1.3	O-RU in the FREERUN state.....	256

11.4.1.4	O-DU in the FREERUN state	256
11.4.1.5	Operation in LOCKED/HOLDOVER state	257
11.4.2	Configurations LLS-C3	258
11.4.2.1	Overview	258
11.4.2.2	M-Plane disconnected	258
11.4.2.3	O-RU in the FREERUN State	258
11.4.2.4	O-DU in the FREERUN state	258
11.4.2.5	Operation in LOCKED/HOLDOVER state	259
11.4.3	Configurations LLS-C4	259
11.4.3.1	Overview	259
11.4.3.2	M-Plane disconnected	259
11.4.3.3	O-RU in the FREERUN state	259
11.4.3.4	O-DU in the FREERUN state	260
11.4.3.5	Operation in LOCKED/HOLDOVER state	260
11.5	S-Plane handling in multiple link scenarios	260
11.6	Announce messages	260
11.7	Elementary procedures	260
11.7.1	PTP time synchronization procedure	260
11.7.2	System frame number calculation from GPS time	261
12	Beamforming functionality	263
12.1	General	263
12.2	Hierarchy of radiation structure in O-RU	263
12.3	Calibration	264
12.4	beamId use for various beamforming methods	264
12.4.1	Overview	264
12.4.2	Predefined-beam beamforming	265
12.4.2.1	Overview	265
12.4.2.2	Beam characteristics	265
12.4.3	Weight-based dynamic beamforming	267
12.4.3.1	Overview	267
12.4.3.2	Optional restriction on sending of beamforming weights	267
12.4.3.3	Weight-based dynamic frequency-domain or time-domain beamforming (not hybrid)	270
12.4.3.4	Weight-based dynamic hybrid beamforming	270
12.4.3.4.1	Overview	270
12.4.3.4.2	Hybrid beamforming with updatable frequency-domain and time-domain weights	270
12.4.3.4.3	Hybrid beamforming with updatable frequency-domain weights and fixed time-domain beams	271
12.4.3.4.4	Hybrid beamforming with Section Type 4 'TIME_DOMAIN_BEAM_CONFIG'	273
12.4.4	Attribute-based dynamic beamforming	274
12.4.5	Channel-information-based beamforming	275
12.5	O-RU antenna model	275
12.5.1	Overview	275
12.5.2	Coordinate systems	275
12.5.3	O-RU antenna model parameters	276
12.5.4	Identification and ordering of array elements	277
12.5.5	Relations between array elements	278
12.5.6	Model usage	278
13	Support of Shared Cell	281
13.1	General	281
13.2	Copy and Combine function	282
13.2.1	DL Copy function	282
13.2.2	UL Combine function	282
13.3	Selective transmission and reception using beamId	284
13.3.1	Introduction	284
13.3.2	Selective transmission function	285
13.3.3	Selective reception function	286
13.4	Multiple Cells Scenario in Cascade mode	287
13.4.1	Introduction	287
13.4.2	Multiple Cell scenario in downlink	287
13.4.3	Multiple Cell scenario in uplink	288

13.5	Delay management for Shared cell.....	290
13.5.1	Introduction.....	290
13.5.2	DL delay management for Shared cell.....	290
13.5.3	UL delay management for Shared cell.....	292
13.5.4	Section Extension 10 support for Shared Cell	294
13.6	S-plane for Shared cell	295
13.6.1	Introduction.....	295
13.6.2	Node behavior at O-RU in FREERUN state.....	295
13.6.2.1	O-DU.....	295
13.6.2.2	O-RU.....	295
13.7	Cascade-FHM mode.....	295
13.7.1	General.....	295
13.7.2	Copy and Combine function	296
13.7.3	Delay management	297
14	Support of Shared O-RU	300
14.1	General	300
14.2	Delay management for Shared O-RU.....	300
14.2.1	General.....	300
14.3	S-Plane for Shared O-RU	300
14.3.1	General.....	300
15	Support of Narrowband-IoT	301
15.1	Overview (informative).....	301
15.2	NB-IoT Downlink enabling methods	301
15.2.1	Extended PRB Grid	301
15.3	NB-IoT Downlink use cases.....	303
15.3.1	Overview	303
15.3.2	In-band mode with 15 kHz subcarrier spacing of host carrier	303
15.3.3	Guard-band mode with 15 kHz subcarrier spacing of host carrier.....	304
15.4	NB-IoT Uplink enabling methods	305
15.4.1	Single Section Single Carrier (SSSC).....	305
15.4.2	Single Section Multiple Carriers (SSMC)	306
15.4.3	Extended PRB Grid	307
15.5	NB-IoT Uplink use cases	308
15.5.1	Overview	308
15.5.2	General case for in-band and guard-band mode	308
15.5.3	Special case of in-band mode and guard-band mode NPUSCH with extended PRB Grid	309
15.5.4	NPRACH without C-plane	309
15.6	subframeId, slotId and symbolId numbering for NB-IoT.....	310
15.6.1	Overview	310
15.6.2	NPUSCH timing	310
15.6.3	NPRACH timing.....	310
Annex A (normative): Compression methods.....		312
A.1	Block floating point compression.....	312
A.1.1	Overview	312
A.1.2	Block floating point compression algorithm.....	312
A.1.3	Block floating point decompression algorithm.....	313
A.2	Block scaling compression	313
A.2.1	Overview	313
A.2.2	Block scaling compression algorithm	314
A.2.3	Block scaling decompression algorithm	314
A.3	μ -Law compression	314
A.3.1	μ -Law compression algorithm	314
A.3.2	μ -Law decompression algorithm	316
A.3.3	μ -law udCompParam and IQ data format	316
A.4	Beamspace compression type I and type II	317
A.4.1	Overview	317
A.4.2	Beamspace compression type I.....	317
A.4.2.1	Overview.....	317
A.4.2.2	Beamspace compression algorithm.....	317

A.4.2.3	Beamspace decompression algorithm	318
A.4.3	Beamspace compression type II.....	319
A.4.3.1	Overview	319
A.4.3.2	Beamspace compression algorithm.....	319
A.4.3.3	Beamspace decompression algorithm	321
A.4.3.4	Kronecker product algorithm	322
A.5	Modulation compression	322
A.6	Selective RE sending.....	325
Annex B (informative): Delay management use cases		326
B.1	Introduction	326
B.2	Latency categories and sub-categories	331
B.3	Example case: Evaluating O-DU / O-RU combinations	336
B.4	Example case: Non-ideal transport O-DU/O-RU categories.....	339
Annex C (informative): Void		340
Annex D (informative): IQ sample and exponent packetization for different bitwidths		341
D.1	IQ sample ordering for big-endian byte order	341
D.2	IQ sample ordering for little-endian byte order.....	346
Annex E (informative): OFDM phase compensation		352
Annex F (normative): Beamforming attributes frame of reference.....		353
Annex G (informative): LAA Algorithms and Examples.....		355
Annex H (informative): S-Plane detailed frequency and phase error budget analysis, and future ITU-T clock types and classes reference		359
H.1	Void.....	359
H.2	Frequency and time error budget analysis.....	359
H.2.1	Overview	359
H.2.2	Network limit guidelines for frequency error budget (LLS-C1 and LLS-C2):	359
H.2.3	Network limit guidelines for time error budget (LLS-C1 and LLS-C2):.....	360
H.2.4	Network limit guidelines for LLS-C1 and LLS-C2 based on the budget.....	362
H.2.5	Network limit guidelines for frequency error budget (LLS-C3):.....	366
H.2.6	Time error budget for network limit (LLS-C3).....	368
H.3	Guidelines for the allowed number of switches:	368
Annex I (informative): Precoding and examples.....		370
I.1	Case 1: Tx Diversity 1-CRS Port Ant0, 1 PRB:.....	370
I.2	Case 2: Tx Diversity 2-CRS Port Ant0, Ant1, and 1 PRB	370
I.3	Case 3: Tx Diversity 4-CRS Port Ant0,1,2,3 and 1 PRB	371
I.4	Case 4: TM3/TM4 3 layers, 4 antenna ports	373
I.5	CRS location assignment.....	375
I.6	First data layer and non-first data layer association using static eAxC_ID and layerId mapping	377
Annex J (informative): Beamforming methods description		378
J.1	Introduction	378
J.2	Predefined-beam beamforming	378
J.3	Weight-based dynamic beamforming.....	378
J.4	Attribute-based dynamic beamforming	379
J.5	Channel-information-based beamforming.....	379
Annex K (informative): Layers of array elements.....		381
K.1	General	381
K.2	Use Case A	381
K.3	Use Case B	382
Annex L (informative): Considerations when operating in non-ideal transport environments ...		383
L.1	General	383
L.2	Deployment scenarios and performance considerations.....	383
L.2.1	Low impact	383

L.2.2	Medium impact	383
L.2.3	High impact	383
L.3	HARQ	384
L.3.1	Introduction.....	384
L.3.2	Synchronous HARQ	384
L.3.3	HARQ interleaving	384
L.3.4	Predictive HARQ	385
L.3.5	Asynchronous HARQ considerations	386
L.4	RACH considerations.....	386
L.4.1	Introduction.....	386
L.4.2	Non-Ideal fronthaul RACH designs.....	387
L.4.2.1	Assumptions.....	387
L.4.2.2	Option 1: Semi-persistent resource allocation.....	387
L.4.2.3	Option 2: RACH retransmission estimation.....	388
L.5	Other latency related considerations.....	389
L.6	Bandwidth limitation considerations	390
Annex M (informative): Use case of selective transmission and reception		392
M.1	General	392
M.2	Selective transmission and reception with non beamforming O-RUs and 1-to-1 mapping of global beamId and local beamId	392
M.3	Selective transmission and reception with beamforming O-RUs and 1-to-N mapping of global beamId and local beamId	393
Change history		395
History		397

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: "*Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards*", which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (<https://ipr.etsi.org/>).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™** and **LTE™** are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM®** and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Mobile Standards Group (MSG).

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The contents of the present document are subject to continuing work within O-RAN and may change following formal O-RAN approval. Should the O-RAN Alliance modify the contents of the present document, it will be re-released by O-RAN with an identifying change of version date and an increase in version number as follows:

version xx.yy.zz

where:

xx: the first digit-group is incremented for all changes of substance, i.e. technical enhancements, corrections, updates, etc. (the initial approved document will have xx=01). Always 2 digits with leading zero if needed.

yy: the second digit-group is incremented when editorial only changes have been incorporated in the document. Always 2 digits with leading zero if needed.

zz: the third digit-group included only in working versions of the document indicating incremental changes during the editing process. External versions never include the third digit-group. Always 2 digits with leading zero if needed.

The present document specifies the control plane, user plane and synchronization plane protocols used over the fronthaul interface linking the O-DU (O-RAN Distributed Unit) with the O-RU (O-RAN Radio Unit) with a Lower Layer Functional Split-7-2x based architecture (explained below). The scope of the present document includes both LTE and NR (5G). A separate document contains the O-RAN M-Plane (management plane) specification.

In the following, "Layer 1" and "Physical Layer" are assumed to be synonymous.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies. In the case of a reference to a 3GPP document (including a GSM document), a non-specific reference implicitly refers to the latest version of that document in Release 15.

Referenced documents which are not found to be publicly available in the expected location might be found at <https://o-ran.org/specifications>.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, O-RAN cannot guarantee their long-term validity.

The following referenced documents are necessary for the application of the present document.

- [1] Void.
- [2] eCPRI Specification V1.0: "Common Public Radio Interface: eCPRI Interface Specification".
- [3] IEEE Std 1588™-2008 "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems".
- [4] 3GPP TS 38.211 Release 15: "NR; Physical channels and modulation".
- [5] Void.
- [6] Void.
- [7] O-RAN ALLIANCE O-RAN.WG4.MP.0: "Management Plane Specification".
- [8] 3GPP TS 38.104 Release 15: "Base Station (BS) radio transmission and reception".
- [9] 3GPP TS 36.104 Release 16: "Base Station (BS) radio transmission and reception".
- [10] IETF RFC 1166: "Internet Numbers".
- [11] IEEE Std 802.1CM™-2018: "IEEE Standard for Local and metropolitan area networks – Time-Sensitive Networking for Fronthaul", 8 June 2018.

- [12] Amendment to IEEE Std 802.1CM™-2018: "IEEE Standard for Local and metropolitan area networks – Time-Sensitive Networking for Fronthaul - Amendment 1: Enhancements to Fronthaul Profiles to Support New Fronthaul Interface, Synchronization, and Syntonization Standards," in IEEE Std 802.1CMde-2020, Oct. 2020.
- [13] Void
- [14] 3GPP TS 36.213: "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures".
- [15] ITU-T Recommendation G.781 (04/2020): "Synchronization layer functions for frequency synchronization based on the physical layer".
- [16] ITU-T Recommendation G.810 (08/1996): "Definitions and terminology for synchronization networks".
- [17] ITU-T Recommendation G.8260 (03/2020): "Definitions and terminology for synchronization in packet networks".
- [18] ITU-T Recommendation G.8261/Y.1361 (2019): Amendment 2 (10/2020) Timing and synchronization aspects in packet networks.
- [19] ITU-T Recommendation G.8262/Y.1362 (2018): Amendment 1 (03/2020) Timing characteristics of a synchronous Ethernet equipment slave clock.
- [20] ITU-T Recommendation G.8262.1/Y.1362 (2019): Amendment 1 (08/2019) Timing characteristics of an enhanced synchronous equipment slave clock.
- [21] ITU-T Recommendation G.8264/Y.1364 (08/2017): Amendment 1 (03/2018) Distribution of timing information through packet networks.
- [22] ITU-T Recommendation G.8271/Y.1366 (03/2020): "Time and phase synchronization aspects of telecommunication networks".
- [23] ITU-T Recommendation G.8271.1/Y.1366.1 (10/2020): "Network limits for time synchronization in packet networks with full timing support from the network".
- [24] ITU-T Recommendation G.8271.2/Y.1366.2 (05/2021): "Network limits for time synchronization in packet networks with partial timing support from the network".
- [25] ITU-T Recommendation G.8272/Y.1367 (11/2018): Amendment 1 (03/2020) Timing characteristics of primary reference time clocks.
- [26] ITU-T Recommendation G.8272.1/Y.1367. (11/2016): Amendment 2 (08/2019) Timing characteristics of enhanced primary reference time clocks.
- [27] ITU-T Recommendation G.8273/Y.1368 (2018): Corrigendum 1 (10/2020) Framework of phase and time clocks.
- [28] ITU-T Recommendation G.8273.2/Y.1368.2 (10/2020): "Timing characteristics of telecom boundary clocks and telecom time slave clocks".
- [29] ITU-T Recommendation G.8273.3/Y.1368.3 (10/2020): Timing characteristics of telecom transparent clocks for use with full timing support from the network.
- [30] ITU-T Recommendation G.8275/Y.1369 (2020): Amendment 1 (05/2021) Architecture and requirements for packet-based time and phase distribution.
- [31] ITU-T Recommendation G8275.1/Y.1369.1 (2020): Amendment 2 (06/2021) Precision time protocol telecom profile for phase/time synchronization with full timing support from the network.
- [32] ITU-T Recommendation G8275.2/Y.1369.2 (2020): Amendment 2 (06/2021) Precision time protocol telecom profile for time/phase synchronization with partial timing support from the network.
- [33] IEEE Std 1588™-2019 "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", November 2019.
- [34] ITU-T Recommendation G.811(09/1997) Amendment 1 (04/2016): "Timing characteristics of primary reference clocks".

- [35] ITU-T Recommendation G.8273.4/Y.1368.4 (03/2020) Amendment 1 (05/2021): Timing characteristics of telecom boundary clocks and telecom time slave clocks for use with partial timing support from the network.
- [36] IEEE Std 802.1Q™-2014 "IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks", 3 November 2014.
- [37] 3GPP TS 38.133 Release 15: "NR; Requirements for support of radio resource management".
- [38] 3GPP TS 36.133 Release 15: "Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management".
- [39] eCPRI Transport Network V1.2 (2018-06-25) "Common Public Radio Interface: Requirements for the eCPRI Transport Network".
- [40] 3GPP TS 36.141 Release 15: "Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) conformance testing".
- [41] 3GPP TS 38.141-1 Release 15: "NR; Base Station (BS) conformance testing Part 1: Conducted conformance testing".
- [42] 3GPP TS 36.211 Release 15: "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation".
- [43] 3GPP TR 38.901: "Study on channel model for frequencies from 0.5 to 100 GHz".
- [44] IEEE Std. 802.2™: "IEEE Standard for Local Area Networks - Logical Link Control".
- [45] IEEE Std. 1914.3™: "IEEE Standard for Radio over Ethernet Encapsulations and Mappings".
- [46] RFC 2460: "Internet Protocol, Version 6 (IPv6) Specification".
- [47] RFC 791: "Internet Protocol".
- [48] RFC 826: "Ethernet Address Resolution Protocol".
- [49] RFC 768: "User Datagram Protocol".
- [50] IEEE Std 145™-1993 (R2004): "IEEE Standard Definitions of Terms for Antennas".
- [51] IEEE Std 802.1X™-2020: "IEEE Standard for Local and Metropolitan Area Networks--Port-Based Network Access Control", February 2020.
- [52] 3GPP TS 38.323 Release 16, V16.6.0 (2021-12): "NR; Packet Data Convergence Protocol (PDCP)".
- [53] IEEE Std 802.3-2018™: "IEEE Standard for Ethernet".
- [54] O-RAN.WG9.XTRP-SYN.0-v03.00: "O-RAN Synchronization Architecture and Solution Specification, Release 03.00".

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, O-RAN cannot guarantee their long-term validity.

The following referenced documents are not necessary for the application of the present document, but they assist the user with regard to a particular subject area.

- [i.1] 3GPP TR 21.905: "Vocabulary for 3GPP Specifications".
- [i.2] 3GPP TR 36.932: "Scenarios and requirements for small cell enhancements for E-UTRA and E-UTRAN".
- [i.3] 3GPP TS 36.331: "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification".

[i.4] 3GPP TR 38.801 (V14.0.0): "Study on new radio access technology: Radio access architecture and interfaces".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the terms given in 3GPP TR 21.905 [i.1] and the following apply. A term defined in the present document takes precedence over the definition of the same term, if any, in 3GPP TR 21.905 [i.1].

beam: beam (of the antenna) is a radiation pattern without restriction to main lobe.

C-Plane: Control Plane: refers specifically to real-time control between O-DU and O-RU, and should not be confused with the UE's control plane.

Cascade mode: Mode of Shared cell which is realized by several O-RUs cascaded in chain.

NOTE: See clause 13.

Category A O-RU: O-RU within which the DL precoding is not done.

NOTE: See Figure 4.2-1.

Category B O-RU: O-RU within which the DL precoding is performed.

NOTE: See Figure 4.2-1

eAxC: extended Antenna-Carrier: a data flow for a single antenna (or spatial stream) for a single carrier in a single sector.

FFO: Fractional Frequency Offset. This is defined as $\Delta f/f_{\text{norm}}$ which is used to describe frequency error, typically on the output of the T-TSC filter in the O-RU.

NOTE: It is the same as the Fractional Frequency Deviation defined in Recommendation ITU-T G.810 [16], clause 4.5.2.

FHM mode: Mode of Shared cell which is realized by FHM and several O-RUs.

NOTE: See clause 13.

GPS Time: GPS Time is a continuous time scale (no leap seconds) defined by the GPS Control segment. It starts at 0^h UTC (midnight) of January 5th to 6th 1980.

NOTE: See IEEE 1588 [3], [33] Annex B for timescales.

Fronthaul: A logical link connecting O-DU and O-RU.

NOTE: Fronthaul transports C-Plane, U-Plane, S-Plane and M-Plane traffic.

Hop: physical link between two S-Plane nodes (where node can be O-DU, switch or O-RU).

NOTE: As defined in IEEE 802.1CM [11], [12].

LAA: Licensed-assisted access: Carrier aggregation with at least one secondary cell operating in the unlicensed spectrum.

LLS: Lower Layer Split: logical interface between O-DU and O-RU when using a lower layer (intra-PHY based) functional split.

LLS-U: Lower Layer Split User Plane: logical interface between O-DU and O-RU when using a lower layer functional split.

LLS-C: Lower Layer Split Control Plane: logical interface between O-DU and O-RU when using a lower layer functional split.

High-PHY: those portions of the PHY layer processing on the O-DU side of the fronthaul interface.

NOTE: This includes forward error correction encode/decode, scrambling, and modulation/demodulation.

Low-PHY: those portions of the PHY layer processing on the O-RU side of the fronthaul interface.

NOTE: This includes FFT/iFFT, digital beamforming, and PRACH extraction and filtering.

M-Plane: Management Plane: refers to non-real-time management operations between the O-DU and the O-RU

NB-IoT: Narrowband Internet of Things which is a Low Power Wide Area Network (LPWAN) radio technology standard specified in 3GPP standards to provide extended coverage such as indoor and underground scenarios, support of massive number of low throughput devices, low delay sensitivity, ultra-low device cost, low device power consumption and optimized network architecture.

O-CU: O-RAN Control Unit – a logical node hosting PDCP, RRC, SDAP and other control functions

O-DU: O-RAN Distributed Unit: a logical node hosting RLC/MAC/High-PHY layers based on a lower layer functional split.

O-RU: O-RAN Radio Unit: a logical node hosting Low-PHY layer and RF processing based on a lower layer functional split.

NOTE: This is similar to 3GPP defined TRP or RRH but more specific in including the Low-PHY layer (FFT/iFFT, PRACH extraction).

PRACH Symbol: A resource in the time domain having the duration of (1/Subcarrier Spacing), following cyclic prefix (excluded) in a PRACH occasion. Cyclic prefix and one or more of such consecutive PRACH Symbols constitute a PRACH preamble or a NPRACH symbol group defined in 3GPP TS 36.211 [42] and 38.211 [4].

PRB: Physical Resource Block. A group of 12 consecutive subcarriers of an OFDM symbol.

Processing element: See clause 7.5 of the M-Plane specification [7] for a definition of processing element.

PRTC: Primary Reference Time Clock as defined in G.8272 [25] and G.8272.1 [26].

NOTE: There are different types of PRTC defined in both G.8272 and G.8272.1. Unless the specific type is indicated, a reference to a PRTC in the present document could include any of these types.

Q<I.F>: denotes a signed two's-complement I+F bit fixed point number with I signed integer bits, and F fractional bits.

Quality Level: Referred to in Recommendation ITU G.781 [15] as "clock source quality level." The clock-source quality-level of a synchronous equipment clock or stand-alone synchronization clock is defined as the grade of clock to which it is ultimately traceable.

RE: Resource Element. One subcarrier of an OFDM symbol.

rx-array: As defined in clause 12.2.

S-Plane: Synchronization Plane: refers to traffic between the O-RU or O-DU and a synchronization controller which is generally an IEEE 1588 [3], [33] Grand Master.

NOTE: Grand Master functionality may be embedded in the O-DU.

Shared Cell: The operation for the same cell by several O-RUs.

NOTE: See clause 11.

Shared O-RU: See clause 14.1 for a definition of Shared O-RU.

Shared O-RU Host: See clause 3.1 of the M-Plane specification [7] for a definition of Shared O-RU Host.

Shared Resource Operator: See clause 3.1 of the M-Plane specification [7] for a definition of Shared Resource Operator.

Slot: a time duration defined in Table 7.5.2.13-3.

NOTE: This is applicable to both LTE and NR. Slot definition in the present document for LTE is different from the one defined in 3GPP TS 36.211 [42], clause 4.

Spatial stream: the data flow on the DL associated with precoded data (may be same as layers or different if there is expansion in the precoding), and on UL associated with the number of outputs from the digital beamforming (sometimes called "beams").

TE(t): Time Error – The difference between the time indicated by a clock or timing signal, and that indicated by a reference clock or timing signal.

NOTE: See Recommendation ITU-T G.810 [16], clause 4.5.13.

TRX: Refers to the specific processing chain in an O-RU associated with D/A or A/D converters.

NOTE: Due to digital beamforming the number of TRXs may exceed the number of spatial streams, and due to analog beamforming, the number of TRXs may be lower than the number of antenna elements.

tx-array: As defined in clause 12.2.

U-Plane: User Plane: refers to IQ sample data transferred between O-DU and O-RU

UNI: User Network Interface as defined by eCPRI network requirement specification

UQ<LF>: denotes an unsigned I+F bit fixed point number with I unsigned integer bits and F fractional bits

3.2 Symbols

For the purposes of the present document, the symbols given in 3GPP TS 38.21 [4] and the following apply:

Δf	Subcarrier spacing
μ	Subcarrier spacing configuration

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in 3GPP TR 21.905 [i.1] and the following apply. An abbreviation defined in the present document takes precedence over the definition of the same abbreviation, if any, in 3GPP TR 21.905 [i.1].

APTS	Assisted Partial Timing Support
ACK/NACK	Acknowledgement and Negative Acknowledgement
BF	Beamforming
BMCA	Best Master Clock Algorithm
C/U-Plane	C-Plane and U-Plane
CA	Carrier Aggregation
CID	Company IDentification
CoMP	Coordinated Multi-Point
CRS	Cell Specific Reference Signal
DAC	Digital Analogue Converter
DIX	Ethernet, named after DEC, Intel and Xerox
DL	Downlink
DRS	Discovery Reference Signal
DSS	Dynamic Spectrum Sharing
EEC	Ethernet Equipment Clock
ESMC	Ethernet Synchronization Messaging Channel
eAxC_ID	eAxC Identifier
eEEC	enhanced Ethernet Equipment Clock
ePRTC	enhanced Primary Reference Time Clock
eNB	evolved NodeB (applies to LTE)
FFO	Fractional Frequency Offset
FHM	Fronthaul Multiplexer
gNB	next generation NodeB (applies to NR)
HARQ	Hybrid Automatic Repeat request
IE	Information element
IQ	In-phase and Quadrature-phase
ITU	International Telecommunication Union
iFFT	inverse Fast Fourier Transformation
LBT	Listen Before Talk
LLC	Logical Link Control
MAC	Media Access Control
MCOT	Maximum Channel Occupancy Time
MCS	Modulation Coding Scheme

MMSE	Minimum Mean Square Error
MTIE	Maximum Time Interval Error
MTU	Maximum Transmission Unit
OTA	Over The Air
OTDOA	Observed Time Delay Of Arrival
OUI	Organizationally Unique Identifier
PCP	Priority Code Point
PDV	Packet Delay Variation
PHY	Physical Layer
PLFS	Physical Layer Frequency Support
PLL	Phase Locked Loop
PPS	Pulse Per Second
PRC	Primary Reference Clock
PRG	Precoding Resource Block Group
PRTC	Primary Reference Time Clock
PTP	Precision Time Protocol
QL	Quality Level
RB	Resource Block
RE	Resource Element
RLC	Radio Link Control
RLC-AM	RLC Acknowledged Mode
RLC-TM	RLC Transparent Mode
RLC-UM	RLC Unacknowledged Mode
RoE	Radio over Ethernet
RX	Receiver
SE	Section Extension
SF	Subframe
SLA	Service Level Agreement
SMO	Service Management and Orchestration
SRO	Shared Resource Operator
SSM	Synchronization Status Message
ST	Section Type
SYNCE	SYNChronous Ethernet
T-BC	Telecom Boundary Clock
T-GM	Telecom Grand Master
T-TC	Telecom Transparent Clock
T-TSC	Telecom Subordinate Clock
TAE	Time Alignment Error
TDEV	Time DEVIation
TLV	Type Length Value
TM	Transmission Mode
TOD	Time of Day
TX	Transmitter
UL	Uplink
URLLC	Ultra Reliable Low Latency Communications
UTC	Universal Time Coordinated
VLAN	Virtual Local Area Network
vRAN	virtualized Radio Access Network

3.4 Document conventions

3.4.1 Terminologies

In the present document, all material is considered normative except where explicitly or implicitly labeled as informative. Labeling as informative includes content outside of tables and figures labeled as "NOTE", content described as "examples", and content carrying the explicit label of "informative" such as some Annexes.

For the purposes of the present document, the following terminologies apply.

ignored Within the present document, it is sometimes stated that certain information elements are to be "ignored" by the receiver (generally for "reserved" fields but in some cases other fields too). In this case, the fields shall be ignored for the purposes of processing at the O-RAN application level, but in the case of certain packet error-checking such as Ethernet-layer CRC or parity calculations, the fields shall be included in the CRC or parity calculations. As well, in the event of packet encryption, the "ignored" fields shall be encrypted along with the other packet payload information.

3.4.2 Fields and bitmasks in messages

For the purposes of the present document, the following conventions apply for the format of messages and data structures within messages.

In accordance with IETF RFC 1166 [10], the left most bit of an octet is the most significant bit (msb) and the right most bit is the least significant bit (lsb). The msb is labelled as 0 and lsb is labelled as 7. This is illustrated by a blue ribbon in tables showing the message format. An example is depicted in Table 3.4.2-1. Note that this bit labelling convention (specifically the blue ribbon header in some tables) for octets is different from the labelling of bits within a field (bracketed bit numbers shown in tables).

To address specific bits within a field, the following notation is used: "X[k]" represents kth bit in a field X with the convention that bit X[0] is the least significant bit of field X and located in the rightmost bit position. Where applicable, a sequence of bits in a field X can be interpreted as an unsigned integer value X_{val} calculated with formula:

$$X_{val} = \sum_{k=0}^{N-1} X[k]2^k$$

where N is number of bits in field X.

Notation "X[a:b]" represents a sequence of bits in field X starting from bit X[a] and ending at bit X[b] inclusive, where $a > b$.

Table 3.4.2-1 below illustrates the format of messages and data structures using this notation.

Table 3.4.2-1: Example of table presenting format of a data structure

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	octet
Y[3:0]				X[11:8]				1	N
X[7:0]								1	N+1

The example data structure presented in Table 3.4.2-1 is interpreted as follows:

- Field Y has 4 bits. Bits Y[3] to Y[0] of field Y are in octet N; Y[3], the most significant bit of field Y is in the most significant bit of octet N.
- Field X has 12 bits. Bits X[11] to X[8] of field X are in octet N; X[8] is in the least significant bit of octet N. Bits X[7] to X[0] of field X are in octet N+1; X[7] is in the most significant bit of octet N+1 and X[0] is in the least significant bit of octet N+1.

This corresponds to a structure that maps every bit as presented in Table 3.4.2-2 below.

Table 3.4.2-2: Interpretation of the example of table presenting format of a data structure

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	octet
Y[3]	Y[2]	Y[1]	Y[0]	X[11]	X[10]	X[9]	X[8]	1	N
X[7]	X[6]	X[5]	X[4]	X[3]	X[2]	X[1]	X[0]	1	N+1

4 Architecture & requirements

4.1 Architectural aspects

4.1.1 Architectural entities

The architecture of eNB or gNB with O-DU and O-RUs is shown in Figure 4.1.1-1. LLS-C and LLS-U provide C-Plane and U-Plane over LLS interface, respectively.

In this architecture, O-DU and O-RU are defined as follows.

O-RAN Distributed Unit (O-DU): a logical node that includes the eNB/gNB functions as listed in clause 4.2 split option 7-2x, excepting those functions allocated exclusively to the O-RU. The O-DU controls the operation of O-RUs.

O-RAN Radio Unit (O-RU): a logical node that includes a subset of the eNB/gNB functions as listed in clause 4.2 split option 7-2x. The real-time aspects of control & user plane communication with the O-RU are controlled by the O-DU.

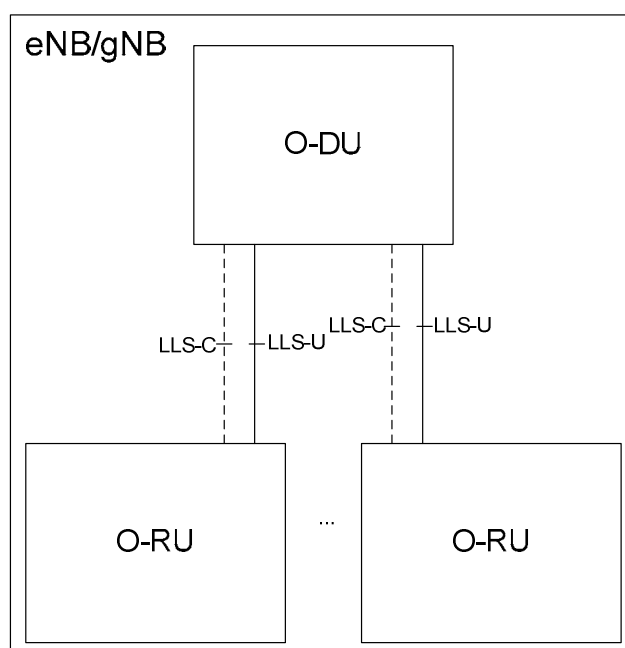


Figure 4.1.1-1: eNB/gNB architecture with O-DU and O-RUs

4.2 Functional split

4.2.1 Introduction

When considering the functional split defining a fronthaul interface there are two competing interests:

- There is a benefit in keeping an O-RU as simple as possible because size, weight, and power draw are primary deciding considerations and the more complex an O-RU, the larger, heavier and more power-hungry the O-RU tends to be.
- There is a benefit in having the interface at a higher level which tends to reduce the interface throughput relative to a lower-level interface – but the higher-level the interface, the more complex the O-RU tends to be.

To resolve this, O-RAN has selected a single split point, known as "7-2x" but allows a variation, with the precoding function to be located either "above" the interface in the O-DU or "below" the interface in the O-RU. For the most part the interface is not affected by this decision, but there are some impacts namely to provide the necessary information to the O-RU to execute the precoding operation. O-RUs within which the precoding is not done (therefore of lower complexity) are called "Category A" O-RUs while O-RUs within which the precoding is done are called "Category B" O-RUs. See Figure 4.2-1 for a depiction of this dual O-RU concept.

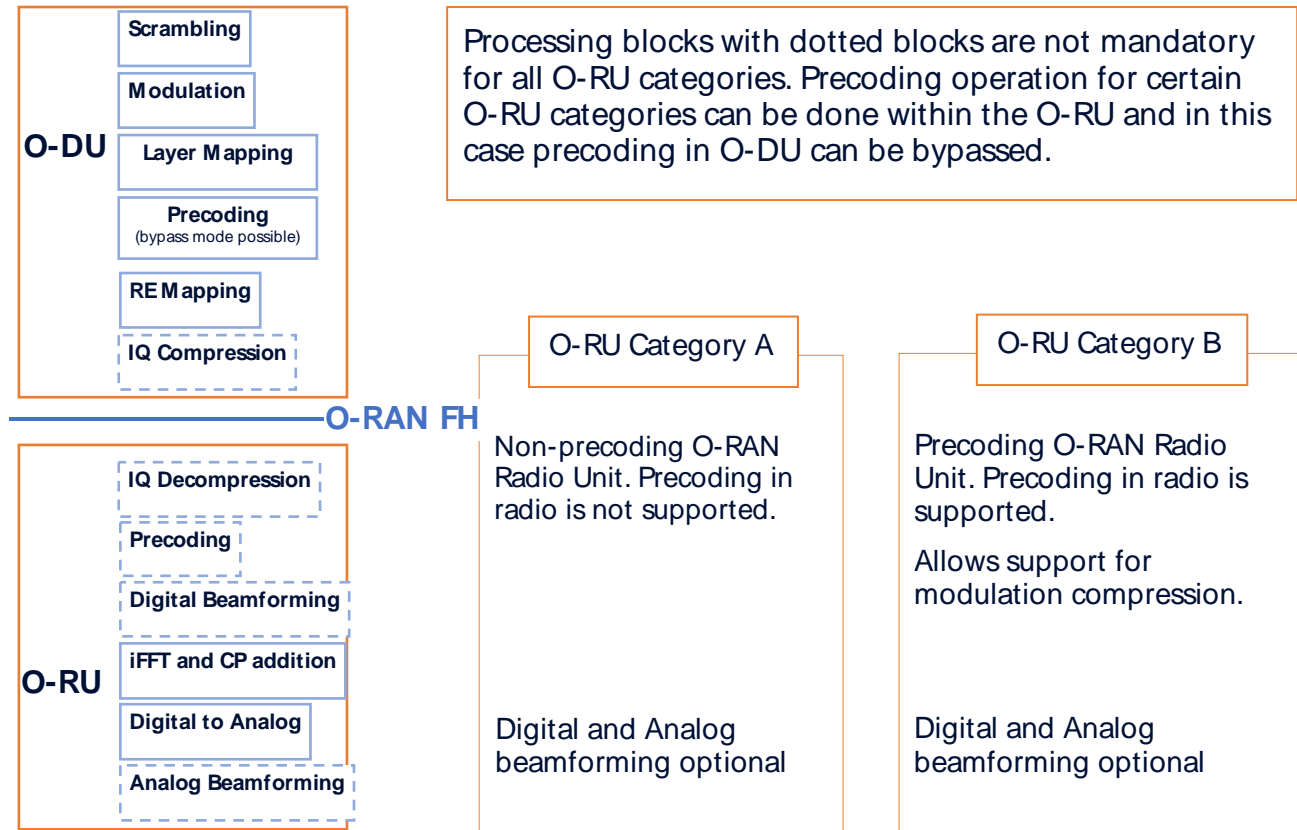


Figure 4.2-1: Split Point and Category A and Category B O-RAN Radio Units

The inclusion of these two O-RU categories has certain implications for the LTE and NR functional splits in both DL and UL which are reflected in the following clauses. In particular, for a Category A O-RU, the precoding shall be executed in the O-DU and any beamforming in the O-RU, if present, would exclude the precoding calculation. For LTE precoding schemes other than Transmit Diversity and Spatial Multiplexing with Large Delay CDD and precoding schemes for NR, the precoding may be included in a digital beamforming processing block within the O-RU for a Category B O-RU (even for analog beamforming O-RUs). To implement Transmit Diversity and Spatial Multiplexing with Large Delay CDD precoding special C-Plane instructions need to be provided to the O-RU from the O-DU; this is described in clause 7.7.3.

4.2.2 Split 7-2x (DL)

4.2.2.1 Description

DL functional split for various physical layer channels is illustrated in Figure 4.2.2.1-1 (LTE Category A O-RUs), Figure 4.2.2.1-2 (LTE Category B O-RUs), Figure 4.2.2.1-3 (NR Category A O-RUs), and Figure 4.2.2.1-4 (NR Category B O-RUs).

When O-DU supports Category A O-RU, O-DU shall support a total number of precoded streams of up to 8. O-DU supporting Category A O-RU may support more than 8 precoded streams. For LTE (e.g., LTE TM9) and NR PDSCH with UE specific reference signals, the DL processing chain specified by 3GPP does not include a precoding operation. The detailed precoding operation referred to in Figure 4.2.2.1-1 and Figure 4.2.2.1-2 for PDSCH with UE specific reference signals is not further described within the present document.

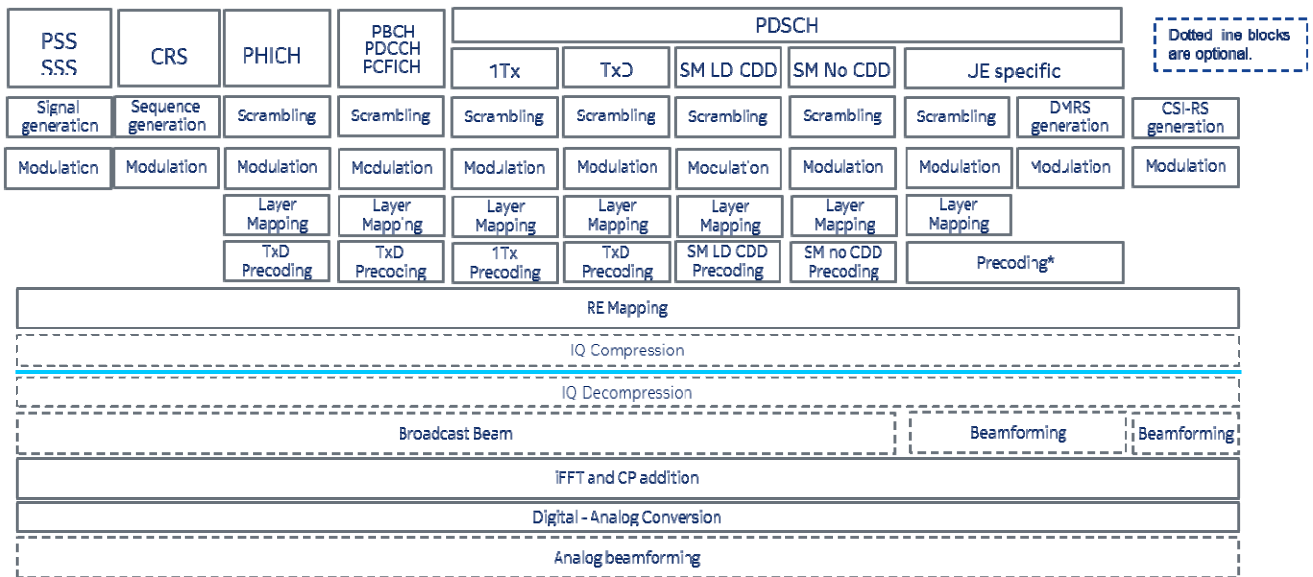


Figure 4.2.2.1-1: Lower layer DL split description, LTE, Category A O-RUs

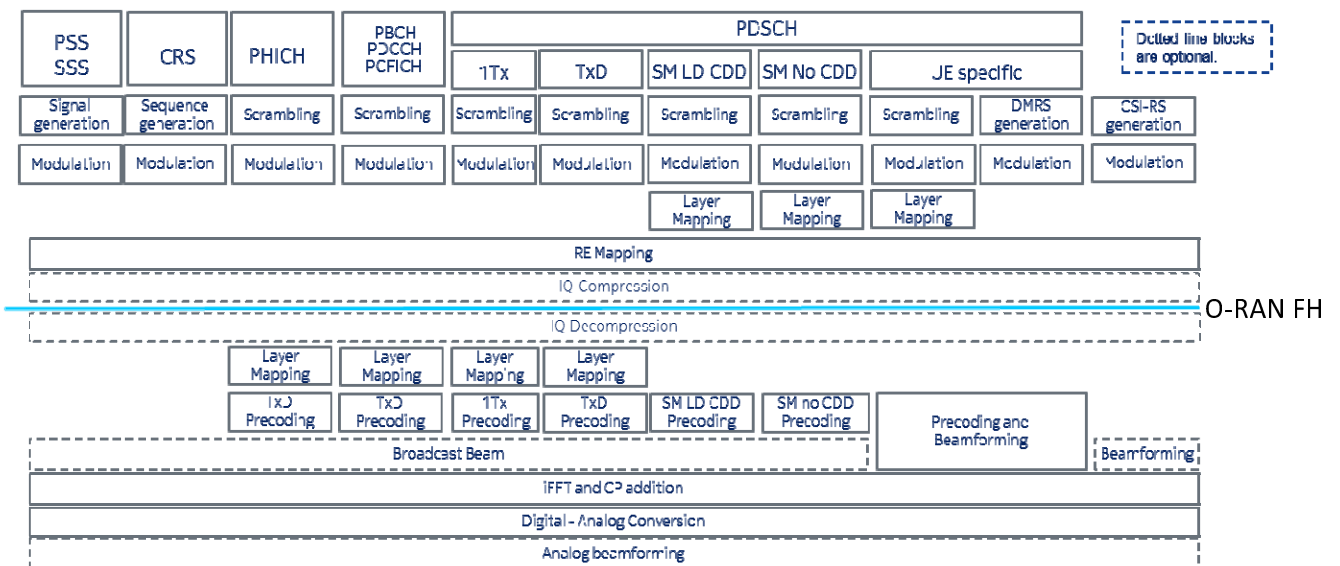
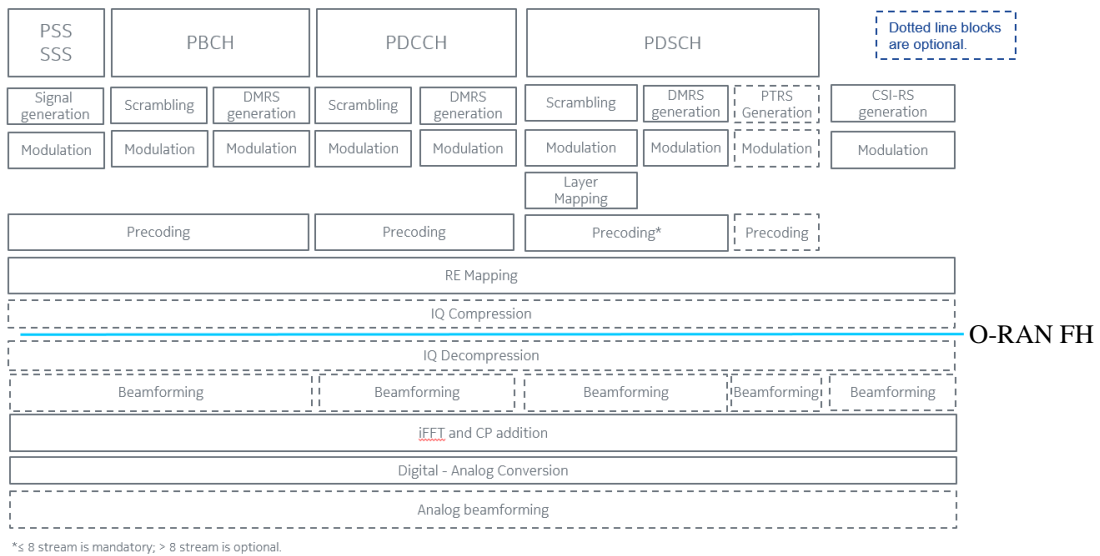


Figure 4.2.2.1-2: Lower layer DL split description, LTE, Category B O-RUs

Figure 4.2.2.1-2 presents DL functional split for Category B O-RU supporting precoding including precoding with Transmit Diversity and Spatial Multiplexing with Large Delay CDD schemes; for Category B O-RU not supporting Transmit Diversity and Spatial Multiplexing with Large Delay CDD schemes the channels precoded with Transmit Diversity or Spatial Multiplexing with Large Delay CDD shall be precoded by O-DU as presented on Figure 4.2.2.1-1 and the other channels shall be precoded on O-RU as presented on Figure 4.2.2.1-2.

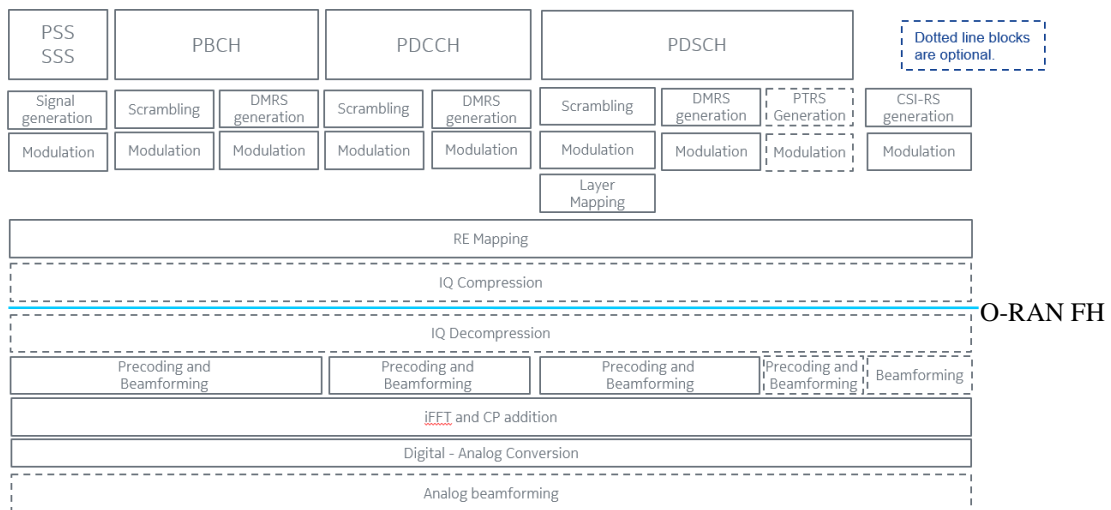
Although not explicitly indicated in Figure 4.2.2.1-2, RE mapping operation is separated into two parts for precoding support at O-RU for transmit diversity and spatial multiplexing modes in Category B O-RU:

- 1) RE mapping to frequency resources is performed at O-DU.
- 2) RE mapping to antenna ports is performed at O-RU after precoding.



NOTE: The above figure illustrates some of the high-level functionalities inside the O-RU and O-DU, and do not bind the internal design of each of the O-RU and O-DU.

Figure 4.2.2.1-3: Lower layer DL split description, NR, Category A Radio



NOTE: The above figure illustrates some of the high-level functionalities inside the O-RU and O-DU, and do not bind the internal design of each of the O-RU and O-DU.

Figure 4.2.2.1-4: Lower layer DL split description, NR, Category B Radio

In the DL, OFDM phase compensation (clause 5.4 of [4]), iFFT, CP addition, and digital beamforming functions reside in the O-RU for Category A and Category B O-RUs. In addition, precoding with schemes other than Transmit Diversity and Spatial Multiplexing with Large Delay CDD reside in the Category B O-RU. Optionally, precoding with Transmit Diversity and Spatial Multiplexing with Large Delay CDD precoding schemes reside in the Category B O-RU supporting these schemes. The PHY functions including resource element mapping, precoding, layer mapping, modulation, scrambling, rate matching and coding not handled by O-RU reside in the O-DU.

NOTE: The option of including some PHY functionality in a radio unit was not included in 3GPP TR 38.801 [i.4], clause 11.1.

4.2.2.2 Benefits and justification (informative)

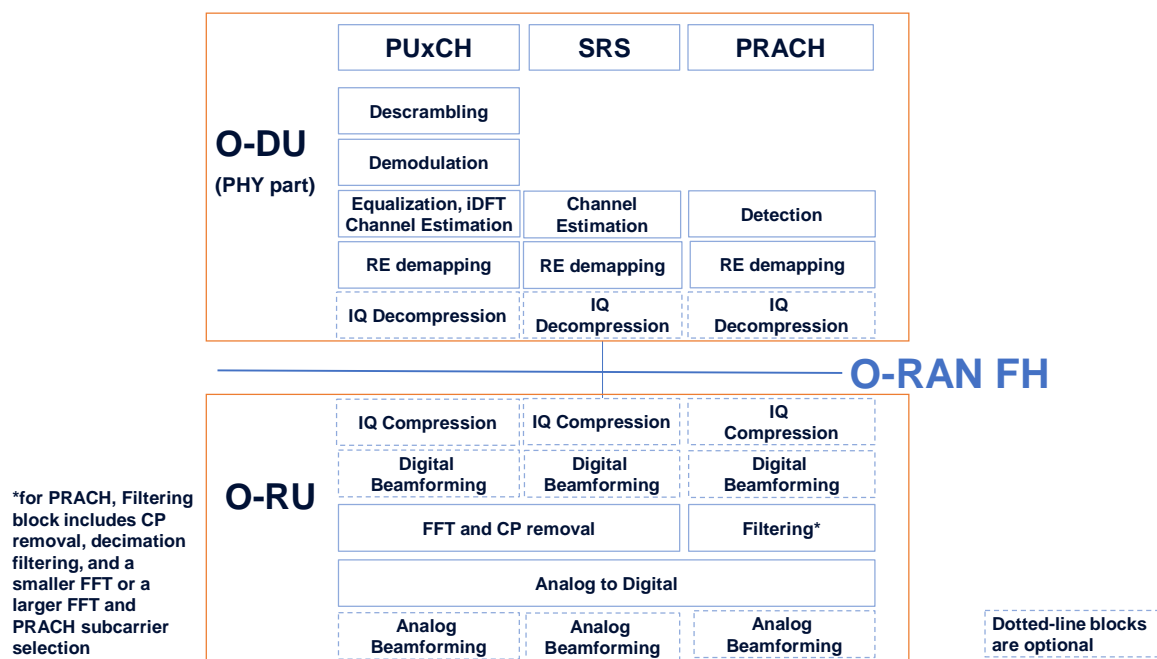
- Interface simplicity: Transfer of user plane data is based on Resource Elements / Physical Resource Blocks, which simplifies the data mapping and limits the required associated control messages

- Transport Bandwidth Scalability: Lower split options (such as splits 7-1 and 8 as defined in 3GPP TR 38.801 [13], clause 11.1) scale based on number of antennas. In contrast, 7-2x interface scales based on "streams", which allows using high number of antennas without higher transport bandwidth. Further, sending only PRBs that contain user data is possible for purpose of reducing transport bandwidth
- Beamforming Support: The same interface design supports different beamforming techniques (digital, analog, hybrid) as well as different beamforming algorithms. Likewise, deployments using only analog beamforming are also possible with the same interface design.
- Interoperability: Less user specific parameters are used at split 7-2x (when compared to higher split options), which simplifies specification.
- Advanced receivers and inter-cell coordination: this option allows implementation of advanced receivers and coordination features, which are also easier to implement and less restricted when most functions are placed at the O-DU. For example, UL CoMP is not possible when the UL upper-PHY processing is in the O-RU.
- Lower O-RU complexity: Less functions at O-RU (when compared to higher split options) allow limiting the number of required real time calculations as well as required memory requirement, especially for Category A O-RUs.
- Future proof-ness: Placing most functions at O-DU will allow introduction of new features via software upgrades without inflicting HW changes at O-RU (e.g., specification changes due to URLLC or new modulation schemes).
- Interface and functions symmetry: If the same interface and split point is used for DL and UL, specification effort is reduced.

4.2.3 Split 7-2x (UL)

4.2.3.1 Description

UL functional split for various physical layer channels and transmission modes are illustrated in Figure 4.2.3.1-1. Likewise, digital beamforming in this context, is the function of antenna port selection or antenna port combining.



NOTE: The above figure illustrates some of the high level functionalities inside the O-RU and O-DU, and do not bind the internal design of each of the O-RU and O-DU.

Figure 4.2.3.1-1: Lower layer UL split description for LTE and NR

In the UL, OFDM phase compensation (for all channels except PRACH when Section Type "3" message is used) according to clause 5.4 of [4], FFT, CP removal and digital beamforming functions reside in the O-RU. The rest of the PHY functions including resource element de-mapping, equalization, de-modulation, de-scrambling, rate de-matching and de-coding reside in the O-DU.

NOTE: The option of including some PHY functionality in a radio unit was not included in 3GPP TR 38.801 [i.4], clause 11.1.

4.2.3.2 Benefits and justification (informative)

The benefits applicable for Option 7-2x for DL (see clause 4.2.2.2) are also applicable for Uplink.

4.3 Data flows

An overall reference of required inputs for some major functional blocks, their granularity and input originating source is compiled in Table 4.3-1 and Table 4.3-2 for DL and UL respectively.

Table 4.3-1: Required information for each functional block (DL)

Function	Required information	Signaling granularity	Information source
Digital beamforming	<ul style="list-style-type: none"> Digital beamforming information Number of TRXs or Beams 	Per RE per beam per slot	Scheduler (O-DU)
iFFT and CP addition	<ul style="list-style-type: none"> FFT size SC spacing CP length 	Per TRX per slot	Scheduler (O-DU)
Analog beamforming	<ul style="list-style-type: none"> Analog beamforming info 	Per frequency band per slot	Scheduler (O-DU)
Listen Before Talk (LBT)	<ul style="list-style-type: none"> LBT parameters 	Per MCOT	Scheduler (O-DU)

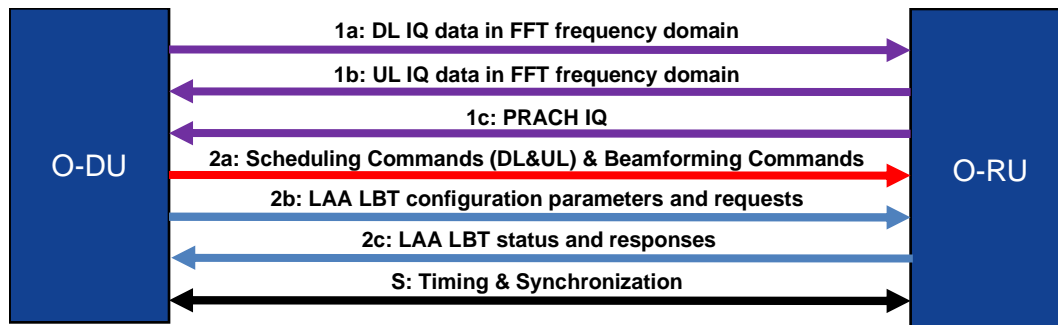
Table 4.3-2: Required Information for each functional block (UL)

Function	Required information	Signaling granularity	Information source
Digital beamforming	<ul style="list-style-type: none"> Digital beamforming information Number of TRXs or Beams 	Per RE per TRX per slot	Scheduler (O-DU)
FFT and CP removal	<ul style="list-style-type: none"> FFT size SC spacing CP length 	Per TRX per slot	Scheduler (O-DU)
Analog beamforming	<ul style="list-style-type: none"> Analog beamforming info 	Per frequency band per slot	Scheduler (O-DU)

In case of lower layer fronthaul based on split option 7-2x for DL and UL, the required external data flows (excluding M-Plane) to exchange information between O-DU and O-RU are categorized as follows. Further, the data flow direction is shown in figure 4.3-1, the data flow contents mapping is described in Table 4.3-3.

- User Plane
 - Data Flow 1a: Flows of IQ Data in FFT frequency domain on DL
 - Data Flow 1b: Flows of IQ Data in FFT frequency domain on UL
 - Data Flow 1c: Flow of PRACH IQ data in FFT frequency domain

- Control Plane
 - Data Flow 2a: Scheduling commands (DL & UL) & Beamforming commands
 - Data Flow 2b: LAA LBT configuration commands and requests
 - Data Flow 2c: LAA LBT status and response messages
- Synchronization Plane
 - Data Flow S: Timing and Synchronization data



NOTE: M-Plane flows are not presented here.

Figure 4.3-1: Lower layer fronthaul data flows

Table 4.3-3: Data flow information mapping

Plane	ID	Name	Contents	Periodicity
U-Plane	1a	DL Frequency Domain IQ Data	DL user data (PDSCH), control channel data (PDCCH, etc.), ...	< slot
	1b	UL Frequency Domain IQ Data	UL user data (PUSCH), control channel data (PUCCH, etc.), ...	< slot
	1c	PRACH Frequency Domain IQ Data	UL PRACH data	< slot
C-Plane	2a	Scheduling Commands (DL&UL) & Beamforming Commands	Scheduling information, FFT size, CP length, Subcarrier spacing, UL PRACH scheduling DL and UL Beamforming commands (e.g., beam index) and scheduling	~ slot
	2b	LAA LBT configuration parameters and requests	LBT Configuration parameters such as lbtHandle, lbtDeferFactor, lbtBackoffCounter, lbtOffset, MCOT, lbtMode, sfnSf, lbtCWconfig_H, lbtCWconfig_T, lbtTrafficClass.	per MCOT/DRS
	2c	LAA LBT status and responses	LBT DL indication parameters such as lbtHandle, lbtResult, initialPartialSFs, bufferError, lbtCWR_Result	
S-Plane	S	Timing and Synchronization	SyncE SSM & IEEE 1588 [3], [33] PTP packets	

4.4 Latency requirements

4.4.1 Introduction

Intra-PHY lower layer fronthaul split has characteristic of a stringent bandwidth and tight latency requirement. This implies use of a special "Fronthaul Service Profile" to be supported by the transport network, and which may differ depending on the operating environment, topology and target use cases. The general concept and latency model is based on eCPRI reference points as specified in eCPRI Specification [2], clause 6.3.1, for delay management definitions (see Figure 4.4.1-1). However, the present document provides additional details for both DL and UL (see Table 4.4.1-1) latency. The delay parameters and how these are determined are explained below; in general it is expected a separate set of timing parameters applies to each O-RU attached to an O-DU. This clause defines multiple approaches which can be used to provide an inter-operable timing solution.

The reference points defined for eCPRI are reflected below in Figure 4.4.1-1. The reference points are:

- O-DU: R1 / R4 – Transmit / Receive interface at O-DU
- O-RU: R2 / R3 – Receive / Transmit interface at O-RU
- Ra: Antenna interface at O-RU

NOTE: When an external antenna is used with a cable imposing a negligible delay, then the O-RU connector to the external antenna can be assumed as Ra.




Figure 4.4.1-1: Definition of reference points for delay management (adapted from [2])

Transmission delay between O-DU and O-RU are specified as T12 (downlink) and T34 (uplink). The transmission delay encompasses only the time from when a bit leaves the sender (R1 / R3) until it is received at the receiver (R2 / R4). In an ethernet transport network, these delays may not be constant due to switching delays (i.e., PDV). To account for this, transport delay shall be considered as a range with upper and lower bounds:

- Downlink transport delay: T12min / T12max
- Uplink transport delay: T34min / T34max

However, fixed timing at Ra is still required. Therefore, Ra is used as a reference point for delay management in the eCPRI model and transmission and reception at the reference points shall be measured relative to Ra, resulting in the following parameters in Table 4.4.1-1.

Table 4.4.1-1: eCPRI O-DU/ O-RU delay model latency parameters (adapted from [2])

		eCPRI	Latency	eCPRI	
				Minimum	Maximum
	O-DU	T1a	Measured from output at O-DU (R1) to transmission over the air.	T1amin	T1amax
	O-RU	T2a	Measured from reception at O-RU (R2) to transmission over the air	T2amin	T2amax
	O-DU	Ta4	Measured from reception at O-RU antenna to reception at O-DU port (R4).	Ta4min	Ta4max
	O-RU	Ta3	Measured from reception at O-RU antenna to output at O-RU port (R3)	Ta3min	Ta3max

A common timescale is used on both ends of the link. The relative time error of the S-plane measurement signals between the O-DU and O-RU, for the purposes of latency requirements management, should be within a limit of 3 μ s ($\pm 1.5 \mu$ s).

- This requirement allows for the measurement of the transport delay (i.e., between O-DU and O-RU UNIs) with a measurement error that is sufficiently lower than the delay that is to be measured (e.g., 10 times lower).
- In case of LLS-C1 and LLS-C2, due to clause 11.3.2 requirements, the relative time error of the S-plane measurement signals between the O-DU and O-RU is within the limit of 3 μ s ($\pm 1.5 \mu$ s).
- In case of LLS-C3 and LLS-C4, the synchronization network should meet G.8271.1 or G.8271.2 network limits in order to meet similar requirements (relative time error of the S-plane measurement signals between the O-DU and O-RU within a limit of 3 μ s).
- The upper bound on the absolute time error requirement at the O-DU S-Plane is dictated by the O-RU's receive window, the delay and PDV in the transport network and the O-DU's internal delays.

4.4.2 Timing parameter relationships

4.4.2.1 General

To ensure proper reception of data at the receiver over the packet interface, there are several relationships between the parameters defined above which shall be met. First it is important to understand the nature of the transmission itself. In either direction (downlink/ uplink) it takes some amount of time for the sender to actually transmit the packets onto the transmission media. However, the amount of data for any interval (e.g., symbol) can vary thus resulting differing transmit times. This transmission time can be affected by several factors including (but not limited to) transport media rate, air interface bandwidth, and amount of data compression. The maximum amount of time allowed for the transmitter to send all data for an interval (Transmission Window) is defined by $T1_{amax} - T1_{amin}$. This is the allowed time, based on transport and O-RU characteristics, and its impacts on O-DU is explained in clause 4.4.2.2.

To account for transport variation and transmission time the receiver implements a reception window. This allows packets containing samples for a specific symbol to be received within the window and still be transmitted at R_a at the required time. The size of the Reception Window shall account for both the maximum transmission time at the sender and the transport variation through the fronthaul network. The result is the first of the delay relationships which shall be met to ensure a working delay solution, see Table 4.4.2.1-1.

$$\text{Reception Window} \geq (\text{Transmission Window} + \text{Transport Variation})$$

Table 4.4.2.1-1: eCPRI Delay Windows

	Reception Window	Transmission Window	Transport Variation
Downlink	$T2_{amax} - T2_{amin}$	$T1_{amax} - T1_{amin}$	$T12_{max} - T12_{min}$
Uplink	$Ta4_{max} - Ta4_{min}$	$Ta3_{max} - Ta3_{min}$	$T34_{max} - T34_{min}$

The position (in time) of the reception/ transmission windows at the O-RU is fixed relative to the air interface. However, the position of the corresponding windows at the O-DU is a function of the O-RU and transport parameters. For guaranteed reception of packets sent from O-DU to O-RU within the O-RU reception window, the following relationships shown in Table 4.4.2.1-2 shall also be met (refer to Annex B for explanation of these inequalities):

Table 4.4.2.1-2: O-DU transmission/ reception window position (in time)

O-DU Timing		Parameter	O-DU Transmit Boundary Relationships
Downlink (Transmit)	No earlier than	$T1_{amax}$	$T1_{amax} \leq (T2_{amax} + T12_{min})$
	No later than	$T1_{amin}$	$T1_{amin} \geq (T2_{amin} + T12_{max})$
Uplink (Receive)	No earlier than	$Ta4_{min}$	$Ta4_{min} \leq (Ta3_{min} + T34_{min})$
	No later than	$Ta4_{max}$	$Ta4_{max} \geq (Ta3_{max} + T34_{max})$

4.4.2.2 O-DU Transmission Window

The U-Plane O-DU transmission window ($T1_{\max} - T1_{\min}$) is defined by the relationships above based on the O-RU reception window and max transport variation. It does not define the exact timing of transmission from the O-DU. Rather, it defines the boundaries that the U-Plane O-DU transmission shall operate within. The window merely represents the mathematical boundaries imposed on the O-DU as a result of the O-RU and Transport constraints. It is possible to define the constraints for any one of O-DU, transport and O-RU based on knowledge of the other two. However, typically the O-RU constraint is predefined based on the equipment, and transport is part of the overall network goals.

The window resulting from the relationships shall be greater than or equal to the actual maximum time required by the O-DU to transmit all data for a symbol ($TX_{\max O-DU}$). That is, the window shall be at least large enough that the O-DU can transmit in the worst case within the window. Where, within the window, the O-DU transmits (e.g., beginning, middle, end) and how much of the window is consumed by the O-DU transmission is a matter of O-DU design.

EXAMPLE: The following downlink example illustrates the concept:

- O-RU parameters: $T2_{\min} = 100 \mu\text{s}$, $T2_{\max} = 260 \mu\text{s}$.
- Transport Parameters (direct fiber of known length): $T12_{\min} = 50 \mu\text{s}$, $T12_{\max} = 51 \mu\text{s}$.

The result indicates an O-DU transmission window as follows:

- $T1_{\max} \leq (260 + 50) = 310 \mu\text{s}$.
- $T1_{\min} \geq (100 + 51) = 151 \mu\text{s}$.

This provides a very large transmit window available to the O-DU. If, for example, the $TX_{\max O-DU}$ is only $30 \mu\text{s}$, then the O-DU can determine where within the window to start its transmission, so long as the transmission completes prior to $T1_{\min}$.

If, however, this same O-DU were paired with an O-RU with smaller reception window (e.g., $T2_{\min} = 100 \mu\text{s}$, $T2_{\max} = 150 \mu\text{s}$) using a transport network with the same $T12_{\min}$, but with $15 \mu\text{s}$ of PDV ($T12_{\max} = 65 \mu\text{s}$), the result is:

- $T1_{\max} \leq (150 + 50) = 200 \mu\text{s}$.
- $T1_{\min} \geq (100 + 65) = 165 \mu\text{s}$.

The delay solution still works ($200 - 165 \geq 30$), but with far less margin, and far less flexibility as to where within the window the O-DU may begin transmission.

The C-Plane O-DU transmission window follows the same concept as the U-Plane O-DU transmit window. However, the C-Plane O-DU transmission window can be larger in size and hence, can start much earlier in time compared to the U-Plane O-DU transmission window.

In certain scenarios such as sending the beamforming weights in real-time for mMIMO application using Section Extension 11, the C-Plane messages need to be spread over time to avoid peak throughput. This use case is shown in Figure 4.4.2.2-1.




Figure 4.4.2.2-1: Example of the benefit of using a longer C-Plane tx/rx window compared to the U-Plane windows

4.4.3 U-Plane/ C-Plane timing

The basic delay parameters above describe the general delay model and characteristics of the fronthaul interface. However, the fronthaul interface is divided into C-Plane and U-Plane parts. The C-Plane shall be available in order to process the corresponding U-Plane packets. To support coordination of C-Plane and U-Plane timing, the fronthaul interface specifies that C-Plane messages shall arrive at the O-RU some amount of time in advance ($T_{cp_adv_dl}$) of the latest possible time the first corresponding U-Plane messages may arrive. As a result, the present document has defined the downlink timing relationships and parameters shown in Figure 4.4.3-1.

NOTE: It is possible for $T_{cp_adv_dl}$ to be set to 0.

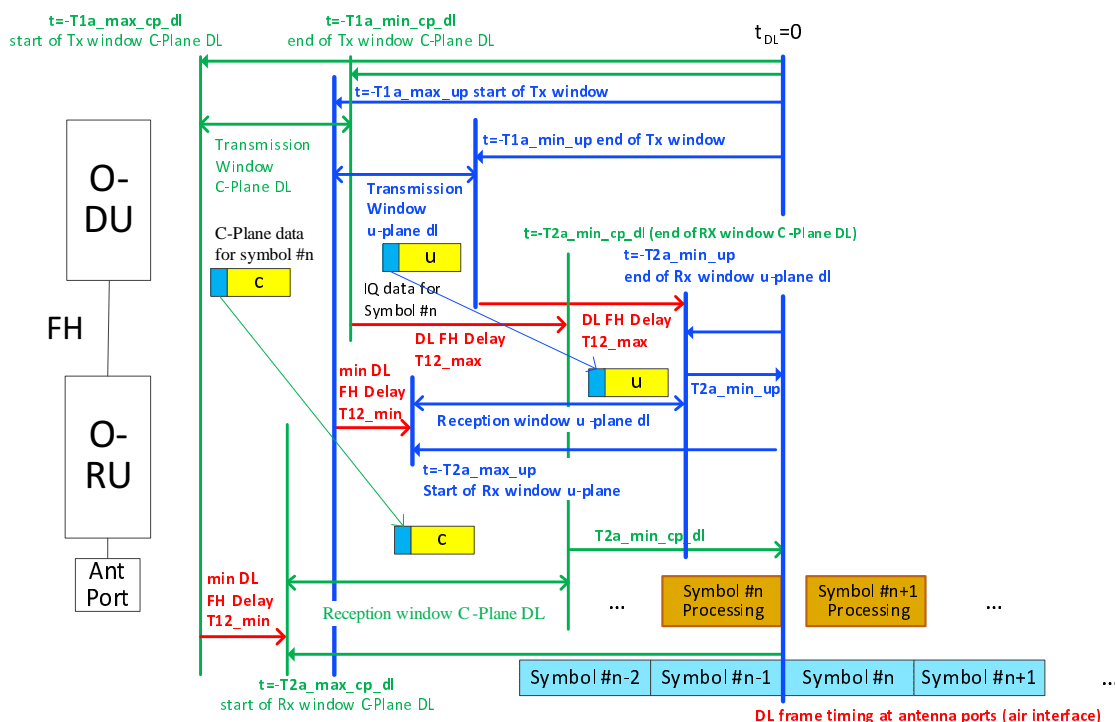


Figure 4.4.3-1: Timing relations per symbol IQ in DL direction (U-Plane and C-Plane)

The present document does not assume that the transmission windows for C-Plane and U-Plane are of the same size. $T_{cp_adv_dl}$ allows for different alignment of the respective transmission/ reception windows as illustrated in the Figure 4.4.3-1. The reference (denoted as " $t_{DL}=0$ ") is the transmission of the earliest IQ sample in time domain within a symbol (including cyclic prefix) which is generated from the IQ data received in a U-Plane messages specific to a symbol identified by symbolId (optionally adjusted by timeOffset).

A downlink C-Plane message may refer to one or more symbols; transmission and reception windows for a downlink C-Plane message referencing multiple symbols are relative to the start of the earliest symbol referenced by the message (identified by startSymbolId; optionally adjusted by timeOffset).

For uplink, the C-Plane is sent from O-DU to O-RU, while the U-Plane is from O-RU to O-DU. Since uplink C-Plane is from O-DU to O-RU, it uses the R1/ R2 reference points, and corresponding timing parameters. However, the uplink C-Plane timing is not characterized relative to the downlink C-Plane or U-Plane timing, but relative to Ra. The resulting fronthaul delay model for non-PRACH channels is shown in Figure 4.4.3-2.

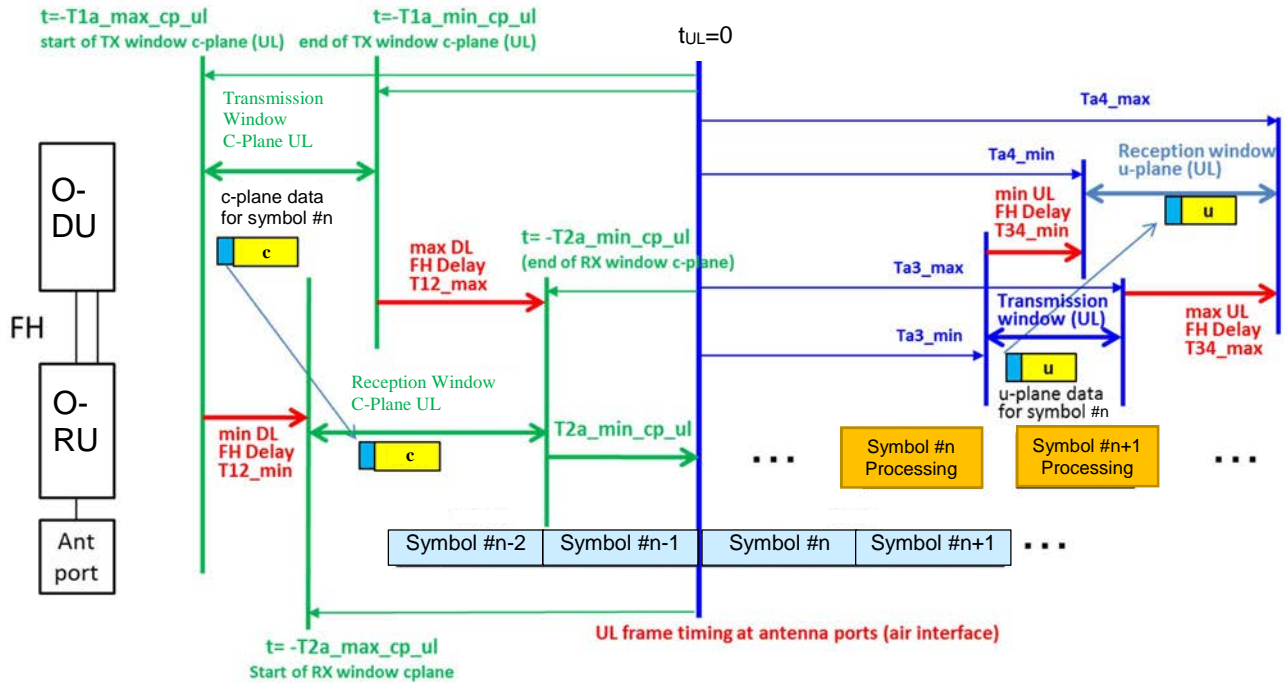


Figure 4.4.3-2: Timing relations per symbol IQ in UL direction (U-Plane and C-Plane)

An uplink C-Plane message not describing PRACH channels can refer to one or more symbols. The reference point for uplink C-Plane message (denoted as " $t_{UL}=0$ ") is the reception of the earliest IQ sample in time domain (start of cyclic prefix) that is specific to the earliest symbol referenced by the message (pointed by startSymbolId; if Section Type 3 message is used, timeOffset will also point to this reference point).

The reference point for uplink U-Plane messages related to non-PRACH channels (denoted as " $t_{UL}=0$ ") is the reception of the earliest IQ sample in time domain within a symbol (start of cyclic prefix) referenced by the message (pointed by symbolId; if Section Type 3 message is used, timeOffset will also point to this reference point).

NOTE: " $t_{UL}=0$ " is not affected by FFT sampling offset configured over M-Plane since configured FFT sampling offset cannot be longer than CP-length.

Due to the characteristics regarding CP-length and symbol-length of PRACH-symbols, PRACH is handled in a separate way. Especially for the long preamble PRACH formats 0...3 (see 3GPP TS 38.211 [4], clause 6.3.3) the length of both the CP and the length of the PRACH Symbol is much larger than corresponding parts for e.g., PUSCH symbols. For the other PRACH formats (A1...C2) the length of the CP could be very long compared to the CP for the PUSCH symbols (up to roughly 14 times longer (C2)). The consequence of either long CP and/or long PRACH Symbol of specific PRACH formats is that the values of $Ta3_{min}$ and $Ta3_{max}$ are increased. Due to this, PRACH CP duration and number of PRACH Symbols are excluded from the values for these parameters provided by the O-RU.

The O-DU is aware of the characteristics of the PRACH format being used (e.g., length of CP, number of PRACH Symbols), and with this information the position and size of the reception window is calculated by the O-DU. The used values for $Ta3_{min}$ and $Ta3_{max}$ are the ones being retrieved from the O-RU for the SCS being used for PRACH.

An uplink C-Plane message describing PRACH can refer to one or more PRACH Symbols depending on the PRACH format. For PRACH formats with multiple PRACH symbols, only a single C-Plane message is used to refer to all PRACH Symbols for a particular PRACH occasion.

The reference point for PRACH C-Plane message is pointed by startSymbolId and:

- Is the reception of the earliest IQ sample in time domain (start of cyclic prefix) that is specific to the PRACH Symbol referenced by the message.
 - This is allowed only for PRACH formats without repetitions and with only one PRACH occasion in a slot, e.g. Format 0.
 - In this case, cpLength is set to a non-zero value in the PRACH C-Plane (Section Type 3) message, which should indicate the actual CP duration.
- Or is the start of the latest symbol timing that starts right at or before the reception of the earliest IQ sample in time domain (after PRACH cyclic prefix) that is specific to the earliest PRACH Symbol referenced by the message.
 - This is allowed for all PRACH formats.
 - In this case, cpLength is set to zero in the PRACH C-Plane (Section Type 3) message.
 - If the SCS value provided by "frameStructure" is equal to or greater than 15 kHz, then the symbol timing used to determine startSymbolId is based on the numerology of the SCS value provided by "frameStructure". Else, if the SCS value provided by "frameStructure" is less than 15 kHz (e.g. for long preamble PRACH formats), then the symbol timing used to determine startSymbolId is based on the numerology of 15 kHz SCS.
 - When **ul-fft-sampling-offset** for PRACH is configured (via M-Plane) to a non-zero value, the timing of the earliest IQ sample is advanced by that value. However, the reference symbol timing that is used to determine the startSymbolId value in the PRACH C-Plane message is unaffected by **ul-fft-sampling-offset**. Therefore, with a non-zero **ul-fft-sampling offset** for PRACH, it is possible that "the latest symbol timing that starts right at or before the reception of the earliest IQ sample in time domain (after PRACH cyclic prefix) that is specific to the earliest PRACH Symbol" may change to an earlier symbol timing compared to the case where **ul-fft-sampling-offset** for PRACH is zero. So, the O-DU needs to take **ul-fft-sampling-offset** for PRACH into account when determining the startSymbolId value. However, it is noted that even if **ul-fft-sampling-offset** for PRACH is configured to a non-zero value, it is expected to be in the order of nanoseconds (i.e., less than 1 microsecond).

The reference point for PRACH U-Plane message is the reception of the earliest IQ sample in time domain within a PRACH Symbol (after PRACH cyclic prefix) referenced by the message.

- For the first PRACH Symbol.
 - If cpLength is set to a non-zero value in the PRACH C-Plane (Section Type 3) message referencing the PRACH Symbol, this reference point for PRACH U-Plane message is pointed by timeOffset (pointing to the start of PRACH cyclic prefix) + cpLength – **ul-fft-sampling-offset** (for PRACH).
 - Else, if cpLength is set to zero in the PRACH C-Plane (Section Type 3) message referencing the PRACH repetition, this reference point for PRACH U-Plane message is pointed by timeOffset (pointing to the start of the PRACH Symbol (after PRACH cyclic prefix)) – **ul-fft-sampling-offset** (for PRACH).
- For subsequent PRACH Symbols (for PRACH formats with multiple PRACH Symbols), this reference point for PRACH U-Plane message is shifted by the PRACH Symbol duration.

Timing relationships for PRACH taking some specific PRACH formats (Format 0, Format 1, Format B4 and Format C0) as examples are shown in Figure 4.4.3-3 through Figure 4.4.3-6.

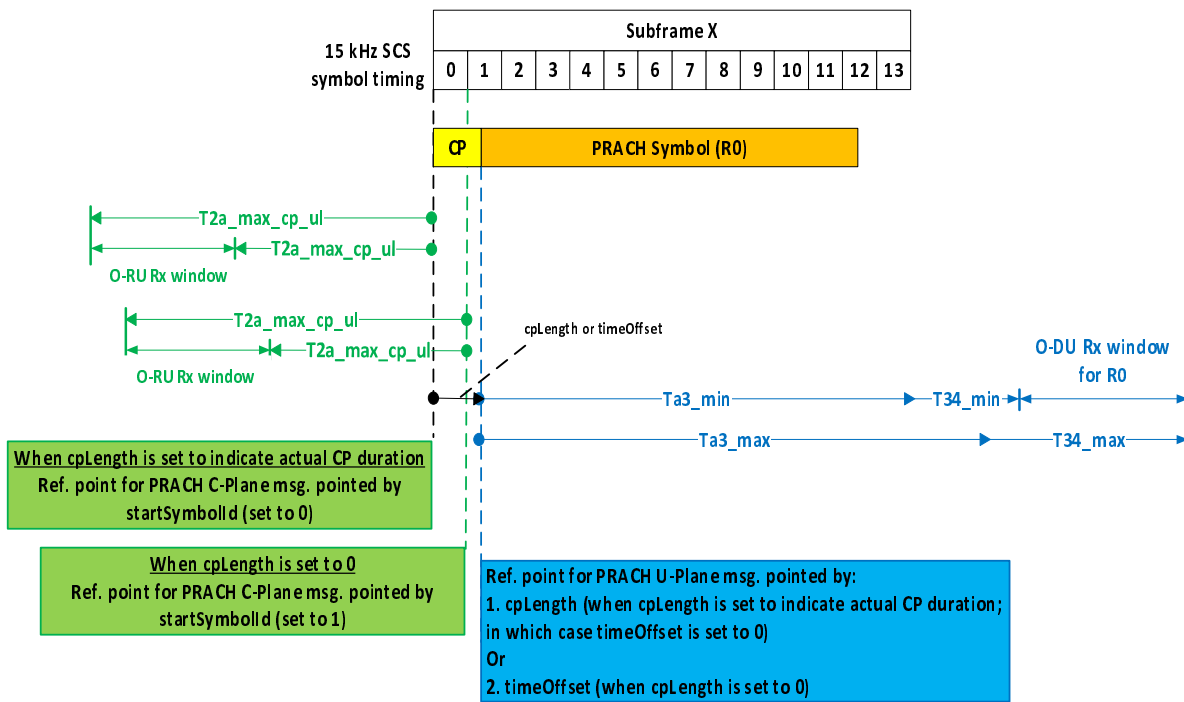


Figure 4.4.3-3: Timing relations for PRACH (Example 1: PRACH Format 0)

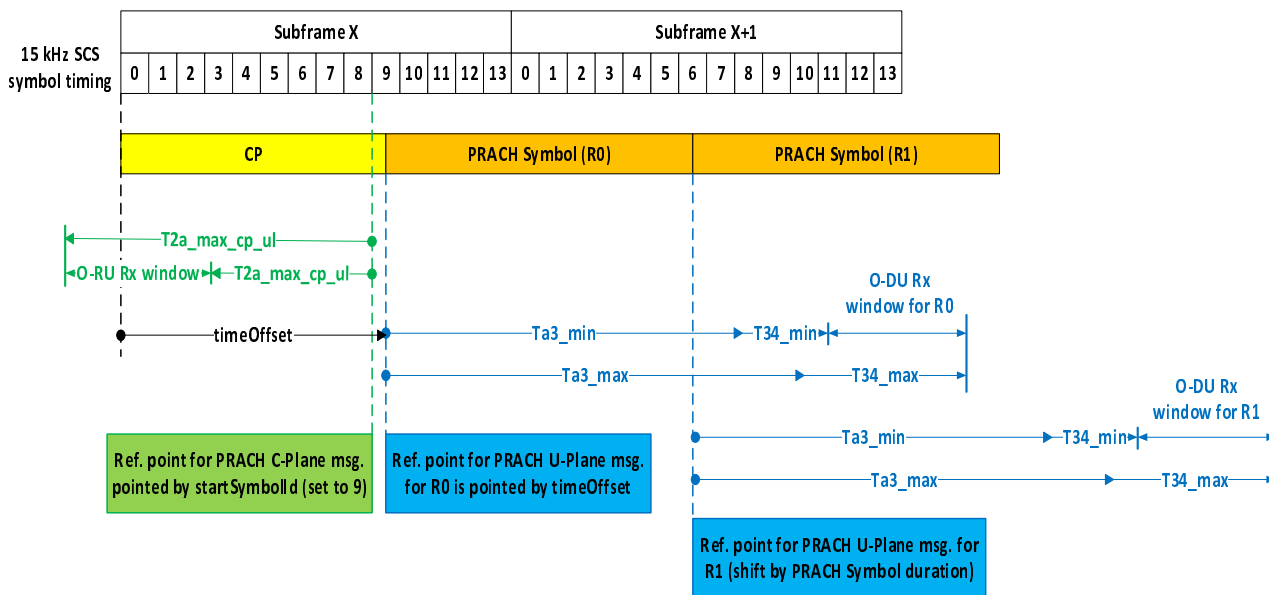


Figure 4.4.3-4: Timing relations for PRACH (Example 2: PRACH Format 1)

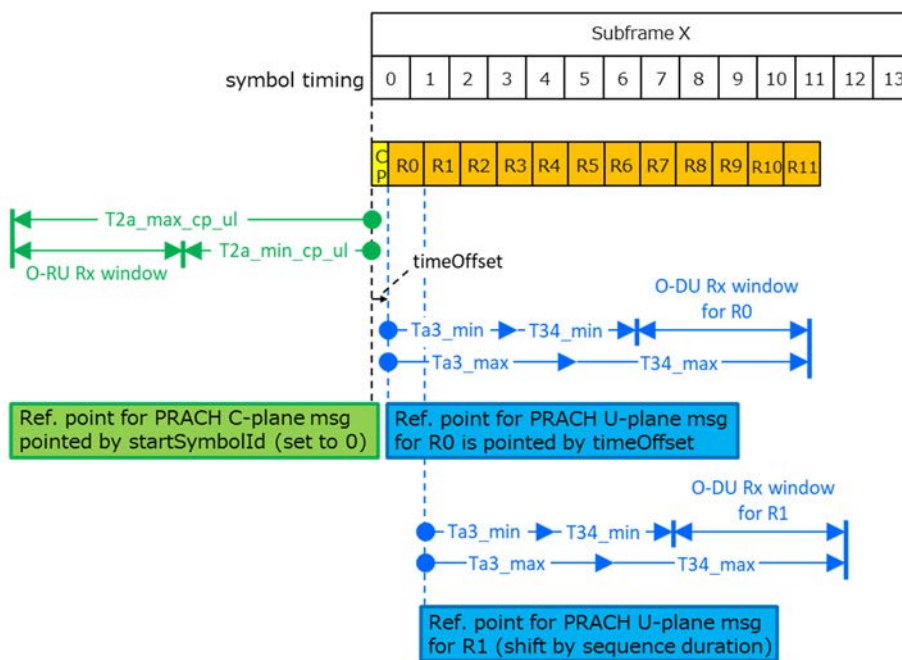


Figure 4.4.3-5: Timing relations for PRACH (Example 3: PRACH Format B4)

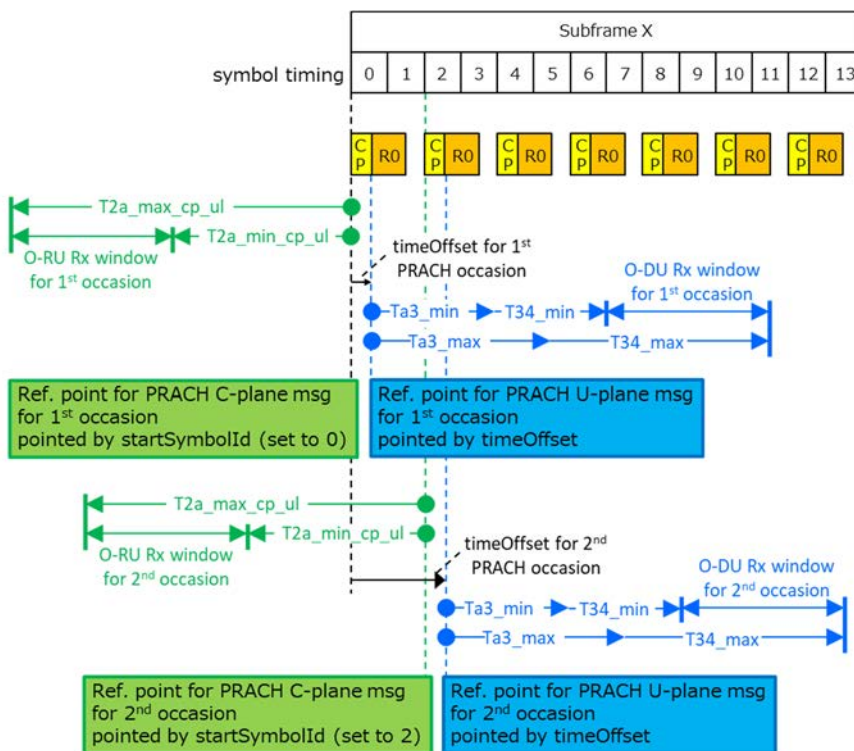


Figure 4.4.3-6: Timing relations for PRACH (Example 4: PRACH Format C0)

The resulting O-RAN delay model parameters are summarized in Table 4.4.3-1. The respective sections within the table for O-RU and O-DU define the delay profile parameters for the equipment.

Table 4.4.3-1: Delay management model parameters

	Parameters	C-Plane		U-Plane	
		DL	UL	DL	UL
Network	T12min	T12_min	T12_min	T12_min	NA
	T12max	T12_max	T12_max	T12_max	NA
	T34min	NA	NA	NA	T34_min
	T34max	NA	NA	NA	T34_max
O-RU	T2amin	T2a_min_cp_dl	T2a_min_cp_ul	T2a_min_up	NA
	T2amax	T2a_max_cp_dl	T2a_max_cp_ul	T2a_max_up	NA
		Tcp_adv_dl	NA	NA	NA
	Ta3min	NA	NA	NA	Ta3_min
	Ta3max	NA	NA	NA	Ta3_max
O-DU	T1amin	T1a_min_cp_dl	T1a_min_cp_ul	T1a_min_up	NA
	T1amax	T1a_max_cp_dl	T1a_max_cp_ul	T1a_max_up	NA
	Ta4min	NA	NA	NA	Ta4_min
	Ta4max	NA	NA	NA	Ta4_max

The relationships between the various delay model parameters are defined below in Table 4.4.3-2 and Table 4.4.3-3:

Table 4.4.3-2: Downlink delay relationships

	Earliest transmission from O-DU	Latest Transmission from O-DU
U-Plane	$T1a_max_up \leq (T2a_max_up + T12_min)$	$T1a_min_up \geq (T2a_min_up + T12_max)$
C-Plane	$T1a_max_cp_dl \leq (T2a_max_cp_dl + T12_min)$	$T1a_min_cp_dl \geq (T2a_min_cp_dl + T12_max)$

Table 4.4.3-3: Uplink delay relationships

	Earliest Reception at O-DU	Latest Reception at O-DU
U-Plane	$Ta4min \leq (Ta3min + T34min)$	$Ta4max \geq (Ta3max + T34max)$
C-Plane	$T1a_min_cp_ul \geq (T2a_min_cp_ul + T12_max)$	$T1a_max_cp_ul \leq (T2a_max_cp_ul + T12_min)$

For the fronthaul interface to operate properly, the transmit and receive windows at the O-DU shall be properly aligned. The O-RU window alignment is always based on Ra. For the O-DU to align transmit and receive windows, it needs the following:

- O-RU Delay Characteristics.
- Transport Network Delay Characteristics.

The accuracy of the reported O-RU delay characteristics shall be 200 ns. The 200 ns value was decided as a compromise between a very high accuracy (tens of nanoseconds) which makes buffering in the O-RU easier and a more modest accuracy (~0.5 microseconds) which is relatively easy for a well-controlled Ethernet network to accomplish. This accuracy applies only to the reported start and end times of the reception and transmission windows relative to Ra. Ra is a fixed reference point as defined by the respective air interface standards and supported features. As an example, if the O-RU reports T2a_min_up of 20 μs, the O-RU may begin processing the data from 19.8 μs to 20.2 μs prior to the symbol over the air transmission time (Ra). However, the O-RU shall still transmit over the air with the precision as defined by the air interface standards.

4.4.4 Computed latency methods

4.4.4.1 Introduction

Clause 4.4.4 addresses supported methods for providing the O-DU with the required information to properly align transmit and receive windows. Table 4.4.4.1-1 lists the supported methods.

Table 4.4.4.1-1: O-DU transmit/ receive window alignment methods

Delay Characteristics	Computed Methods	
	Defined Transport	Measured Transport
O-RU Delay	O-RU specific values available to O-DU	O-RU specific values available to O-DU
Transport Delay	Pre-defined Min/ Max	Min – Measured; PDVmax - SLA

The goal of delay management for fronthaul interfaces is to ensure that the transmission and reception windows at the O-DU are properly aligned to support the O-RU and transport network characteristics. With the computed methods, the O-DU shall compute the required transmit and receive windows based on O-RU delay and transport network characteristics. The O-DU shall then determine if the O-DU can accommodate the required transmit and receive windows. If so, it shall set its transmit and receive windows (based on O-DU design) to the most appropriate settings. As an example, an O-DU can have a fixed set of transmit and receive windows. The O-DU may then select the fixed settings that best fit the transport latency (T12, T34) and O-RU delay characteristics for the connected O-RUs.

The O-DU should raise notification if it is not capable of accommodating the required transmit and receive windows.

Based on the downlink delay relationships defined in Table 4.4.3-2, following conditions apply for the placement of the transmit window at the O-DU to support the reception window at the O-RU:

- The O-DU shall transmit DL U-Plane messages:
 - Early enough (i.e. before T1a_min_up) to ensure it is received at O-RU before T2a_min_up
 - Late enough (i.e. after T1a_max_up) to ensure it is received at O-RU after T2a_max_up
- Similarly, the O-DU shall transmit C-Plane messages within their respective windows (T1a_min_cp_dl and T1a_max_cp_dl for downlink C-Plane messages, and T1a_min_cp_ul and T1a_max_cp_ul for uplink C-Plane messages)

Based on the uplink delay relationships defined in Table 4.4.3-3, the following conditions apply for alignment of the receive window at the O-DU to support the transmission window at the O-RU,:

- The O-DU shall set its receive window for UL U-Plane messages:
 - Early enough (i.e. Ta4_min or before) to ensure UL U-Plane messages transmitted by O-RU after Ta3_min is received
 - Late enough (i.e. Ta4_max or after) to ensure UL U-Plane messages transmitted by O-RU before Ta3_max is received

The O-DU constraints for delay management to operate properly are given in Table 4.4.4.1-2:

Table 4.4.4.1-2: O-DU constraints

O-DU	Constraint		
	Parameter	Description	Relationship
Transmit Window (U-Plane)	TXmax _{O-DU}	maximum transmit window needed by an O-DU	$(T1a_max_up - T1a_min_up) \geq TXmax_{O-DU}$
	T1a_max_up _{O-DU}	Maximum T1a_max_up supported by an O-DU	$T1a_max_up_{O-DU} \geq T1a_max_up$
Receive Window	RXmax _{O-DU}	Maximum receive window supported by an O-DU	$RXmax_{O-DU} \geq (Ta4_max - Ta4_min)$
	Ta4_max _{O-DU}	Maximum uplink latency relative to Ra supported by an O-DU	$Ta4_max_{O-DU} \geq Ta4_max$
Transmit Window (C-Plane)	TXma _{XO-DU-C-Plane}	maximum transmit window needed by an O-DU	$(T1a_max_cp_dl - T1a_min_cp_dl) \geq TXma_{XO-DU-C-Plane}$

4.4.4.2 Fronthaul timing domain

With the computed delay approach, O-DU delay parameters apply to a timing domain. A timing domain is defined as the set of O-DU and O-RU Physical Fronthaul Network Interface ports to which the computed O-DU delay parameters apply. The O-DU timing domain can encompass a single O-DU port, or can encompass multiple O-DU ports, where the same O-DU timing parameters apply to all O-DU ports in the domain.

Similarly, each O-DU port in the domain can support C-Plane and U-Plane traffic to 1 or more O-RU ports. Therefore, a fronthaul timing domain can encompass one or more O-DU Physical Fronthaul Interface Ports and one or more O-RU Physical Fronthaul Interface Ports.

The O-DU shall compute transmit and receive windows with the single set of delay parameters to encompass the worst case of Transport and O-RU delay characteristics within the domain. The following definitions are used:

- Transport Network Parameters for timing domain

$T12_min_{ij}$: $T12_min$ between O-DU Physical Fronthaul Interface Port (i) and O-RU Physical Fronthaul Interface Port (j)

$T34_min_{ij}$: $T34_min$ between O-DU Physical Fronthaul Interface Port (i) and O-RU Physical Fronthaul Interface Port (j)

Table 4.4.4.2-1 lists delay parameters for O-RU Port (j) within a timing domain.

NOTE: It is assumed that if a timing domain has multiple $T12_min$ and $T34_min$ values, then the transport delay characteristics are measured. For measured transport delay it is not possible to measure the maximum delay. $T12_max$ and $T34_max$ in this case is computed by adding a pre-defined (e.g. via SLA) worst case variation (PDVmax) to the corresponding transport minimum delay values.

Table 4.4.4.2-1: Radio delay parameters for O-RU Port "j" within timing domain

	Downlink U-Plane	Uplink U-Plane	Downlink C-Plane	Uplink C-Plane
Minimum	$T2a_min_up_j$	$Ta3_min_j$	$T2a_min_cp_dl_j$	$T2a_min_cp_ul_j$
Maximum	$T2a_max_up_j$	$Ta3_max_j$	$T2a_max_cp_dl_j$	$T2a_max_cp_ul_j$
NOTE:	Tcp_adv_dl may be different across RUs within a domain, therefore $T2a_min_cp_dl$ and $T2a_max_cp_dl$ shall be used to determine downlink C-Plane window.			

The resulting parameters to be used for determining the O-DU window for the timing domain are shown in Table 4.4.4.2-2 and Table 4.4.4.2-3.

Table 4.4.4.2-2: Transport network parameters for timing domain

	Downlink	Uplink
Minimum	$T12_min = \text{MIN}(T12_min_{ij})$	$T34_min = \text{MIN}(T34_min_{ij})$
Maximum	$T12_max = \text{MAX}(T12_min_{ij}) + \text{PDVmax}$	$T34_max = \text{MAX}(T34_min_{ij}) + \text{PDVmax}$

Table 4.4.4.2-3: O-RU delay parameters for timing domain

		U-Plane	C-Plane
Downlink	Minimum	$T2a_min_up = \text{MAX}(T2a_min_up_j)$	$T2a_min_cp_dl = \text{MAX}(T2a_min_cp_dl_j)$
	Maximum	$T2a_max_up = \text{MIN}(T2a_max_up_j)$	$T2a_max_cp_dl = \text{MIN}(T2a_max_cp_dl_j)$
Uplink	Minimum	$Ta3_min = \text{MIN}(Ta3_min_j)$	$T2a_min_cp_ul = \text{MAX}(T2a_min_cp_ul_j)$
	Maximum	$Ta3_max = \text{MAX}(Ta3_max_j)$	$T2a_max_cp_ul = \text{MIN}(T2a_max_cp_ul_j)$

The tables defined above are used to derive the corresponding delay parameters to be used in the timing relationships defined in Table 4.4.3-2 and Table 4.4.3-3 to support the computed methods defined below.

With computed delay approach, there are two situations that are considered:

- **Defined Transport Method:** the network delay is pre-defined (usually by the network operator) and the delays are computed based on that definition. The definition generally considers the maximum network latency, with an assumption that a smaller delay can more easily be accommodated.
- **Measured Transport method:** the network delay is estimated based on actual measurements of packet latency on the downlink and uplink (sometimes a symmetric delay is assumed but this is not necessary).

Regardless of whether the Defined Transport Method or Measured Transport Method is used, the calculations are mainly the same. The next two clauses describe in more detail these two methods.

4.4.4.3 Defined Transport method

With the Defined Transport method, the O-DU transmit and receive windows are determined based on pre-defined transport network characteristics, and the delay characteristics of the O-RUs within the timing domain. For this approach to work, the O-DU shall obtain the O-RU delay characteristics as defined in Table 4.4.4.2-3 for each O-RU in the timing domain. The O-DU can adapt its transmit and receive windows to accommodate the O-RU delay characteristics if O-DU has that capability. The adaptation of the O-DU transmission and reception windows can be limited based on the O-DU design. Table 4.4.4.3-1 details the O-DU and O-RU delay profile parameters.

Table 4.4.4.3-1: DL and UL delay profiles

	O-RU delay profile	O-DU delay profile
Downlink	T2a_min_up, T2a_max_up, T2a_min_cp_dl, T2a_max_cp_dl, Tcp_adv_dl	T1a_max_up _{O-DU} , TXmax _{O-DU}
Uplink	Ta3_min, Ta3_max, T2a_min_cp_ul, T2a_max_cp_ul	Ta4_max _{O-DU} , RXmax _{O-DU}

Details on how O-RU parameters are obtained by the O-DU are specified in the fronthaul M-Plane specification [7]. Optionally, an O-RU can adapt its delay profile information, especially for uplink, based on O-DU delay profile and transport delay (T12_min, T34_min). If this option is supported by the O-RU, then the O-DU shall provide its delay profile as well as the T12_min to the O-RU. The exchange of this information is specified in the fronthaul M-Plane specification [7]. With this option, the O-RU receives the O-DU delay profile before sending its delay profile to the O-DU. The O-DU is unaware that the O-RU delay profile information has been modified. It simply uses the delay profile provided by the O-RU as it would without this option.

Delay characteristics for an O-RU can vary based on air interface properties. Air interface properties which may be used as the basis for supporting different delay characteristics are limited to:

- Channel Bandwidth
- SCS

A set of delay characteristics which applies to a combination of the above properties is referred to as a delay profile. For each supported combination of the above properties, the O-RU shall identify a delay profile. Multiple combinations of the above properties can utilize the same O-RU delay profile. These delay profiles are O-RU specific and are not within the scope of the present document.

When calculating the O-DU transmit and receive window for a timing domain, the O-DU shall use the delay profile applicable for each O-RU based on the air interface properties used by the O-RU in the specific network configuration. If any of the O-DU constraints in Table 4.4.4.1-2 is violated, it cannot be assured for the O-DU to adjust transmit and reception windows to meet the required transmit or receive windows.

The O-RU and O-DU can have multiple delay profiles depending on the design as shown in Table 4.4.4.2-3. Using the delay parameters specified for the corresponding profile for each O-RU in the timing domain, the uplink and downlink O-RU parameters to use for the timing domain can be determined as specified in clause 4.4.4.2. The transport delay parameters are constant for this method and specified as part of the network. The delay relationships defined in Table 4.4.3-2 and Table 4.4.3-3 can then be applied to determine the O-DU transmit and receive delay parameters. Finally, the O-DU uplink and downlink constraints as defined in Table 4.4.4.1-2 will need to be ensured, by appropriate fronthaul transport network delay dimensioning and selection of O-RU and O-DU equipment, for the calculated windows to be supported.

4.4.4.4 Measured Transport method

With the Measured Transport method, the O-DU transmit and receive windows are determined based on the delay characteristics of the O-RUs as defined in clause 4.4.4.3, and measured transport delays between all O-DU ports and O-RU ports in the timing domain. The O-DU shall measure the transport delay(s) for all O-RU/O-DU Physical Fronthaul Network Interface port pairs in the timing domain and adapt its transmit and receive windows to accommodate the measured transport delay characteristics.

If the measured transport method is to be used, the O-DU shall use the One-Way Delay Measurement messages as defined in the eCPRI specification [2], clause 3.2.4.6. The message format is shown in Table 4.4.4.4-1. Some of the benefits of using the eCPRI approach are:

- Measurement consistency between the C-Plane and U-Plane is achieved by using a single processing element (see clause 3.1 for the definition of a processing element) and associated C/U-Plane endpoint address for the measurement.
- eCPRI One-Way Delay Measurement allows for varying packet sizes to better simulate real traffic.
- eCPRI approach allows for separate T12 and T34 measurements.

NOTE 1: For this delay measurement, the use of eCPRI One-Way Delay protocol is specified because no other specified transport mechanism supports one-way delay measurements – this is true even if RoE is used as the C-Plane and U-Plane transport mechanism.

Table 4.4.4.4-1: eCPRI One-Way Delay Measurement message

One-Way Delay Measurement (Type 5)									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ecpriVersion				ecpriReserved			ecpriConcatenation	1	Octet 1
ecpriMessage = 5								1	Octet 2
ecpriPayload								2	Octet 3
Measurement ID								1	Octet 5
Action Type								1	Octet 6
TimeStamp (seconds)								6	Octet 7
TimeStamp (nanoseconds)								4	Octet 13
Compensation value (nanoseconds)								8	Octet 17
Dummy bytes								L	Octet 25
									Octet M

The O-DU is the initiator for measuring T12 and optionally T34 according to Figure 4.4.4.4-1 (T12) and Figure 4.4.4.4-2 (T34).

eCPRI Specification [2] defines two methods for measuring the One-Way-Delay, a 1-Step and a 2-Step method. With the 1-Step method the sender of the Request message includes the values of t_1 and t_{CV1} in the request. With the 2-Step method these values are included in a Follow-Up message. Detailed information of this procedure is specified in [2], clause 3.2.4.6.

Usage of either 1-Step or 2-Step method for the T34 measurement is controlled by M-Plane O-RU capability parameters **one-step-t34-supported** and **two-step-t34-supported**. When the O-RU supports Measured Transport Method (eCPRI message 5) the O-RU shall support at least one of the methods for measurement of T34. When the two parameters are reported by the O-RU in the capability report, the O-RU shall support both 1-Step and 2-Step method of the T12 measurement. If the O-DU will initiate a T34 measurement, the O-DU shall support both 1-Step and 2-Step method of the T34 measurement. When the two parameters are not reported by the O-RU in the capability report, but measured transport method is supported by the O-RU, the O-DU may probe the O-RU capabilities by using either 1-Step or 2-Step method when performing a T12 or T34 measurement.

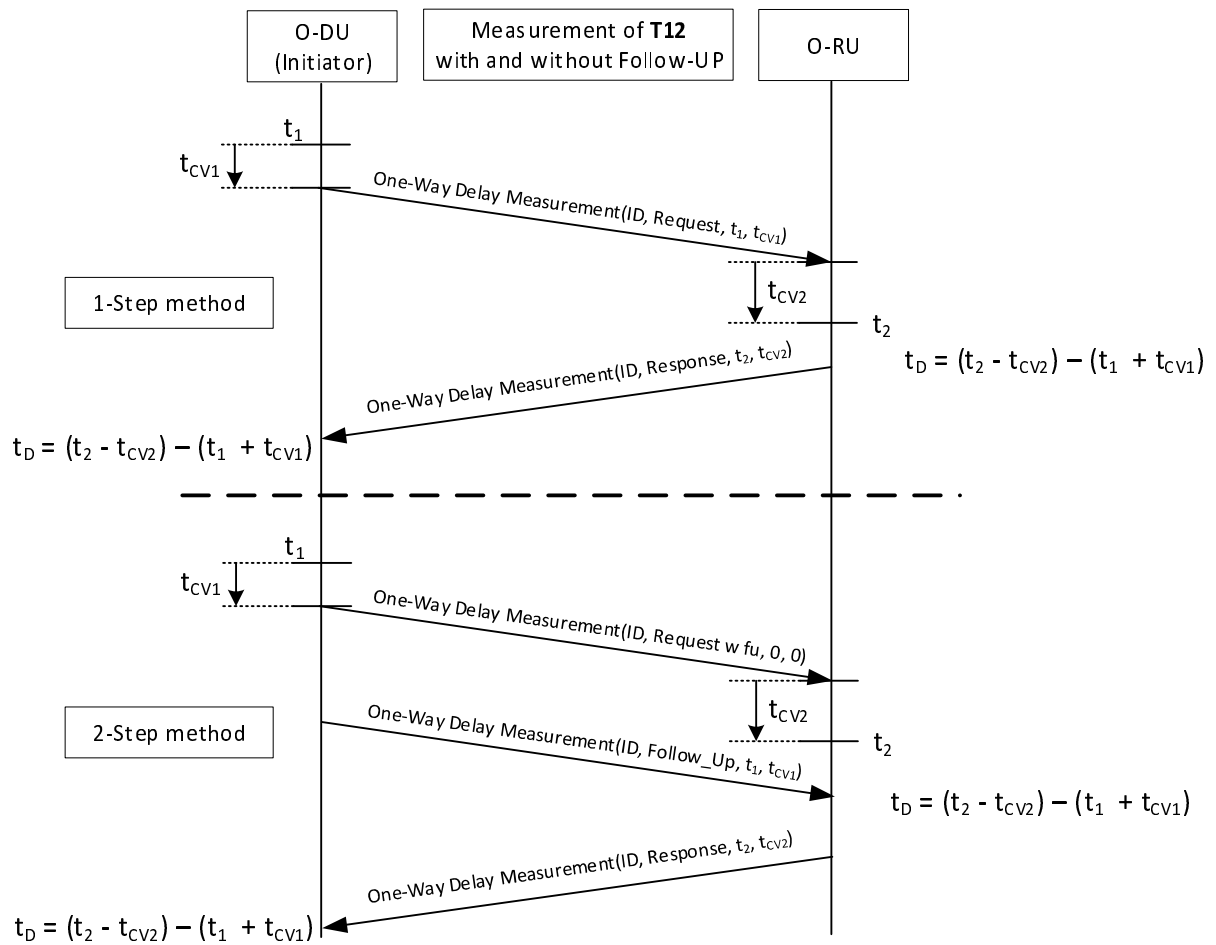


Figure 4.4.4.4-1: eCPRI One-Way Delay T12 Measurement procedure

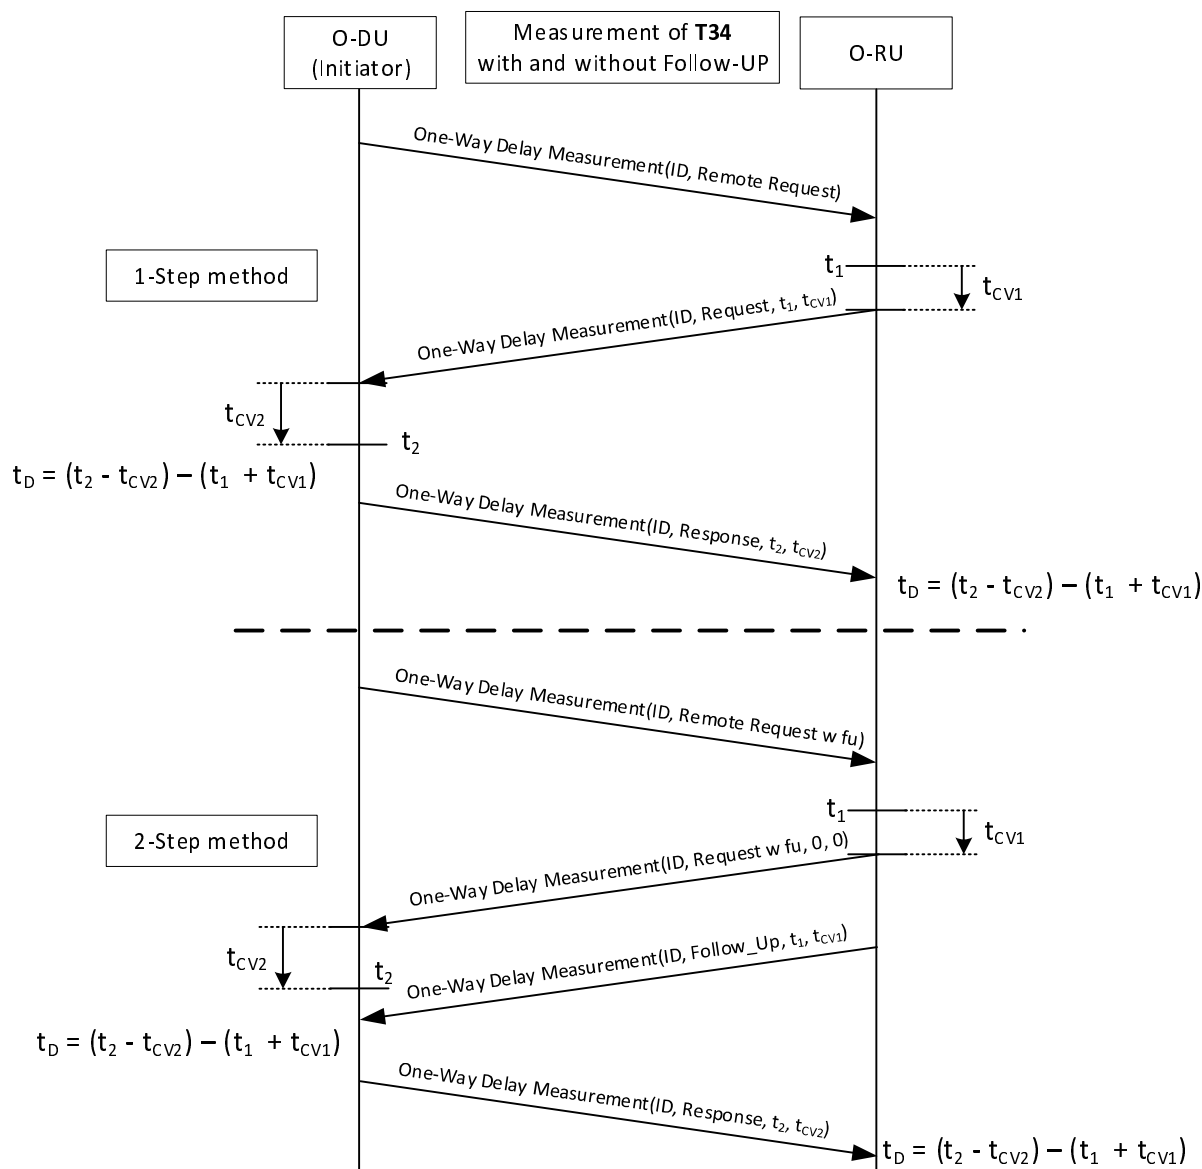


Figure 4.4.4.4-2: eCPRI One-Way Delay T34 Measurement procedure

When measuring the one-way delay, the goal is to determine the minimum transport delay. Since any individual packet can experience different delays through the network due to PDV, it is necessary to perform multiple measurements. Since the desired value is the minimum (T12_min or T34_min), the O-DU should use the minimum delay measurement among the various measurements to estimate the minimum transport delay. The O-DU may use the estimated T12_min value as the T34_min value.

NOTE 2: There is no theoretical upper bound for transport delay. The maximum values are determined based on the measured minimum values plus a maximum allowed transport variation (PDVmax) defined as part of the SLA for the transport network.

The O-DU shall measure the delay corresponding to each O-DU and O-RU port pair within the timing domain. The O-DU shall then use the per port pair minimum delay as per Table 4.4.4.2-1 to provide the transport parameters to be used in the delay relationships specified in Table 4.4.3-2 and Table 4.4.3-3. The resulting delay configuration will meet the O-DU constraints defined in Table 4.4.4.1-2 provided that appropriate fronthaul transport network delay dimensioning and selection of O-RU and O-DU equipment are made.

Since the transport delay is necessary to establish correct timing, when using the measured transport method, the O-DU shall perform this measurement initially before sending of C/U-Plane traffic. The measurement should be performed with packets using the same L2 CoS Priority or L3 DSCP Code as the U-Plane data uses, to ensure the measurement accurately reflects the U-Plane transport latency. The O-DU may also do this periodically to verify delay if desired.

Also, if an error is detected by traffic counts used for error detection defined in clause 4.4, it can be necessary to re-measure the delay and update the O-DU or O-RU timing based on the new measurements.

4.4.5 Latency categories for O-DU with dynamic timing advance

As an aide to selection of equipment for use in specific transport network use cases, O-DU and O-RU are categorized based on delay capabilities. Equipment is categorized with a Category [AAAA-ZZ] and Sub-category [.00-.1000]. The category can be used to determine the maximum T_{12max} and T_{34max} which the equipment can support. Categories are determined as follows:

- $T_{1a_max_{O-DU}} \geq (T_{12max} + T_{2a_min_up})$
- $T_{a4_max_{O-DU}} \geq (T_{34max} + T_{a3_max_up})$

By definition $T_{1a_max_up_{O-DU}}$ is the earliest that the O-DU can begin transmission. This results in that the earliest that the O-DU can be guaranteed to complete transmission is $T_{1a_max_up_{O-DU}} - TX_{max_{O-DU}}$. The result is that:

- $(T_{1a_max_up_{O-DU}} - TX_{max_{O-DU}} - T_{2a_min_up}) \geq T_{12max}$

Similarly, for uplink $T_{a4_max_up_{O-DU}}$ is the latest that the O-DU may receive any packet from the O-RU. Therefore:

- $(T_{a4_max_up_{O-DU}} - T_{a3_max_up}) \geq T_{34max}$

Both of these are illustrated in the **Figure 4.4.5-1** below:




Figure 4.4.5-1: Uplink and downlink timing parameter relationship

Using these inequalities, O-DU and O-RU can be assigned categories based on the constraints shown in Table 4.4.5-1:

Table 4.4.5-1: Maximum transport latency constraints

Parameters	Description	O-DU	Constraint O-RU Processing	Transport
Downlink	Max Transport	$T_{1a_max_up_{O-DU}} - TX_{max_{O-DU}}$	$T_{2a_min_up}$	$T_{12_max} = T_{1a_max_up_{O-DU}} - TX_{max_{O-DU}} - T_{2a_min_up}$
Uplink	Max Transport	$T_{a4_max_up_{O-DU}}$	$T_{a3_max_up}$	$T_{34_max} = T_{a4_max_up_{O-DU}} - T_{a3_max_up}$

The maximum range that can be supported is based on the minimum value of T_{12_max} and T_{34_max} from above. This requires separate categorization for UL and DL for each O-DU and O-RU. Therefore, an uplink category pair (AX) and downlink category pair (BY) should be evaluated for each pairing of O-DU and O-RU. The minimum distance (T_{12_max} , T_{34_max}) represented by the corresponding pairs (AX, BY) then identifies the maximum transport latency which can be supported by the pair.

Each endpoint has a specific value for the characteristics defined in Table 4.4.5-1. In order to categorize endpoints, the endpoints are grouped based on the range in which the maximum transport characteristic falls. The category identification then means that the maximum transport characteristic for the endpoint falls somewhere between the maximum and minimum values for the category.

Since the category identifies that the maximum latency supported for the endpoint falls somewhere within the category range, the resulting category pair (AX) defines a range of $T12_max$ and $T34_max$ for the combination. The combination can be used to identify a maximum $T12_max$ and $T34_max$, and minimum $T12_min$ and $T34_min$ that the combination can support. The category pair identifies a best and worst case for the maximum range the combination can support. The actual maximum range supported by a specific O-DU and O-RU pair with a given category combination will fall somewhere within the range. That is, for a given O-DU with category A and O-RU with category X, the $T12_max$ and $T34_max$ supported by the pair is defined as:

- $Latency_min_{AX} \leq (T12_max = T34_max) \leq Latency_max_{AX}$

Additionally, the paired combination can be mapped to supported use cases. Annex B.2 contains the category definitions, $Latency_min$, $Latency_max$, and use case tables.

The categories are only useful for identifying the maximum possible range ($T12_max$, $T34_max$) values. To be able to usefully identify equipment for a given use-case, it is also necessary to be able to identify $T12_min$ and $T34_min$ values. This can be defined in terms of variation, where variation indicates how much lower $T12_min$ and $T34_min$ can be relative to $T12_max$ and $T34_max$ respectively. That is:

- $DL\ Variation = T12_max - T12_min$
- $UL\ Variation = T34_max - T34_min$

An additional level of categorization is required to address the transport delay variation (e.g. $T12_{max} - T12_{min}$) that a given equipment type pair can support. This is addressed using sub-categories. The total variation that can be supported is a combination of the receive window at the receiver, and the transmit variation. However, the receive window will likely be much larger than the transmit variation. Each endpoint is assigned to both an uplink and downlink category. The O-DU DL sub-category is based on transmit variation, while the O-DU UL sub-category is based on receive window. For UL, the roles are reversed, and the O-RU DL sub-category is based on receive window, and the UL sub-category is based on transmit variation.

As a result, sub-categories are defined based on the role for the endpoint in the given transmission direction (UL, DL).

Downlink variation ($T_{variation_DL}$) is defined in Figure 4.4.5-2 as follows:

$$T_{variation_DL} \leq (T2a_max_up - T2a_min_up - TXmax_{O-DU})$$

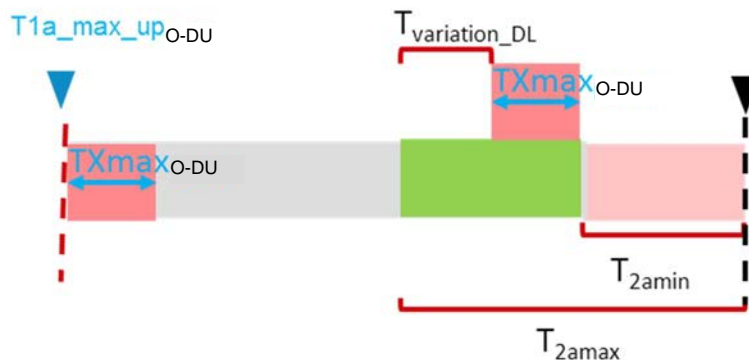


Figure 4.4.5-2: Downlink transport variation

Uplink variation ($T_{variation_UL}$) is defined in Figure 4.4.5-3 as follows:

$$T_{variation_UL} \leq (RXmax_{O-DU} - (Ta3_max - Ta3_min))$$




Figure 4.4.5-3: Uplink transport variation

Using the respective constraints, the variability can be defined based on O-DU and O-RU delay characteristics as shown in Table 4.4.5-2:

Table 4.4.5-2: Maximum transport latency variation constraints

Parameters	Description	O-DU	Constraint O-RU Processing	Transport
Downlink	Max Variation	$T_{x\max_{O-DU}}$	$T_{2a_max_up} - T_{2a_min_up}$	$T_{12_max} - T_{12_min} = (T_{2a_max_up} - T_{2a_min_up}) - T_{x\max_{O-DU}}$
Uplink	Max Variation	$RX_{\max_{O-DU}}$	T_{a3_max}	$T_{34_max} - T_{34_min} = RX_{\max_{O-DU}} - (T_{a3_max_up} - T_{a3_min_up})$

Sub-categories for each direction are defined using the endpoint characteristics. Similar to categories, sub-categories indicate that the endpoint has a variation characteristic that falls somewhere in the specified sub-category range. Since transmit variability and receive windows are expected to have different ranges, separate tables are created. However, since sub-categories are numeric [#####], to simplify these values are assigned in a similar way. The ##### value represents the time in 10's of μs for the given endpoint property.

For example, if an O-RU with a category X has a $T_{2a_min} - T_{2a_max} = 163 \mu\text{s}$, the sub-category for the O-RU is [.16]. Making the full category for the O-RU X.16 (the "." is not a decimal point but a delimiter only). By using this approach, the receiver sub-category and transmitter sub-category can be directly used to determine the dynamic range in either direction:

TX sub-category: .06

RX sub-category: .16

Resulting $T_{\text{variability}} = 16 \times 10 - 6 \times 10 = 100 \mu\text{s}$. This provides a lower bound on the maximum dynamic range supported in one direction. The minimum of the UL and DL $T_{\text{variability}}$ can be used to determine the maximum dynamic range for the combination (e.g. maximum difference between T_{12_max} and T_{12_min} that can be supported).

Annex B.3 contains the sub-category definitions and resulting combination dynamic range use case tables.

With dynamic timing advance, since the O-DU can set its earliest transmission time based on the measured T_{12_max} , the full range of variability is available regardless of the T_{12_max} . For example, if an O-DU/ O-RU combination has 100 μsec worth of variability, and a maximum range T_{12_max} of 250 μsec , the combination can support the measured $T_{12_max} (\leq 250 \mu\text{sec})$ and T_{12_min} of $T_{12_max} - 100 \mu\text{sec}$. So, if the measured T_{12_max} is 200 μsec , the combination can support T_{12_max} of 200 μsec and T_{12_min} of 100 μsec .

4.4.6 Latency categories for O-DU with fixed timing advance

The same category concepts apply for O-DU which support fixed transmit and receive windows. The earliest transmit window and latest receive window timing defines the maximum range that the O-DU can support. $T_{1a_max_up_{O-DU}}$ for the O-DU is equal to the $T_{1a_max_up}$ of the earliest fixed transmit window. Similarly, $T_{a4_max_{O-DU}}$ is equal to the T_{a4_max} of the latest receive window. The O-DU category is then assigned based on this value.

Fixed transmit and receive windows do however impact variability. The fixed transmit time results in sending packets earlier than absolutely needed to achieve the actual T_{12_max} (i.e. the T_{12_max} value that is either measured or configured by the service provider). Figure 4.4.6-1 illustrates this concept:




Figure 4.4.6-1: Fixed transmit time illustration

Since the TX_{DL_Start} is earlier than required for T_{12_max} by T_{1a_offset} , the latest packets arrive at least T_{1a_offset} before $T_{2a_min_up}$. This results in T_{1a_offset} worth of O-RU receive buffer being always used, implying that the O-RU needs additional buffer to compensate T_{1a_offset} . Since the TX_{DL_Start} is fixed, the T_{12min} is also fixed at:

$$T_{12min} = TX_{DL_Start} - (T_{2a_max_up} - TX_{max_O_DU})$$

Since T_{12min} is unaffected, the resulting $T_{variation}$ is reduced by T_{1a_offset} . To support T_{12min} , O-RU need to provide $T_{1a_offset} + T_{12max} - T_{12min} + TX_{max_O_DU}$

When $T_{12_max} + T_{2a_min_up} = T_{1a_max_up_O_DU}$ (e.g. at maximum range), then $T_{1a_offset} = 0$. But as the actual T_{12_max} decreases, the amount of receive buffer consumed by T_{1a_offset} increases by the corresponding amount.

The same applies for UL. The result is that the T_{34_min} is defined relative to the T_{34_max} for the combination, not relative to the measured T_{34_max} .

4.4.7 Non Delay-managed U-Plane traffic

While most U-Plane data shall meet the delay constraints described above, there can be certain types of U-Plane traffic for which the reception windows may not be applicable. Not applying the transmission and reception window constraints to such data allows the transmission of the data over the fronthaul interface to be spread to reduce the peak bandwidth required by the interface.

Figure 4.4.7-1 uses SRS to illustrate an example non delay-managed U-Plane traffic.




Figure 4.4.7-1: SRS as non delay-managed U-Plane traffic

In the example, a large amount of SRS data is received during a special slot on the air interface. By allowing the SRS data to be transmitted from the O-RU to the O-DU outside of the normal transmit window, and instead allowing it to use otherwise un-used fronthaul interface bandwidth, the required peak fronthaul bandwidth requirement can be reduced.

Non delay-managed U-Plane traffic flows shall be uniquely identified from delay managed U-Plane traffic flows. The identification of U-Plane traffic flows using eAxC identifiers is described in clause 5.4, and the eAxC shall be used to indicate delay-managed versus non delay-managed traffic (e.g. some eAxC or Pcid-identified packets may be used to carry delay-managed U-Plane traffic while other eAxC or Pcid-identified packets may be used to carry non delay-managed U-Plane traffic). Non delay-managed U-Plane traffic flows shall not use the same eAxC identifiers as delay managed U-Plane traffic flows. Different RU_Port_IDs shall be allocated within the eAxC_ID for non delay-managed U-Plane flows from those used for delay managed U-Plane traffic flows in the same direction (UL, DL). RU_Port_IDs used in UL for non delay-managed traffic need not be the same as those used in the DL.

The sender may send non delay-managed traffic without regard to the transmission window for the corresponding air interface symbol. Similarly, the receiver need not discard received non delay-managed traffic received outside the normal reception window for the corresponding air interface symbol. However, the air interface symbol time shall still be contained within the U-Plane packets header for reference.

NOTE: The status of traffic flows being delay-managed versus non delay-managed (which is managed within the application) is independent of any transport prioritization that may be configured.

4.5 Reception window monitoring

The reception window can be used for the purpose of error detection. The reception windows at O-RU and O-DU absorb delay variation in the fronthaul. Although its maximum size may vary depending on implementation, the position of windows at O-DU may be adjustable via parametrization as described in clause 4.3.

Data flow error detection during fronthaul operation may be monitored and reported via counters based on receipt of packets relative to the reception window. The receiving node monitors the actual reception timing and generates counters regarding the status and whether packets were successfully received in proper timing.

The following reception window performance counters shall be supported:

- Data was received on time (within reception window) – no action needed
- Data was received too early – link latency re-measurement may be needed

- Data was received too late – link latency re-measurement may be needed

The above counters and other performance monitoring counters may be implemented as defined in clause 9.1. Based on those counts the system may adjust its operation. The specific reaction will be implementation-specific and depend on such considerations as the number of instances needed before justifying sending an alarm, the ability of the system to measure link latencies, or other design-dependent judgments. Additional details regarding the counts, alarms, and fault handling are outside of the scope of the present document.

Dynamic adaptation of windows at Tx side to handle larger delay values may be employed. For instance, fronthaul distance may be extended by handling a portion (or most) of the fixed delay component at Tx buffer side by adapting the transmission timing based on information provided by RX side. In other words, rather than accommodating a latency range of zero to "max", it is possible to instead accommodate a longer latency of "offset" to "max+offset" (longer latency but same range). This is described in detail in clause 4.4.

4.6 Transmission windows

4.6.1 Normal transmission

O-RU and O-DU send C-Plane and U-Plane messages at unspecified time within appropriate transmission window (see clause 4.4.2 for general description of transmission windows). This default O-RU transmission behavior in uplink may be altered by configuration over M-Plane or with Section Extension 18 in C-Plane message as described in clauses 4.6.2, 4.6.3 and 4.6.4.

4.6.2 Uniformly distributed transmission

If O-RU supports uniformly distributed transmission feature, then O-RU may be configured (over M-Plane or with Section Extension 18 in C-Plane message) to uniformly distributed transmission of UL U-Plane messages over the transmission window (**uniformly-distributed-transmission-supported** in M-Plane).

The definition of "uniformly distributed" shall be that the total number of packets should as evenly as possible be spread out in time over the transmission window when they are sent on the interface.

The transmission window shall be divided into N sub-windows where N is the total number of packets to be transmitted during the transmission window. The N packets shall be sent one packet per sub-window. The transmission shall start at a random time within the sub-window.

The above description of when packets are sent can be disturbed by other traffic being transmitted by the O-RU. If that is the case the main object with 'uniform distribution' is that the traffic load shall be evenly distributed in time as much as possible.

In many cases transmission windows will be overlapping between different user data flows (different eAxCs or different section IDs etc.) When this is the case the same "rule" as described above is still valid, that is within the specific sub-window, the data packet should be sent at a random time. If several packets (due to different flows) are to be sent within the same sub-window, each packet shall be sent at a random time. If the transmission time for the different packets collide then the packets are buffered in any order and sent when link capacity is available.

4.6.3 Ordered transmission

If O-RU supports ordered transmission, the O-RU may be configured (over M-Plane) to ordered transmission of UL U-Plane messages within transmission window (**ordered-transmission-supported** in M-Plane). If O-RU is configured to order transmission of a set of eAxC then O-RU reports how eAxC are grouped into ordering groups (G) and relative order of eAxCs within each group (R) (**transmission-order and transmission-order-group** in M-Plane).

Assuming

A and B be eAxCs configured to order transmission,

G_A and G_B be ordering groups of eAxC A and B respectively,

R_A and R_B be relative ordering values of eAxC A and B respectively.

M_A and M_B be messages of eAxC A and B respectively that carry data related to the same symbol (the symbol is identified by the same frame number and the same symbol number within the frame and the same duration as derived from SCS) and transmitted in transmission windows W_A and W_B respectively.

If A and B belong to the same ordering group ($G_A = G_B$), and relative order of A is less than relative order of B ($R_A < R_B$), and start of transmission window W_A is same as start of transmission window W_B , and end of transmission window W_A is same as end of transmission window W_B , then O-RU transmits message M_A before message M_B . Otherwise, O-RU transmits message M_A and M_B in unspecified order. Therefore, order of U-Plane messages of eAxCs of different ordering groups and eAxCs of the same relative order (within a group) is unspecified. Messages transmitted in different transmission windows are ordered only as required to transmit them within transmission windows.

4.6.4 Scheduled transmission

If O-RU supports scheduled transmission, the O-RU may be configured (over M-Plane or with Section Extension 18 in C-Plane message) to shift (delay) and/or resize UL U-Plane messages transmission window.

O-RU may be commanded to shift the transmission window by given offset (delay) expressed as number of symbols. Depending on the O-RU capability the eAxC specific window offset (O) and window size (S) may be provided over:

- 1) M-Plane (**transmission-window-offset** and **transmission-window-size** in M-Plane) or
- 2) C-Plane (in Section Extension 18).

The window offset (O) and size (S) are specific to each eAxC and depending on O-RU the offset and size may be:

- 1) specific to all REs in a symbol within eAxC
This is applicable to O and S values provided over M-Plane and C-Plane.
- 2) specific to REs selected by section descriptions in a symbol within an eAxC.
This is applicable only to O and S provided in C-Plane. In this case, the same values of O and S shall be used for every Section Extension 18 description referring to REs in the same PRB using different values of reMask.

A C-Plane message can address multiple eAxCs if Section Extension 7 is used or mapping between of an eAxC_ID to multiple eAxCs is configured with M-Plane.

The time at which the O-RU UL U-Plane message is transmitted depends on whether the eAxC is configured for delay managed or non delay-managed traffic:

- If eAxC is configured for delay managed traffic and is also configured (via C-Plane or M-Plane) to apply transmission window offset (O) and window size (S), then the O-RU shall transmit the UL U-Plane message with data specific to REs in symbol N (N is the number of the symbol within a frame), in a window that starts at the same time as UL U-Plane normal transmission window corresponding to symbol N+O and has a duration as configured by window size setting. O-RU reports per endpoint max number of symbols and max number of PRBs that can be buffered (**max-buffered-prbs** and **max-buffered-symbols** in M-Plane). The O-DU shall not request from the O-RU UL a reception and transmission window configuration that exceeds O-RU buffering capacity. For calculation of buffering load, O-DU shall assume that the O-RU allocates its buffer at the time reference for Ta3_min (typically start of symbol) and the buffer is released at the end of transmission window.
- If eAxC is configured for non delay-managed traffic and is also configured (via C-Plane or M-Plane) to apply transmission window offset (O) and window size (S), then the O-RU shall transmit the UL U-Plane message with data specific REs in to symbol N (N is the number of the symbol within a frame) in a window that starts after of UL U-Plane normal transmission window corresponding to symbol N+O. In this case only the start of transmission window is defined – the end of the transmission window is unspecified. O-RU may report per endpoint max number of symbols and max number of PRBs that can be buffered (**max-buffered-prbs** and **max-buffered-symbols** in M-Plane). O-DU can request from O-RU UL reception and transmission window configuration that exceeds O-RU buffering capacity and O-RU shall handle it on best effort basis. For calculation of buffering load, O-DU shall assume that O-RU allocates buffer at the time reference for Ta3_min (typically start of symbol) and the buffer is not released before at the start of transmission window. U-Plane message transmission of eAxC configured for non delay-managed traffic shall be handled by O-RU on best effort basis and transmission is not guaranteed.

If O-RU does not report, over M-Plane, support for independent U-Plane transmission window control (feature INDEPENDENT-TRANSMISSION-WINDOW-CONTROL in M-Plane), then all endpoints with transmission window control enabled, handling the same carrier type and SCS, shall be configured (via M-Plane or C-Plane) with parameter values resulting in transmission windows that coincide between the endpoints. That is, for every symbol N, the effective transmission window for symbol N shall start at the same time and end at the same time for all the endpoints handling same carrier type and SCS. This restriction applies only to endpoints that have transmission control enabled over M-Plane.

4.7 O-RU external antenna delay handling

4.7.1 Introduction

The latency model of clause 4.4 assumes that the antenna delay is negligible compared to the O-RU's internal delay or is known by the O-RU. Figure 4.4.1-1 defines the $T2a$ and $Ta3$ values as the total processing time between Fronthaul interface input/output ports (R2 and R3) and the antenna interface point (Ra).

Figure 4.7-1 defines the reference points for delay management where the O-RU external antenna delay parameters are defined. $T2a$ and $Ta3$ will have a new definition when external antenna delays are introduced. $T2a$ is the O-RU internal delay between R2 and Rd and $Ta3$ is the O-RU internal delay between Ru and R3.




Figure 4.7-1: Definition of reference points including external antenna delay

Table 4.7.1-1: eCPRI O-RU additional delay latency parameters

	Parameter	Latency	Minimum	Maximum
Downlink	Tda	Timing difference between the output of DL signal at the antenna connector (Rd) of O-RU and the transmission over the air (Ra).	-	-
Uplink	Tau	Timing difference between the reception over the air (Ra) and the input of UL signal at the antenna connector (Ru) of O-RU.	-	-

NOTE: The latency model of clause 4.4 does not allow to adjust timing to compensate transmission delays between antenna ports of the O-RU and the external antenna. To handle external delays between O-RU and antenna reference point, a method is specified in clause 4.7.2 which has minimal impact on the O-DU implementation

4.7.2 Minimal O-DU impact method (defined transport method)

4.7.2.1 Description

When using Defined Transport Method as described in clause 4.4.4.3 and non-negligible external delays between the O-RU and antenna are present, this method will have a minimal impact on the O-DU implementation. O-DU shall configure the O-RU with Tda and Tau via the M-Plane [7]. Instead of taking Tda and Tau under consideration when calculating the transmission and reception windows on the O-DU, the external delay is taken from the fronthaul delay ($T12$, $T34$) timing budget.

When using this method, the values configured to O-RU for **t12-max** and **t34-max** shall maintain the same value as if there is no external delay, i.e. $\tau = T_{da} = 0$ (zero). It will be the responsibility of the operator to ensure that the actual fronthaul latency between R1 and R2 respectively R3 and R4 shall not exceed $(\mathbf{t12-max} - T_{da})$ and $(\mathbf{t34-max} - \tau)$.

This method will keep the transmission and reception windows on the O-DU at the same positions in time with or without the external delays. The O-RU shall transmit the downlink signal at reference point Rd at $(t_{DL}=0 - T_{da})$. The O-RU shall start the processing of the uplink signal received at the reference point Ru at $(t_{UL}=0 + \tau)$.

For downlink C-Plane and U-Plane, the O-RU shall shift its reception windows T_{da} earlier in time. For uplink C-Plane the reception window shall not be shifted. The transmission window for uplink U-Plane data will be shifted by τ on the O-RU but since the external delay is taken from the fronthaul latency budget the uplink U-Plane messages arrive within the reception window on the O-DU. The sizes for reception and transmission windows remain unchanged on the O-DU and on the O-RU.

NOTE: In comparison with latency model of clause 4.4, this method can lead to changes in the O-DU's implementation of its O-RU M-Plane configuration flows.

4.7.2.2 Minimal O-DU impact method - Example

This clause shows an example of the timing relations for downlink and uplink when using the method described in clause 4.7.2.1.

Table 4.7.2.2-1 contains values for all relevant parameters used in this example.

Table 4.7.2.2-1: Parameters for example

Parameter	Latency ($T_{da} = \tau = 0$) [μs]	Latency ($T_{da} = \tau = 25$) [μs]
Physical T12_max, T34_max (NOTE 1)	100	75
T12_max, T34_max	100	
T12_min, T34_min	0	
T2a_min_up_dl	100	
T2a_max_up_dl	300	
T1a_min_up_dl	= $T2a_min_up_dl + T12_max = 100 + 100 = 200$	
T1a_max_up_dl	= $T2a_max_up_dl = 300$	
T2a_min_cp_dl	250	
T2a_max_cp_dl	600	
T1a_min_cp_dl	= $T2a_min_cp_dl + T12_max = 250 + 100 = 350$	
T1a_max_cp_dl	= $T2a_max_cp_dl = 600$	
T2a_min_cp_ul	100	
T2a_max_cp_ul	500	
T1a_min_cp_ul	= $T2a_min_cp_ul + T12_max = 100 + 100 = 200$	
T1a_max_cp_ul	= $T2a_max_cp_ul = 500$	
Ta3_min_up	50	
Ta3_max_up	150	
Ta4_min_up	= $Ta3_min_up = 50$	
Ta4_max_up	= $Ta3_max_up + T34_max = 150 + 100 = 250$	
NOTE:	The configured T12_max and T34_max values are 100 μs . Due to the external delay of 25 μs the actual physical T12_max and T34_max are decreased by 25 μs according to the method description. The physical max fronthaul latency is thus 75 μs .	

Below are figures similar to Figure 4.4.3-1 and Figure 4.4.3-2, showing the timing relations both with and without the usage of the "Minimal O-DU Impact"-method. Figure 4.7.2.2-1 and Figure 4.7.2.2-3 assumes zero external delay both for downlink and uplink, the purpose of these two figures are to show that the relationships stated in Table 4.4.3-2 and Table 4.4.3-3: are still valid when applying the method in clause 4.7.2.1, i.e. the transmission and reception windows are unchanged on the O-DU.

The blue lines and arrows of the figures apply to U-Plane messages, and the green lines and arrows apply to C-Plane messages.

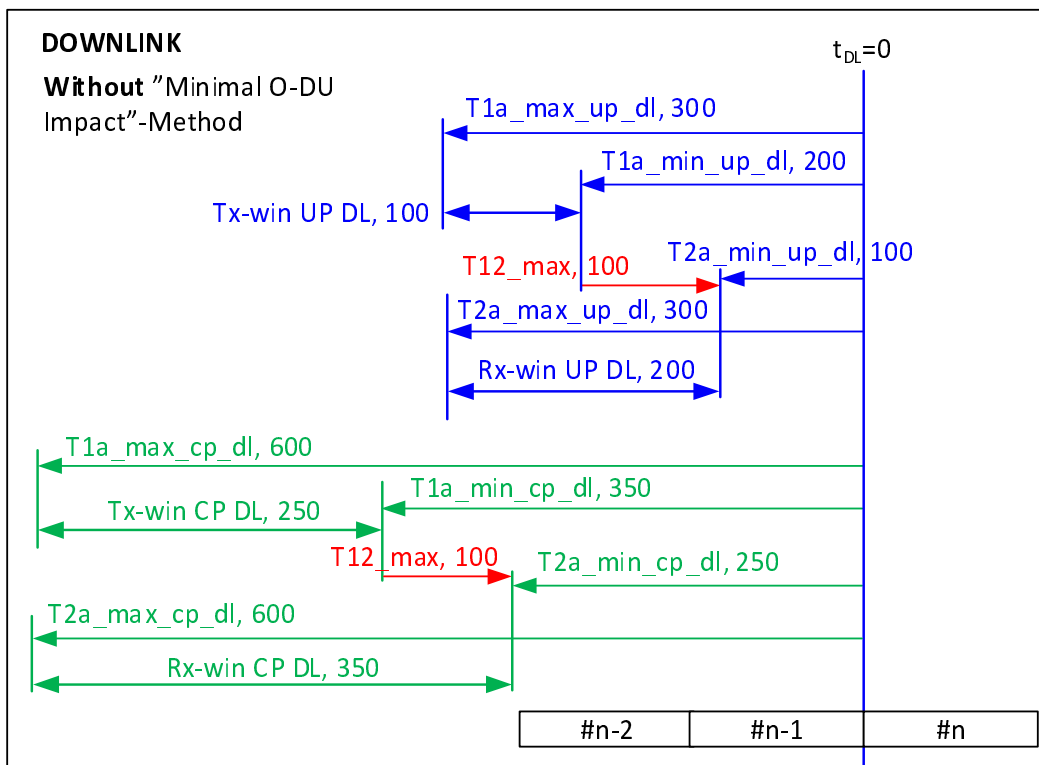


Figure 4.7.2.2-1: Downlink without using minimal O-DU impact method

Figure 4.7.2.2-1 shows the timing relations for downlink without any external delays. T12_max in this example is 100 μs, both as configured value and as the actual physical maximum fronthaul latency.

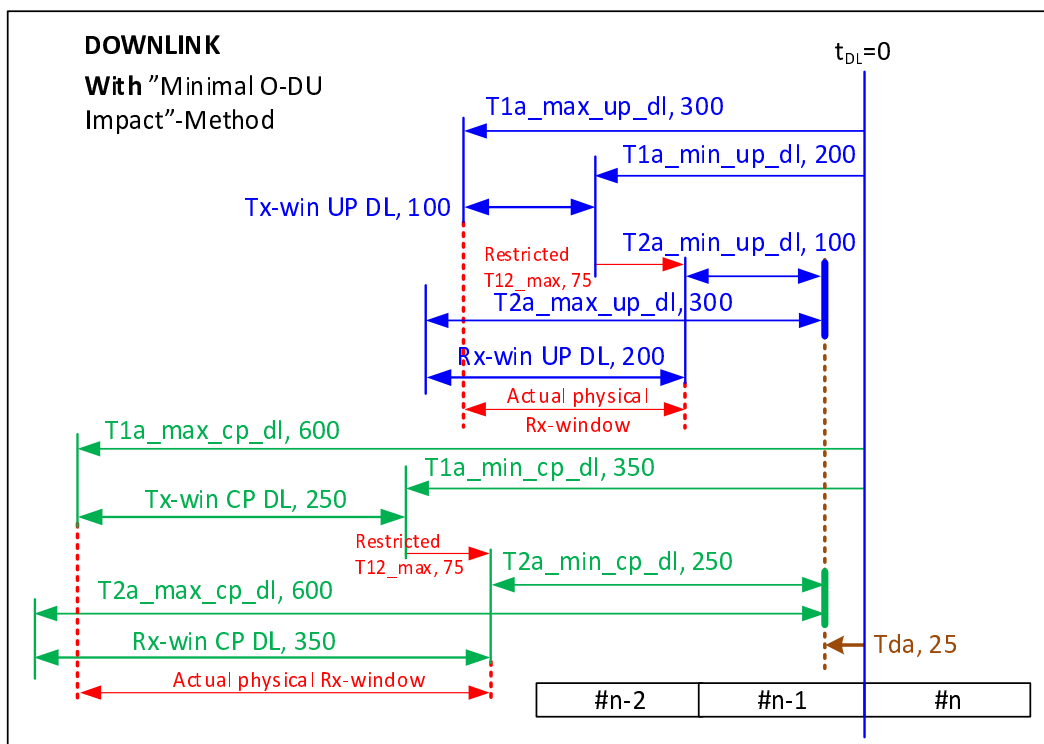


Figure 4.7.2.2-2: Downlink with usage of minimal O-DU impact method

Figure 4.7.2.2-2 shows the timing relations for downlink when an external delay of 25 μs is present and the method described in clause 4.7.2.1 is used. As shown, the transmission windows on the O-DU are the same as the ones in Figure 4.7.2.2-1. The shifting of the reception windows on the O-RU is shown, i.e. these windows start T_{da} μs earlier compared to when not using the method in clause 4.7.2.1, the sizes of the reception windows is the same. The red arrows and text show the actual reception windows on the O-RU when restricting the actual physical fronthaul latency from 100 μs to 75 μs but still configuring 100 μs to the system. As can be seen, there now is a part at the start of the reception windows that should be "un-used". While it is not expected by the O-RU to receive C-Plane or U-Plane messages during these periods the O-RU will process these messages.

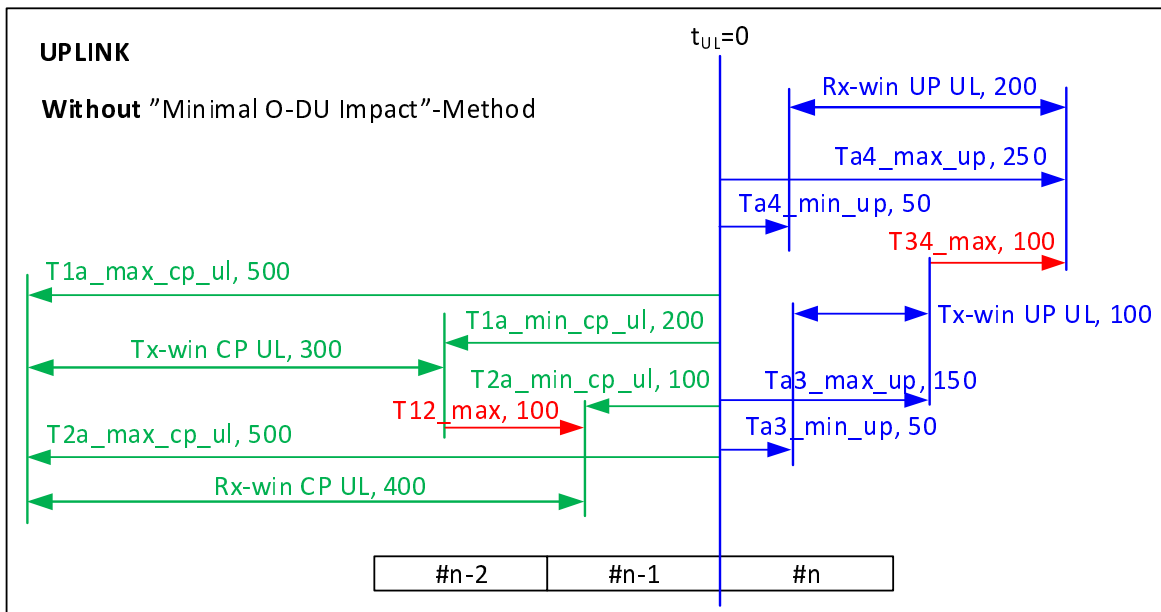


Figure 4.7.2.2-3: Uplink without using 4.7.2

Figure 4.7.2.2-3 shows the timing relations for uplink without any external delays. T_{12_max} and T_{34_max} in this example are 100 μs , both as configured value and as the actual physical maximum fronthaul latency.

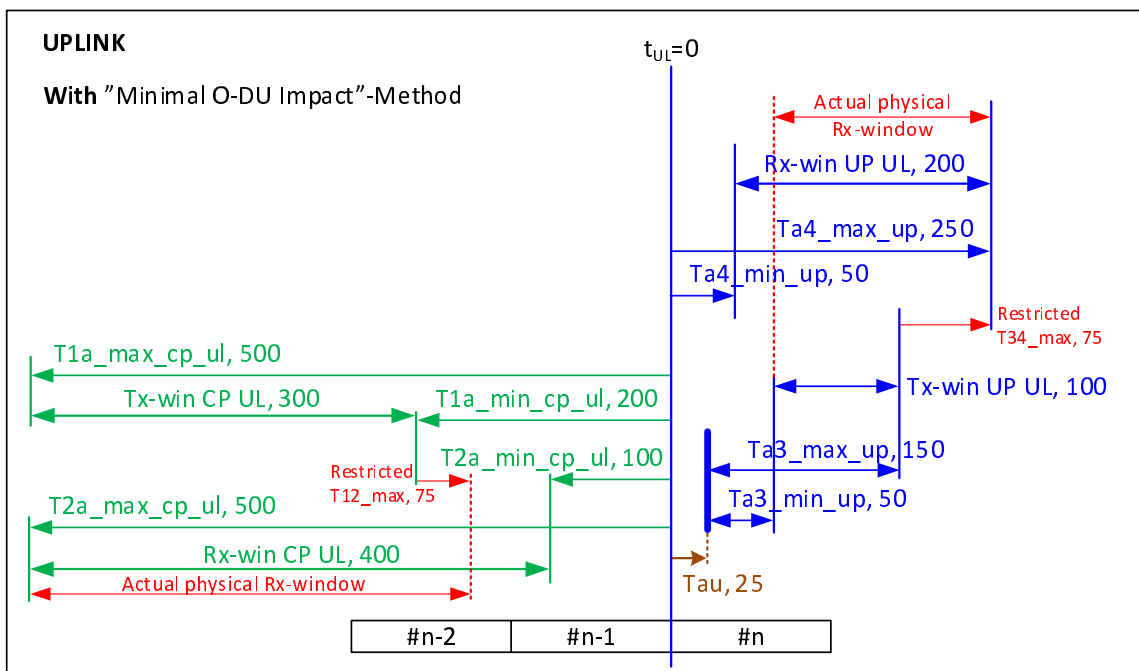


Figure 4.7.2.2-4: Uplink with usage of 4.7.2

Figure 4.7.2.2-4 shows the timing relations for uplink when an external delay of 25 μs is present and the method described in clause 4.7.2.1 is used. As shown, the transmission window on the O-DU are the same as the ones in Figure 4.7.2.2-3. The reception window on the O-RU for C-Plane messages is not shifted for uplink traffic. The size of the reception window is the same. The red arrows and text show the actual reception windows on the O-RU and on the O-DU when restricting the actual physical fronthaul latency from 100 μs to 75 μs but still configuring 100 μs to the system. As can be seen there is a part at the end of the reception window on the O-RU that will be "un-used". While it is not expected by the O-RU to receive C-Plane messages during this period the O-RU will process these messages. On the O-DU there will be a part at the start of the reception window that will be "un-used". While it is not expected by the O-DU to receive U-Plane messages during this period the O-DU will process these messages. The transmission window on the O-RU for the U-Plane messages are shifted τ later in time due to the external latency between reference points Ra and Ru.

4.8 UL frame timing to DL frame

4.8.1 Overview

UL radio frame has a timing advance offset to DL radio frame with the same frame number at UE antenna port. $N_{TA\text{ offset}}$ is a default timing offset parameter applied to UE, unless another value is configured in the cell, which is specified in 3GPP TS 36.211 [42] clause 8.1 and 38.211 [4] clause 4.3.1. **n-ta-offset** is a M-Plane parameter used to control the timing offset of UL radio frame with respect to the DL radio frame with the same frame number at O-RU antenna. Timing relation between UL radio frame and DL radio frame which is configured by **n-ta-offset** is described in [Figure 4.8.1-1](#).

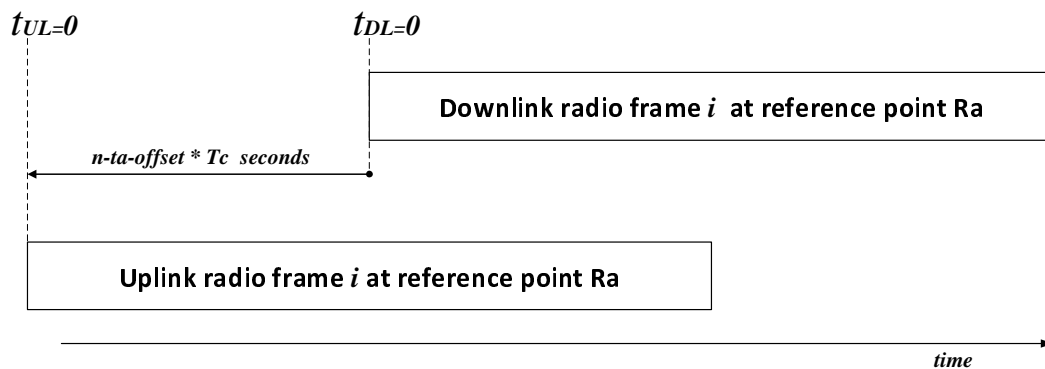


Figure 4.8.1-1 : Uplink-Downlink Timing Relation

Optionally, **t-da-offset** and **t-au-offset** M-Plane parameters control radio timing on each direction at O-RU's antenna connector for external-antenna-delay-handling O-RU and contribute to the timing offset of UL frame to DL frame. Refer to clause [4.7.1](#) for those parameters. Timing relation configured by **t-da-offset** and **t-au-offset** on top of **n-ta-offset** is described in [Figure 4.8.1-2](#).

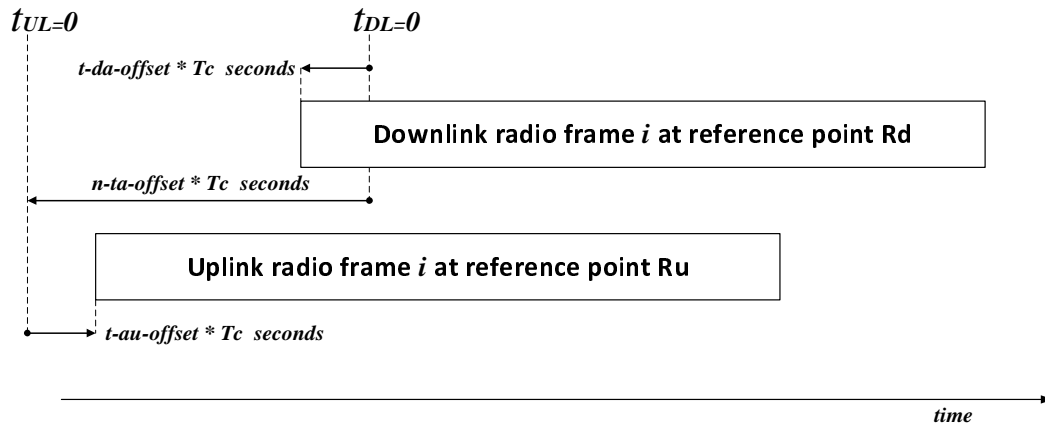


Figure 4.8.1-2 : Uplink-Downlink Timing Relation for O-RU handling external antenna delay

4.8.2 Carrier types and default values of n-ta-offset

Default **n-ta-offset** values are described in [Table 4.8.2-1](#) according to carrier type. This value is configurable up to 65535 (~33.3 μ s) with T_c unit herein T_c is 1/1.96608 GHz, specified in 3GPP TS 38.211 [4] clause 4.1. The time offset between UL frame and DL frame with the same frame number is calculated as **n-ta-offset** $\times T_c$.

Table 4.8.2-1: Default values of n-ta-offset

Carrier type and frequency range	$N_{TA\ offset}$ (Unit: T_c) (NOTE 4)	Default n-ta-offset (Unit: T_c) (NOTE 5)
LTE FDD (NOTE 1) NB-IoT FDD (NOTE 2) NR FR1 FDD with LTE-NR and/or NB-IoT-NR coexistence case (NOTE 3)	0	0
LTE TDD (NOTE 1) NB-IoT TDD (NOTE 2) NR FR1 TDD with LTE-NR and/or NB-IoT-NR coexistence case (NOTE 3)	39936	39936
NR FR1 FDD or TDD band with neither LTE-NR nor NB-IoT-NR coexistence case (NOTE 3)	25600	25600
NR FR2 (NOTE 3)	13792	13792
NOTE 1: See 3GPP TS 36.211 [42] clause 8.1 NOTE 2: See 3GPP TS 36.211 [42] clause 10.0.1 NOTE 3: See 3GPP TS 38.133 [37] Release 16 Table 7.1.2-2 NOTE 4: Timing value specified in 3GPP specification according radio technology and its use case. NOTE 5: Default value of M-Plane parameter n-ta-offset to be set by O-DU.		

5 Transport & Protocol Architecture

5.1 Transport encapsulation types

5.1.1 Ethernet encapsulation

Ethernet can be used as transport mechanism for both U-Plane and C-Plane. In this case, messages are transmitted over standard Ethernet frames (see Figure 5.1.1-1). The supported frame format is Ethernet II/DIX frame with type interpretation of the type length field. The Length-interpretation and multiplexing data with LLC (IEEE 802.2 [44] LLC standard) is not supported. A globally administered unicast MAC address is assigned per Ethernet interface. The Ethernet termination feature supports MAC Client Data field sizes of up to 1500 bytes for basic frames, and up to 9000 bytes for jumbo frames.

To ensure QoS is supported, 3-bit Priority Code Point (PCP) field as defined in IEEE 802.1Q [36] shall be supported meaning the capability for VLAN tagging for the traffic shall be supported and M-Plane configurable (but using VLAN tags may be omitted for certain network configurations e.g., LLS-C1). Priority tagged frames (i.e., frames tagged with VLAN ID=0) are always received, as required by standard. Transmitting priority tagged frames is optional.

The MAC address resolution for U-Plane and C-Plane shall be performed via an M-Plane procedure specified in [7] regardless of whether the destined node supports IP or not.

The O-DU and O-RU shall support Intermediate L2 switch MAC address learning to avoid flooding U-Plane and/or C-Plane (not necessary for topology LLS-C1); this is done as a result of test Ethernet frame exchanges during the M-Plane procedure "Ethernet connectivity checking" as specified in [7].

Further, both the transport header and payload are contained within the Ethernet data field. For this encapsulation, either the eCPRI Ethertype or the IEEE 1914.3 [45] Ethertype shall be used.

Preamble (8 Bytes)	Destination MAC Address (6 Bytes)	Source MAC Address (6 Bytes)	VLAN Tag (4 Bytes)	Type/Length (Ethertype) (2 Bytes)	Payload (46...1500 Bytes)	FCS (4 Bytes)	IFG (12 Bytes)
-----------------------	---	------------------------------------	-----------------------	---	------------------------------	------------------	-------------------

Figure 5.1.1-1: Native Ethernet frame with VLAN

5.1.2 IP/UDP encapsulation

IP/UDP can be used as transport mechanism for both U-Plane and C-Plane. In this case, IP version 4 shall be supported according to RFC791 (Internet Protocol), and/or IPv6 as per IETF RFC 2460 [46] (both IPv4 and IPv6 are optional, see Table 10.2-1). The encapsulation mechanism is identified by "IPv4" or "IPv6" Ethertype (see Figure 5.1.2-1 and Figure 5.1.2-2). O-RAN equipment shall not send IPv4 packets with any IPv4 header option field or IPv6 packets with "Next Header" value that is an extension header or "Next Header = 59". If the IPv4 header option field or IPv6 packets with "Next Header" value that is an extension header or "Next Header = 59" is detected in the received packet, the packet shall be silently discarded at the receiving equipment (though the existence of the discarded packet may be logged). For IPv6, if the value in "Next Header" is not an extension header and value is not equal to "59", it indicates that the next item in the header is upper-layer header and the packets shall be processed by the receiver (refer to IETF RFC 2460 [46]).

IP based applications can be bound to virtual or physical interface addresses. Each configured physical and logical (i.e. VLAN-) interface shall be configured with a separate subnet. As an option, Alias IP addresses (i.e. more than one IP address on the same interface) can be supported as well. If supported, they can be assigned to logical and physical interfaces as well as to virtual interfaces (loopback interfaces).

The C/U-Plane packets shall not use IPv4 or IPv6 fragmentation due to stringent bandwidth and tight latency requirements (such received packets shall be silently discarded). The minimum MTU that is configurable shall be as per the relevant standard, i.e., IETF RFC 791 [47] for IPv4 and IETF RFC 2460 [46] for IPv6.

The resolution of MAC addresses and the flooding traffic reduction through L2 switch mac address learning shall be done via ARP according to IETF RFC 826 [48] only when IPv4 is supported by the transmitting and receiving nodes.

When using IP/UDP encapsulation, the UDP destination port field shall identify the encapsulated protocol. The UDP destination port field shall be set during initialization via M-Plane configuration. When the eCPRI Specification or IEEE 1914.3 [45] define their respective UDP destination ports these will become the default for M-Plane configuration.

On egress, the sending node shall set the checksum for transmitted UDP datagrams as per the relevant standards:

- For IPv4 UDP datagrams the field can be filled with a proper value to mark that the checksum is valid or to 0 to mark that the checksum is not calculated (see IETF RFC 768 [49]).
- For IPv6 UDP datagrams the field shall be filled with a proper value (see IETF RFC 2460 [46]).

On ingress, the receiving node may validate UDP checksums if the checksum field contains a proper value.

Preamble (8 Bytes)	Destination MAC Address (6 Bytes)	Source MAC Address (6 Bytes)	VLAN Tag (4 Bytes)	Type/Length (EtherType) (2 Bytes)	IPv4 (20 Bytes)	UDP (8 Bytes)	Payload (1...1472 Bytes)	FCS (4 Bytes)	IFG (12 Bytes)
-----------------------	---	------------------------------------	-----------------------	---	--------------------	------------------	-----------------------------	------------------	-------------------

Figure 5.1.2-1: Native IPv4 packet with VLAN

Preamble (8 Bytes)	Destination MAC Address (6 Bytes)	Source MAC Address (6 Bytes)	VLAN Tag (4 Bytes)	Type/Length (EtherType) (2 Bytes)	IPv6 (40 Bytes)	UDP (8 Bytes)	Payload (1...1452 Bytes)	FCS (4 Bytes)	IFG (12 Bytes)
-----------------------	---	------------------------------------	-----------------------	---	--------------------	------------------	-----------------------------	------------------	-------------------

Figure 5.1.2-2: Native IPv6 packet with VLAN

While the numbers in Figure 5.1.1-1, Figure 5.1.2-1 and Figure 5.1.2-2 imply standard Ethernet packets, use of Jumbo frames (up to 9000 bytes in length) is not precluded in the present document.

5.1.3 Transport headers

5.1.3.1 Introduction

O-RAN allows for multiple different transport headers within the Ethernet payload to further describe how the application data is to be handled in the C-Plane and U-Plane. In each case the transport header is 8 bytes in length and provides basic data routing capabilities, including description of the data flow type, sending and reception port identifiers, ability to support concatenation of multiple application messages in a single Ethernet packet, and sequence numbering. The following clauses describe the possible transport headers.

5.1.3.2 eCPRI transport header

5.1.3.2.1 Overview

The definition of the eCPRI transport header is shown in Table 5.1.3.2.1-1 below.

Table 5.1.3.2.1-1: eCPRI transport header field definitions

Section Type: any									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ecpriVersion				ecpriReserved			ecpriConcatenation	1	Octet 1
ecpriMessage								1	Octet 2
ecpriPayload								2	Octet 3
ecpriRtcid / ecpriPcid								2	Octet 5
ecpriSeqid								2	Octet 7

Each field within the eCPRI Transport Header is further described in the following clauses.

5.1.3.2.2 ecpriVersion (eCPRI protocol revision)

Description: This parameter indicates the eCPRI protocol version. This parameter is part of the eCPRI common header.

Value range: {0001b=eCPRI version 1.0, 1.1, 1.2 and 2.0, where the interpretation of the eCPRI message shall follow the eCPRI specification versions up to 2.0; 0000b and 0010b-1111b=Reserved for future eCPRI protocol revisions}.

Type: unsigned integer.

Field length: 4 bits.

Default Value: 0001b (eCPRI version 1.0, 1.1, 1.2 and 2.0).

5.1.3.2.3 ecpriReserved (eCPRI reserved)

Description: This parameter is reserved for eCPRI future use. This parameter is part of the eCPRI common header.

Value range: {001b-111b=Reserved}.

Type: unsigned integer.

Field length: 3 bits.

Default Value: 000b (reserved fields should always be set to all zeros).

5.1.3.2.4 ecpriConcatenation (eCPRI concatenation indicator)

Description: This parameter indicates when eCPRI concatenation is in use (allowing multiple eCPRI messages in a single Ethernet payload). This parameter is part of the eCPRI common header.

Value range: {0b=No concatenation, 1b=Concatenation}.

Type: binary bit.

Field length: 1 bit.

Default Value: 0b (no concatenation).

5.1.3.2.5 ecpriMessage (eCPRI message type)

Description: This parameter indicates the type of service conveyed by the message type. This parameter is part of the eCPRI common header. In the present document, only values "0000 0000b" and "0000 0010b" and "0000 0101b" are used.

Value range:

0000 0000b = IQ data message

0000 0010b = Real-time control data message

0000 0101b = transport network delay measurement message (see clause 4.4.4.4 for full message format)

other values not recognized within the present document.

Type: unsigned integer.

Field length: 8 bits.

Valid Values: 0x0 (U-Plane data) or 0x2 (C-Plane data) or 0x5 (network delay measurement messages).

5.1.3.2.6 ecpriPayload (eCPRI payload size)

Description: This parameter is the size in bytes of the payload part of the corresponding eCPRI message. It does not include any padding bytes following the eCPRI message. The maximum supported payload size is $2^{16}-1$, but the actual

size may be further limited by the maximum payload size of the underlying transport network. This parameter is part of the eCPRI common header. See Figure 5.1.3.2.6-1 for a pictorial representation.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

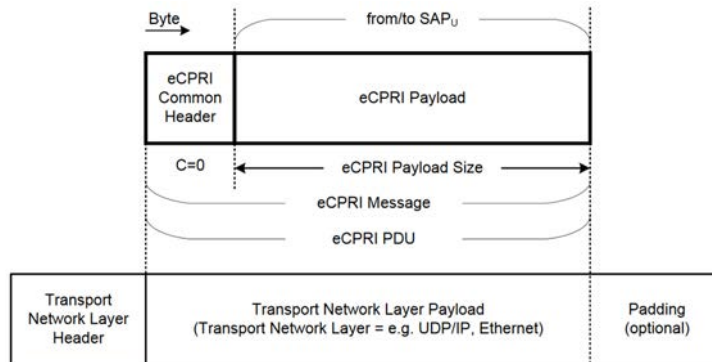


Figure 5.1.3.2.6-1: eCPRI payload - no concatenation (Ref. [2])

5.1.3.2.7 ecpriRtcid / ecpriPcid (real time control data / IQ data transfer message series identifier)

Description: This parameter is an eAxC identifier (eAxC_ID) and identifies the specific data flow associated with each C-Plane (ecpriRtcid) or U-Plane (ecpriPcid) message. It is the analog of CPRI's "AxC" (antenna-carrier) value so is designated here as "eAxC" ("e" for "extended" to accommodate multiple bands and multiple component carriers). Multiple O-DU processors may contribute to a single eAxC.

An endpoint may support mixed numerologies by indicating the capability via an M-Plane parameter; and a Section Type 3 message can be used to select one of the numerologies (frameStructure) from the capability list of possible numerologies for that endpoint. In this case, a single eAxC_ID may be used for channels with different characteristics (frameStructure, cpLength, timeOffset, freqOffset), e.g. channels with different numerology and PRACH can share same eAxC_ID. Alternatively, an endpoint may be simpler, supporting a single numerology by declaring non-support of mixed numerology capability via the M-Plane parameter. In this case, a Section Type 3 message can be used to convey different values for parameters except for the frameStructure parameter. In this case a unique eAxC_ID would be used for each mixed numerology channel (frameStructure), i.e. channels with different numerologies will use different eAxC_IDs. An endpoint may declare non-support for Section Type 3, in which case it is suitable for channels like PDSCH/PUSCH.

The eAxC_ID value assigned to an O-RU's endpoint is used to identify the endpoint within O-RU and shall be unique among all endpoints of O-RU of the same direction (Tx or Rx). An O-RU's endpoint could in principle be associated with multiple physical and virtual fronthaul interfaces (Ethernet ports and VLANs), but an eAxC_ID value shall be unique and cannot be used to address different endpoints. The association of an O-RU's endpoint with multiple fronthaul interfaces is subject to O-RU design restrictions reported by O-RU via the M-Plane – interfaces not associated with a given endpoint shall not be configured via M-Plane for transferring C-Plane and U-Plane message to the endpoint. Also, fronthaul interfaces not configured via M-Plane shall not be used for transferring C-Plane and U-Plane messages to the endpoint. In other words, the O-DU shall obey the O-RU restrictions as conveyed on the M-Plane.

The O-DU may decide to use same eAxC_ID value to address endpoints on different O-RUs, but eAxC_ID values assigned to one O-RU shall be unique within the O-RU's endpoints of same direction. eAxC_ID values assigned to one FHM shall be unique in the same direction. There can be O-RUs which do not support, within a same eAxC_ID, independent sequence checkers for C-Plane messages describing U-Plane DL and C-Plane messages describing U-Plane UL, which is specified in clause 5.1.3.2.7. This O-RU capability limitation can be signaled by O-RUs and be interpreted by O-DUs from M-Plane specification v03.00. If either or both of the O-RU and O-DU only supports previous versions of the M-Plane specification, then the network operators and vendors intending to work with such O-RUs will need to ensure that the O-DU can interpret the O-RU capability limitation in other non-standardized ways, i.e., through off-line discussions. To interoperate with such O-RUs, O-DUs should avoid requiring at the O-RU, within a same eAxC_ID, independent sequence checkers for C-Plane messages describing U-Plane DL and C-Plane messages describing U-Plane UL, either by:

- using different eAxC_IDs for them; or
- if the O-DU chooses to use the same eAxC_ID, using a shared sequence generator for C-Plane messages describing U-Plane DL and C-Plane messages describing U-Plane UL as shown in Figure 5.1.3.2.8-2 and described in clause 5.1.3.2.8.

However, support of such behavior in order to interoperate with such O-RUs is optional for the O-DUs, and incompatible operation is likely to result in data outages and spurious error reports.

The present document does not define behavior regarding transferring C-Plane and U-Plane messages to or from an endpoint over multiple fronthaul interfaces at the same time. For example, in UL, if an endpoint were to be assigned to multiple fronthaul ports, there is no way to instruct the endpoint which fronthaul port to use for any given message. In addition, in DL there is no accommodation for multiple delay windows per eAxC_ID when multiple links are used. Therefore, an endpoint shall be assigned exactly one fronthaul interface via the M-Plane, but different endpoints may use different fronthaul interfaces (ports) so the O-RU as a whole may use multiple fronthaul ports.

eAxC_ID is specific to eCPRI Message Type = 2 (C-Plane) and eCPRI Message Type = 0 (U-Plane) messages. In the present document, one eAxC contains only one spatial stream (i.e., one beam per subcarrier) at a time. When precoding in the O-RU, then each eAxC contains one layer at a time, except for TxD (LTE TM2) when a single eAxC (single ecpriRtcid and ecpriPcid) represents all TxD layers.

eAxC_ID subfields: One eAxC identifier (eAxC_ID) comprises a band and sector identifier (BandSector_ID), a component-carrier identifier (CC_ID) a spatial stream identifier (RU_Port_ID) and a Distributed Unit identifier (DU_Port_ID).

Bit allocation is subdivided as follows:

- DU_Port_ID: Used to differentiate processing units at O-DU (e.g., different baseband cards). It can be expected the O-DU will assign these bits, and the O-RU will attach the same value to the UL U-Plane messages carrying the same sectionId data.
- BandSector_ID: Aggregated cell identifier (distinguishes bands and sectors supported by the O-RU).
- CC_ID: distinguishes Component Carrier supported by the O-RU.
- RU_Port_ID: designates logical flows such as data layers or spatial streams, and logical flows such as separate numerologies (e.g. PRACH) or signaling channels requiring special antenna assignments such as SRS.

The assignment of the DU_port_ID, BandSector_ID, CC_ID, and RU_Port_ID as part of the eAxC_ID is done solely by the O-DU via the M-Plane. Furthermore, the O-RU doesn't need an explicit definition of any bit-level allocation within any of the four fields of the eAxC_ID.

The bitwidth of each of the above fields is variable and set via M-Plane messaging. This is to allow flexibility given it can be expected that not all fields will simultaneously need their maximum range for any given O-RU. It can be expected the M-Plane message will configure the O-RU and O-DU with the appropriate bitwidth of each of the four fields, and the SMO that does the actual assignment will assure all 16 bits are allocated (with or without padding).

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b = eAxC_ID}.

Bit allocations: See Table 5.1.3.2.7-1.

Table 5.1.3.2.7-1: Bit allocations

0 (msb)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15 (lsb)	Number of Octets
DU_Port_ID				BandSector_ID				CC_ID				RU_Port_ID				2

Type: unsigned integer (concatenated bit fields).

Field length: 16 bits.

5.1.3.2.8 ecpriSeqid (message identifier)

Description: This parameter provides unique message identification and ordering on two different levels. The first octet of this parameter is the Sequence ID, which is used to identify ordering of messages within an eAxC message stream. The Sequence ID field increments and wraps independently for each U-Plane eAxC DL, U-Plane eAxC UL, C-Plane eAxC DL, and C-Plane eAxC UL, even if they share the same eAxC_ID. The Sequence ID is used to verify that all messages are received and to reorder messages that are received out of order. The second octet of this parameter is the Subsequence ID. The Subsequence ID is used to verify ordering and implement reordering when radio-transport-level (eCPRI or IEEE-1914.3) fragmentation occurs. Radio-transport (eCPRI or IEEE-1914.3) fragmentation is a method of splitting U-Plane messages containing one or more sections whose length exceeds the maximum packet or message length of the underlying protocol. The Subsequence ID field consists of a 7 bit Subsequence counter and a single bit field, called E-bit. The Subsequence number increments starting from zero for each fragment of a U-Plane message. The E bit is used to indicate the last message of the radio-transport level fragments. It shall be set to zero except for the last message of the U-Plane fragment. In the case of C-Plane messages radio-transport fragmentation is not allowed, therefore the Subsequence ID shall be set to zero, and the E bit set to one. See clause 5.5 for a description of the fragmentation process.

As an alternative to radio-transport-level fragmentation, application fragmentation can be implemented. In this case the application can take the responsibility to ensure all transport messages are not too long (fit within the necessary transport payload size). When this "application layer fragmentation" is used, the subsequence identifier shall always be set to "0", and the E-bit set to "1" (see clause 5.5).

Figure 5.1.3.2.8-1 shows how the sequence numbers are intended to be generated and checked by C-Plane and U-Plane messaging.

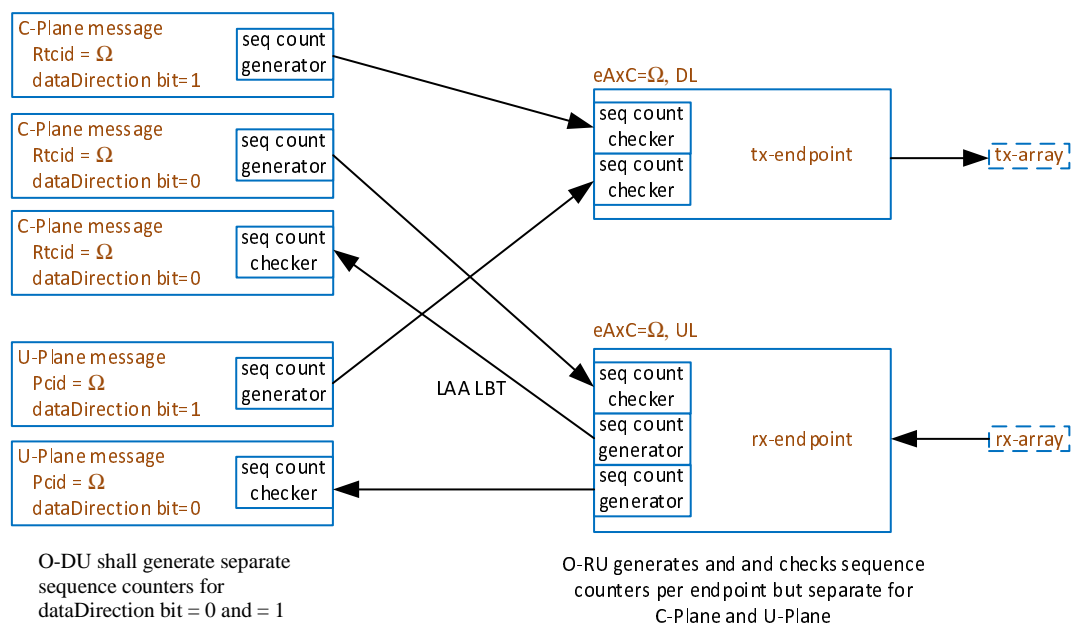


Figure 5.1.3.2.8-1: Sequence count generation and checking for C-Plane and U-Plane in O-DU and O-RU

Implications of Figure 5.1.3.2.8-1 are as follows:

- Sequence ID is unique per eAxC (ecpriRtcid or ecpriPcid): different eAxC values have their own Sequence ID generation
- Sequence ID is generated by the fronthaul interface transmitter in either the O-DU or O-RU
- The sequence generator and checker for a C-Plane message describing U-Plane DL is independent of the sequence generator and checker for a C-Plane message describing U-Plane UL.

Table 5.1.3.2.8-1 summarizes this situation consistent with Figure 5.1.3.2.8-1:

Table 5.1.3.2.8-1: Description of shared or independent sequence generators

direction of message flow	type of flow	U-Plane data direction	Same/different eAxC between DL & UL	shared/independent sequence generator
O-DU → O-RU	C-Plane	DL (dataDirection bit =1)	same or different	independent
O-DU → O-RU	C-Plane	UL (dataDirection bit =0)	same or different	independent
O-RU → O-DU	C-Plane	UL (dataDirection bit =0)	same or different	independent
O-DU → O-RU	U-Plane	DL (dataDirection bit =1)	same or different	independent
O-RU → O-DU	U-Plane	UL (dataDirection bit =0)	same or different	independent

As described in clause 5.1.3.2.7, there are O-RUs which do not support, within a same eAxC ID, independent sequence checkers for C-Plane messages describing U-Plane DL and C-Plane messages describing U-Plane UL. As also described in clause 5.1.3.2.7, to interoperate with such O-RUs, O-DUs may choose, among other options, to use the same eAxC ID but with a shared sequence generator for C-Plane messages describing U-Plane DL and C-Plane messages describing U-Plane UL as described in Figure 5.1.3.2.8-2 and Table 5.1.3.2.8-2 below.

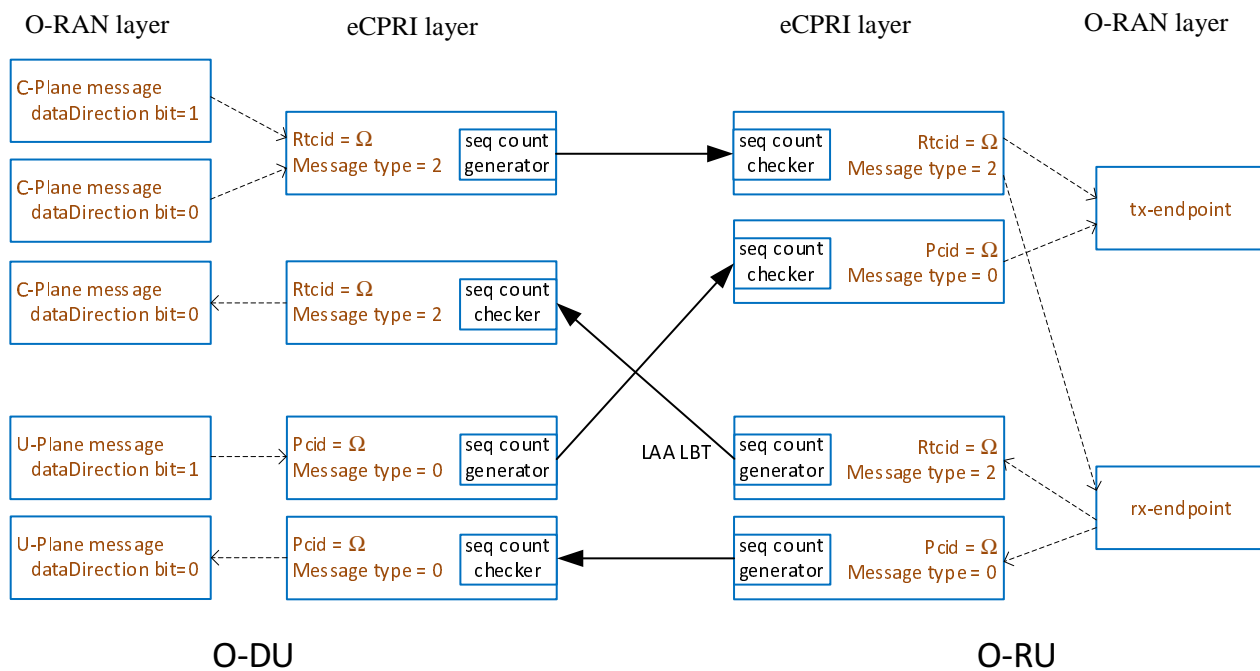


Figure 5.1.3.2.8-2: Optional case of shared sequence generation and checking for C-Plane and U-Plane in O-DU and O-RU

Table 5.1.3.2.8-2: Optional case of description of shared sequence per eAxC_ID generators

direction of message flow	type of flow	U-Plane data direction	Same/different eAxC between DL & UL	shared/independent sequence generator
O-DU → O-RU	C-Plane	DL (dataDirection bit =1)	same	Shared
O-DU → O-RU	C-Plane	UL (dataDirection bit =0)	same	Shared
O-RU → O-DU	C-Plane	UL (dataDirection bit =0)	same or different	Independent
O-DU → O-RU	U-Plane	DL (dataDirection bit =1)	same or different	Independent
O-RU → O-DU	U-Plane	UL (dataDirection bit =0)	same or different	Independent

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}

Bit allocations: See Table 5.1.3.2.8-3

Table 5.1.3.2.8-3: Bit allocations

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
Sequence ID								1	Octet 1
E bit	Subsequence ID							1	Octet 2

Type: unsigned integer (concatenated bit fields)

Field length: Sequence ID: 8 bits; subsequence ID: 7 bits; E-bit: 1 bit

5.1.3.3 IEEE 1914.3 transport header

5.1.3.3.1 Overview

As an alternative to eCPRI as a transport header, IEEE 1914.3 [45] may be used. The definition of the IEEE 1914.3 transport header is shown in Table 5.1.3.3.1-1 below.

Table 5.1.3.3.1-1: IEEE 1914.3 transport header field definitions

Section Type: any									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
RoEsubType								1	Octet 1
RoEflowId								1	Octet 2
RoElength								2	Octet 3
RoEorderInfo								4	Octet 5

Each field within the IEEE 1914.3 Transport Header is further described in the following clauses.

5.1.3.3.2 RoEsubType (sub type / message type)

Description: This field indicates the payload type within the IEEE 1914.3 Standard for Radio over Ethernet Encapsulations and Mappings (RoE) subType range. RoE allows RoE subTypes in the range 128 to 191 to be mapped to external organizations and companies using a subtype mapping table (below). This table has two fields. The OUI/CID field uses a unique Company ID (CID) value. For the present document, the OUI/CID assigned by IEEE to xRAN (0xFAEB6E) shall be used. The payload structure mapping field assigns the RoE subtype to O-RAN message types. When the payload structure mapping field is set to 0xFFFF, this indicates that that RoE subtype is not mapped to an O-RAN message type.

A given O-RAN implementation may choose to have fixed (default) mappings, or alternatively, it may choose to configure the subType mapping table using control packets as described in IEEE 1914.3 [45].

Table 5.1.3.3.2-1 indicates the use by O-RAN of the IEEE-defined xRAN subtype mapping; future discussions with IEEE are needed to convert these to O-RAN subtype mapping.

Table 5.1.3.3.2-1: RoE subtype mapping

RoE subType	OUI/CID subType mapping table (mapSubtype)	
	OUI/CID Mapping (3 bytes) bit39 <-----	Payload structure mapping (2 bytes) -----> bit 0
128	xRAN=0xFAEB6E	0x0001, IQ (No concatenation)
129	xRAN=0xFAEB6E	0x0002, IQ (With concatenation)
130	xRAN=0xFAEB6E	0x0003, Ctrl (No concatenation)
131	xRAN=0xFAEB6E	0x0004, Ctrl (With concatenation)
132 to 191	xRAN=0XXXXXX (do not care)	0xFFFF (IEEE1914.3 [45] default), unused/unmapped by xRAN in the present document.

Value range: 128 to 191. Default values are shown below.

Type: unsigned integer.

Field length: 8 bits.

Valid Values: 128 (U-Plane, no concat), 129 (U-Plane, with concat), 130 (C-Plane, no concat), 131 (C-Plane, with concat).

5.1.3.3.3 RoEflowID (flow identifier)

Description: The RoEflowID is a mechanism which can identify specific flows between end-points. RoEflowID, 0xFF is reserved for RoE control packets. O-RAN has no current use for this field.

Value range: 0 – 0xFE.

Type: unsigned integer.

Field length: 8 bits.

Description: This field is currently unused.

5.1.3.3.4 RoElength (length)

Description: This field is the size in bytes of the payload part of the message. The payload length field value shall be the total number of octets following the O-RAN common header. It does not include the Ethernet FCS or following bytes.

Value range: 0 – 0xFFFF.

Type: unsigned integer.

Field length: 16 bits.

5.1.3.3.5 RoEorderInfo (order information)

Description: This field is split into seven sub-fields as shown in Table 5.1.3.3.5-1.

Table 5.1.3.3.5-1: Layout of RoEOrderInfo field

Field	Length	Note
DU_Port_ID	16 bits	Used to differentiate processing units at O-DU (e.g., different baseband cards). It can be expected the O-DU will assign these bits, and the O-RU will attach the same value to the UL U-Plane messages carrying the same sectionId data. See sub-clause 5.1.3.2.7 for further information.
BandSector_ID		Aggregated cell identifier (distinguishes bands and sectors supported by the O-RU). See sub-clause 5.1.3.2.7 for further information.
CC_ID		Distinguishes Component Carrier supported by the O-RU. See sub-clause 5.1.3.2.7 for further information.
RU_Port_ID		Used to differentiate spatial streams or beams on the O-RU. See sub-clause 5.1.3.2.7 for further information.
Sequence_ID	8 bits	Unique message ordering sequence. See sub-clause 5.1.3.2.8 for further information.
E_Bit	1 bit	Marks the last message pertaining to the section. See sub-clause 5.1.3.2.8 for further information.
Subsequence_ID	7 bits	Unique message ordering sub-sequence. See sub-clause 5.1.3.2.8 for further information.

Value range: 0 to 0xFFFF FFFF.

Type: unsigned integer (concatenated bit fields).

Field length: 32 bits.

5.2 Protocol architecture

5.2.1 C-Plane

Figure 5.2.1-1 depicts the protocol stack for C-Plane. Data can be optionally transmitted over IP Layer 3 if supported by the transmitting and receiving nodes.




Figure 5.2.1-1: C-Plane protocol structure

5.2.2 U-Plane

Figure 5.2.2-1 depicts the protocol stack for U-Plane. Data can be optionally transmitted over IP Layer 3 if supported by the transmitting and receiving nodes.




Figure 5.2.2-1: U-Plane protocol structure

5.2.3 S-Plane

Frequency and time synchronization of O-DUs and O-RUs via Ethernet use Synchronous Ethernet and IEEE 1588 [3], [33] Precision Time Protocol (PTP). Transport of PTP directly over L2 Ethernet (Recommendation ITU-T G.8275.1 [31] full timing on-path support) is assumed in the present document, whilst transport of PTP over UDP/IP (Recommendation ITU-T G.8275.2 [32] partial timing support from the network) is also possible albeit with unassured synchronization performance. Security mechanisms will not be used for this Plane. The protocol stack for PTP and SyncE (Synchronous Ethernet) over L2 Ethernet is depicted in Figure 5.2.3-1.




Figure 5.2.3-1: S-Plane protocol structure

5.3 Quality of service

The LLS interface needs to support the ability to distinguish between data flows with different QoS requirements. Configurable priority levels (via the M-Plane) for traffic prioritization of flows shall be supported on each node on the network path. Default values for the respective O-RAN planes are indicated as shown in Table 5.3-1. Priority marking per packet shall be used in each protocol layer except for protocols that do not allow such tagging e.g., Recommendation ITU-T G.8264 [21] for SyncE and Recommendation ITU-T G.8275.1 [31] for full timing support PTP. For operation at Layer 2, prioritization is performed by specifying a configurable value for the PCP tag in the IEEE 802.1Q VLAN header on the outgoing traffic (see IEEE Standard 802.1Q [36], clauses 6.9.3 and 9).

Table 5.3-1: Quality of service classes

Plane	L2 CoS Priority (range 0-7)	L3 DSCP Code	Preemption(NOTE 1)
S-Plane	G.8264 [21]: N/A G.8275.1 [31]: N/A G.8275.2 [32]: Default 7	G.8264: N/A G.8275.1: N/A G.8275.2: EF (Expedited Forwarding)	non-preemptable
U-Plane (NOTES 2,3)	Default: 7	EF (Expedited Forwarding)	non-preemptable
C-Plane (NOTES 2,3)	Default: 7	EF (Expedited Forwarding)	non-preemptable
M-Plane (NOTE 3)	Default: 2	AF2x (Assured Forwarding)	preemptable
Other traffic	Default: 1	BE (Best Effort)	preemptable
NOTE 1: Not all networks will support pre-emption, so this only applies to networks supporting preemption			
NOTE 2: Applies also to LBP [7]			
NOTE 3: Applies also to ARP, if used, where IPv4 is used			

For U-Plane separate priorities can be supported based on *ecpriRtcid* / *ecpriPcid* (or the equivalent fields in RoE packet headers). This allows for different prioritization of some channels (e.g. PRACH/SRS), or services (URLLC). This requires the ability to configure priority assignments to *ecpriRtcid* / *ecpriPcid* via M-Plane messages. The default U-Plane priority can be applied to flows not specifically configured with a different value via the M-Plane. If a node does not support configurable *ecpriRtcid* / *ecpriPcid* priorities, then all U-Plane traffic will use the default value.

QoS failure (dropped packets) may be detected using the sequence numbers that make up part of the transport header. Loss of a packet in the U-Plane will impact only a specific symbol. Loss of a C-Plane packet is more serious, with the possibility of impacting an entire slot's worth of data. A broader impact is also possible with C-Plane packet loss: if downloaded beamforming weights, beamforming attributes, or channel information is lost then a *beamId* may be incorrectly associated with an "old" set of beamforming weights thereby possibly impacting multiple slots. However, packet loss can be detected via the sequence number so a recovery action may be taken to re-send messages that may have a persistent impact. The C-Plane packet loss can also be more precisely detected via the ACK/NACK procedure (as described in clause 7.2.8).

5.4 Data flow identification

Differentiation between a combined C/U-Plane data flow for a single eAxC_ID and Management Plane data flow traffic can be achieved using the following options:

- Data flow separation based on TCP/UDP (applicable when layer 3 transport is used for the C/U-Plane).
- Data flow separation based on VLAN (applicable when layer 2 or layer 3 is used for the C/U-Plane transport):
 - VLAN ID for the combined U-Plane and C-Plane data flow for a given eAxC_ID is configured via the M-Plane.
- Data flow separation based on usage of different MAC addresses (applicable when Layer 2 is used for the combined C/U Plane transport):

EXAMPLE: One MAC address used for the combined C/U-Plane data flow and a second MAC address used for the M-Plane data flow or additional MAC addresses used for the U-Plane for baseband load sharing purposes.
- Data flow separation based on different EtherTypes (applicable when Layer 2 is used for C/U Plane transport).

The U-Plane application also needs to uniquely associate different data flows (e.g. spatial streams) each with a unique C/U Plane endpoint address. This can be achieved in an O-RU using the eAxC identifier, and in the O-DU using the eAxC identifier in combination with transport-based endpoint identifiers to differentiate O-RUs. In addition, O-RU data flows can be switched/routed to different O-DUs (or different O-DU ports or O-DU processors) according to the transport-based identifiers associated with an eAxC_ID (referred to as processing-elements in the WG4 M-Plane Specification) to allow frames/packets to be switched/routed by network equipment with no visibility of the eAxC values carried in the eCPRI/IEEE 1914.3 transport header (see clause 5.1.3).

Different transport identifiers (e.g. UDP/IP, VLAN, MAC) can be used based on whether layer 2 or layer 3 transport is used for the C/U-Plane:

- Separation of the combined C/U-Plane data flow using UDP-port identifiers (applicable when layer 3 transport is used for the C/U-Plane).
- Separation of the combined C/U-Plane data flow using VLAN identities (applicable when layer 2 or layer 3 is used for the C/U-Plane transport).
- Separation of the combined C/U-Plane data flow using different MAC addresses (applicable when Layer 2 is used for C/U Plane transport).

5.5 Fragmentation

5.5.1 Introduction

Fragmentation shall be applied in case data (U-Plane or C-Plane data) with Ethernet transport overheads to be transferred exceed maximum transmission unit (MTU) of the network. The present document allows two methods for fragmentation, application layer fragmentation and transport layer fragmentation:

5.5.2 Application layer fragmentation

- Application creates C-Plane or U-Plane messages, which when including overheads can fit to MTU requirements set by network
- Sequence ID: Sequence ID increases for every message, E=1, Subsequence ID=0

An example of application fragmentation is provided in Figure 5.5.2-1, wherein a long data section is split between multiple U-Plane messages. In the example it is a single data section (0xABC) so described as such in the C-Plane, and in the U-Plane the sectionId is cited multiple times in the multiple U-Plane messages but with differing and consecutive groups of PRBs as indicated by the startPrbu and numPrbu fields. The I and Q word width in the example is 14 bits, which for such a large data section requires multiple U-Plane messages. The ecpriRtcid / ecpriPcid (0xDEF3) is shown with a random example value, as is the sequence ID in the C-Plane (0x99) and independently in the U-Plane (0x40-0x48).

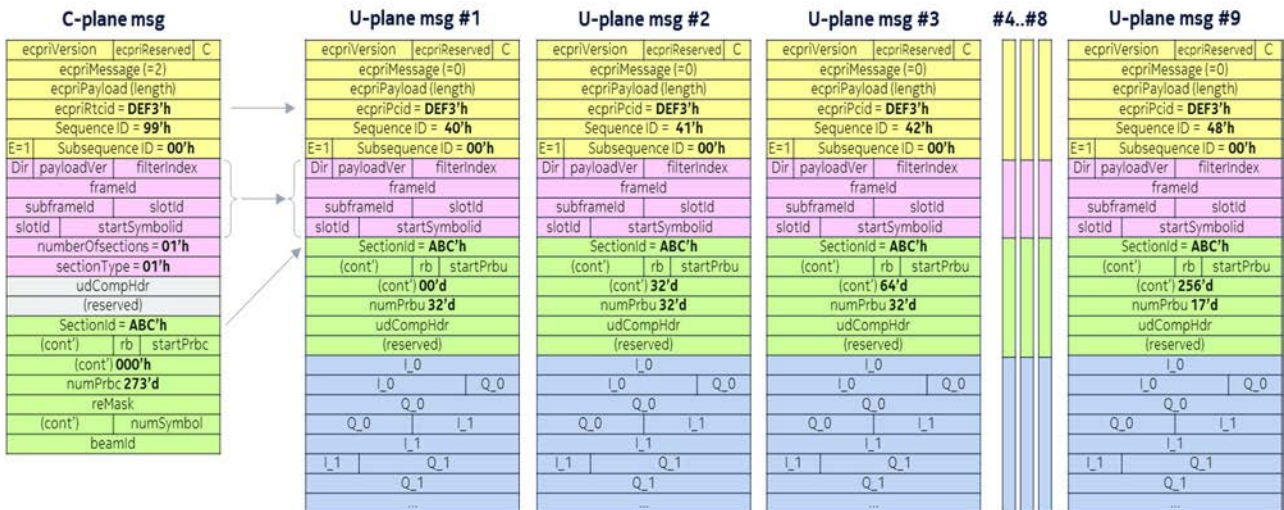


Figure 5.5.2-1: Example of application-level fragmentation of U-Plane messages

5.5.3 Radio transport layer (eCPRI or IEEE-1914.3) fragmentation

- Application creates U-Plane messages, which when including overhead may exceed MTU requirements set by network
- Radio transport layer splits message which may contain more than one section into pieces such that the fragments with overheads fit to MTU requirements set by network.
- Sequence ID: Sequence ID remains same for all the fragments. Subsequence ID starts from 0 for the first fragment and counts up for each fragment. Last fragment flagged with E=1, others E=0. This is shown in Table 5.5.3-1.

Table 5.5.3-1: Example of sequence numbers usage

Application Layer Fragmentation			
eAxC	SeqId	E	Sub-SeqId
Ω	0	1	0
Ω	1	1	0
Ω	2	1	0
Ω	3	1	0
Ω	4	1	0
Ω	5	1	0
Ω	6	1	0
Ω	7	1	0
Ω	8	1	0
Ω	9	1	0

eCPRI Layer Fragmentation			
eAxC	SeqId	E	Sub-SeqId
Ω	0	0	0
Ω	0	1	1
Ω	1	0	0
Ω	1	0	1
Ω	1	0	2
Ω	1	1	3
Ω	2	0	0
Ω	2	0	1
Ω	2	1	2
Ω	3	1	0

5.5.4 Fragmentation considerations

Application layer fragmentation should be applied so that maximum size standard IEEE 802.3 [53] Ethernet frames can be used. (Jumbo frames, if used, increase the maximum MTU size).

- In case of L2 only solution, application layer maximum transmission unit size is standard IEEE 802.3 Ethernet frame payload size (1500 bytes) – transport overhead (8 bytes) = 1492 bytes (or larger for Jumbo frames)

When a U-Plane section payload is fragmented due to large number of PRBs, it will be divided into multiple groups of PRBs such that each group (including the application headers) can fit to the MTU requirements. Each group of PRBs shall have its respective control fields including the same section ID from the C-Plane and startPrbu and numPrbu to identify its contents (see Figure 5.5.2-1). When an application layer message exceeds the MTU size set via the M-Plane configuration, radio-transport layer fragmentation can be applied at the transport layer and has the advantage of allowing application layer messages to fit whatever is the MTU size without the application having to know about transport-level details.

6 Security

6.1 General

Security requirements are depicted in Table 6.1-1. The CUS-Plane has sensitivity to latency. Security controls should be selected with consideration of latency requirements.

Table 6.1-1: Security requirements for User Plane, Control Plane, and Synchronization Plane

Plane	Authentication and Authorization (at interface level)	Integrity (protection from packet modifications and injection)	Confidentiality (encryption protection)	Availability (including performance degradation)	Remarks
C-Plane	IEEE 802.1X supplicant functionality for NAC [51] (NOTE 2)	Not currently specified	Not currently specified	IEEE 802.1X supplicant functionality for NAC [51] (NOTE 2)	
U-Plane	IEEE 802.1X supplicant functionality for NAC [51] (NOTE 2)	No additional protection mechanisms are specified	No additional protection mechanisms are specified	IEEE 802.1X supplicant functionality for NAC [51] (NOTE 2)	3GPP Control Plane and User Plane messages that are transported via the Open Fronthaul U-Plane (LLS-UP) are confidentiality and integrity protected by the Packet Data Convergence Protocol (PDCP) [52].
S-Plane	IEEE 802.1X supplicant functionality for NAC [51] (NOTE 2)	Not currently specified	Not currently specified	Optional Grand Master Clock redundancy (NOTE 1). IEEE 802.1X supplicant functionality for NAC [51] (NOTE 2)	
NOTE 1: Principles for Grand Master Clock redundancy are provided in Annex P.2.3.1 of IEEE 1588 [33] and in the applicable PTP profile (e.g., ITU-T G.8275.1 [31]). Guidelines are provided in Annex G of O-RAN Synchronization Architecture and Solution Specification [54]					
NOTE 2: Please refer to Table 10.2-1 for support of IEEE 802.1X supplicant functionality by the O-DU and the O-RU.					

7 C-Plane Protocol

7.1 General

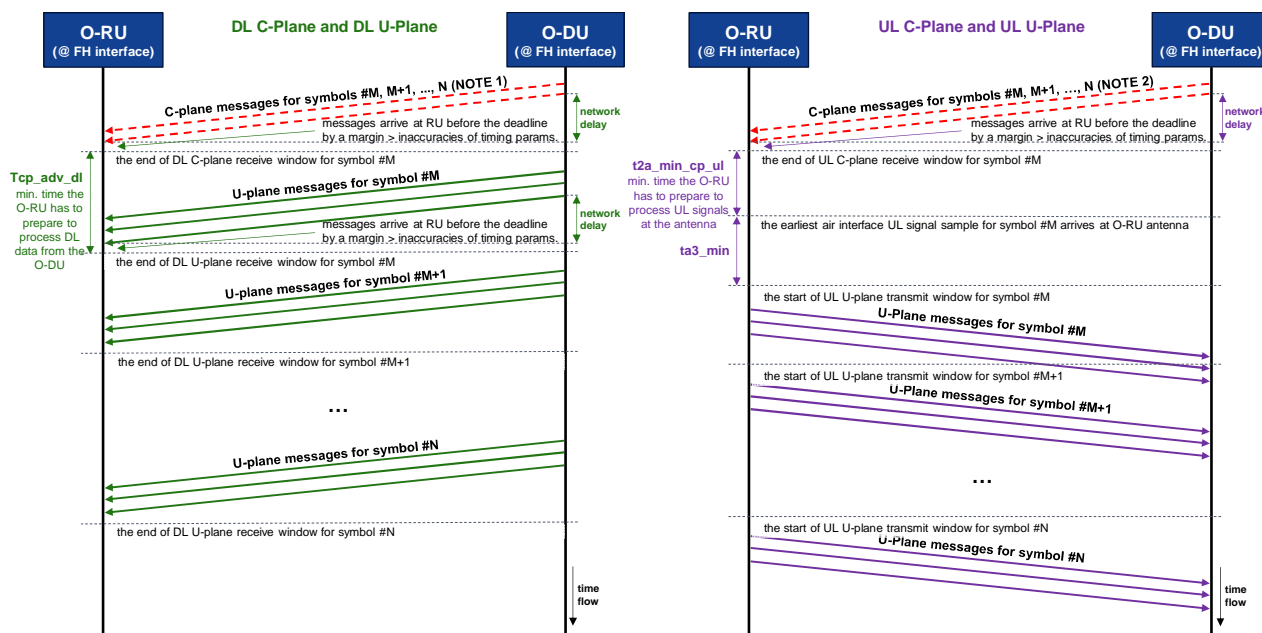
7.1.1 C-Plane transport

Either eCPRI or IEEE 1914.3 shall be used as an encapsulation mechanism for the Control Plane messages. It is assumed that a different data flow is used other than the U-Plane channel. Further, C-Plane messages are not concatenated with U-Plane messages within same Ethernet frame.

7.2 Elementary procedures

7.2.1 Scheduling and beamforming commands transfer procedure

C-Plane messages shall be exchanged between O-DU and O-RU. The main purpose of these messages is to transmit data-associated control information required for processing of user data (e.g., scheduling and beamforming commands) if such information is not provided via M-Plane (see clause 8.2.2 for more details). Messages shall be sent separately for DL related commands and UL related commands (see Figure 7.2.1-1). See also Figure 7.2.5.2-1 for the special case of LTE LAA wherein there are UL as well as DL C-Plane message flows. For purpose of increased flexibility, C-Plane messages may be sent either jointly or separately depending on the channel for which information is conveyed. For example, PUCCH and PUSCH may be bundled or not bundled into a single C-Plane message depending on implementation.



NOTE 1: A DL C-Plane message describing multiple symbols shall arrive at O-RU at least Tcp_adv_dl before the end of DL U-Plane receive window for symbol **startSymbolId** (the earliest symbol described by the message).

NOTE 2: A UL C-Plane message describing multiple symbols shall arrive at O-RU at least $t2a_min_cp_ul$ before the earliest air interface UL signal sample of symbol **startSymbolId** (the earliest symbol described by the message) arrives at O-RU antenna.

Figure 7.2.1-1: Scheduling and beamforming commands transfer procedure

7.2.2 Symbol numbering and duration

For the NR case, data symbol numbering in a slot shall be from zero to thirteen (for extended CP, eleven). For LTE the symbol numbering shall be mapped to the NR numbering as follows:

LTE with normal prefix maps to NR with same μ and normal prefix:

LTE symbol l_{LTE} of slot $2i$ is mapped to NR slot i , symbol $l' = l_{LTE}$

LTE symbol l_{LTE} of slot $2i + 1$ is mapped to NR slot i , symbol $l' = 7 + l_{LTE}$

LTE with extended prefix the following applies:

LTE symbol l_{LTE} of slot $2i$ is mapped to NR slot i , symbol $l' = l_{LTE}$

LTE symbol l_{LTE} of slot $2i + 1$ is mapped to NR slot i , symbol $l' = 6 + l_{LTE}$

The symbol duration can be determined from the SCS provided by "frameStructure" and knowledge of whether LTE or NR is supported; 3GPP TS 36.211 [42] clause 5.2.1 (UL) and clause 6.2.1 (DL), and 3GPP TS 38.211 [4] clause 4.3 provide the necessary formulas.

When the SCS value provided by "frameStructure" is less than 15 kHz (e.g., for long preamble PRACH formats), numerology for symbolId shall be based on 15 kHz SCS.

7.2.3 Mixed numerology and PRACH handling

7.2.3.1 Overview

5G NR defines the Physical Resource Block (PRB) such that the number of subcarriers per PRB is the same (i.e., 12) for all numerologies. However, subcarrier spacing may differ, resulting in mixed numerology used in time and frequency domains. Likewise, mixed numerologies may be applied to both DL and UL respectively. An example is shown in Figure 7.2.3.1-1.




Figure 7.2.3.1-1: Mixed numerology in time and frequency domain

7.2.3.2 Frequency indexing across multiple numerologies

RB indexing (i.e. the physical resource block) is always dependent on the used numerology. For PRACH and SSB the relation is with the physical channel. In case of PRACH, indexing of RBs inside of each PRACH frequency block follows 3GPP specifications and PRACH subcarrier extraction shall be performed in the O-DU. In particular, the first RE of the first RB addressed inside section Id shall correspond with the first guard tone used at the lower edge of the PRACH frequency block based on the PRACH SCS. The O-RU, using the filter index and SCS of corresponding PUSCH, knows the PRACH subcarrier position and guard tones and can then adjust its filtering parameters (e.g. filter center, passband). For the SSB case, its own separate frequency offset is assigned.

Figure 7.2.3.2-1 shows a PRACH example illustrating startPrbc, numPrbc, and freqOffset. Specifically,

- freqOffset indicate the location of lowest RE's center in the lowest RB defined by frameStructure, with respect to center-of-channel-bandwidth.
- startPrbc and numPrbc in the Section Type 3 indicate the PRB in the RB grid defined by corresponding freqOffset and frameStructure.

- If `freqOffset` refers to the lowest RE's center in the lowest RB of the PRACH block (as shown in Figure 7.2.3.2-1), then `startPrbc` shall be set to 0.
- If `freqOffset` refers to lower than the start of the PRACH block, then `startPrbc` shall take a non-zero value, where indexing starts from the PRB referred to by `freqOffset`.
- Note that `numPrbc` may exclude PRBs with all-guard-tones at the upper end of the PRACH frequency block. (Hence, in the example in Figure 7.2.3.2-1, `numPrbc` could be set to 71 or 72).

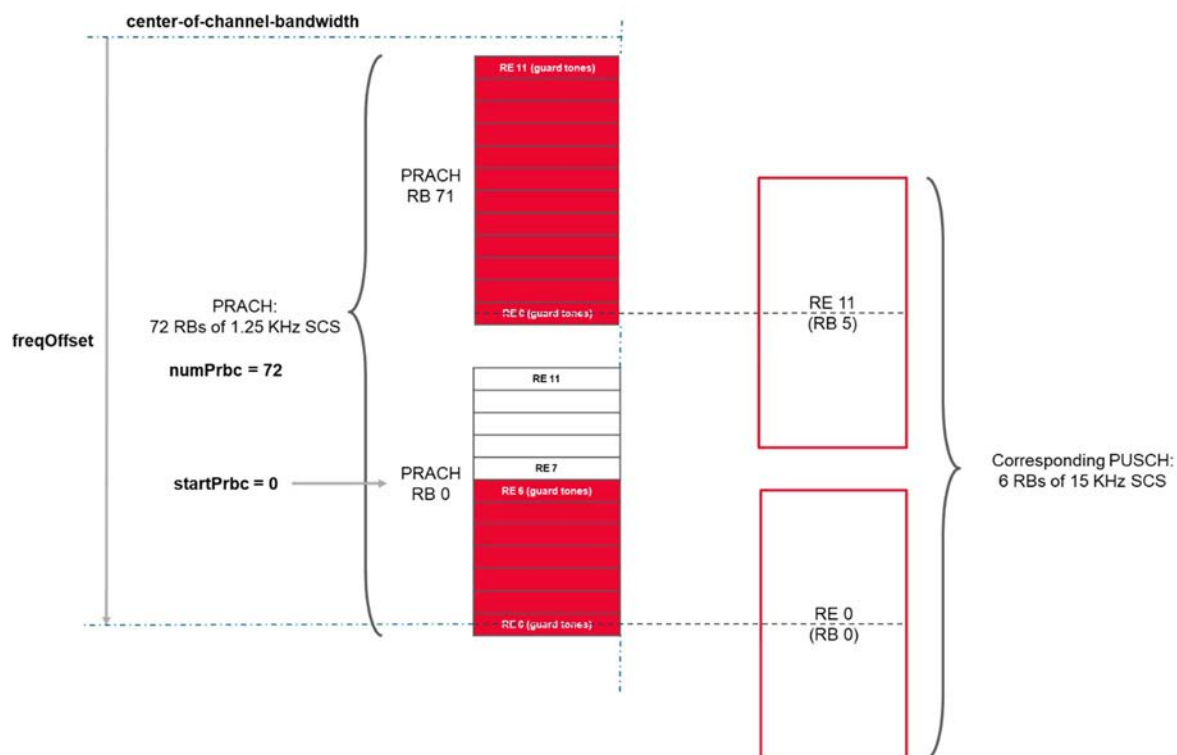



Figure 7.2.3.2-1: startPrbc, numPrbc, and freqOffset Illustration - PRACH example

For each numerology (and PRACH/SSB), the `freqOffset` IE shall determine the frequency offset between center of channel bandwidth (configured via M-Plane) and center of subcarrier corresponding to RE#0 of RB#0 with resolution of half the SCS of the respective numerology. This concept is depicted in Figure 7.2.3.2-2, Figure 7.2.3.2-3 and Figure 7.2.3.2-4.

- The center of channel bandwidth (component carrier center frequency in Hz) is the common reference to all numerologies and PRACH/SSB channels. The center of channel bandwidth shall be configured at carrier setup over M-Plane.
- Frequency offset resolution of half the SCS value allows center of channel bandwidth to be aligned with an RE center, or an RE edge.



NOTE 1: Each frequency block is either requested with separate section descriptions in one ST 3 C-plane message, as separate section descriptions in multiple ST 3 C-plane messages or as a single section description covering multiple frequency blocks in one ST 3 C-plane message.

NOTE 2: In this example, 5G NR 100MHz carrier PUSCH SCS is 30 kHz, PRACH SCS is 1.25 kHz.

NOTE 3: In 5G NR, PRACH frequency blocks are continuous.

NOTE 4: In the example, guard tones are included.

Figure 7.2.3.2-2: RB mapping– PRACH example

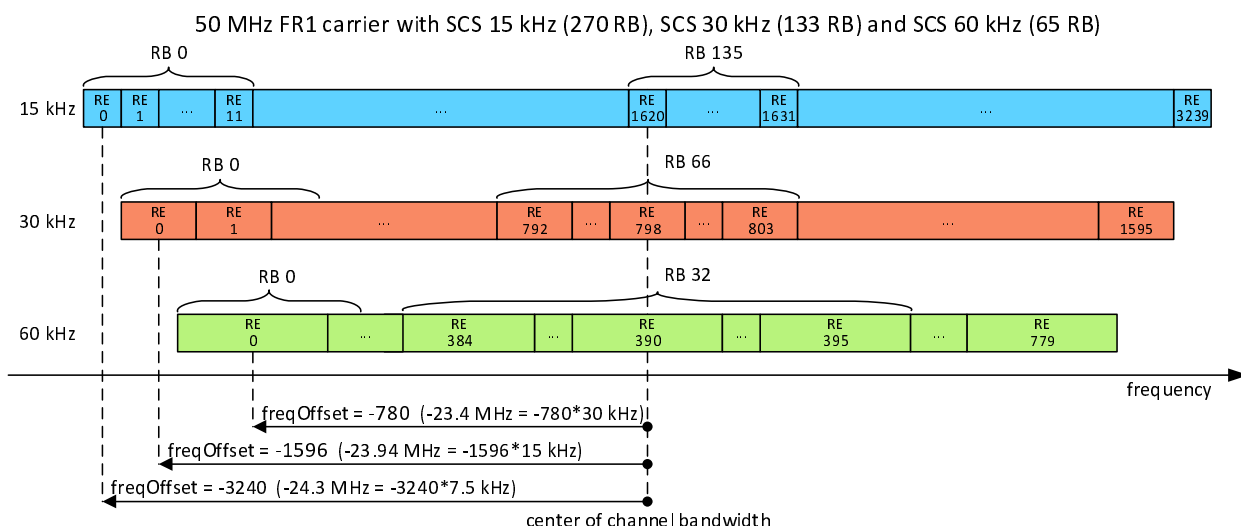


Figure 7.2.3.2-3: RB mapping and support of mixed numerologies - example with mixed numerologies




Figure 7.2.3.2-4: RB mapping and support of mixed numerologies - example SSB

7.2.3.3 Slot indexing with mixed numerologies

For a given frequency range FR1 or FR2, the highest possible numerology supported by the O-RU shall be used as the common reference (separately for FR1 and FR2) for the start of the slot identified by slotId. UL and DL shall use the same reference numerology for slotId. If the highest numerology supported by the O-RU allows both normal and extended CP then normal CP shall be used as reference. Note that the O-RU shall advance slots with extended CP against the reference. The symbol duration and position in time is calculated from the μ value (SCS from frameStructure in message field or configured via M-Plane) and the slotId field in the C/U-Plane message. The value of the sectionId field in C/U-Plane messages addressed per eAxC shall be unique per slot identified by slotId value.

For example, in FR1 case, the maximum slot numbers per subframe is four, when the maximum supported SCS by the O-RU is 60 kHz

- $\mu=0 \rightarrow$ slot #0; $\mu=1 \rightarrow$ slot #0 & #2; $\mu=2 \rightarrow$ slot #0, #1, #2, #3 (see Table 7.2.3.3-1)

Table 7.2.3.3-1: slotId indexing (example 1)

FR 1						FR 2					
$\mu=2$ 60 kHz (Highest supported FR 1 SCS)		$\mu=1$ 30 kHz		$\mu=0$ 15 kHz		$\mu=4$ 240 kHz (Highest supported FR 2 SCS)		$\mu=3$ 120 kHz (FR2)		$\mu=2$ 60 kHz (FR2)	
slotId	symbolId	slotId	symbolId	slotId	symbolId	slotId	symbolId	slotId	symbolId	slotId	symbolId
0	0	0	0	0	0	0	0	0	0	0	0
	1						1		1		
	2						1		2		
	...						2		...		
	11						2		11		
	12						2		12		
1	0	2	11	2	1	1	0	4	0	4	0
	1						11		1		1
	2						12		2		2
	...						13	
	11						13		11		11
	12						13		12		12
2	0	2	0	2	0	8	0	8	0
	1						1		1		1
	2						11		2		2
	...						12	
	11						13		11		11
	12						13		12		12
3	0	2	11	2	12	12	0	8	0	8	0
	1						11		1		1
	2						12		2		2
	...						13	
	11						13		11		11
	12						13		12		12
0	0	0	0	0	0	0	0	14	0	12	0
	1						1		1		1
	2						11		2		2
	...						12	
	11						13		11		11
	12						13		12		12
0	0	0	0	0	0	0	0	14	0	12	0
	1						1		1		1
	2						11		2		2
	...						12	
	11						13		11		11
	12						13		12		12

Table 7.2.3.3-2: slotId indexing (FR1 + FR2 example 2)

FR 2						FR 1					
$\mu=4$ 240 kHz (Highest supported FR 2 SCS)		$\mu=3$ 120 kHz		$\mu=2$ 60 kHz		$\mu=2$ 60 kHz (Highest supported FR 1 SCS))		$\mu=1$ 30 kHz		$\mu=0$ 15 kHz	
slotId	symbolId	slotId	symbolId	slotId	symbolId	slotId	symbolId	slotId	symbolId	slotId	symbolId
0	0	0	0	0	0	0	0	0	0	0	0
	...		1		...		1				
1	0	0	12	0	1	0	1	0	0	0	0
	...		13		...		2				
2	0	2	0	4	0	1	0	...	2	...	1
	...		1		...		11				
3	0	2	12	4	1	1	1	...	2	...	1
	...		13		...		12				
4	0	4	0	4	0	1	0	11	12	...	2
	...		1		...		1				
5	0	4	12	4	1	1	1	11	12	...	2
	...		13		...		2				
6	0	6	0	8	0	2	0	2	13
	...		1		...		11				
7	0	6	12	8	1	2	1	2	13
	...		13		...		12				
8	0	8	0	8	0	2	0	2	13
	...		1		...		1				
9	0	8	12	8	1	2	1	2	13
	...		13		...		2				
...
14	0	14	0	12	0	3	0	12	13	...	12
	...		1		...		11				
15	12	14	12	12	1	3	1	12	13	...	13
	...		13		...		12				
0	0	0	12	0	0	0	0	0	0	0	0
	...		1		...		1				
...

Numerology for slotId shall be based on the highest possible numerology supported by the O-RU. For example:

- If the O-RU supports only FR1 SCSs i.e., 15 kHz, 30 kHz and 60 kHz, it shall be as per the left hand side table of Table 7.2.3.3-1;
- If the O-RU supports only FR2 SCSs i.e., 60 kHz, 120 kHz and 240 kHz, it shall be as per the right hand side table of Table 7.2.3.3-1;
- If the O-RU supports both FR1 SCSs i.e., 15 kHz, 30 kHz, 60 kHz and FR2 SCSs i.e., 60 kHz, 120 kHz and 240 kHz, it shall be as per Table 7.2.3.3-2;

7.2.3.4 PRACH formats with multiple PRACH Symbols

Certain PRACH formats have PRACH preambles constructed from multiple repetitions of a PRACH Symbol, with the Cyclic Prefix (CP) used only with the first PRACH Symbol. Therefore, the O-RU shall be informed how to correctly execute CP extraction and FFT. This shall be achieved by sending a single control message spanning over multiple symbols which reduces the number of C-Plane messages and data sections required.

EXAMPLE: In Figure 7.2.3.4-1, number of symbols = 4, CP length = 0, and time offset duration is adjusted by an equivalent time value of 1152 samples to compensate for setting CP length =0.

Optionally O-RU can advertise the list of the supported formats for the specific endpoint. If the O-RU chooses not to, it means O-RU supports all PRACH formats and the O-DU may select whatever format that is 3GPP compliant with the selected SCS. If the O-RU implements earlier releases of the M-Plane specification which does not support such reporting, then the negotiation of PRACH formats to use shall happen offline, i.e. operators and/or vendors shall ensure that the intended PRACH formats are supported by both O-RU and O-DU.

The example in Figure 7.2.3.4-1 depicts the scenario with format A2 and 30 kHz SCS, in which CP extraction occurs only once, yet there are four associated FFT operations.




Figure 7.2.3.4-1: Example of PRACH format A2

7.2.4 DL precoding configuration parameters and indications

This clause addresses the method of specifying the precoding operation in the Category B O-RU. For Category A O-RU, if there is precoding in the DL processing chain, the precoding shall be implemented in the O-DU so the descriptions in this clause are not relevant. Annex I contains more information and examples regarding this precoding. Support for Transmit Diversity and Spatial Multiplexing with Large Delay CDD precoding schemes is optional for Category B O-RU. If O-RU does not support those precoding schemes, O-DU shall handle precoding of channels using those schemes on O-DU side as defined for Category A O-RU case.

O-DU: 12 REs are generally sent on the interface from the O-DU to the O-RU but for certain IQ formats (selective RE sending as described in Annex A.1), fewer REs may be sent. In this case the missing REs would be considered by the O-RU to be equal to zero in both I and Q.

For Single TX:

For layer mapping at O-RU, complex-valued modulation symbols mapped in a sequence starting with $d^{(q)}(0), \dots, d^{(q)}(M_{\text{symb}}^{(q)} - 1)$ to frequency REs (k,l) and are packed into a PRB. A single eAxC is used for this purpose.

- For CRS mapping at O-RU, all CRS REs belong to a single Tx antenna and are mapped to frequency REs (k,l) for one layer and are packed into a PRB for transmission and are unpacked at O-RU (see Annex I for details).

For Transmit Diversity:

- For layer mapping at O-RU, complex-valued modulation symbols mapped in a sequence starting with $d^{(q)}(0), \dots, d^{(q)}(M_{\text{symb}}^{(q)} - 1)$ to frequency REs (k,l) and are packed into a PRB. A single eAxC is used for this purpose.

- For CRS mapping at O-RU, all CRS RE's are mapped to frequency REs (k,l) for all layers and are packed into a PRB for transmission and are unpacked at O-RU (see Annex I for details).

For Spatial Multiplexing with Large Delay CDD:

- At the O-DU, layer mapped symbols for each layer v , mapped in sequence starting with $x(i) = [x^{(0)}(i) \dots x^{(v-1)}(i)]^T$ to frequency REs (k,l) are packed into each PRB (see Annex I).
- For precoding at the O-RU, different eAxCs are used for each layer (See Annex I).
- For CRS mapping at the O-RU, all CRS RE's for each layer are packed into a PRB for transmission and are unpacked at the O-RU (see Annex I for details).
- All C-Plane message parameters are kept the same for precoding purposes.

For other LTE precoding schemes and NR

- Precoding in the O-RU may be implemented in various ways that are vendor-defined and vendor-specific, especially for LTE TM7-TM10 and NR as there is no 3GPP-mandated precoding operation for these cases.

O-RU: From the C-Plane precoding Section Extension, the O-RU determine the transmission scheme using txScheme field.

For Single TX:

- At the O-RU, input modulated symbols $d^{(q)}(0), \dots, d^{(q)}(M_{\text{symb}}^{(q)} - 1)$ shall be unpacked and used to perform layer mapping, precoding and antenna port mapping for single tx.
- For transmission on a single antenna port, a single layer shall be used, $v=1$, and the mapping is defined as $x^{(0)}(i) = d^{(0)}(i)$ with $M_{\text{symb}}^{\text{layer}} = M_{\text{symb}}^{(0)}$.
- For single tx transmission on a single antenna port, precoding is defined by $y^{(p)}(i) = x^{(0)}(i)$ where $p \in \{0, 4, 5, 7, 8, 11, 13, 107, 108, 109, 110\}$ is the number of the single antenna port used for transmission of the physical channel and $i = 0, 1, \dots, M_{\text{symb}}^{\text{ap}} - 1$, $M_{\text{symb}}^{\text{ap}} = M_{\text{symb}}^{\text{layer}}$.
- For antenna port mapping $p=\{0\}$, each $y(i) = [y^{(p)}(i)]^T$ RE shall go to antenna port $y_p(i)$ after antenna port mapping.
- Since the PRB contains CRS sequences for one antenna port, the RE should extract the CRS RE's using crsSymbolNumber, crsReMask and crsShift (see Annex I for details) and are mapped to the appropriate RE position.

For Transmit Diversity:

- At the O-RU, input modulated symbols $d^{(q)}(0), \dots, d^{(q)}(M_{\text{symb}}^{(q)} - 1)$ shall be unpacked and used to perform layer mapping, precoding and antenna port mapping.
- The appropriate precoder shall be selected based on number of layers and antenna ports.
- For antenna port mapping $p=\{0..N\}$, each $y(i) = [y^{(p)}(i)]^T$ RE shall go to each antenna port $y_p(i)$ after antenna port mapping.
- Since the PRB contains CRS sequences for N antenna ports, the RE should extract the CRS RE's using crsSymbolNumber, crsReMask and crsShift (see Annex I for details) which are mapped to the appropriate RE position and rest of the REs are populated with zero data.

For Spatial Multiplexing with Large Delay CDD:

- Input layer mapped symbols $x(i) = [x^{(0)}(i) \dots x^{(v-1)}(i)]^T$ shall be used to perform precoding at the O-RU based on numLayers, layerID.
- The O-RU shall change the precoder per RE automatically based on the number of antenna ports and number of layers.

- After precoding, for antenna port mapping $p=\{0..N\}$, each $y(i)=[y^{(p)}(i)]^T$ RE shall go to each antenna port $y_p(i)$ after antenna port mapping.
- Since all PRBs contain CRS sequences for N antenna ports, then based on the layerID (layer 0) the O-RU shall extract CRS sequence using `crsSymbolNumber`, `crsReMask` and `crsShift` (see Annex I) for CRS mapping to each of the antenna ports using the `reMask` bit field; the CRS REs from other layers can be ignored.

For other LTE precoding schemes and NR

- One way to implement precoding in the O-RU is via the `beamId` values, wherein a `beamId` points to a beamforming vector that also implements the precoding operation.

7.2.5 LAA commands transfer procedure

7.2.5.1 Introduction

This procedure shall be used to exchange C-Plane messages between O-DU and O-RU. The main purpose of these messages is to support LAA feature in the O-RU/O-DU. See Annex G for more details on the LAA message flow.

7.2.5.2 LBT procedure overview

The LBT procedure shall be used to configure the O-RU with the parameters needed to do LBT prior to PDSCH or DRS transmission OTA. The O-RU shall report the LBT process outcome (either success or failure) in the indication message.

In contrast to licensed spectrum, where the O-RU may continuously send data or reference signals OTA, in unlicensed spectrum, the O-RU can only send discontinuous bursts of data (of length MCOT) or periodic DRS signals. If the O-RU cannot manage a congestion window by itself, the O-DU shall send a `LBT_DL_CONFIG` request message prior to each OTA transmission on the LAA SCell. This includes both DRS and PDSCH (i.e., MCOT burst) transmissions. For every `LBT_DL_CONFIG` request message, the O-RU shall reply with the `LBT_DL` indication messages which include the LBT outcome and the status of the subframe (transmitted/dropped) (see Figure 7.2.5.2-1).

If the O-RU can manage a congestion window by itself, the O-DU does not need to send `LBT_DL_CONFIG` request message to start the LBT procedure at the O-RU. Because the O-RU can determine the start timing of LBT procedure, the requirements of the O-DU for the LBT procedure at the O-RU is a delivery of needed information to adjust a congestion window managed by the O-RU. To transmit this information, the O-DU shall send `LBT_CWCONFIG_REQ` message, and the O-RU shall notify the O-DU by `LBT_CWCONFIG_RSP` message to indicate that the `LBT_CWCONFIG_REQ` message is successfully received or not. Also, the O-RU can adjust its congestion window based on the information on the `LBT_CWCONFIG_REQ` message. For every `LBT_CWCONFIG_REQ` request message, the O-RU shall reply with the `LBT_DL` indication message(s) which include the LBT outcome and the status of the subframe (transmitted/dropped).

For the DRS signal, initially, the O-DU shall configure the O-RU with the DRS parameters (DMTC period and DMTC offset) via the M-Plane.

Before every DRS period, the O-DU shall send the `LBT_DL_CONFIG` request message and the DRS signal. The O-RU shall do the LBT procedure and transmit the DRS signal OTA in case of LBT success. The O-RU then shall send the `LBT_DL` indication.

The DRS signal can be transmitted over the fronthaul interface similar to other reference signals such as the PSS/SSS/CRS.

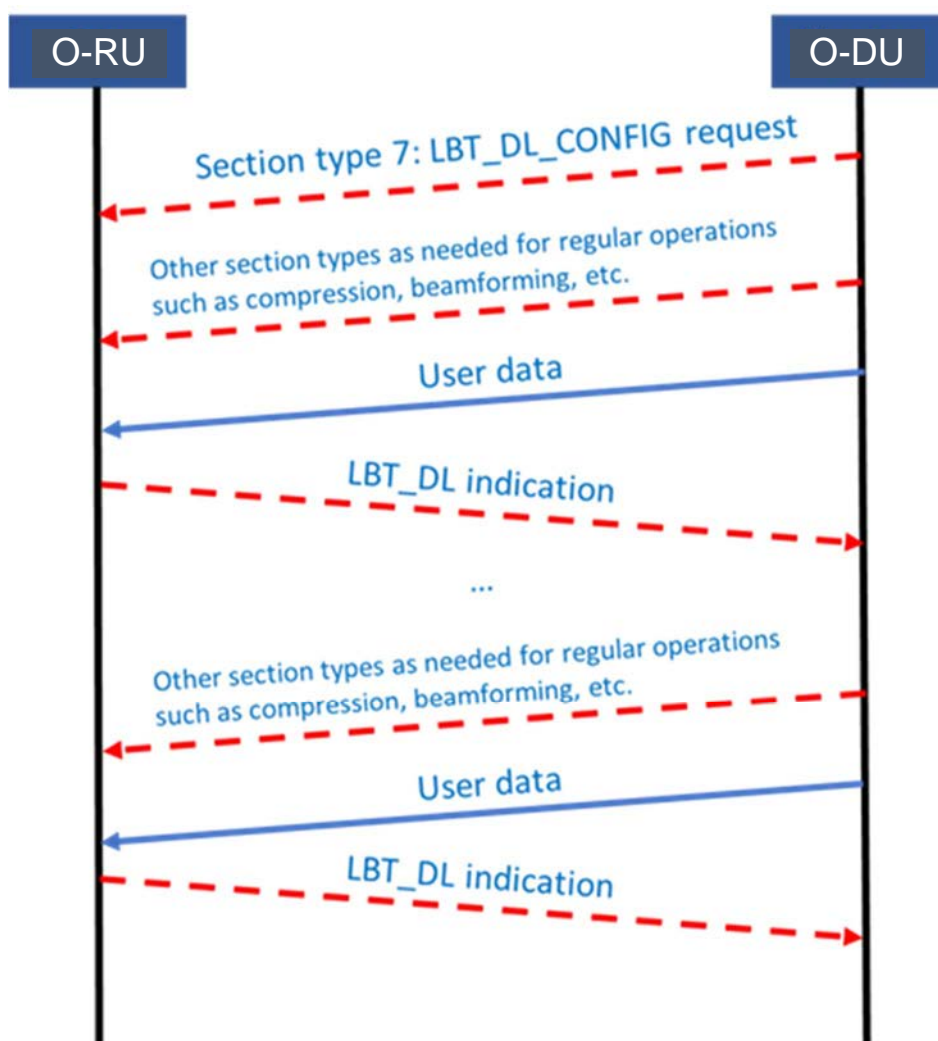


Figure 7.2.5.2-1: LBT message flow

7.2.5.3 Definitions

- D_{OW} : Maximum one-way fronthaul latency.
- P_{DU} : Maximum O-DU processing time (reading an upstream O-RU C-Plane message, MAC processing, sending the downstream U-Plane data).
- P_{RU} : Maximum O-RU processing time (sending an upstream C-Plane message, receiving downstream U-Plane data, transmitting U-Plane data OTA).
- B_{RU} : O-RU Buffer depth (for LAA) ceiled to an integer number of OFDM symbols in microseconds (up to 140 symbols, for example (i.e., MCOT)). This is equivalent to the minimum amount of data that shall be buffered at the O-RU. This parameter could be less than or equal to the actual O-RU buffer size communicated via the M-Plane.
- Expired symbol: Symbol where the current time at the O-RU exceeds its target transmission time (i.e., SFN/SF) OTA.
- Normal (partially-filled) SF assumption: T_{start} is the start of the transmission time, which occurs at the first SF (slot) boundary the O-RU encounters after starting the LBT CAT 4 process and after finishing the minimum sensing duration. Initially known to the O-DU (since it sends the LBT start time and the LBT parameters to the O-RU). It can be updated in real-time based on the LBT outcome and the communication between the O-DU and O-RU. Here, a "slot" follows the LTE definition of containing seven symbols.

- LBT_DL indication Threshold (LBT_{th}): The latest time the O-DU is expected to receive the LBT_DL indication message from the O-RU.

$$LBT_{th} = T_{start} + (P_{RU} + D_{OW} + P_{DU})$$

- CU Transmission time threshold (CUT_{th}): The time at which the O-DU is required to send the data to the O-RU to be transmitted OTA at time T_{start} .

$$CUT_{th} = T_{start} - (P_{RU} + D_{OW} + P_{DU})$$

- x : Minimum time between any two potential start transmission times.

$$x = \begin{cases} 1 \text{ ms}, & \text{Normal SF} \\ 0.5 \text{ ms}, & \text{Partially - filled SF} \end{cases}$$

7.2.5.4 General guidelines for the LAA-procedure

7.2.5.4.1 PDSCH Transmission

- O-DU should avoid buffer overflow or underflow at the O-RU:
 - O-DU should only send B_{RU} worth of data to the O-RU for every single transmission
 - O-DU should plan for the data to be received at the O-RU only P_{RU} before the actual OTA transmission time at the O-RU
- O-RU should avoid buffer overflow or underflow at the O-RU:
 - O-RU should flush its buffer (by dropping any expired symbols) as soon as any symbol becomes expired (i.e., current time is larger than the symbol's scheduled time)
 - O-RU should immediately send a success LBT indication to the O-DU once the channel is acquired.
 - The O-RU should send an error message to the O-DU if data received is larger than its local buffer
 - The O-RU should send a subframe drop or transmission message to the O-DU when buffered subframe is dropped because scheduled time is passed, or is transmitted after LBT success.
- The O-RU should have a buffer that satisfies the following equation: $B_{RU} \geq 2 \times D_{OW} + P_{DU} + P_{RU}$
- LBT CAT 4 considerations:
 - The O-DU can configure the O-RU (via the M-Plane) with the threshold on the LBT CAT 4 duration (e.g., 8 ms). Once this threshold has exceeded, the O-RU sends a failure LBT indication to the O-DU, which in return sends back a new LBT config request. The O-RU can then restart the LBT CAT 4 process.
- The data signal may be received at the O-RU before or after the LBT_PDSCH_REQ
- The O-RU should generate the reservation signal locally whenever needed

7.2.5.4.2 DRS Transmission

- Recall:
 - a. PSS/SSS being part of a DRS may occur outside subframe 0 and 5
 - b. CRS/CSI-RS/PSS/SSS do not vary with subframe number but are kept unchanged across subframes 0–4 (call it DRS_v1) and 5–9 (DRS_v2).
- Assumption: DRS OTA transmission starts at the SF boundaries.
- The DRS signal may be received before or after the LBT_DRS_REQ.

- LBT_DRS_REQ is sent once per DRS window (e.g., DMTC offset = 6 ms is conveyed to the O-RU via the M-Plane).
- DRS signal is sent every SF in the 6 SFs of the DRS window until LBT succeeds.
- $DRS_{th,1} = SF0 - (P_{RU} + D_{OW} + P_{DU})$.
- $DRS_{th,2} = SF5 - (P_{RU} + D_{OW} + P_{DU})$.
- SF0: The start of SF0, SF1, SF2, SF3, or SF4.
- SF5: The start of SF5, SF6, SF7, SF8, or SF9.
- At $DRS_{th,1}$, O-DU sends DRS_v1 (i.e., to be transmitted within SFs 0-4).
- At $DRS_{th,2}$, O-DU sends DRS_v2 (i.e., to be transmitted within SFs 5-9).

NOTE: LBT indication with failure outcome is sent for every sensing period.

7.2.5.4.3 Congestion window information transmission

- O-DU should send information regarding a congestion window adjustment to O-RU:
 - HARQ feedback information for the reference subframe and number of TB are included.
- O-RU should adjust its managed congestion window value and notify to O-DU the packet reception status:
 - O-RU should adjust its congestion window value based on received information.
 - O-RU should immediately send a success LBT_CWCONFIG_RSP to the O-DU when the message is successfully received. If not, O-RU should send a fail LBT_CWCONFIG_RSP.

7.2.6 Dynamic Spectrum Sharing (DSS)

7.2.6.1 Overview

Dynamic Spectrum sharing allows different technologies – 4G LTE and 5G NR in the present document, to share same frequency carrier dynamically, i.e., the O-DU may allocate one or more set(s) of frequency-time resources (or PRBs) to LTE or NR dynamically. The O-RU shall support this feature for the O-DU to take advantage of such an allocation. The present document supports DSS via using dedicated eAxC_ID (i.e., endpoints) for LTE and NR or via using Section Extension 9 for DSS.

7.2.6.2 Dynamic Spectrum Sharing (DSS) via dedicated endpoints

When using dedicated endpoints for DSS, dedicated endpoints shall be configured in the O-RU for LTE and NR (hence dedicated eAxC_ID) to implement DSS. For example, when an O-DU allocates a set of PRBs for LTE, the O-DU shall use an eAxC_ID dedicated for LTE (assigned to a carrier identified as an LTE carrier) to address the LTE endpoint; at a different time when the O-DU allocates the same set of PRBs for NR, the O-DU shall use an eAxC_ID dedicated for NR (assigned to a carrier identified as an NR carrier) to address the NR endpoint in the O-RU.

When using this approach, a carrier will be assigned to LTE or NR, so when both LTE and NR are to be supported at the same time ("DSS"), two carriers with overlapping frequency ranges are to be identified, one being LTE and the other being NR. It will be up to the O-RU to overlap the two carriers such that the same frequency range can support both NR and LTE, and it will be up to the O-DU to assure that no RE is assigned to both LTE and NR at the same time.

7.2.6.3 Dynamic Spectrum Sharing (DSS) via Section Extension 9 for DSS

Using Section Extension 9 for DSS, the O-DU shall use a single eAxC_ID (and hence a single endpoint) for conveying LTE and NR PRB allocations if the O-RU is capable of handling both LTE and NR processing using one endpoint (i.e., the O-RU indicates support for Section Extension 9). A single carrier shall be allocated to either LTE, NR or both, and

supporting DSS in this manner means the carrier (to which the endpoint is assigned) shall be designated as supporting both. The O-DU shall indicate whether the PRB allocation information is applicable to LTE or to NR via Section Extension 9.

7.2.7 Channel information based beamforming

In the channel information based beamforming method, the O-DU shall provide channel information per UE periodically (generally less often than every slot) using Section Type 6 C-Plane message(s) and then on a slot-by-slot basis the O-DU shall provide scheduling information using Section Type 5 C-Plane message which the O-RU shall use along with the channel information to calculate the proper beamforming weights for the specific slot with its co-scheduled UEs. To reduce peak bandwidth, the transmission/ reception window constraints shall not apply to a Section Type 6 message. The O-RU is expected to use the latest possible available channel information for a ueId for channel information based beamforming operation i.e., the O-RU uses channel information for a ueId available to the O-RU at the end of receive window of a Section Type 5 message. The O-RU should store updated channel information for a ueId when it is received, to be used when that ueId is scheduled in future.

7.2.8 ACK/NACK Procedure for C-Plane Messages

If an O-RU reports its capability of supporting both Section Extension 22 and Section Type 8, then the O-DU may append a Section Extension 22 with an ackNackReqId to the section description to request for ACK/NACK feedback from the O-RU. SE 22 can be used with all Section Types and Section Extensions sent from O-DU to O-RU, except with Section Type 7. The O-DU shall not request O-RU to send from an endpoint more ACK/NACK feedbacks in a symbol than endpoint's capability reported by O-RU in M-plane parameter **max-ack-nack-per-symbol**; that is, number of SE 22 instances in all C-plane messages for one endpoint that address the same symbol with frameId, subframeId, slotId and startSymbolId shall not exceed the value of **max-ack-nack-per-symbol**. The O-RU may also report the capability of **nack-supported**; if it is reported as 'false', the O-DU shall assume that NACK feedback can not be reported by the O-RU and that only ACK feedback can be reported. The O-RU may also report the capability of max number of Section Extension 22 (ACK/NACK request) supported in a C-Plane message by **max-num-se22-per-cplane-message**. The O-DU shall not include more Section Extension 22 in a C-Plane message than the reported value.

If the O-RU receives one or more section description(s) with Section Extension 22, it shall send an ACK/NACK feedback C-Plane message (Section Type 8) to the O-DU, which contains the ACK feedback including ackId(s) for the section description(s) correctly received and correctly processed and not dropped. If the O-RU supports sending NACK feedback (nack-supported = true) and Section Extension 22 is present within the section description containing an O-RAN protocol error, then the O-RU shall drop the section description and send nackId(s) for the section description to the O-DU using Section Type 8. O-RU may send one or more ackIds and/or nackIds in one ST 8 message. The O-RU shall send ACK/NACK messages with frameId, subframeId slotId and symbolId set to values of frameId, subframeId, slotId and startSymbolId respectively received in the C-plane message with the Section Extension 22.

The O-RU shall send ACK/NACK messages within the ACK transmit window, and the O-DU is expected to receive ACK/NACK messages within the ACK receive window. The ACK transmit/receive window is similar with UL U-Plane transmit/receive window as defined in clause 4.4.3, specifically as Figure 4.4.3-2, with the difference that the parameters Ta3_min_ack and Ta3_max_ack can be negative values, i.e. Ta3_min_ack takes the place of Ta3_min and Ta3_max_ack takes the place of Ta3_max.

The O-DU populates a unique identifier for a section description in ackNackReqId in Section Extension 22 for the O-RU to use to report the ACK/NACK status. The O-RU shall indicate the ackNackReqId of correctly decoded section description in ackId field. The O-RU shall indicate the ackNackReqId of the incorrectly decoded section description in the nackId field of the ACK/NACK feedback message, if the O-RU supports sending NACK feedback.

NOTE: Section descriptions after an erroneous and dropped section description might be impacted e.g., dropping of a symInc flag, and it is expected an O-DU, knowing a section description was dropped via receiving a NACK, will understand the impact on subsequent section descriptions.

By the end of ACK receive window for a symbol M (where M is the value of startSymbolId in the C-plane message which carries SE 22), if the O-DU receives the ACK/NACK feedback message which contains the NACK (NACK means one of the nackId(s) equals to the ackNackReqId in Section Extension 22) for a section description, then the O-DU may interpret that there is an error in the section description sent for symbol M, and may retransmit the section description by reconstructing the information. If neither NACK nor ACK (ACK means one of the ackId(s) equals to the ackNackReqId in Section Extension 22) for a section description is received by the O-DU by the end of the ACK receive window for symbol M, then the O-DU may interpret that either the C-plane message containing section description with SE 22, for symbol M was lost or section descriptions within the C-Plane message sent for symbol M

was dropped before Section Extension 22 is decoded, and the O-DU may choose to resend the lost section description and if so, shall appropriately update any required fields to ensure correct reception at the O-RU

In addition, if the O-DU receives the NACK message for C-Plane message(s) from O-RU before sending the corresponding U-Plane message(s), it may skip sending the U-Plane message(s) to reduce the fronthaul traffic. [Figure 7.2.8-1](#) shows an example of this ACK/NACK procedure.

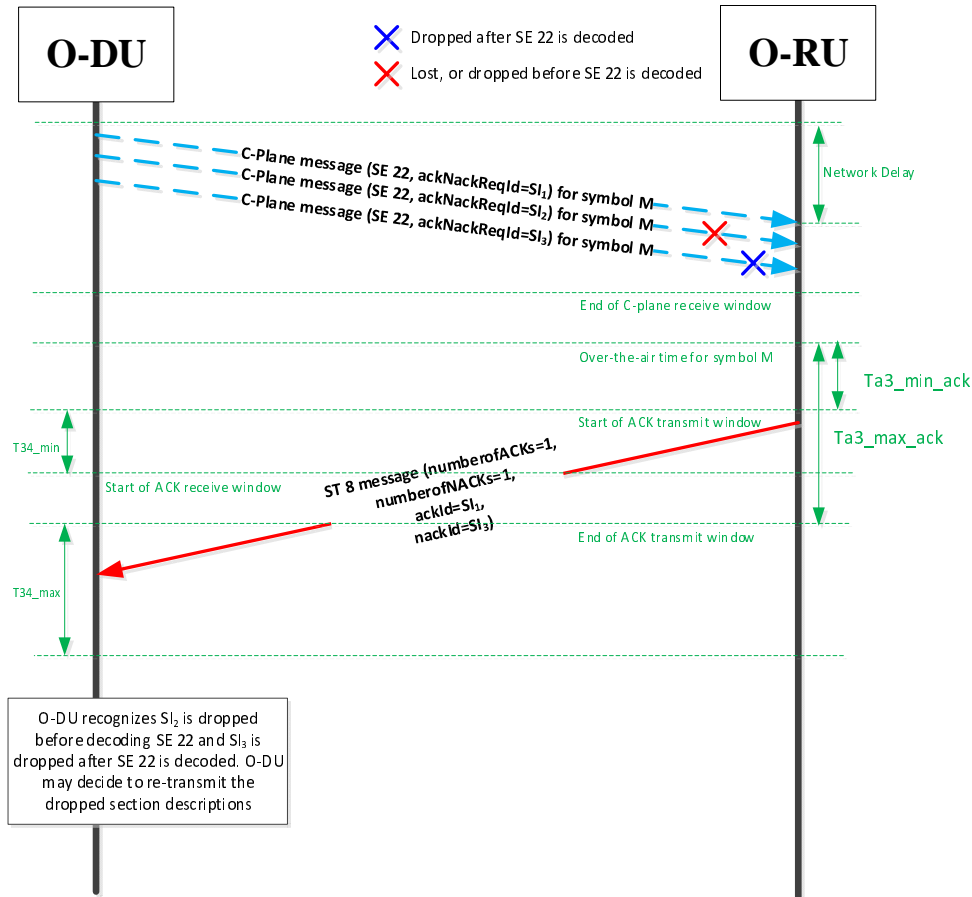


Figure 7.2.8-1: Example of ACK/NACK procedure

7.2.9 Section Type 4: Slot Configuration Control

7.2.9.1 Overview

Section Type 4 C-Plane message is used to send Slot Level configuration from the O-DU to the O-RU which may apply to all endpoints belonging to a carrier, all endpoints associated with a tx-array or rx-array or all the endpoints in the O-RU. Each Section Type 4 C-Plane message comprises of Section Type 4 common header followed by one or more Section Type 4 command type headers. Both Section Type 4 header and each command type shall be 4-byte aligned. 4-byte alignment shall also be applicable when multiple commands are present in a Section Type 4 C-Plane message. Each command type shall comprise of a common header specifying `st4CmdType`, `st4CmdLen`, `numSlots` for which the configuration specified by the command applies, and `ackNackReqId` field. The common part of the command header is followed by command configuration-specific fields.

NOTE: Loss of ST 4 message especially with `numSlots` other than 1 (see clause 7.5.3.40) can have a lasting negative impact on the O-RU operation.

7.2.9.2 Section Type 4 commands

7.2.9.2.1 TIME_DOMAIN_BEAM_CONFIG (Time-domain beamforming command)

This command is defined to send time-domain beamforming information from the O-DU to the O-RU. This command may be used by the O-DU for sending time-domain beamforming information to the O-RU and may also be used in the case of hybrid beamforming to send time-domain beamforming information separate from frequency-domain beamforming information. Refer to clause 12.4.3.4.4 to understand how redundancy is avoided when this command is used with Section Type 4.

This command provides the flexibility for sending time-domain beamforming weights. Here are some of the variants of this command configuration:

- Specify time-domain beam group with associated time-domain beam number and beamforming weights per time-domain beam number. Refer to Table 7.4.6-3.
- Specify time-domain beam group with time-domain beam number vector. Refer to Table 7.4.6-4.
- Specify time-domain beam group only. Refer to Table 7.4.6-5.

Because the values of beamforming weights never expire, the value of numSlots should be set to zero for this command type and the O-RU shall apply the beamforming configuration until the configuration is changed by a subsequent message.

The C-Plane message with Section Type 4 TIME_DOMAIN_BEAM_CONFIG command shall apply to the direction specified by dataDirection bit in the Section Header.

7.2.9.2.2 TDD_CONFIG_PATTERN (TDD configuration command)

This command may be used to send TDD UL, DL, and guard symbol configuration as an alternative to using the dataDirection bit in the C-plane message. It enables support of semi-static and dynamic TDD use cases. Using this command enables the O-RU to know TDD configuration for an entire slot and also a set of consecutive slots (if 'numSlots > 0' refer to clause 7.5.3.40). If 'numSlots' in the Section Type 4 header is set to the value of zero, TDD_CONFIG_PATTERN sent in the respective 'slotID' shall last till the next Section Type 4 message with modified TDD_CONFIG_PATTERN is sent to the O-RU. If 'numSlots' is set to a non-zero value, the TDD configuration shall only last for the value specified by numSlots.

If Section Type 4 C-Plane message with TDD_CONFIG_PATTERN command is followed by C-Plane message using Section Type 0, 1, 3, 5, 6 which specifies symbol direction using 'dataDirection' bit, O-DU shall set 'dataDirection' bit direction to align with TDD configuration specified in TDD_CONFIG_PATTERN. Refer to clause 12.4.3.4.4 for details.

The C-Plane message with Section Type 4 TDD_CONFIG_PATTERN command shall apply to both UL and DL (as specified in Table 7.5.2.17-1) regardless of the direction specified by the dataDirection bit in the section header.

7.2.9.3 ACK/NACK support

If the O-DU requests the O-RU to provide ACK/NACK for a given command and if the O-RU supports Section Type 8 C-Plane message, O-DU shall set 'ackNackReqId' to a non-zero value. Supporting ACK/NACK for Section Type 4 is similar to using SE 22 in conjunction with Section Type 8 with two differences:

- 1) Instead of using SE 22 to request Section Type 8 C-Plane message from O-RU, the Section Type 4 message command common header has the field 'ackNackReqId' to request for Section Type 8 message from the O-RU.
- 2) The granularity of requesting ACK/NACK Section Type 8 C-Plane message from O-RU is per Section Type 4 command instead of per section header when SE 22 is used with ST 0, 1, 3, 5, 6.

7.2.9.4 O-RU handling

If the O-RU supports Section Type 4, support for this Section Type and associated command type shall be indicated by the O-RU on a per endpoint basis using M-Plane parameters for Section Type 4. Each Section Type 4 'cmdScope' value shall be associated with a bitmask 'st4-reception-mask' defined in M-Plane. The C-Plane reception window for Section

Type 4 C-Plane messages shall be determined from the bits set to '1' in the 'st4-reception-mask'. The value of 'startSymbolId' shall be equal to the symbol number set to '1' in the M-Plane bitmask 'st4-reception-mask'. If the O-RU receives a Section Type 4 message which is not aligned with 'st4-reception-mask', the O-RU shall drop the Section Type 4 message. If a Section Type 4 command is invoked more than once within a slot for the same 'cmdScope', the value of 'startSymbolId' for each invocation shall be equal to the corresponding symbol set to '1' in per command scope 'st4-reception-mask'.

7.2.9.5 Latency requirements

To ensure adequate time for processing a Section Type 4 message with a given 'cmdScope' value, the O-DU shall send the message to the Section Type 4 supporting endpoint with the earliest reception window (largest value 'T2a_min_cp_ul/dl') of all those endpoints impacted by the message. In the case of multiple O-DUs sharing RF circuits within a single O-RU, a specific O-DU may need to send the Section Type 4 message to assure the rx-window is sufficiently early to allow all endpoints to correctly process the message.

7.3 Elements for the C-Plane protocol

7.3.1 Section Types

C-Plane messages shall be encapsulated using a two-layered header approach. The first layer consists of an eCPRI common header or IEEE 1914.3 common header, including corresponding fields used to indicate the message type, while the second layer is an application layer including necessary fields for control and synchronization. Within the application layer, a "section" defines the characteristics of U-Plane data to be transferred or received from a beam with one pattern id. In general, the transmission header, application header, and sections shall be aligned on 4-byte boundaries (but see the exception for Section Type 6 below) and are transmitted in "network byte order" meaning the most significant byte of a multi-byte parameter is transmitted first.

When the optional "little endian byte order" is chosen via M-Plane, the beamforming weights (Section Type 1,3 or 4, and Section Extension 1, 11 or 19) and ciIsample/ciQsample (Section Type 6) fields shall use little endian byte order to transmit the complex numbers. Annex D.2 shows little endian byte order format for various IQ data bitwidths.

Table 7.3.1-1 describes the Section Types that are supported within the C-Plane.

Table 7.3.1-1: Section Types

Section Type	Target Scenario	Remarks
0	Unused Resource Blocks or symbols in Downlink or Uplink	Indicates to the O-RU that certain Resource Blocks or symbols will not be used (idle periods, guard periods). Likewise, there are no associated U-Plane messages containing IQ data for this Section Type. The purpose is to inform the O-RU that transmissions may be halted during the specified idle interval for e.g. power-savings or to provide an interval for calibration.
1	Most DL/UL radio channels (NOTE)	Here "most" refers to channels not requiring time or frequency offsets such as are needed for mixed-numerology channels
2	Reserved for future use	
3	PRACH and mixed-numerology channels (NOTE)	Channels requiring time or frequency offsets or different-than-nominal SCS values
4	Slot Configuration Control	Slot configuration for multiple eAxC_IDs with one or multiple Section Type 4 configuration commands
5	UE scheduling information (ueld assignment to section)	Provides scheduling information for uelds
6	Channel information	Sends UE-specific channel information from the O-DU to the O-RU
7	LAA	Messages communicated between O-DU and the O-RU in both directions to configure LBT for PDSCH/DRS transmission and to report the LBT outcome.
8	ACK/NACK Feedback	Sent from the O-RU to the O-DU, providing ACK/NACK feedback for section descriptions in C-Plane messages
9-255	Reserved for future use	
NOTE:	When PRACH having same numerology as other UL channel, Section Type 1 can alternatively be used by O-DU for PRACH signaling. In this case, O-RU is not expected to perform any PRACH specific processing.	

By default, Section Type 6 is 1-byte aligned instead of 4-byte-aligned. In the present document, the O-RU may optionally support a boolean flag "st6-4byte-alignment-required" to indicate support for 4-byte aligned Section Type 6 as part of M-Plane O-RU capability exchange. If O-RU sets "st6-4byte-alignment-required=true", the O-DU shall provide Section Type 6 messages according to [Table 7.4.8-1](#) with the octets shown as "<padding with zeros to the next 4-byte boundary>" included.

However, if O-RU sets boolean flag "st6-4byte-alignment-required=false" or the O-DU does not receive this boolean flag, then the O-DU shall provide Section Type 6 messages according to [Table 7.4.8-1](#) with the octets shown as "<padding with zeros to the next 4-byte boundary>" not included.

The parameter value true or false for "st6-4byte-alignment-required" applies to all Section Type 6 messages that O-DU may send to the O-RU.

The network operator should ensure that the O-DU in the network is prepared to conform to the possible presence of an O-RU's M-Plane capability flag "st6-4byte-alignment-required", or the network operator should ensure that the O-RUs connected to an O-DU that is not prepared to conform to the O-RU's 4-byte aligned section type 6 capability indicated by "st6-4byte-alignment-required" flag will properly process 1-byte-aligned Section Type 6 messages.

7.3.2 Section Extensions

Within a data section description, the section header may include an "extension flag" which indicates parameters that apply to the section beyond those within the section header (the extension flag is often adjacent to the beamId or ueId fields). The presence of this extension flag indicates that following the header, a Section Extension is present. Within the Section Extension is another extension flag, adjacent to the "extType" field, which indicates that a second (or third, etc.) extension is present. In this way any number of Section Extensions may be included within a data section. This provides an extensibility for section parameters without the need to continually redefine the section header or create new Section Types to accommodate future fronthaul specification needs.

The Section Extension takes the form reminiscent of "TLV": there is a type field specifically "extType" which is a 7-bit field describing the specific extension type, then a one-byte (sometimes 2 bytes) length field specifically "extLen" detailing how many 4-byte words are contained in the extension (minimum of 1 including the "extType" and "extLen" fields), and then some number of parameters that depend on the "extType" value for their definition. In all cases the Section Extension shall be an integer number of (4-byte) words in length.

Specific valid values of "extType", their meanings and their associated parameters are detailed in clause 7.7.

7.4 Section Type elements

7.4.1 Overview

A common frame format shall be used for C-Plane messages, consisting of a transport layer (see clause 5.1.3) and an application layer. The application layer is within the transport layer payload and shall consist of a common header for time reference, followed by information and parameters dependent and specific to the Section Type in use. Multiple sets of section data of the same Section Type value can be lined up one after another within the payload. To minimize packet rate over the interface, transmitter should fill messages with as many subsequent sections (with or without sequential section IDs) as possible. However, sets of section data of different Section Type values shall be sent via separate messages (i.e. different values of Section Type shall not be mixed within a single C-Plane message payload). When using Section Types, the formats shown in following clauses shall be used.

7.4.2 Section Type 0 elements

Section Type 0 is used for indicating idle or guard periods from O-DU to O-RU; see Table 7.4.2-1

- Common Header Fields:
 - **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
 - **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
 - **filterIndex** (filter index) field: 4 bits.
 - **frameId** (frame identifier) field: 8 bits.

- **subframeId** (subframe identifier) field: 4 bits.
 - **slotID** (slot identifier) field: 6 bits.
 - **startSymbolId** (start symbol id) field: 6 bits.
 - **numberOfSections** (number of sections) field: 8 bits.
 - **sectionType** (Section Type) field: 8 bits:
 - value = 0 shall be set.
 - **timeOffset** (time offset) field: 16 bits.
 - **frameStructure** (frame structure) field: 8 bits.
 - **cpLength** (cyclic prefix length) field: 16 bits.
 - **reserved** (reserved for future use) field: 8 bits.
- Section Fields:
- **sectionId** (section identifier) field: 12 bits.
 - **rb** (resource block indicator) field: 1 bit.
 - **symInc** (symbol number increment command) field: 1 bit.
 - **startPrbc** (starting PRB of data section description) field: 10 bits.
 - **numPrbc** (number of contiguous PRBs per data section description) field: 8 bits.
 - **reMask** (resource element mask) field: 12 bits.
 - **numSymbol** (number of symbols) field: 4 bits.
 - **ef** (extension flag) field: 1 bit.
 - **reserved** (reserved for future use) field: 15 bits.

Table 7.4.2-1: Scheduling and beamforming commands frame format (Section Type 0)

Section Type 0: idle / guard periods										
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
transport header, see clause 5.1.3									8	Octet 1
dataDirection	payloadVersion			filterIndex				1	Octet 9	
frameId									1	Octet 10
subframeId			slotId						1	Octet 11
slotId		startSymbolId							1	Octet 12
numberOfSections									1	Octet 13
sectionType = 0									1	Octet 14
timeOffset									2	Octet 15
frameStructure									1	Octet 17
cpLength									2	Octet 18
Reserved									1	Octet 20
sectionId									1	Octet 21
sectionId			rb	symInc	startPrbc				1	Octet 22
startPrbc									1	Octet 23
numPrbc									1	Octet 24
reMask[11:4]									1	Octet 25
reMask[3:0]			numSymbol						1	Octet 26
ef	reserved (7 bits)								1	Octet 27
reserved (8 bits)									1	Octet 28
Section Extensions as indicated by "ef" if any									var	Octet 29
...										
sectionId									1	Octet N
sectionId			rb	symInc	startPrbc				1	N+1
startPrbc									1	N+2
numPrbc									1	N+3
reMask[11:4]									1	N+4
reMask[3:0]			numSymbol						1	N+5
ef	reserved (7 bits)								1	N+6
reserved (8 bits)									1	N+7
Section Extensions as indicated by "ef" if any									var	N+8

NOTE: Shading: yellow is transport header, pink is radio application header, others are repeated sections.

7.4.3 Section Type 1 elements

Section Type 1 is used for most of the Downlink and Uplink radio channels – some channels especially PRACH and mixed-numerology channels may need more information elements contained in other Section Types. However, if Section Type 1 is used for PRACH channel, it should be processed in the same manner as a non-PRACH channel is processed in O-RU; see Table 7.4.3-1

○ Common Header Fields:

- **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
- **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
- **filterIndex** (filter index) field: 4 bits.
- **frameId** (frame identifier) field: 8 bits.
- **subframeId** (subframe identifier) field: 4 bits.
- **slotId** (slot identifier) field: 6 bits.
- **startSymbolId** (start symbol id) field: 6 bits.
- **numberOfSections** (number of sections) field: 8 bits.

- **sectionType** (Section Type) field: 8 bits:
 - value = 1 shall be set.
 - **udCompHdr** (user data compression header) field: 8 bits.
 - **reserved** (reserved for future use) field: 8 bits.
- Section Fields:
- **sectionId** (section identifier) field: 12 bits.
 - **rb** (resource block identifier) field: 1 bit.
 - **symInc** (symbol number increment command) field: 1 bit.
 - **startPrbc** (starting PRB of data section description) field: 10 bits.
 - **numPrbc** (number of contiguous PRBs per data section description) field: 8 bits.
 - **reMask** (resource element mask) field: 12 bits.
 - **numSymbol** (number of symbols) field: 4 bits.
 - **ef** (extension flag) field: 1 bit.
 - **beamId** (beam identifier) field: 15 bits.

Table 7.4.3-1: Scheduling and beamforming commands frame format (Section Type 1)

Section Type 1: DL/UL control msgs											
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes			
transport header, see clause 5.1.3									8	Octet 1	
dataDirection	payloadVersion			filterIndex						1	Octet 9
frameId									1	Octet 10	
subframeId				slotId						1	Octet 11
slotId		startSymbolId							1	Octet 12	
numberOfSections									1	Octet 13	
sectionType = 1									1	Octet 14	
udCompHdr									1	Octet 15	
reserved									1	Octet 16	
sectionId									1	Octet 17	
sectionId			rb	symInc	startPrbc				1	Octet 18	
startPrbc									1	Octet 19	
numPrbc									1	Octet 20	
reMask[11:4]									1	Octet 21	
reMask[3:0]			numSymbol						1	Octet 22	
ef = 1	beamId[14:8]								1	Octet 23	
beamId[7:0]									1	Octet 24	
Section Extensions as indicated by "ef"									var	Octet 25	
...											
sectionId									1	Octet N	
sectionId			rb	symInc	startPrbc				1	N+1	
startPrbc									1	N+2	
numPrbc									1	N+3	
reMask[11:4]									1	N+4	
reMask[3:0]			numSymbol						1	N+5	
ef = 0	beamId[14:8]								1	N+6	
beamId[7:0]									1	N+7	
Section Extensions as indicated by "ef"									var	N+8	

NOTE: Shading: yellow is transport header, pink is radio application header, others are repeated sections

7.4.4 reserved (reserved for future use)

7.4.5 Section Type 3 elements

Section Type 3 is used for PRACH and mixed-numerology channels; see Table 7.4.5-1:

- Common Header Fields:
 - **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
 - **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
 - **filterIndex** (filter index) field: 4 bits.
 - **frameId** (frame identifier) field: 8 bits.
 - **subframeId** (subframe identifier) field: 4 bits.
 - **slotID** (slot identifier) field: 6 bits.
 - **startSymbolId** (start symbol identifier) field: 6 bits.
 - **numberOfSections** (number of sections) field: 8 bits.
 - **sectionType** (Section Type) field: 8 bits.
 - value = 3 shall be set.
 - **timeOffset** (time offset) field: 16 bits.
 - **frameStructure** (frame structure) field: 8 bits.
 - **cpLength** (cyclic prefix length) field: 16 bits.
 - **udCompHdr** (user data compression header) field: 8 bits.
- Section Fields:
 - **sectionId** (section identifier) field: 12 bits.
 - **rb** (resource block identifier) field: 1 bit.
 - **symInc** (symbol number increment command) field: 1 bit.
 - **startPrbc** (starting PRB of data section description) field: 10 bits.
 - **numPrbc** (number of contiguous PRBs per data section description) field: 8 bits.
 - **reMask** (resource element mask) field: 12 bits.
 - **numSymbol** (number of symbols) field: 4 bits.
 - **ef** (extension flag) field: 1 bit.
 - **beamId** (beam identifier) field: 15 bits.
 - **freqOffset** (frequency offset) field: 24 bits.
 - **reserved** (reserved for future use) field: 8 bits.

Table 7.4.5-1: Scheduling and beamforming commands frame format (Section Type 3)

Section Type 3: PRACH & mixed-numerology									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
transport header, see clause 5.1.3								8	Octet 1
dataDirection	payloadVersion			filterIndex				1	Octet 9
frameId								1	Octet 10
subframeId				slotId				1	Octet 11
slotId		startSymbolId						1	Octet 12
numberOfSections								1	Octet 13
sectionType = 3								1	Octet 14
timeOffset								2	Octet 15
frameStructure								1	Octet 17
cpLength								2	Octet 18
udCompHdr								1	Octet 20
sectionId								1	Octet 21
sectionId			rb	symInc	startPrbc			1	Octet 22
startPrbc								1	Octet 23
numPrbc								1	Octet 24
reMask[11:4]								1	Octet 25
reMask[3:0]			numSymbol					1	Octet 26
ef	beamId[14:8]							1	Octet 27
beamId[7:0]								1	Octet 28
freqOffset								3	Octet 29
reserved (8 bits)								1	Octet 32
Section Extensions as indicated by "ef"								var	Octet 33
...									
sectionId								1	Octet N
sectionId			rb	symInc	startPrbc			1	N+1
startPrbc								1	N+2
numPrbc								1	N+3
reMask[11:4]								1	N+4
reMask[3:0]			numSymbol					1	N+5
ef	beamId[14:8]							1	N+6
beamId[7:0]								1	N+7
freqOffset								3	N+8
reserved (8 bits)								1	N+11
Section Extensions as indicated by "ef"								var	N+12

NOTE: shading: yellow is transport header, pink is radio application header, others are repeated sections

7.4.6 Section Type 4 elements

Section Type 4 (Slot level configuration); see Table 7.4.6-1, Table 7.4.6-2:

○ Common Header Fields:

- **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
- **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
- **scs** (filter index) field: 4 bits.
- **frameId** (frame identifier) field: 8 bits.
- **subframeId** (subframe identifier) field: 4 bits.
- **slotID** (slot identifier) field: 6 bits.

- **startSymbolId** (start symbol identifier) field: 6 bits.
 - **reserved** (reserved for future use) field: 4 bits.
 - **cmdScope** (scope of Section Type 4 command) field: 4 bits.
 - **sectionType** (Section Type) field: 8 bits:
 - value = 4 shall be set.
 - **numberOfST4Cmds** (number of sections) field: 8 bits.
 - **reserved** (for future use): 8 bits.
- **st4CmdType** = 1 shall be set for **st4CmdType** : TIME_DOMAIN_BEAM_CONFIG; see Table 7.4.6-3, Table 7.4.6-4, Table 7.4.6-5:
- **st4CmdType** (TIME_DOMAIN_BEAM_CONFIG command type): 8 bits.
 - **st4CmdLen** (length of section type 4 command): 16 bits.
 - **numSlots** (number of slots for which the command configuration shall be valid): 8 bits.
 - **ackNackReqId** (ACK/NACK request id for a Section Type 8 command which shall be used for ackId and nackId in Section Type 8): 16 bits.
 - **reserved** (reserved for future use) field, required for 4-byte alignment: 2 bytes.
 - **symbolMask** (symbol to which time-domain beamforming config shall apply. Symbol bits before 'startSymbolId' in Section Type 4 common header should be set to 0 by O-DU and shall be ignored by O-RU): 14 bits.
 - **disableTDBFNs** ('1' = only 'tdBeamGrp' is included; '0' = 'tdBeamNum' with associated time-domain beamforming groups are included): 1 bit
 - **tdBeamGrp** (time-domain beam group, uniquely identifying a vector of length **K**, of time-domain beam numbers): 15 bits.
 - **bfwCompHdr** (beamforming weight compression header. Present only if disableTDBWs=0, else the field is a reserved field): 8 bits.
 - **reserved** (reserved for future use) field, required for 4-byte alignment: 3 bytes.
 - **disableTDBFWs** ('1' = only tdBeamNum are included; '0' = tdBeamNum with associated time-domain beamforming weights are included): 1 bit.
 - **tdBeamNum** (time-domain beam number associated with tdBeamGrp $N = K$): 15 bits.
 - **bfwCompParam** (Beamforming weight compression parameter, present only if disableTDBFWs=0): variable .
 - **bfWI** (time-domain beamforming weight in-phase value): variable.
 - **bfWQ** (time-domain beamforming weight quadrature-phase value): variable.
- Padding rules: 4-byte boundary for each set of weights when weights are present (disableTDBFWs = 0), mandatory padding to a 4-byte boundary only after the last beamId when weights are absent (disableTDBFWs = 1).
- **st4CmdType** = 2 shall be set for **st4CmdType**: TDD_CONFIG_PATTERN; see Table 7.4.6-6:
- **st4CmdType** (TDD_CONFIG_PATTERN command type): 8 bits.
 - **st4CmdLen** (length of section type 4 command): 16 bits.
 - **numSlots** (number of slots for which the command configuration shall be valid): 8 bits.
 - **ackNackReqId** (ACK/NACK request id for a Section Type 8 command which shall be used for ackId and nackId in Section Type 8): 16 bits.

- **reserved** (reserved for future use) field, required for 4-byte alignment: 2 bytes.
- **dirPattern** (Tx/Rx direction symbol mask, 1 = DL symbol, 0 = UL symbol): 14 bits.
- **guardPattern** (guard bits for Tx->Rx switch, 1=guard symbol, 0 = non-guard symbol, either UL or DL per dirPattern): 14 bits.

Table 7.4.6-1: Slot message frame format (Section Type 4)

Section Type 4: Slot-level Config											
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes			
transport header, see clause 5.1.3								8	Octet 1		
dataDirection	payloadVersion			scs						1	Octet 9
frameId								1	Octet 10		
subframeId				slotId						1	Octet 11
slotId		startSymbolId								1	Octet 12
reserved				cmdScope						1	Octet 13
sectionType = 4								1	Octet 14		
numberOfST4Cmds								1	Octet 15		
reserved (8 bits)								1	Octet 16		

Table 7.4.6-2: Section Type 4 Command common header format

Section Type 4: Command common header format									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
st4CmdType = 1								1	Octet 17
st4CmdLen								1	Octet 18
st4CmdLen								1	Octet 19
numSlots								1	Octet 20
ackNackReqId								2	Octet 21
reserved								2	Octet 23

Table 7.4.6-3: Slot level command type TIME_DOMAIN_BEAM_CONFIG (st4CmdType = 1) with disableTDBFNs = 0 and disableTDBFWs = 0

st4CmdType 'TIME_DOMAIN_BEAM_CONFIG'									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
transport header, see clause 5.1.3								8	Octet 1
Section Type 4 section header								8	Octet 9
Section Type 4 common part of the command header								8	Octet 17
reserved		symbolMask[13:8]						1	Octet 25
		symbolMask[7:0]						1	Octet 26
disableTDBFNs	tdBeamGrp[14:8] (for tdBeamGrp 1)							1	Octet 27
		tdBeamGrp[7:0] (for tdBeamGrp 1)						1	Octet 28
		bfwCompHdr						1	Octet 29
		reserved						3	Octet 30
disableTDBFWs	tdBeamNum[14:8] (for tdBeamNum 1)							1	Octet 31
		tdBeamNum[7:0] (for tdBeamNum 1)						1	Octet 32
		bfwCompParam (for tdBeamNum 1)						var	
		bfl (for tdBeamNum 1, for TRX 0)						var	
		bflQ (for tdBeamNum 1, for TRX 0)						var	
		remaining time-domain beamforming weights bfl and bflQ (for tdBeamNum 1, for $(\frac{K}{K'}) - 1$ TRXs)						var	
		zero pad to a 4-byte boundary						var	
...									
disableTDBFWs	tdBeamNum[14:8] ((for tdBeamNum K')							1	
		tdBeamNum[7:0] (for tdBeamNum K')						1	
		bflCompParam						var	
		bfl (TRX 0)						var	
		bflQ (TRX 0)						var	
		remaining time-domain beamforming weights bfl and bflQ (for tdBeamNum K th , for $(\frac{K}{K'}) - 1$ TRXs)						var	
		zero pad to a 4-byte boundary						var	

Table 7.4.6-4: Slot level command type TIME_DOMAIN_BEAM_CONFIG (st4CmdType = 1) with disableTDBFNs = 0 and disableTDBFWs = 1

st4CmdType 'TIME_DOMAIN_BEAM_CONFIG'									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
transport header, see clause 5.1.3								8	Octet 1
Section Type 4 section header								8	Octet 9
Section Type 4 common part of the command header								8	Octet 17
reserved		symbolMask[13:8]						1	Octet 25
		symbolMask[7:0]						1	Octet 26
disableTDBFNs	tdBeamGrp[14:8] (for tdBeamGrp 1)							1	Octet 27
		tdBeamGrp[7:0] (for tdBeamGrp 1)						1	Octet 28
		bflCompHdr						1	Octet 29
		reserved						1	Octet 30
disableTDBFWs	tdBeamNum[14:8] (for tdBeamNum 1)							1	Octet 33
		tdBeamNum[7:0] (for tdBeamNum 1)						1	Octet 34
...									
disableTDBFWs	tdBeamNum[14:8] (for tdBeamNum K')							1	
		tdBeamNum[7:0] (for tdBeamNum K')						1	
		zero pad to a 4-byte boundary						var	

Table 7.4.6-5: Slot level command type TIME_DOMAIN_BEAM_CONFIG (st4CmdType =1) with disableTDBFNs = 1

st4CmdType 'TIME_DOMAIN_BEAM_CONFIG'									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
transport header, see clause 5.1.3								8	Octet 1
Section Type 4 section header								8	Octet 9
Section Type 4 common part of the command header								8	Octet 17
reserved		symbolMask[13:8]						1	Octet 25
		symbolMask[7:0]						1	Octet 26
disableTDBFNs	tdBeamGrp[14:8] (for tdBeamGrp 1)						1	Octet 27	
		tdBeamGrp[7:0] (for tdBeamGrp 1)						1	Octet 28
		bfwCompHdr						1	Octet 29
		reserved						3	Octet 30

Table 7.4.6-6: Slot level command type TDD_CONFIG_PATTERN (st4CmdType =2)

st4CmdType 'TDD_CONFIG_PATTERN'									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
transport header, see clause 5.1.3								8	Octet 1
Section Type 4 section header								8	Octet 9
Section Type 4 common part of the command header								8	Octet 17
reserved		dirPattern[13:8]						1	Octet 25
		dirPattern[7:0]						1	Octet 26
reserved		guardPattern[13:8]						1	Octet 27
		guardPattern[7:0]						1	Octet 28

7.4.7 Section Type 5 elements

Section Type 5 is used for UE scheduling information; see Table 7.4.7-1:

- Common Header Fields:
 - **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
 - **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
 - **filterIndex** (filter index) field: 4 bits.
 - **frameId** (frame identifier) field: 8 bits.
 - **subframeId** (subframe identifier) field: 4 bits.
 - **slotID** (slot identifier) field: 6 bits.
 - **startSymbolId** (start symbol identifier) field: 6 bits.
 - **numberOfsections** (number of sections) field: 8 bits.
 - **sectionType** (Section Type) field: 8 bits:
 - value = 5 shall be set.
 - **udCompHdr** (user data compression header) field: 8 bits.
 - **reserved** (reserved for future use) field: 8 bits.
- Section Fields:
 - **sectionId** (section identifier) field: 12 bits.
 - **rb** (resource block identifier) field: 1 bit.

- **symInc** (symbol number increment command) field: 1 bit.
- **startPrbc** (starting PRB of data section description) field: 10 bits.
- **numPrbc** (number of contiguous PRBs per data section description) field: 8 bits.
- **reMask** (resource element mask) field: 12 bits.
- **numSymbol** (number of symbols) field: 4 bits.
- **ef** (extension flag) field: 1 bit.
- **ueId** field: 15 bits.

Table 7.4.7-1: UE scheduling information frame format (Section Type 5)

Section Type 5: UE scheduling information conveyance									# of bytes	
0 (msb)	1	2	3	4	5	6	7 (lsb)			
transport header, see clause 5.1.3									8	Octet 1
dataDirection	payloadVersion			filterIndex					1	Octet 9
frameId									1	Octet 10
subframeId				slotId					1	Octet 11
slotId		startSymbolId							1	Octet 12
numberOfSections									1	Octet 13
sectionType = 5									1	Octet 14
udCompHdr									1	Octet 15
reserved									1	Octet 16
sectionId									1	Octet 17
sectionId			rb	symInc	startPrbc				1	Octet 18
startPrbc									1	Octet 19
numPrbc									1	Octet 20
reMask[11:4]									1	Octet 21
reMask[3:0]			numSymbol						1	Octet 22
ef	ueId[14:8]								1	Octet 23
ueId[7:0]									1	Octet 24
Section Extensions as indicated by "ef"									var	Octet 25
...										
sectionId									1	Octet N
sectionId			rb	symInc	startPrbc				1	N+1
startPrbc									1	N+2
numPrbc									1	N+3
reMask[11:4]									1	N+4
reMask[3:0]			numSymbol						1	N+5
ef	ueId[14:8]								1	N+6
ueId[7:0]									1	N+7
Section Extensions as indicated by "ef"									var	N+8

NOTE: Shading: yellow is transport header, pink is radio application header, others are repeated sections

7.4.8 Section Type 6 elements

Section Type 6 is used for sending channel information for a specific ueId; see Table 7.4.8-1:

- Common Header Fields:
 - **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
 - **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
 - **filterIndex** (filter index) field: 4 bits.

- **frameId** (frame identifier) field: 8 bits.
- **subframeId** (subframe identifier) field: 4 bits.
- **slotID** (slot identifier) field: 6 bits.
- **startSymbolId** (start symbol identifier) field: 6 bits.
- **numberOfSections** (number of sections) field: 8 bits.
- **sectionType** (Section Type) field: 8 bits:
 - value = 6 shall be set.
- **numberOfUEs** (number of UE-specific channel information data sets) field: 8 bits.
- **ciCompHdr** (channel information compression header) field: 8 bits.
- Section Fields:
 - **ef** (extension flag) field: 1 bit.
 - **ueId** field: 15 bits.
 - **regularizationFactor** (regularization factor used for MMSE reception) field: 16 bits.
 - **reserved** (reserved for future use) field: 4 bits.
 - **rb** (resource block identifier) field: 1 bit.
 - **symInc** (symbol number increment command) field: 1 bit.
 - **startPrbc** (starting PRB of data section description) field: 10 bits.
 - **numPrbc** (number of contiguous PRBs per data section description) field: 8 bits.
 - **ciCompParam** (channel information compression parameter): 0 or 8 bits.
 - **ciIsample** (channel information value, in-phase sample) field: 1-16 bits.
 - **ciQsample** (channel information value, quadrature sample) field: 1-16 bits.

Table 7.4.8-1: UE channel information frame format (Section Type 6)

Section Type 6: channel information conveyance											
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes			
transport header, see clause 5.1.3									8	Octet 1	
dataDirection	payloadVersion			filterIndex						1	Octet 9
frameId									1	Octet 10	
subframeId				slotId						1	Octet 11
slotId		startSymbolId							1	Octet 12	
numberOfSections									1	Octet 13	
sectionType = 6									1	Octet 14	
numberOfUEs									1	Octet 15	
ciCompHdr									1	Octet 16	
ef	ueld[14:8]								1	Octet 17	
ueld[7:0]									1	Octet 18	
regularizationFactor									2	Octet 19	
reserved			rb	symInc	startPrbc				1	Octet 21	
startPrbc									1	Octet 22	
numPrbc									1	Octet 23	
ciCompParam (for the first PRB or all PRBs of the first UE, not always present)									var	Octet 24	
ciSample (first PRB, first antenna)									var		
ciQsample (first PRB, first antenna)									var		
ciSample (first PRB, second antenna)									var		
ciQsample (first PRB, second antenna)									var		
...											
ciSample (first PRB, last antenna)									var		
ciQsample (first PRB, last antenna)									var		
...											
ciCompParam (for the last PRB, not always present)									var		
ciSample (last PRB, last antenna)									var		
ciQsample (last PRB, last antenna)									var		
<padding with zeros to the next 4-byte boundary> may not be present (NOTE 2)									var		
Section Extensions as indicated by "ef"									var		
...											
ef	ueld[14:8]								1	Octet N	
ueld[7:0]									1	N+1	
regularizationFactor									2	N+2	
Reserved			rb	symInc	startPrbc				1	N+4	
startPrbc									1	N+5	
numPrbc									1	N+6	
ciCompParam (for the first PRB or all PRBs of the last UE, not always present)									var	N+7	
ciSample (first PRB, first antenna)									var		
ciQsample (first PRB, first antenna)									var		
ciSample (first PRB, second antenna)									var		
ciQsample (first PRB, second antenna)									var		
...											
ciSample (first PRB, last antenna)									var		
ciQsample (first PRB, last antenna)									var		
...											
ciCompParam (for the last PRB, not always present)											
ciSample (last PRB, last antenna)									var		
ciQsample (last PRB, last antenna)									var		
<padding with zeros to the next 4-byte boundary> may not be present (NOTE 2)									var		
Section Extensions as indicated by "ef"									var		

NOTE 1: shading: yellow is transport header, pink is radio application header, others are repeated sections.
NOTE 2: "may not be present" depends on O-RU boolean flag "st6-4byte-alignment-required".
NOTE 3: The remainder of N divided by 4 is 1 if 4-byte alignment is to be used.

7.4.9 Section Type 7 elements

Section Type 7 is used to support LAA; see Table 7.4.9-1:

- Common Header Fields:
 - **reserved** (reserved for future use) field: 1 bit.
 - **payloadVersion** (payload version) field: 3 bits:
 - value = 1 shall be set (1st protocol version for payload and time reference format).
 - **reserved** (reserved for future use) field: 4 bits.
 - **frameId** (frame identifier) field: 8 bits.
 - **subframeId** (subframe identifier) field: 4 bits.
 - **slotID** (slot identifier) field: 6 bits.
 - **reserved** (reserved for future use) field: 14 bits.
 - **sectionType** (Section Type) field: 8 bits:
 - value = 7 shall be set.
 - **reserved** (reserved for future use) field: 16 bits.
- Section Fields:
 - **laaMsgType (LAA message type)** field: 4 bits.
 - **laaMsgLen** field: 4 bits.
- laaMsgType = 0 shall be set for LBT_DL_CONFIG.request: LBT_PDSCH_REQ; see Table 7.4.9-2:
 - **lbtHandle** (An opaque handling returned in LBT_PDSCH_RSP) field: 16 bits.
 - **lbtOffset** (LBT start time in microseconds from the beginning of the subframe scheduled by this message) field: 10 bits.
 - **lbtMode** (LBT process type) field: 2 bits.
 - **reserved** (reserved for future use) field: 1 bit.
 - **lbtDeferFactor** (Defer factor in sensing slots as described in 3GPP TS 36.213 [14], clause 15.1.1) field: 3 bits.
 - **lbtBackoffCounter** (LBT backoff counter in sensing slots as described in 3GPP TS 36.213 [14], clause 15.1.1) field: 10 bits.
 - **MCOT** (LTE TXOP duration in subframes) field: 4 bits.
 - **reserved** (reserved for future use) field: 10 bits.
- laaMsgType = 1 shall be set for LBT_DL_CONFIG.request: LBT_DRS_REQ; see Table 7.4.9-3:
 - **lbtHandle** (An opaque handling returned in LBT_DRS_RSP) field: 16 bits.
 - **lbtOffset** (LBT start time in microseconds from the beginning of the subframe scheduled by this message) field: 10 bits.
 - **lbtMode** (LBT process type) field: 2 bits.
 - **reserved** (reserved for future use) field: 28 bits.
- laaMsgType = 2 shall be set for LBT_DL.indication: LBT_PDSCH_RSP; see Table 7.4.9-4:
 - **lbtHandle** (An opaque handling returned in LBT_PDSCH_RSP) field: 16 bits.

- **lbtPdschRes** (LBT result of SFN/SF) field: 2 bits.
 - **inParSF** (Indicates whether the initial SF in the LBT process is full or partial) field: 1 bit.
 - **sfStatus** (subframe status) field: 1 bit.
 - **sfnSf** (SFN/SF of subframe which is dropped or successfully transmitted at O-RU) field: 12 bits.
 - **reserved** (reserved for future use) field: 24 bits.
- laaMsgType = 3 shall be set for LBT_DL.indication: LBT_DRS_RSP; see Table 7.4.9-5:
 - **lbtHandle** (An opaque handling returned in LBT_DRS_RSP) field: 16 bits.
 - **lbtDrsRes** (LBT result of SFN/SF) field: 1 bit.
 - **reserved** (reserved for future use) field: 7 bits.
 - laaMsgType = 4 shall be set for LBT_buffer error: LBT_Buffer_Error; see Table 7.4.9-6:
 - **lbtHandle** (An opaque handling returned in LBT_DRS_RSP) field: 16 bits.
 - **lbtBufErr** (LBT buffer error type) field: 1 bit.
 - **reserved** (reserved for future use) field: 7 bits.
 - laaMsgType = 5 shall be set for LBT_DL_CONFIG.request: LBT_CWCONFIG_REQ; see Table 7.4.9-7:
 - **lbtHandle** (An opaque handle returned in LBT_CWCONFIG_RSP) field: 16 bits.
 - **lbtCWConfig_H** (HARQ feedback information regarding number of NACK about reference subframe as described in 3GPP TS 36.213 [14], clause 15.1.3) field: 8 bits.
 - **lbtCWConfig_T** (number of TB to manage congestion window as described in 3GPP TS 36.213 [14], clause 15.1.3) field: 8 bits.
 - **lbtMode** (LBT process type) field: 2 bits.
 - **lbtTrafficClass** (LBT traffic class) field: 3 bits.
 - **reserved** (reserved for future use) field: 19 bits.
 - laaMsgType = 6 shall be set for LBT_DL.indication: LBT_CWCONFIG_RSP; see Table 7.4.9-8:
 - **lbtHandle** (An opaque handling returned in LBT_CWCONFIG_RSP) field: 16 bits.
 - **lbtCWR_Rst** (Notification LBT_CWCONFIG_REQ message successful or not) field: 1 bit.
 - **reserved** (reserved for future use) field: 7 bits.

Table 7.4.9-1: LAA Message, O-DU to O-RU -OR- O-RU to O-DU (Section Type 7)

Section Type 7: LAA Message, O-DU to O-RU or O-RU to O-DU										
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
transport header, see clause 5.1.3								8	Octet 1	
reserved	payloadVersion			reserved					1	Octet 9
frameId								1	Octet 10	
subframeId				slotId				1	Octet 11	
slotId	reserved							1	Octet 12	
reserved								1	Octet 13	
sectionType = 7								1	Octet 14	
reserved								1	Octet 15	
reserved								1	Octet 16	
laaMsgType				laaMsgLen				1	Octet 17	
Payload (see below) plus padding to 32-bit boundary								3 or 7	Octet 18+	
NOTE: frameId, subframeId and slotId serve to provide a time stamp on the LAA LBT C-Plane message and do not describe any U-Plane data associated with the Section Type 7 C-Plane message.										

Table 7.4.9-2: LAA Request Message, O-DU to O-RU LBT_PDSCH_REQ

Section Type 7: LAA Request Message, O-DU to O-RU LBT_PDSCH_REQ									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
laaMsgType = 0000b (LBT_PDSCH_REQ)				laaMsgLen = 2 (2 words)				1	Octet 17
lbtHandle								2	Octet 18
lbtOffset[9:2]								1	Octet 20
lbtOffset[1:0]	lbtMode	reserved		lbtDeferFactor				1	Octet 21
lbtBckoffCounter[9:2]								1	Octet 22
lbtBckoffCounter[1:0]	MCOT				reserved			1	Octet 23
reserved								1	Octet 24

Table 7.4.9-3: LAA Request Message, O-DU to O-RU LBT_DRS_REQ

Section Type 7: LAA Request Message, O-DU to O-RU LBT_DRS_REQ									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
laaMsgType = 0001b (LBT_DRS_REQ)				laaMsgLen = 2 (2 words)				1	Octet 17
lbtHandle								2	Octet 18
lbtOffset[9:2]								1	Octet 20
lbtOffset[1:0]	lbtMode	reserved						1	Octet 21
reserved								1	Octet 22
reserved								1	Octet 23
reserved								1	Octet 24

Table 7.4.9-4: LAA Response Message, O-RU to O-DU LBT_PDSCH_RSP

Section Type 7: LAA Response Message, O-RU to O-DU LBT_PDSCH_RSP									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
laaMsgType = 0010b (LBT_PDSCH_RSP)				laaMsgLen = 2 (2 words)				1	Octet 17
lbtHandle									
lbtPdschRes		inParSF	sfStatus	sfnSf[11:8]				1	Octet 20
sfnSf[7:0]									
reserved									
reserved									
reserved									

Table 7.4.9-5: LAA Response Message, O-RU to O-DU LBT_DRS_RSP

Section Type 7: LAA Response Message, O-RU to O-DU LBT_DRS_RSP										
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
laaMsgType = 0011b (LBT_DRS_RSP)				laaMsgLen = 1 (1 word)				1	Octet N	
lbtHandle										
lbtDrsRes	reserved								1	N+3

Table 7.4.9-6: LAA Response Message, O-RU to O-DU LBT_Buffer_Error

Section Type 7: LAA Response Message, O-RU to O-DU LBT_Buffer_Error										
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
laaMsgType = 0100b (LBT_Buffer_Error)				laaMsgLen = 1 (1 word)				1	Octet N	
lbtHandle										
lbtBufErr	reserved								1	N+3

Table 7.4.9-7: LAA Request Message, O-DU to O-RU LBT_CWCONFIG_REQ

Section Type 7: LAA Request Message, O-DU to O-RU LBT_CWCONFIG_REQ									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
laaMsgType = 0101b (LBT_CWCONFIG_REQ)				laaMsgLen = 2 (2 words)				1	Octet N
lbtHandle									
lbtCWConfig_H									
lbtCWConfig_T									
lbtMode	lbtTrafficClass			reserved				1	N+5
reserved									

Table 7.4.9-8: LAA Request Message, O-RU to O-DU LBT_CWCONFIG_RSP

Section Type 7: LAA Request Message, O-RU to O-DU LBT_CWCONFIG_RSP										
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
laaMsgType = 0110b (LBT_CWCONFIG_RSP)				laaMsgLen = 1 (1 word)				1	Octet N	
lbtHandle										
lbtCWR_Rst	reserved								1	N+2

7.4.10 Section Type 8 elements

Section Type 8 is used for sending ACK/NACK from O-RU to O-DU; see Table 7.4.10-1:

- Common Header Fields
 - **reserved** (reserved for future use) field: 1 bit
 - **payloadVersion** (payload version) field: 3 bits
 - value = 1 shall be set (1st protocol version for payload and time reference format)
 - **reserved** (reserved for future use) field: 4 bits
 - **frameId** (frame identifier) field: 8 bits
 - **subframeId** (subframe identifier) field: 4 bits
 - **slotId** (slot identifier) field: 6 bits
 - **symbolId** (symbolId identifier) field: 6 bits
 - **reserved** (reserved for future use) field: 8 bits
 - **sectionType** (Section Type) field: 8 bits
 - value = 8 shall be set
 - **numberOfAcks** (number of ACKs) field: 8 bits
 - **numberOfNacks** (number of NACKs) field: 8 bits
- Section Fields
 - **ackId** (ACK Id) field: 16 bits
 - **nackId** (NACK Id) field: 16 bits

Table 7.4.10-1: ACK Message, O-RU to O-DU (Section Type 8)

Section Type 8 : ACK Message, O-RU to O-DU									
0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
transport header, see clause 5.1.3								8	Octet 1
reserved	payloadVersion			reserved				1	Octet 9
frameId								1	Octet 10
subframeId				slotId				1	Octet 11
slotId		symbolId						1	Octet 12
reserved								1	Octet 13
sectionType = 8								1	Octet 14
numberOfAcks								1	Octet 15
numberOfNacks								1	Octet 16
1st ackId								2	Octet 17
2nd ackId								2	Octet 19
...									
nth ackId (NOTE)								2	Octet 15+2*n
1st nackId								2	Octet 17+2*n
2nd nackId								2	Octet 19+2*n
...									
mth nackId (NOTE)								2	Octet 15+2*(m+n)
Padding to align with 32-bit boundary									

NOTE: n is the value in the field of numberOfAcks and m is the value in the field of numberOfNacks.

7.5 Coding of Section Type IEs

7.5.1 Transport layer

See clause 5.1.3 for transport header information element details.

7.5.2 Application layer, common

7.5.2.1 dataDirection (data direction (gNB Tx/Rx))

Description: This parameter indicates the gNB data direction.

Value range: {0b=Rx (i.e., UL), 1b=Tx (i.e., DL)}.

Type: binary bit.

Field length: 1 bit.

7.5.2.2 payloadVersion (payload version)

Description: This parameter defines the payload protocol version valid for the following IEs in the application layer. In the present document payloadVersion=001b shall be used.

Value range: {000b-111b=Payload version}.

Type: unsigned integer.

Field length: 3 bits.

Default Value: 001b (version 1 assumed).

7.5.2.3 filterIndex (filter index)

Description: This parameter defines an index to the channel filter to be used between IQ data and air interface, both in DL and UL. For most physical channels filterIndex =0000b may be used which indexes the standard channel filter, e.g., 100 MHz channel filter for 100MHz nominal carrier bandwidth. Another use case is PRACH in UL, where different filter indices can be used for different PRACH formats, assuming that before FFT processing of PRACH data there is a separate PRACH filter or PRACH filter in addition to the standard channel filter in UL. For PRACH, when Section Type 3 message is used, there is typically also a frequency offset (see freqOffset) applied before the PRACH filter. filterIndex coding shall conform to Table 7.5.2.3-2.

- Filter index shall be commanded from the O-DU to the O-RU. However, it is not mandatory to command special filters, and filter index = 0000b is also allowed for PRACH.
- When using filter indices corresponding to PRACH, the first RE of the first PRB addressed inside the section Id shall correspond with the first guard tone used at the lower edge of the PRACH frequency block as specified in 3GPP TS 36.211 [42], clause 5.7 and TS 38.211 [4], clause 6.3.3.2. In case of fragmentation of one data section over multiple U-Plane messages, this is applicable to the first fragmented section.
- Since different PRACH formats may be assigned to the same filterIndex value, optionally O-RU can notify by M-Plane the specific supported formats (grouped into PRACH format-groups) in o-ran-uplane-conf.yang module on a per-endpoint basis.
- O-DU shall set filterIndex as 1000b for NPUSCH and 0100b for NPRACH as specified in Table 7.5.2.3-1 in ST 3 C-plane message when SSSC or SSMC method is used for NB-IoT carriers for UL direction as specified in clause 15.4.1 and 15.4.2.

Value range: {0000b-1111b}.

Table 7.5.2.3-2: Filter index

Value of IE "filter Index"	Usage	PRACH preamble formats	Minimum filter pass band
0000b=0x0	standard channel filter	N/A	
0001b=0x1	UL filter for PRACH preamble formats	LTE-0, LTE-1, LTE-2, LTE-3, NR-0, NR-1, NR-2	839 x 1.25 kHz = 1048.75 kHz
0010b=0x2		NR-3	839 x 5 kHz = 4195 kHz
0011b=0x3		NR-A1, NR-A2, NR-A3, NR-B1, NR-B2, NR-B3, NR-B4, NR-C0, NR-C2	139 x Δf^{RA} (See SCS in Table 7.5.2.13-3)
0110b=0x6			1151 x 15 kHz = 17265 kHz
0111b=0x7			571 x 30 kHz = 17130 kHz
0100b=0x4	UL filter for NPRACH	LTE-NB0, LTE-NB1 LTE-NB0-a, LTE-NB1-a LTE-NB2	48 x 3.75 kHz = 180 kHz and 144 x 1.25 kHz = 180 kHz for SSSC, or numPrbc x 12 x SCS for SSMC (NOTE)
0101b=0x5	UL filter for PRACH preamble formats	LTE-4	139 x 7.5 kHz = 1042.5 kHz
1000b=0x8	UL filter for NPUSCH	NA	180 kHz for SSSC, or numPrbc x 12 x SCS for SSMC (NOTE)
1001b...1111b	Reserved		

NOTE: Refer to clause 15.4.1 and 15.4.2 for SSSC and SSMC operation.

Type: unsigned integer.

Field length: 4 bits.

Default Value: 0000b (no special filter).

7.5.2.4 frameld (frame identifier)

Description: This parameter is a counter for 10 ms frames (wrapping period 2.56 seconds), specifically frameId = frame number modulo 256. When the SCS is less than 15 kHz in e.g. PRACH and NB-IoT, then the frameId shall take the value of the first symbol in a multi-symbol C-Plane section description.

Value range: {0000 0000b-1111 1111b}.

Type: unsigned integer.

Field length: 8 bits.

7.5.2.5 subframeld (subframe identifier)

Description: This parameter is a counter for 1 ms sub-frames within a 10ms frame. When the SCS is less than 15 kHz in e.g. PRACH and NB-IoT, then the subframeId shall take the value of the first symbol in a multi-symbol C-Plane section description.

Value range: {0000b-1111b}.

Type: unsigned integer.

Field length: 4 bits.

7.5.2.6 slotld (slot identifier)

Description: This parameter is the slot number within a 1ms sub-frame. All slots in one sub-frame shall be counted by this parameter, slotId running from 0 to $N_{\text{slot}}-1$. If a C-Plane message is addressing symbols of SCS less than 15 kHz,

the slot duration exceeds duration of a sub-frame ($N_{\text{slot}} = 1$) and slotId shall be set to zero. In the present document the maximum $N_{\text{slot}}=16$, All other values of the 6 bits are reserved for future use.

Value range: {00 0000b-00 1111b=slotID, 01 0000b-11 1111b=Reserved}.

Type: unsigned integer.

Field length: 6 bits.

7.5.2.7 startSymbolId (start symbol identifier)

Description: This parameter identifies the symbol number (within a slot) of the earliest symbol, to which the information of this message is applicable. When the SCS value provided by "frameStructure" in the C-Plane message is less than 15 kHz, startSymbolId shall be based on the frame structure of SCS 15 kHz, specified in clause 7.2.2. When SSSC or SSMC method is used (see clauses 15.4.1 and 15.4.2), startSymbolId shall be set as per clause 15.6.

NOTE: If the C-Plane message containing this parameter includes a Section Extension (e.g., SE 6, SE 12) or command type (Section Type 4) containing a mask for symbols (e.g., symbolMask, dirPattern, guardPattern), all bits in the mask before startSymbolId are expected to be set to zero by the O-DU.

Value range: {00 0000b-11 1111b}.

Type: unsigned integer.

Field length: 6 bits.

7.5.2.8 numberOfSections (number of sections)

Description: This parameter indicates the number of data section descriptions (separate citations of section ID even for multiple citations of the same sectionId) included in this C-Plane message.

Value range: {0000 0000b-1111 1111b}.

Type: unsigned integer.

Field length: 8 bits.

7.5.2.9 sectionType (Section Type)

Description: This parameter determines the characteristics of U-Plane data to be transferred or received from a beam with one pattern id.

Value range: {0000 0000b-1111 1111b=Section Type, see Table 7.3.1-1 for Section Type meanings}.

Type: unsigned integer.

Field length: 8 bits.

7.5.2.10 udCompHdr (user data compression header)

See clause 8.3.3.13 for the description of this parameter.

The udCompHdr information is provided on the U-Plane, instructing the O-RU (on DL) and O-DU (on UL) how to interpret and decompress the received U-Plane data. For UL U-Plane data compression, the O-DU shall instruct the O-RU via udCompHdr in a C-Plane message. This means the udCompHdr in a C-Plane message only has meaning for UL U-Plane data (dataDirection = 0b) and should be set to 0x00 when dataDirection = 1b (DL data), and the receiving O-RU shall ignore the received udCompHdr value. When static data format and compression is used (see clause 8.1.2) then udCompHdr is not needed in any C-Plane message and should be set to 0x00, and the receiving O-RU shall ignore the received udCompHdr value.

The O-DU shall not send different values of udCompHdr for the same data section, even if it is sent over separate C-Plane messages discriminated by different reMask values. Only a single compression method per data section is supported (see clause 7.5.3.1).

7.5.2.11 numberOfUEs (number Of UEs)

Description: This parameter applies to Section Type 6 messages and indicates the number of UEs (for which channel information is provided) are included in the message.

NOTE: This allows the parser to determine when the last UE's data has been parsed.

Value range: {0000 0000b-1111 1111b}.

Type: unsigned integer.

Field length: 8 bits.

7.5.2.12 timeOffset (time offset)

Description: This parameter defines the time_offset from the start of the slot to the start of the Cyclic Prefix (CP) in number of samples T_s ($=1/30.72\text{MHz}$ as specified in 3GPP TS38.211 [4], clause 4.1). The value shall be less than the slot length. The slot length is based on the SCS on which the startSymbolId numerology is based (see Table 7.2.3.3-1 and Table 7.2.3.3-2 in clause 7.2.3.3). For the mixed numerology case (non-PRACH), timeOffset points to the same timing pointed by startSymbolId. For PRACH, refer to clause 4.4.3.

$$\text{time_offset} = \text{timeOffset} \times t_s$$

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.5.2.13 frameStructure (frame structure)

Description: This parameter defines the frame structure. The first 4 bits define the FFT/iFFT size being used for all IQ data processing related to this message. The second 4 bits define the sub carrier spacing as well as the number of slots per 1ms sub-frame according to 3GPP TS 38.211 [4], taking for completeness also 3GPP TS 36.211 [42] into account. The parameter $\mu = 0, \dots, 4$ from 3GPP TS 38.211 [4] is extended to apply for PRACH processing. Table 7.5.2.13-1 shall be used for the frameStructure bit allocation, while Table 7.5.2.13-2 shall be used for the FFT size field definitions and Table 7.5.2.13-3 shall be used for the sub-carrier spacing field definitions.

NOTE: The parameter "FFTSize" does not dictate the O-RU's method of time-to-frequency conversion.

Value range: {0000 0000b-1111 1111b}

Bit allocations

Table 7.5.2.13-1: frameStructure bit allocation

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
FFT Size				μ (Subcarrier spacing)				1	Octet 1

Table 7.5.2.13-2: FFT size

Value of IE "FFT_size"	FFT/iFFT size
0000b	Reserved (no FFT/iFFT processing)
0001b...0011b	Reserved
0100b	16
0101b	32
0110b	64
0111b	128
1000b	256
1001b	512
1010b	1024
1011b	2048
1100b	4096
1101b	1536
1110b	3072
1111b	Reserved

Table 7.5.2.13-3: Sub-Carrier spacing

Value of IE "SCS"	3GPP " μ "	Subcarrier spacing Δf	Number of slots per 1 ms sub-frame: N_{slot}	Slot length
0000b	0	15 kHz	1	1 ms
0001b	1	30 kHz	2	500 μ s
0010b	2	60 kHz	4	250 μ s
0011b	3	120 kHz	8	125 μ s
0100b	4	240 kHz	16	62.5 μ s
0101b...1011b	NA	Reserved	Reserved	Reserved
1100b	NA	1.25 kHz	1	1 ms
1101b	NA	3.75 kHz (LTE-specific)	1	1 ms
1110b	NA	5 kHz	1	1 ms
1111b	NA	7.5 kHz (LTE-specific)	1	1 ms

Type: unsigned integer (concatenated bit fields).

Field length: 8 bits.

7.5.2.14 cpLength (cyclic prefix length)

Description: This parameter defines the length CP_length of the Cyclic Prefix (CP) as follows, based on T_s (= 1/30.72 MHz as specified in 3GPP TS 38.211 [4], clause 4.1):

$$\text{CP_length} = \text{cpLength} \times T_s$$

The cpLength parameter is used with Section Types 0 and 3, and is applicable to all symbols addressed by the C-Plane message. In addition to the values defined in 3GPP, the value 0 shall be supported. The slot start time reference differentiation between normal CP and extended CP shall be based on cpLength and frameStructure. If cpLength is zero then differentiation between normal CP and extended CP time reference shall be based on cp-type parameter conveyed in the M-Plane. With Section Type 1, the CP properties shall be configured over M-Plane and the cyclic prefix shall have the following properties:

- **cp-type:** Normal or extended.
- **cp-length:** same unit as cpLength defined in this clause; used for symbol 0 for NR & LTE, and symbol $7 \times 2^\mu$ for NR (μ as defined in Table 7.5.2.13-3 with "NA" for μ shall be replaced by "0").
- **cp-length-other:** same unit as cpLength defined in this clause; used for other symbols than those using cp-length.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.5.2.15 ciCompHdr (channel information compression header)

Description: This parameter defines the compression method and IQ bit width for the channel information.

If the O-RU supports static channel information compression through M-Plane configuration, and static compression is configured, this field shall be set to a value corresponding to the compression method configured via M-Plane. If the O-RU supports dynamic channel information compression, then O-DU may configure dynamic compression via M-Plane with or without providing a list of compression methods that will be used dynamically (i.e., a subset of methods supported by O-RU). If dynamic compression is configured via M-Plane without providing a list of compression methods, this field shall be set to a value corresponding to one of dynamic compression methods supported by O-RU. If dynamic compression is configured with a list of compression methods, then this field shall be set to a value corresponding to one of compression methods in the list. If the ciCompHdr parameter value in a C-Plane message and M-Plane configuration are not compatible, then the O-RU may discard the message or process the message as per M-Plane configuration or as per ciCompHdr value (this may result in O-RU incorrectly decoding channel information).

Value range: {0000 0000b-1111 0001b}

Bit allocations: Table 7.5.2.15-1 shall be used for the ciCompHdr bit allocation, Table 7.5.2.15-2 shall be used to encode the cilqWidth value, Table 7.5.2.15-3 shall be used for the ciCompMeth definition, and Table 7.5.2.15-4 shall be used for the ciCompOpt definition..

Table 7.5.2.15-1: ciCompHdr bit allocation

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
cilqWidth				ciCompMeth			ciCompOpt	1	Octet 1

Table 7.5.2.15-2: cilqWidth definition

cilqWidth	Bit width of each I and each Q
0000-1111b	value of cilqWidth except a value of zero means 16 bits e.g. cilqWidth = 0000b means I and Q are each 16 bits wide; e.g. cilqWidth = 0001b means I and Q are each 1 bit wide; e.g. cilqWidth = 1111b means I and Q are each 15 bits wide

Table 7.5.2.15-3: ciCompMeth definition

ciCompMeth	compression method	cilqWidth meaning
000b	no compression	bitwidth of each uncompressed I and Q value
001b	block floating point,	bitwidth of each I and Q mantissa value
010b	block scaling	bitwidth of each I and Q scaled value
011b	u-law	bitwidth of each compressed I and Q value
100b –111b	reserved for future methods	depends on the specific compression method

Table 7.5.2.15-4: ciCompOpt definition

ciCompOpt	Compression Option
0b	compression per UE, one ciCompParam exists before the I/Q value of each UE
1b	compression per PRB, one ciCompParam exists before the I/Q value of each PRB
NOTE 1: ciCompOpt parameter is applicable only when compression is used i.e., ciCompMeth set to value other than 000b.	
NOTE 2: When channel information compression is supported, O-RU shall support both compression per UE (ciCompOpt=0) and compression per PRB (ciCompOpt=1).	

Type: unsigned integer (concatenated bit fields).

Field length: 8 bits.

Default Value: 0000 0000b (no compression, 16-bit I and Q).

7.5.2.16 scs (sub-carrier spacing)

Description: This parameter specifies the subcarrier spacing for the slot to which Section Type 4 slot level message shall apply. This parameter is used by the O-RU to derive reception window based on the scs for mixed numerology endpoints. Refer to Table 7.5.2.13-3: Sub-Carrier spacing.

Value range: Same as Table 7.5.2.13-3

Type: unsigned integer.

Field length: 4 bits.

7.5.2.17 cmdScope (command scope)

Description: Section Type 4 C-Plane messages are structured to apply to a group of endpoints. The group of endpoints to which a given 'st4CmdType' shall apply is derived from eAxC_ID value in the Section Type 4 C-Plane message transport header, the dataDirection bit in the Section Type 4 header, and the value of this parameter 'cmdScope' in Section Type 4 section header. An endpoint's support for Section Type 4 shall mean that the endpoint (eAxC_ID in the C-Plane transport header) can directly receive and execute the Section Type 4 command and all endpoints in the command scope are impacted by the Section Type 4 command regardless of their support for Section Type 4.

e.g., If a group of eAxC_IDs: 1 to 20 are assigned to tx-array-0 in the O-RU, a Section Type 4 command with transport header eAxC_ID=1, 'dataDirection = 1', cmdScope = 'ARRAY-COMMAND (0000b)' shall apply to all eAxC_IDs: 1 to 20 associated with tx-array-0.

The defined value of command scope is defined in table 7.5.2.17-1.

Table 7.5.2.17-1: command scope values and description

cmdScope	Description
ARRAY-COMMAND	<p>The command applies to:</p> <ul style="list-style-type: none"> all endpoints associated with the array associated with the endpoint receiving the ST 4 message and, all endpoints associated with the arrays having shared or co-located array elements with the array associated with the endpoint receiving the ST 4 message. <p>If the command type applies to both UL and DL then the command applies to endpoints of both RX and TX directions regardless of the direction of endpoint receiving the ST 4 message. Otherwise, the command applies to endpoints of the same direction as the of endpoint receiving the ST 4 message.</p>
CARRIER-COMMAND	<p>The command applies to all endpoints associated with the array carrier associated with the endpoint receiving the ST 4 message.</p>
O-RU-COMMAND	<p>If the command type applies to both UL and DL then the command applies to endpoints of both RX and TX directions regardless of the direction of endpoint receiving the ST 4 message. Otherwise the command applies to endpoints of the same direction as the of endpoint receiving the ST 4 message. In the case of multiple O-DUs sharing RF circuits in a single O-RU, the O-DUs are expected to cooperate to ensure the impact of the Section Type 4 command is appropriate for all affected O-DUs.</p>

NOTE: To handle shared/co-located array elements of [tr]x-arrays, if an endpoint assigned to a large array is addressed by ST 4 command, any sub-arrays of the large array are impacted without having to send the ST 4 command explicitly to endpoints assigned to the sub-arrays. If an endpoint is assigned to a sub-array addressed by a ST 4 command, only array-elements within the sub-array are impacted, and no other array-elements in the sub-array's larger array.

Value range: {0000b – 1111b}.

- 0000b: ARRAY-COMMAND.
- 0001b: CARRIER-COMMAND.

- 0010b: O-RU-COMMAND
- 0011b – 1111b: reserved

Type: unsigned integer.

Field length: 4 bits.

7.5.2.18 numberOfST4Cmds (Number of Section Type 4 commands)

Description: Number of Section Type 4 commands which are included in the Section Type 4 C-Plane message; the minimum value is one because there shall always be at least one Section Type 4 command in a Section Type 4 message.

Value range: {0000 0001b – 1111 1111b}.

0000 0000b - reserved

Type: unsigned integer.

Field length: 8 bits.

7.5.3 Application layer, sections

7.5.3.1 sectionId (section identifier)

Description: If C-Plane and U-Plane coupling via sectionId (see clause 7.8.1.2) is used then this parameter identifies individual data sections that are described by data section descriptions within the C-Plane message. The purpose of the sectionId is to map U-Plane data sections to the corresponding C-Plane message (and Section Types) associated with the data. Two or more C-Plane data section descriptions with same sectionId may be cited corresponding to a single U-Plane data section containing a combined payload for both citations (e.g., for supporting mixed CSI RS and PDSCH). This case is applicable when usage of reMask is complimentary (or orthogonal) and different beam directions (i.e. beamIds) are given the resource elements. Data sections are specific to an eAxC so a different eAxC may have independently defined data sections (e.g., have different ranges of PRBs contained within), unless otherwise specified (e.g., as in SE 10).

sectionId attributes:

- sectionIds are specific to a slot, so sectionId values may be "reused" for each slot, and the sectionId value for one slot has no specified relation to the sectionId value for a different slot.
- All data in section descriptions with same sectionId value shall have same rb, startPrbc, numPrbc, udCompHdr, and numSymbol IE fields' content.
- An upper bound on the max number of section IDs that can be addressed per eAxC or per sets of eAxCs, per symbol and per slot, for DL and for UL respectively, shall be conveyed via M-Plane messaging as part of the O-RU capabilities description.
- The sectionId cited in a C-Plane message shall have the same value as the sectionId in the corresponding U-Plane message for the given data section as defined by the frameId, subFrameId, slotId, startSymbolId fields and range of relevant PRBs as indicated by the totality of the specified startPrb(c/u) and numPrb(c/u) fields.

If C-Plane and U-Plane coupling via sectionId (see clause 7.8.1.2) is not used, then the sectionId value is not used for identification of data sections in U-Plane messages and is not used for identification of data section descriptions in C-Plane messages. In this case above rules for uniqueness of sectionId value within slot and restriction on rb, startPrbc, numPrbc, udCompHdr, and numSymbol IE fields' content do not apply. Sender shall send sectionId = 4095. The only exception to this is "Coupling via Frequency and Time with Priorities (Optimized)" where a unique value of sectionId is still used for identification of the highest priority data section descriptions in C-Plane messages (clause 7.8.1.5, restriction 5). The corresponding U-Plane messages shall still use sectionId value as 4095 in this case.

Value range: {0000 0000 0000b-1111 1111 1111b}.

Type: unsigned integer.

Field length: 12 bits.

7.5.3.2 rb (resource block indicator)

Description: This parameter indicates if every RB is used or every other RB is used. The starting RB is defined by startPrbc and total number of used RBs is defined by numPrbc.

EXAMPLE: RB=1, startPrb=1, numPrb=3, then the PRBs used are 1, 3, and 5.

If numPrbc=0 (i.e., all PRBs), then the sending node shall set the rb value to zero, and the receiving node shall ignore whatever rb value is received and assume it is zero.

If section description includes Section Extension 6 or 12 then the sender shall set the parameter rb to zero. Refer to clause 7.7.6.1, requirement #3 for an exception with respect to Section Extension 6.

Value range: {0b=every RB used; 1b=every other RB used}.

Type: binary bit.

Field length: 1 bit.

Default Value: 0b (every RB used).

7.5.3.3 symInc (symbol number increment command)

Description: If the section description does not include Section Extension 6, 12 and 19, this parameter is used to indicate which symbol number is relevant to the given section description. If the section description includes Section Extension 6, 12 or 19 then set of symbols relevant for the section description shall be given in symbolMask (see clauses 7.7.6.4, 7.7.12.3 and 7.7.19.6). In this case, symInc and numSymbol do not affect the set of symbols relevant for the section description but can affect the following section descriptions if they do not include Section Extension 6, 12 and 19. Regardless of the presence of Section Extension 6, 12 and 19, for each C-Plane message a symbol number shall be maintained and starts with the value of startSymbolId. The same value shall be used for each section in the message as long as symInc is zero. When symInc is one, the maintained symbol number shall be incremented to the next symbol, and that new symbol number shall be used for that section and each subsequent section until the symInc bit is again detected to be one. In the case of a multiple-symbol data section (numSymbol > 1), the new symbol number shall be the one after the last symbol in the data section. In this manner, multiple symbols may be handled by a single C-Plane message.

Additional considerations regarding use of symInc in the C-Plane:

- a) SymInc may be used when different data section descriptions have a different number of symbols under certain conditions namely that the data section descriptions in the C-Plane message are carefully arranged. The arrangement shall be such that the following procedure as applied by the O-RU works correctly:

symbol[s] - symbol addressed by section description $s=1 \dots N$ (considers section descriptions in one C-Plane message, $s=1$ identifies the first section description) shall be calculated as follows:

Assuming symInc[s] and numSymbols[s] are values of corresponding fields of section description s
For sake of simplicity, symbol[0] = startSymbolId, symInc[0]=0 and numSymbols[0]=1
for $s=1 \dots N$ (all section descriptions in a message)

If symInc[s] = 0 then symbol[s] = symbol[s-1]

Else symbol[s] = symbol[s-1] + numSymbols[s-1]

- b) SymInc is specific to a data section description and, if any of Section Extension 6, 12 and 19 are not present, SymInc is affecting range of symbols described by the section description. When C-Plane and U-Plane coupling via sectionId (see clause 7.8.1.2) is used and a data section is referenced by multiple data section descriptions (e.g. with different reMask values) if the value of symInc is to be set to 1, then only the first invocation of the sectionId shall have symInc=1 and all other invocations of the sectionId shall have symInc=0 in the same C-Plane message. This assures that the above rule can be applied (in the same C-Plane message only) while assuring that all invocations of the data section descriptions will refer to the same symbols, which shall be the case whether the invocations are in the same C-Plane message or in separate C-Plane messages.

Use of symInc in the U-Plane is independent of the use of symInc in the C-Plane. U-Plane messages include data for a single symbol (to avoid degrading the latency performance of the U-Plane). For this reason, use of symInc=1 shall be prohibited in the U-Plane.

Value range: {0b=use the current symbol number; 1b=increment the current symbol number and use that}.

Type: binary bit.

Field length: 1 bit.

Default Value: 0b (do not increment the current symbol number).

7.5.3.4 startPrbc (starting PRB of data section description)

Description: In absence of Section Extensions 6, 12 and 13 in the section description, startPrbc parameter conveys the first (lowest frequency) PRB described by the section description. The interpretation of startPrbc is affected by presence of Section Extensions 6, 12 and 13; see clause 7.7.6, 7.7.12 and 7.7.13 for more details.

Section description shall address resource elements without ambiguity: a resource element addressed by a section description shall not be addressed by other section description, with exception for C-Plane and U-Plane coupling via frequency and time with priorities which resolves ambiguity by differences in priorities (see clause 7.8.1.4 and 7.8.1.5). If coupling of C-Plane and U-Plane via sectionId value is used a PRB addressed by a section description shall not be addressed by a section description with a different value of sectionId.

freqOffset affects the frequency span for specific range of PRB numbers. Therefore "shall address resource elements without ambiguity" shall consider the value of freqOffset.

Center frequency F_{RE} in Hz of a subcarrier corresponding to RE #k (k = 0 ... 11) in PRB identified by startPrbc is defined by:

$$F_{RE} = \text{center_of_channel_bandwidth} + \text{frequency_offset} + \text{startPrbc} \times 12 \times \Delta f + k \times \Delta f + \text{dc_skip}$$

where:

center_of_channel_bandwidth in Hz is configured in M-Plane,

frequency_offset is calculated from freqOffset field (see clause 7.5.3.11) if freqOffset is present in the C-Plane message or calculated from M-Plane parameter offset-to-absolute-frequency-center (in this case frequency_offset = offset-to-absolute-frequency-center $\times \Delta f \times 0.5$ where Δf is frequency in Hz corresponding to subcarrier spacing configured via M-Plane) otherwise,

Δf is frequency in Hz corresponding to the subcarrier spacing configured in frameStructure (see clause 7.5.2.13) if frameStructure is present in the message, or to the subcarrier spacing configured via M-Plane otherwise,

dc_skip depends on carrier type:

$$\begin{aligned} &\text{if carrier is of LTE DL type and frequency_offset} + \text{startPrbc} \times 12 \times \Delta f + k \times \Delta f \geq 0 \text{ then dc_skip} = \Delta f \\ &\text{else dc_skip} = 0. \end{aligned}$$

The above formulation is intended to result in resource grid position in frequency compatible with 3GPP requirements, defined in TS 36.211 [42], clause 5.2.1 (UL) and clause 6.2.1 (DL), and 3GPP TS 38.211 [4], clause 4.4.2.

Value range: {00 0000 0000b-11 1111 1111b}.

Type: unsigned integer.

Field length: 10 bits.

7.5.3.5 reMask (resource element mask)

Description: This parameter defines the Resource Element (RE) mask within a PRB. Each bit setting in the reMask indicates if the section control is applicable to the RE sent in U-Plane messages (0=not applicable; 1=applicable). MSB indicates the value for the RE of the lowest frequency in a PRB.

Different REs in a PRB may be referenced by different data section descriptions. This is restricted to data section descriptions having the same sectionId but different reMask if C-Plane and U-Plane coupling via sectionId is used (see clause 7.8.1.2). If C-Plane and U-Plane coupling method independent of sectionId (see clauses 7.8.1.3 or 7.8.1.4 or 7.8.1.5) is used, then data sections descriptions referencing different REs in a PRB shall have different reMask values but sectionId is not relevant. In addition, if C-Plane and U-Plane coupling via frequency and time with priorities (see clause 7.8.1.4) is used a RE in a PRB may be referenced by more than one data section description (in terms of reMask)

but only configuration carried with data section description of highest priority is applicable (see clauses 7.8.1.4 and 7.7.6.4).

The maximum number of different reMask values that may be applied to a PRB is an O-RU characteristic that is conveyed from the O-RU to the O-DU via the M-Plane. If any RE in a PRB is never pointed to by an reMask (but other REs in that PRB are), the "missing" RE should be set to zero in the U-Plane, and no beamforming ID or other processing should be applied to the "missing" RE. An RE shall not be referenced more than once in a data section.

Value range: {0000 0000 0000b-1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 12 bits.

Default Value: 1111 1111 1111b (all REs in the block applicable).

7.5.3.6 numPrbc (number of contiguous PRBs per data section description)

Description: In absence of Section Extensions 6, 12, 13, and 21 in the section description and the M-Plane configuration **ci-prb-group-size**, the numPrbc parameter conveys the number of PRBs described by the section description. The interpretation of numPrbc is affected by presence of Section Extensions 6, 12, 13 and 21, and the M-Plane configuration **ci-prb-group-size**; see clause 7.7.6, 7.7.12, 7.7.13, 7.5.3.13, and 7.7.21 for more details. Presence of M-Plane configuration **ci-prb-group-size** only affects the interpretation of numPrbc parameter within Section Type 6.

Section description shall address resource elements without ambiguity: a resource element addressed by a section description shall not be addressed by other section description, with exception for C-Plane and U-Plane coupling via frequency and time with priorities which resolves ambiguity by differences in priorities (see clause 7.8.1.4 and 7.8.1.5). If coupling of C-Plane and U-Plane via sectionId value is used a PRB addressed by a section description shall not be addressed by a section description with a different value of sectionId.

Value range: {0000 0001b-1111 1111b, 0000 0000b = all PRBs in the specified SCS and carrier bandwidth}.

Value 0000 0000b is reserved for NR cases wherein the total number of PRBs may be more than 255. In this case, the sending node shall set the startPrbc value to zero and the receiving node shall ignore whatever startPrbc value is received and assume it is zero. For all other cases a non-zero value of numPrbc shall be used.

If the number of PRBs that needs to be sent is more than 255 but is not all of the PRBs in the specified SCS and carrier bandwidth, then multiple data sections describing smaller number of PRBs are needed.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.7 numSymbol (number of symbols)

Description: In section description without Section Extension 6, 12 and 19, this parameter defines number of symbols, or number of PRACH Symbols in a PRACH occasion in the case of PRACH or number of PRACH Symbols in NPRACH symbol group in the case of NPRACH, to which the section control is applicable. At minimum, the section control shall be applicable to at least one symbol. However, possible optimizations could allow for several (up to 14) symbols, if e.g., all 14 symbols use the same beamId.

In section description with Section Extension 6, 12 or 19, the set of symbols referred by the description is conveyed with symbolBitmask field and numSymbol parameter does not reflect number of symbols referred by the description. In this case, the value of numSymbol can be zero or non-zero number of symbols as it affects the set of symbols referred by the following section descriptions that do not include Section Extension 6, 12 and 19 and have symInc flag set to 1.

Refer to clause 7.5.3.3 for more details.

Value range: {0001b-1110b=number of symbols, 0000b=reserved (can be used in section description with Section Extension 6 or 12), 1111b=reserved}.

Type: unsigned integer.

Field length: 4 bits.

7.5.3.8 ef (extension flag)

Description: This parameter is used to indicate if this section has any Section Extensions as described in clause 7.3.2 included in the message

Value range: {0b=no Section Extensions; 1b=one or more Section Extensions are included in this section}.

Type: binary bit.

Field length: 1 bit.

Default Value: 0b (no Section Extensions).

7.5.3.9 beamId (beam identifier)

Description: This parameter defines the beam pattern to be applied to the U-Plane data. beamId = 0 means no beamforming operation will be performed. No beamforming operation implies that the O-RU shall not apply any phase or amplitude weights to the U-Plane data and that the resulting RF signal will be applied to all antenna elements in the group equally.

The beamId encodes the beamforming to be done on the O-RU. This beamforming may be digital, analog or both ("hybrid beamforming") and the beamId provides all the information necessary for the O-RU to select the correct beam (or weight table or beam attributes from which to create a beam). It is intended that the beamId be global for the O-RU meaning there are 32767 possible beams shared within the O-RU for all Rtcids/Pcids and shared between UL and DL (beamId=0x0000 is reserved for no beamforming). The specific mapping of beamId to weight table, beam attributes, directionality, beam adjacency or any other beam designator is specific to the O-RU design and shall be conveyed via M-Plane from the O-RU to O-DU upon startup.

An upper bound on the number of beamIds that can be addressed per eAxC or per sets of eAxCs, per symbol and per slot, for DL and for UL respectively, shall be conveyed via M-Plane messaging as part of the O-RU capabilities description. In addition to this, for O-RUs with C-Plane message processing limits, additional limits to restrict the number of highest priority sections per C-Plane message on top of eAxC limits shall be applied (clause 7.8.2).

Value range: {000 0000 0000 0001b-111 1111 1111 1111b; 000 0000 0000 0000b means no BF to be done}.

Type: unsigned integer.

Field length: 15 bits.

Default Value: 0000 0000 0000 0000b (no beamforming).

7.5.3.10 ueId

Description: This parameter provides a logical identifier for the set of channel information associated with a spatial stream of a UE sent via Section Type 6. This shall be used to support channel information sending from the O-DU to the O-RU. This is just a label and the specific value has no meaning regarding types of UEs that may be supported within the system.

If the O-RU reports the capability of 'non-scheduled-ueId-supported' = TRUE for an endpoint, and the O-DU configures 'non-scheduled-ueId-enabled' = TRUE, and if the O-DU sets ueId to 0x7FFF in SE 10 with beamGroupType set to 10b when used with ST 5, the O-RU shall interpret that the PRBs in the section description are not scheduled for the specific eAxC_ID (refer to clause 7.9.13 for more details).

Value range: {000 0000 0000 0000b-111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 15 bits.

7.5.3.11 freqOffset (frequency offset)

Description: This parameter defines the frequency offset with respect to the carrier center frequency before additional filtering (e.g. for PRACH) and FFT processing (in UL) in steps of one half the subcarrier spacings Δf . The frequency

offset shall be consistent per data section (one cannot use reMask to allow different frequency offsets for different REs in the PRBs).

Frequency offset in Hz is calculated as:

$$\text{frequency_offset} = \text{freqOffset} \times \Delta f \times 0.5$$

where Δf is frequency in Hz corresponding to subcarrier spacing provided in frameStructure (see clause 7.5.2.13).

Frequency span resulting from frameStructure, freqOffset, startPrbc, numPrbc, and rb shall not exceed channel bandwidth configured for eAxC over M-Plane.

Value range: { 0x000000=no offset,
0x000001 – 0x7FFFFFF = positive frequency offset,
0x800000 – 0xFFFFFFFF = negative frequency offset }.

Type: signed integer.

Field length: 24 bits.

7.5.3.12 regularizationFactor (regularization factor)

Description: This parameter provides a signed value to support MMSE operation within the O-RU when beamforming weights are supported in the O-RU, so related to Section Type 6.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.5.3.13 cilsample, ciQsample (channel information I and Q values)

Description: These values are the channel information complex values relayed from the O-DU to the O-RU, related to Section Type 6. In absence of Section Extensions 21 and the M-Plane configuration **ci-prb-group-size**, the order of transmission shall be first Prbc for the first antenna to the last antenna, then second Prbc for the first to last antenna, and so on until reaching the last Prbc for the first to last antenna. In the presence of Section Extension 21 or the M-Plane configuration **ci-prb-group-size**, the order of transmission shall be first PRB group for the first antenna to the last antenna, then second PRB group for the first to last antenna, and so on until reaching the last PRB group for the first to last antenna. Each PRB group for one antenna shall include one pair of ciIsample/ciQsample. The bit-width per I and Q value shall be variable (determined by M-Plane configuration **channel-information-bitwidth** if channel information compression is not configured, and determined by ciCompHdr field when either static or dynamic channel information compression is configured) so after the very last Q value, some number of padding (set to zero) bits may be inserted to get to the next byte boundary. When the optional "little endian byte order" is chosen via M-Plane, the ciIsample/ciQsample shall use little endian byte order to transmit the complex numbers. Refer to Annex D.2 for details.

If an O-RU supports receiving and processing channel information with PRB group size greater than one PRB (indicated via M-Plane) and if O-RU is configured via M-Plane to receive and process channel information with PRB group size greater than one PRB, then O-RU shall apply every pair of ciIsample/ciQsample for one PRB group received in Section Type 6 to **ci-prb-group-size** number of continuous PRBs start from index $\text{startPrbc} + i \times \text{ci-prb-group-size}$, where **ci-prb-group-size** is configured via M-Plane, and index i ranges from 0 to numPrbc-1. If M-Plane configuration **ci-prb-group-size** exists, the numPrbc in Section Type 6 indicates the number of PRB groups per antenna and is used to determine the length of the Section Type 6 message, and the actual number of PRBs to which Section Type 6 is applied is determined by the method described above. Also, numPrbc shall not be set to 0 when **ci-prb-group-size** is configured. If $\text{startPrbc} + i \times \text{ci-prb-group-size}$ is larger than the maximum available PRBs configured via M-Plane (**number-of-prb**), the channel information shall only be applied to those PRBs within the maximum available PRBs.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b} for each I and Q value.

Type: signed integer.

Field length: variable bit-width per I and Q value.

7.5.3.14 laaMsgType (LAA message type)

Description: This parameter defines the LAA message type being conveyed within the Section Type 7 C-Plane message. For LAA, the C-Plane messages may flow from the O-DU to the O-RU (as usual) or from the O-RU to the O-DU. All of these messages relate to the Listen Before Talk (LBT) LAA operation and provide a "handshake" between the O-DU and the O-RU to manage the LBT operation. The laaMsgType definitions shall conform to Table 7.5.3.14-1.

Value range: {0001b-1111b}.

Table 7.5.3.14-1: laaMsgType definition

laaMsgType	IssMsgType definition	IssMsgType meaning
0000b	LBT_PDSCH_REQ	O-DU to O-RU request to obtain a PDSCH channel
0001b	LBT_DRS_REQ	O-DU to O-RU request to obtain the channel and send DRS
0010b	LBT_PDSCH_RSP	O-RU to O-DU response, channel acq success or failure
0011b	LBT_DRS_RSP	O-RU to O-DU response, DRS sending success or failure
0100b	LBT_Buffer_Error	O-RU to O-DU response, reporting buffer overflow
0101b	LBT_CWCONFIG_REQ	O-DU to O-RU request, congestion window configuration
0110b	LBT_CWCONFIG_RSP	O-RU to O-DU response, congestion window config. response
0100b – 1111b	reserved for future methods	

Type: unsigned integer.

Field length: 4 bits.

7.5.3.15 laaMsgLen (LAA message length)

Description: This parameter defines number of 32-bit words in the LAA section, including the word containing the IssMsgLen parameter. Zero is a reserved value.

Value range: {0001b-1111b=number of 32-bit words in the section from 1 to 16 words (4 to 64 bytes)
0000b is a reserved value}.

Type: unsigned integer.

Field length: 4 bits.

7.5.3.16 lbtHandle

Description: This parameter provides a label that is included in the configuration request message (e.g., LBT_PDSCH_REQ, LBT_DRS_REQ) transmitted from the O-DU to the O-RU and returned in the corresponding response message (e.g., LBT_PDSCH_RSP, LBT_DRS_RSP).

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.5.3.17 lbtDeferFactor (listen-before-talk defer factor)

Description: Defer factor in sensing slots as described in 3GPP TS 36.213 [14] clause 15.1.1. This parameter is used for LBT CAT 4 and can take one of three values: {1, 3, 7} based on the priority class. Four priority classes are defined in 3GPP TS 36.213 [14] clause 15.1.1.

Value range: {001b, 011b, 111b} or {1, 3, 7} in decimal.

Type: unsigned integer.

Field length: 3 bits.

7.5.3.18 lbtBackoffCounter (listen-before-talk backoff counter)

Description: LBT backoff counter in sensing slots as described in 3GPP TS 36.213 [14] clause 15.1.1.

Value range: {00 0000 0000b – 11 1111 1111b} (0-1023 decimal).

Type: unsigned integer.

Field length: 10 bits.

7.5.3.19 lbtOffset (listen-before-talk offset)

Description: LBT start time in microseconds from the beginning of the subframe scheduled by this message.

Value range: {00 0000 0000b – 11 1110 0111b} or {0 – 999} in decimal.

Type: unsigned integer.

Field length: 10 bits.

7.5.3.20 MCOT (maximum channel occupancy time)

Description: LTE TXOP duration in subframes as described in 3GPP TS 36.213 [14] clause 15.1.1. The maximum values for this parameter are {2, 3, 8, 10} based on the priority class. Four priority classes are defined in 3GPP TS 36.213 [14] clause 15.1.1.

Value range: {1-10} in decimal.

Type: unsigned integer.

Field length: 4 bits.

7.5.3.21 lbtMode (LBT mode)

Description: Part of multi-carrier support. Indicates whether full or partial LBT process is carried out (multi carrier mode B according to 3GPP TS 36.213 [14] clause 15.1.5.2).

00b = full LBT (regular LBT, sending reservation signal until the beginning of the SF/slot).

01b = Partial LBT (looking back 25 µsec prior to transmission as indicated in 3GPP TS 36.213 [14] clause 15.1.5.2).

10b = Partial LBT (looking back 34 µsec prior to transmission as indicated in 3GPP TS 36.213 [14] clause 15.1.5.2).

11b = full LBT and stop (regular LBT, without sending reservation signal; O-RU senses the spectrum for the defer factor + a sensing slot right before OTA when the O-DU is ready to transmit data as indicated in 3GPP TS 36.213 [14] clause 15.1.1. i.e., right before the SF/slot boundary).

Value range: {00b - 11b}.

Type: unsigned integer.

Field length: 2 bits.

7.5.3.22 lbtPdschRes (LBT PDSCH result)

Description: LBT result of SFN/SF:

- 00b = not sensing – indicates that the O-RU is transmitting data.
- 01b = currently sensing – indicates the O-RU has not yet acquired the channel.
- 10b = success – indicates that the channel was successfully acquired.
- 11b = Failure – indicates expiration of the LBT timer. The LBT process should be reset.

Value range: {00b - 11b}.

Type: unsigned integer.

Field length: 2 bits.

7.5.3.23 sfStatus (subframe status)

Description: indicates whether the subframe was dropped or transmitted:

- 0 – subframe was dropped.
- 1 – subframe was transmitted.

Value range: {0,1}.

Type: binary bit.

Field length: 1 bit.

7.5.3.24 lbtDrsRes (LBT DRS result)

Description: LBT result of SFN/SF:

- 0 – SUCCESS – indicates that DRS is sent.
- 1 – FAILURE – indicates that DRS is not sent.

Value range: {0,1}.

Type: binary bit.

Field length: 1 bit.

7.5.3.25 initialPartialSF (Initial partial SF)

Description: Indicates whether the initial SF in the LBT process is full or partial:

- 0 – full SF (two slots, 14 symbols).
- 1 – partial SF (only second slot, last 7 symbols).

Value range: {0, 1}.

Type: binary bit.

Field length: 1 bit.

7.5.3.26 lbtBufErr (LBT buffer error)

Description: Indicates when an LBT buffer overflow has occurred:

- 0 – reserved.
- 1 – buffer overflow – data received at O-RU is larger than the available buffer size.

Value range: {0, 1}.

Type: binary bit.

7.5.3.27 sfnSf (SFN/SF end)

Description: SFN/SF which is dropped by O-RU because of time expired or successfully transmitted.

Value range: SFN: {0000 0000b – 1111 1111b} or {0 – 255} in decimal, SF: {0000b – 1001b} or {0 – 9} in decimal.

Type: unsigned integer.

Field length: 12 bits.

7.5.3.28 lbtCWConfig_H (HARQ parameters for congestion window management)

Description: Total number of HARQ NACK feedback messages received for the lbtCWConfig_T transport blocks transmitted over LAA cell(s) to be used for adjusting the congestion window.

Value range: {0000 0000b – 1111 1111b} or {0 – 255} in decimal.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.29 lbtCWConfig_T (TB parameters for congestion window management)

Description: Total number of Transport Blocks (TB) transmitted over the LAA cell to be used for adjusting the congestion window.

Value range: {0000 0000b – 1111 1111b} or {0 – 255} in decimal.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.30 lbtTrafficClass (traffic class priority for congestion window management)

Description: Channel access priority class as defined in TS 36.213 [14].

Value range: {000 - 111} or {0 - 7} in decimal:

1 - 4: traffic class priority.

0, 5 - 7: reserved.

Type: unsigned integer.

Field length: 3 bits.

7.5.3.31 lbtCWR_Rst (notification about packet reception successful or not)

Description: Notifies to O-DU whether the O-RU receives LBT_CWCONFIG_REQ message successfully or not:

- 0 – SUCCESS – indicates successful reception of LBT_CWCONFIG_REQ.
- 1 – FAILURE – indicates failure of receiving LBT_CWCONFIG_REQ.

Value range: {0, 1}.

Type: binary bit.

Field length: 1 bit.

7.5.3.32 ciCompParam (channel information compression parameter)

Description: This parameter applies to the compression method specified by the associated ciCompMeth value. If ciCompOpt (a subfield of ciCompHdr) is 0, this parameter applies to the following vector of ciIsample, ciQsample for all PRBs of a specific UE. If ciCompOpt (a subfield of ciCompHdr) is 1, this parameter applies to the following vector of ciIsample, ciQsample for all antennas of a specific PRB. ciCompParam shall be configured according to Table 7.5.3.32-1.

Value range: {0000 0000b-1111 1111b}.

Bit allocations:

Table 7.5.3.32-1: ciCompParam bit allocations for various values of ciCompMeth

ciCompMeth	0 (msb)	1	2	3	4	5	6	7 (lsb)	compParam size
000b = no compression	absent								0 octet
001b = block floating point	reserved (set to all zeros)				Exponent (unsigned)				1 octet
010b = block scaling	blockScaler (unsigned, 1 integer bit, 7 fractional bits)								1 octet
011b = μ -law	compBitWidth				compShift				1 octet
100b –111b	reserved (set to all zeros)								? octet

Type: unsigned integer (concatenated bit fields).

Field length: zero for ciCompMeth value 000b, 8 bits for ciCompMeth values 001b, 010b and 011b; other ciCompMeth values may imply other lengths but will always be an integer number of bytes.

7.5.3.33 reserved (reserved for future use)

Description: This parameter is reserved for future use. Transmitter shall send value "0", while receiver shall ignore the value received.

Value range: {variable}.

Type: variable.

Field length: variable.

7.5.3.34 numberOfAcks (number of ACKs)

Description: This parameter is used to indicate the number of ACKs included in the ACK/NACK message for one eAxC Id.

Value range: {0000 0000b-1111 1111b}.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.35 numberOfNacks (number of NACKs)

Description: This parameter is used to indicate the number of NACKs included in the ACK/NACK message for one eAxC Id.

Value range: {0000 0000b-1111 1111b}.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.36 ackId (ACK Id)

Description: This parameter is used to indicate a correctly received section description which carries a Section Extension 22. The ackId in the Section Type 8 message shall be set to the ackNackReqId in the Section Extension 22.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.5.3.37 nackId (NACK Id)

Description: This parameter is used to indicate a section description with Section Extension 22 is received but dropped after Section Extension 22 is decoded. The nackId in the Section Type 8 message shall be set to the ackNackReqId in the Section Extension 22.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.5.3.38 st4CmdType (Section Type 4 command type)

Description: This parameter shall be used specifically with Section Type 4 to specify the command type. Value of this parameter specifies a unique command type value which specifies a slot level configuration to be applied to single/multiple eAxC_IDs. Section Type 4 cannot be used without a Section Type 4 command.

Value range: {0000 0000b-1111 1111b}.

ST4CmdType	ST4CmdType definition	ST4CmdType meaning
0000 0001b	TIME_DOMAIN_BEAM_CONFIG	Command to send time-domain beam group uniquely identifying time-domain beam number vector and/or associated time-domain beamforming weights
0000 0010b	TDD_CONFIG_PATTERN	Command specifying symbol direction: UL/DL and the guard symbols (DL->UL switch), meant for dynamic TDD configurations use cases.
0000 0000b, 0000 0011b – 1111 1111b	reserved for future command types	

Type: unsigned integer.

Field length: 8 bits.

7.5.3.39 st4CmdLen (Section Type 4 command length)

Description: This parameter provides the length of the section type 4 command in units of 32-bit (or 4-byte) words. The value zero is reserved, as there is always at least one word in the Section Type 4 command. If multiple Section Type 4 commands are present in a C-Plane Section Type 4 message, this value shall be used to specify the length of each command independently

Value range: {all zeros – all ones} – from one word to 255 words (8 bits). The value zero is reserved.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.40 numSlots (number of slots)

Description: This parameter shall specify the number of contiguous slots for which a given Section Type 4 command is applicable. If the command configuration is expected to be applicable until changed, the value of this parameter shall be set to zero.

Command configuration specified using a given 'st4CmdType' shall apply to the group of endpoints associated with the 'st4CmdScope' in the section header for the duration specified by 'numSlots'. Any new Section Type 4 command having the same 'st4CmdType' and the same 'st4CmdScope' shall be applicable only after the expiration of duration equal to 'numSlots'.

Value range: {0000 0000b – 1111 1111b}.

0000 0000b – command configuration shall last until updated by the next command of the same 'st4CmdType'.

Type: unsigned integer.

Field length: 8 bits.

7.5.3.41 ackNackReqId (ACK/NACK Request ID)

Description: refer to clause 7.7.22.2 for more details. This parameter is a field in the common header part of each Section Type 4 command. This field enables the O-DU to request the O-RU for ACK/NACK feedback using a Section Type 8 C-Plane message for a specific Section Type 4 command. The value of this parameter shall be set to zero for two specific cases.

Case 1: If an O-RU reports its capability of NOT supporting Section Type 8.

Case 2: If an O-RU reports its capability of supporting Section Type 8, but the O-DU chooses not to request for Section Type 8 ACK/NACK feedback for the Section Type 4 command.

Value range: {0000 0000 0000 0000b – 1111 1111 1111 1111b}.

0000 0000 0000 0000b – No Section Type 8 ACK/NACK C-Plane message expected from O-RU for the 'st4CmdType'. The value of zero for 'ackNackReqId' shall be interpreted as no request for ACK/NACK ST 8 message from O-RU when used in Section Type 4 C-Plane message commands.

Type: unsigned integer.

Field length: 16 bits.

7.5.3.42 symbolMask (TIME_DOMAIN_BEAM_CONFIG symbol mask)

Description: refer to clause 7.7.6.4 for the format of this parameter. This parameter applies to Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG. This parameter specifies the symbols to which time-domain beamforming config specified in the command shall apply. Bits in symbolMask corresponding to symbols earlier than symbol identified by startSymbolId shall be set to 0 by the O-DU.

7.5.3.43 disableTDBFNs (disable sending tdBeamNum)

Description: This parameter is used to enable/disable sending of beamforming numbers 'tdBeamNum' as part of Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG. Disable sending time-domain beamforming numbers implies that once the tdBeamGrp and associated time-domain beam number vector mapping is established in the O-RU, the O-DU may choose to send only tdBeamGrp for associated 'tdBeamNum' vectors. This value shall be set to '1b' only when the mapping between the time-domain beamforming weights and [tdBeamGrp, tdBeamNum] has been established by sending tdBeamGrp, tdBeamNum and per tdBeamNum time-domain beamforming weight in an earlier ST 4 message. Refer to Table 7.4.6-5.

Value range: {0b = time-domain beam numbers included in TIME_DOMAIN_BEAM_CONFIG command.

1b = time-domain beam numbers excluded in TIME_DOMAIN_BEAM_CONFIG command}.

Type: binary bit.

Field length: 1 bit.

7.5.3.44 tdBeamGrp (time-domain beam group identity)

Description: This parameter is used for Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG. This parameter uniquely identifies the time-domain beam number vector (refer to time-domain beam-number (#) in Figure 12.4.3.4.3-1). Each tdBeamGrp is a unique identifier of a vector of size equal to K' i.e., K' 'tdBeamNum(s)'. Only one tdBeamGrp shall be specified with one instance of 'TIME_DOMAIN_BEAM_CONFIG' which shall apply to 'symbolMask' specified in this command header.

Value range: {000 0000 0000 0000b -111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 15 bits.

7.5.3.45 bfwCompHdr (beamforming weight compression header)

Description: refer to clause 7.7.1.2.

This parameter defines the compression method and IQ bit width for the beamforming weights when specifying beamforming weights in Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG. The time-domain beamforming weight compression method may not be the same as the frequency-domain beamforming weight compression method, the compression can be completely independent and can use any of the methods supported by O-RU.

7.5.3.46 disableTDBFWs (disable sending time-domain beamforming weights)

Description: This parameter is used to enable/disable sending of beamforming weights as part of Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG. Disabling sending time-domain beamforming weights allows an operation such that only 'tdBeamNum' associated with 'tdBeamGrp' is sent in the command and not the actual time-domain beamforming weights. This value shall be set to '1b' only when the mapping between the tdBeamNum and time-domain beamforming weights is established by sending tdBeamNum and associated time-domain beamforming weights in an earlier ST 4 message. Refer to Table 7.4.6-4.

Value range: {0b = time-domain beamforming weights included in TIME_DOMAIN_BEAM_CONFIG command.
1b = time-domain beamforming weights excluded in TIME_DOMAIN_BEAM_CONFIG command}.

Type: binary bit.

Field length: 1 bit.

7.5.3.47 bfwCompParam (beamforming weight compression parameter)

Description: refer to clause 7.7.1.3.

Applies to Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG.

7.5.3.48 bfwI, bfwQ (beamforming weights I and Q values)

Description: refer to clause 7.7.1.4 for bfwI and 7.7.1.5 for bfwQ.

Applies to Section Type 4 command type: TIME_DOMAIN_BEAM_CONFIG.

7.5.3.49 dirPattern (symbol data direction (gNB Tx/Rx) pattern)

Description: This parameter is a bitmask for symbols in a slot where each bit of bitmask specifies the symbol's data direction. Bits in dirPattern corresponding to the symbols that are guard symbols shall be ignored by the O-RU. Bits in dirPattern corresponding to symbols earlier than the symbol identified by startSymbolId shall be set to 0 by the O-DU and ignored by the O-RU.

Value of 1 = DL symbol, Value of 0 = UL symbol. Applies to Section Type 4 command type: TDD_CONFIG_PATTERN.

LSB of dirPattern indicates symbol zero (the first symbol in a slot).

Value range: {00 0000 0000 0000b-11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

7.5.3.50 guardPattern (guard pattern bitmask)

Description: This parameter is a bitmask for symbols in a slot used in conjunction with the dirPattern bitmask. Each bit of the bitmask associated with a given symbol specifies a guard symbol for gNB DL to UL switch. Bits in guardPattern corresponding to symbols earlier than the symbol identified by startSymbolId shall be set to 0 by the O-DU and ignored by the O-RU.

Value of 1 = guard symbol, 0 = non-guard symbol.

LSB of guardPattern indicates symbol zero (the first symbol in a slot).

Value range: {00 0000 0000 0000b-11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

7.6 Section Extension elements

7.6.1 Overview

Table 7.6.1-1 lists all the Section Extension parameters defined within the C-Plane:

Table 7.6.1-1: Section Extension commands

extType	meaning	extLen	extension parameters	octets	meaning
0	reserved	1 (1 word)	reserved reserved	1 1	for future use for future use
1	beamforming weights	var	bfwCompHdr bfwCompParam bfwI (for TRX 0) bfwQ (for TRX 0) ... bfwI (for last TRX) bfwQ (for last TRX)	1 1 var var var var	IqWidth(3:0) compMeth(3:0) depends on compr. method beamforming weight I value beamforming weight Q value beamforming weight I value beamforming weight Q value
2	beamforming attributes	var	bfaCompHdr bfAzPt bfZePt bfAz3dd bfZe3dd bfAzSl bfZeSl	2 var var var var 3b 3b	BF attributes compr. header BF azimuth pointing param BF zenith pointing param BF azimuth beamwidth param BF zenith beamwidth param BF azimuth sidelobe param BF zenith sidelobe param
3	DL Precoding configuration parameters and indications	var (3 or 4 words)	codebookIndex layerId txScheme numLayers crsReMask crsSymNum crsShift beamIdAP1 beamIdAP2 beamIdAP3	1 4b 4b 4b 12b 4b 1b 15b 15b 15b	precoder codebook layer ID for DL Tx transmission scheme number of layers in DL Tx CRS RE Mask CRS symbol number CRS shift command Beam ID, Antenna Port 1 Beam ID, Antenna Port 2 Beam ID, Antenna Port 3
4	modulation compr. params	1 (1 word)	csf modCompScaler	1b 15b	constellation shift flag mod. compr. scale value
5	modulation compression additional scaling parameters	var	mcScaleReMask csf mcScaleOffset	12b 1b 15b	Position of same scaling bits constellation shift flag added mod. compr. scale values
6	Non-contiguous PRB allocation	2 (2 words)	repetition rbgSize rbgMask priority symbolMask	1b 3b 28b 2b 14b	repetition flag number of PRBs in the group mask of RBGs in the symbol(s) priority of section description mask of symbols in the slot
7	Multiple-eAxC designation	1	eAxCmask	16b	eAxC mask
8	regularization factor	1	regularizationFactor	16b	regularization factor
9	Dynamic Spectrum Sharing parameters	1	technology	1	interface name
10	Multiple ports grouping	var	beamGroupType numPortc beamId (or ueId)	2b 6b 15b	type of beam grouping the number of ports beam ID (or UE ID)

extType	meaning	extLen	extension parameters	octets	meaning
11	Flexible BF weights	var	disableBFWs RAD numBundPrb bfwCompHdr bfwCompParam for bundle 0 beamId (for PRB bundle 0) bfwI (for TRX 0, bundle 0) bfwQ (for TRX 0, bundle 0) ... bfwI (for last TRX, bundle 0) bfwQ (for last TRX bundle 0) ... bfwCompParam for last bundle beamId (for last PRB bundle) bfwI (for TRX 0, last bundle) bfwQ (for TRX 0, last bundle) ... bfwI (for last TRX, last bundle) bfwQ (for last TRX & bundle)	1b 1b 1 1 var 2 var var ... var var ... var 2 var var ... var var	disable beamforming weights Reset After (PRB) Discontinuity Number of bundled PRBs bitWidth(3:0) compMeth(3:0) depends on compr. method Beam ID BF weight I value for bundle 0 BF weight Q value for bundle 0 ... BF weight I value for bundle 0 BF weight Q value for bundle 0 ... depends on compr. method Beam ID BF weight I value for last bundle BF weight Q value for last bundle ... BF weight I value for last bundle BF weight Q value for last bundle
12	non-contiguous PRB allocation with frequency ranges	var	priority, symbolMask offStartPrb(1) numPrb(1) ... offStartPrb(R-1) numPrb(R-1)	2b 14b 8b 8b ... 8b 8b	priority of section description mask of symbols in the slot offset to start of PRB range #1 number of PRBs in the range #1 ... offset to start of PRB range number of PRBs in the range
13	frequency hopping	var	nextSymbolId(1) nextStartPrbc (1) ... nextSymbolId (R-1), nextStartPrbc (R-1)	4b 10b ... 4b 10b	start symbol of hop #1 start PRB for hop #1 ... start symbol of hop #R-1 start PRB for hop #R-1
14	Null-layer Info. for ueId-based beamforming	var	nullLayerInd	1	Nulling-layer indication
15	Mixed-numerology Info. for ueId-based beamforming	var	frameStructure freqOffset cpLength	1 3 2	FFT size, mu (SCS) Frequency offset Cyclic prefix length
16	Antenna mapping for UE channel Info based UL beamforming	var	antMask ... antMask	8 var	Bitmask of Max. 64 antenna Bitmask of Max. 64 antenna
17	User port group indication	var	numUeId	4b	Number of ueIds per user
18	Uplink Transmission Management	2 (2 words)	transmissionWindowOffset transmissionWindowSize toT	16b 14b 2b	transmission window offset transmission window size type of transmission

extType	meaning	extLen	extension parameters	octets	meaning
19	Compact multiple port beamforming information	var	disableBFWs repetition priority numPortc symbolMask bfwCompHdr ... portReMask(for port 1) portSymbolMask(for port 1) ... bfwCompParam(for port 1) ... beamId(for port 1) bfwI (for port 1 and TRX 0) ... bfwQ (for port 1 and TRX 0) ... portReMask(for last port) portSymbolMask(for last port) bfwCompParam(for last port) ... beamId(for last port) bfwI (for last port and TRX 0) ... bfwQ (for last port and TRX 0) ...	1b 1b 2b 6b 14b 8b ... 12b 14b ... var ... 2 var ... var ... 12b 1b 1b 3b 28b 14b ... 10b ... 10b 12b 1b 1b 3b 28b	disable BF weights repetition flag priority of section description num ports resource symbol bitmask BF weight compression header ... RE bitmask for port 1 Symbol bitmask for port 1 ... BF weight compression parameter for port 1 ... beam identifier for port 1 BF weight in-phase value for port 1 TRX0 ... BF weight quadrature-phase value port 1 TRX0 ... RE bitmask for last port Symbol bitmask last port BF weight compression parameter for last port ... beamId for last port BF weight in-phase value for last port and TRX0 ... BF weight quadrature-phase value last port TRX0
20	Dedicated puncturing	var	numPuncPatterns symbolMask(1) ... startPuncPrb(1) ... numPuncPrb(1) puncReMask(1) rb(1) rbgIncl(1) rbgSize(1) rbgMask(1) symbolMask(last) ... startPuncPrb(last) ... numPuncPrb(last) puncReMask(last) rb(last) rbgIncl(last) rbgSize(last) rbgMask(last)	8b 14b ... 10b ... 10b 12b 1b 1b 3b 28b 14b ... 10b ... 10b 12b 1b 1b 3b 28b	number of puncturing patterns first puncturing pattern symbol mask ... first PRB to which puncturing pattern applies ... number of contiguous PRBs to which first puncturing pattern applies ... first puncturing pattern RE mask ... first RB indicator first rbg included flag first rbg size first rbg bitmask last puncturing pattern symbol mask ... last PRB to which puncturing pattern applies ... number of contiguous PRBs to which last puncturing pattern applies ... last puncturing pattern RE mask last RB indicator last rbg included flag last rbg size last rbg bitmask
21	Variable PRB group size for Channel Information	1 (1 word)	ciPrbGroupSize	1	channel information PRB group size
22	ACK/NACK request	1 (1 word)	ackNackReqId	2	used to identify the section description for which ACK/NACK feedback is requested

extType	meaning	extLen	extension parameters	octets	meaning
23	Multiple symbol modulation compression parameters	var	numSymPrbPattern	4b	number of symbol and resource block pattern
			symMask	14b	symbol mask part of symPrbPattern
			prbPattern	4b	physical resource block pattern part of symPrbPattern
			numMcScaleOffset	4b	number of modulation compression scaling value per symPrbPattern
			mcScaleReMask	12b	modulation compression power scale RE mask
			csf	1b	constellation shift flag
24-127	reserved	1 (1 word)	reserved reserved	1 1	for future use for future use

7.6.2 Common parameters

7.6.2.1 extType (extension type)

Description: This parameter provides the extension type which provides additional parameters specific to the subject data extension. An O-RU or O-DU receiving a "reserved" Section Extension shall ignore the extension and all parameters contained within it.

Value range: {all zeros – all ones}. See Table 7.6.1-1 for values.

Type: unsigned integer.

Field length: 7 bits.

7.6.2.2 ef (extension flag)

Description: This parameter is used to indicate if there is another extension present (ef=1) or this is the last extension (ef=0).

Value range: {0b=no more extensions; 1b=one or more extensions are included after this one}.

Type: binary bit.

Field length: 1 bit.

7.6.2.3 extLen (extension length)

Description: This parameter provides the length of the Section Extension in units of 32-bit (or 4-byte) words. The value zero is reserved, as there is always at least one word in the extension (the word containing the extType and extLen).

Value range: {all zeros – all ones} – from one word to 255 words (8 bits) or from one word to 65535 words (16 bits). The value zero is reserved.

Type: unsigned integer.

Field length: 8 bits for all Section Extensions except for Section Extension 11, 19 and 20 for which the field length is 16 bits.

7.7 Coding of Section Extension IEs

7.7.1 SE 1: Beamforming weights

7.7.1.1 Overview

Section Extension 1 applies to the sending of beamforming weights from the O-DU to the O-RU. When Section Extension 1 is present in a section description, O-RU shall associate the set of beamforming weights conveyed in Section Extension 1 with the beamId conveyed in the section description. In subsequent C-Plane messages, the same beamId can be invoked without sending the beamforming weights again.

NOTE: This allows downloaded weights to have "persistence" which saves DL throughput by not requiring sending of weights multiple times.

The parameter K is defined in clause 12.5.3. Refer to Annex D.1 for details on the usual weight byte order. The optional "little endian byte order" is applied to bfwI/bfwQ fields if configured via M-Plane. Refer to Annex D.2 for details of little endian byte order. This Section Extension applies only to Section Types 1 and 3. The structure of Section Extension 1 is presented in Table 7.7.1.1-1.

Table 7.7.1.1-1: Format of Section Extension 1 (beamforming weights)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x1							1	Octet N
extLen								1	Octet N+1
bfwCompHdr								1	Octet N+2
bfwCompParam								var	Octet N+3
bfwI (for TRX 0)								var	
bfwQ (for TRX0)								var	
remaining beamforming weights bfwI and bfwQ up to K TRXs								var	
zero pad to 4-byte boundary								var	

7.7.1.2 bfwCompHdr (beamforming weight compression header)

Description: This parameter defines the compression method and IQ bit width for the beamforming weights in the specific section in the C-Plane message. In this way each set of weights may employ a separate compression method. For the block compression methods, the block size is the entire vector of beamforming weights, not some subset of them.

If the O-RU supports static beamforming weight compression through M-Plane configuration, and static compression is configured, this field shall be set to a value corresponding to the compression method configured via M-Plane. If the O-RU supports dynamic beamforming weight compression, then O-DU may configure dynamic compression via M-Plane with or without providing a list of compression methods that will be used dynamically (i.e., a subset of methods supported by O-RU). If dynamic compression is configured via M-Plane without providing a list of compression methods, this field shall be set to a value corresponding to one of dynamic compression methods supported by O-RU. If dynamic compression is configured with a list of compression methods, then this field shall be set to a value corresponding to one of compression methods in the list. If bfwCompHdr parameter value in a C-Plane message and M-Plane configuration are not compatible, then O-RU may discard the message or process the message as per M-Plane configuration or as per bfwCompHdr value.

Value range: {0000 0000b-1111 1111b}

Bit allocations: see Table 7.7.1.2-1

Table 7.7.1.2-1: bfwCompHdr definition

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
bfwIqWidth (see Table 7.7.1.2-2)				bfwCompMeth (see Table 7.7.1.2-3)			1	Octet 1	

Table 7.7.1.2-2: bfwIqWidth definition

bfwIqWidth	Bit width of each I and each Q
0000-1111b	Value of bfwIqWidth except a value of zero means 16 bits. EXAMPLE 1: bfwIqWidth = 0000b means I and Q are each 16 bits wide. EXAMPLE 2: bfwIqWidth = 0001b means I and Q are each 1 bit wide. EXAMPLE 3: bfwIqWidth = 1111b means I and Q are each 15 bits wide.

Table 7.7.1.2-3: bfwCompMeth definition

bfwCompMeth	compression method	bfwIqWidth meaning
0000b	no compression	bitwidth of each uncompressed I and Q value
0001b	block floating point	bitwidth of each I and Q mantissa value
0010b	block scaling	bitwidth of each I and Q scaled value
0011b	μ -law	bitwidth of each compressed I and Q value
0100b	beamspace compression type I	bitwidth of each beamspace I and Q coefficient
0101b	beamspace compression type II	bitwidth of each beamspace I and Q coefficient
0110b – 1111b	reserved for future methods	N.A.

Type: unsigned integer (concatenated bit fields).

Field length: 8 bits.

Default Value: 0000 0000b (no compression, 16-bit I and Q).

7.7.1.3 bfwCompParam (beamforming weight compression parameter)

Description: This parameter applies to the compression method specified by the associated bfwCompMeth value.

Value range: {0000 0000b-1111 1111b}.

Bit allocations: As presented in Table 7.7.1.3-1, allocation of bits in bfwCompParam depends on value of subfield bfwCompMeth of field bfwCompHdr described in clause 7.7.1.2.

Table 7.7.1.3-1: bfwCompParam definition

bfwCompMeth	0 (msb)	1	2	3	4	5	6	7 (lsb)	bfwCompParam size
0000b = no compression	absent								0 octets
0001b = block floating point	reserved (set to all zeros)				exponent (unsigned)				1 octet
0010b = block scaling	blockScaler (unsigned, 1 integer bit, 7 fractional bits)								1 octet
0011b = μ -law	compBitWidth				compShift				1 octet
0100b = beamspace I	activeBeamspaceCoefficientMask								ceil(K/8) octets (see NOTE 1)
	blockScaler (unsigned, 1 integer bit, 7 fractional bits)								
0101b = beamspace II	activeBeamspaceCoefficientMask								ceil(K/8) octets (see NOTE 1)
	reserved (set to all zeros)				exponent (unsigned) (see NOTE 2)				
NOTE 1: K is the number of elements in uncompressed beamforming weight vector (see clause 12.5.3 and Annex J). K is O-RU-specific and is calculated from parameters describing tx-array or rx-array, conveyed from the O-RU to the O-DU via the M-Plane as part of the initialization procedure.									
NOTE 2: For bfwCompMeth value 0101b (beamspace II), block floating point compression is used for the beamspace coefficients.									

Type: unsigned integer (concatenated bit fields).

Field length: an integer number of bytes; the length is given in column "bfwCompParam size" in Table 7.7.1.3-1.

7.7.1.4 bfwl (beamforming weight in-phase value)

Description: This parameter is the in-phase beamforming weight value. Refer to Annex D.1 for details on the usual weight byte order. The optional "little endian byte order" can be configured via M-Plane; see Annex D.2 for details on byte order. The maximum number of weights in the Section Extension is O-RU-specific and will be conveyed from the O-RU to the O-DU via the M-Plane as part of the initialization procedure. For bfwCompMeth indicating beamspace I and beamspace II the number of weights in the Section Extension is determined by activeBeamspaceCoefficientMask, for other bfwCompMeth values the number of weights in the Section Extension is always the maximum.

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 1-16 bits.

7.7.1.5 bfwQ (beamforming weight quadrature value)

Description: This parameter is the quadrature beamforming weight value. Refer to Annex D.1 for details on the usual weight byte order. The optional "little endian byte order" can be configured via M-Plane; refer to Annex D.2 for detail on byte order. The maximum number of weights in the Section Extension is O-RU-specific and will be conveyed from the O-RU to the O-DU via the M-Plane as part of the initialization procedure. For bfwCompMeth indicating beamspace I and beamspace II the number of weights in the Section Extension is determined by activeBeamspaceCoefficientMask, for other bfwCompMeth values the number of weights in the Section Extension is always the maximum.

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 1-16 bits.

7.7.2 SE 2: Beamforming attributes

7.7.2.1 Overview

Section Extension 2 applies only to Section Types 1 and 3. Section Extension 2 applies to the sending of beamforming attributes from the O-DU to the O-RU which is described in Annex J.

The Table 7.7.2.1-1 shows the format of Section Extension 2.

Table 7.7.2.1-1: Section Extension 2 Data Format

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x02							1	Octet N
extLen								1	Octet N+1
bfaCompHdr								2	Octet N+2
bfAzPt								var	Octet N+3
bfZePt								var	
bfAz3dd								var	
bfZe3dd								var	
zero-padding	bfAzSl				bfZeSl			1	
zero padding to achieve 4-byte alignment as needed									

7.7.2.2 bfaCompHdr (beamforming attributes compression header)

Description: This parameter defines the bit width for the beamforming attributes extension parameters. In this way each set of beamforming attributes may employ a different bit width.

Value range: {0000 0000 0000b-1111 1111 1111b}

Bit allocations: Table 7.7.2.2-1 shows bit allocations within bfaCompHdr.

Table 7.7.2.2-1: bfaCompHdr definition

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
Reserved		bfAzPtWidth (see Table 7.7.2.2-2)			bfZePtWidth (see Table 7.7.2.2-3)			1	Octet 1
Reserved		bfAz3ddWidth (see Table 7.7.2.2-4)			bfZe3ddWidth (see Table 7.7.2.2-5)			1	Octet 2

Table 7.7.2.2-2: bfAzPtWidth definition

bfAzPtWidth	Bit width of bfAzPt
000-111b	value of bfAzPtWidth

Table 7.7.2.2-3: bfZePtWidth definition

bfZePtWidth	Bit width of bfZePt
000-111b	value of bfZePtWidth

Table 7.7.2.2-4: bfAz3ddWidth definition

bfAz3ddWidth	Bit width of bfAz3dd
000-111b	value of bfAz3ddWidth

Table 7.7.2.2-5: bfZe3ddWidth definition

bfZe3ddWidth	Bit width of bfZe3dd
000-111b	value of bfZe3ddWidth

For each of the four bitwidth values in this parameter (bfAzPtWidth, bfZePtWidth, bfAz3ddWidth, and bfZe3ddWidth) the following mapping shall be used:

000b = no bits, the field is not applicable (e.g., O-RU does not support it) or the default value shall be used.

001b = 2-bit bitwidth

010b = 3-bit bitwidth

011b = 4-bit bitwidth

100b = 5-bit bitwidth

101b = 6-bit bitwidth

110b = 7-bit bitwidth

111b = 8-bit bitwidth (this is the highest bitwidth anticipated to be needed)

Type: unsigned integer (concatenated bit fields).

Field length: 16 bits (4 bits are reserved).

Default Value: 0011 1111 0011 1111b (8-bit azimuth and zenith pointing angle and 8-bit azimuth and zenith beamwidth).

7.7.2.3 bfAzPt (beamforming azimuth pointing parameter)

Description: This parameter is the azimuth beamforming pointing angle in degrees. The valid range of values is O-RU-specific and is conveyed from the O-RU to the O-DU as part of the initialization procedure via the M-Plane.

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 0-8 bits.

7.7.2.4 bfZePt (beamforming zenith pointing parameter)

Description: This parameter is the zenith beamforming pointing angle in degrees. The valid range of values is O-RU-specific and is conveyed from the O-RU to the O-DU as part of the initialization procedure via the M-Plane.

Value range: {all zeros – all ones}.

Type: unsigned integer.

Field length: 0-8 bits.

7.7.2.5 bfAz3dd (beamforming azimuth beamwidth parameter)

Description: This parameter is the azimuth beamforming beamwidth in degrees. The valid range of values is O-RU-specific and is conveyed from the O-RU to the O-DU as part of the initialization procedure via the M-Plane. The value (000b) corresponds to the minimum valid beamwidth.

Value range: {all zeros – all ones}.

Type: unsigned integer.

Field length: 0-8 bits.

7.7.2.6 bfZe3dd (beamforming zenith beamwidth parameter)

Description: This parameter is the zenith beamforming beamwidth in degrees. The valid range of values is O-RU-specific and is conveyed from the O-RU to the O-DU as part of the initialization procedure via the M-Plane. The value (000b) corresponds to the minimum valid beamwidth.

Value range: {all zeros – all ones}.

Type: unsigned integer.

Field length: 0-8 bits.

7.7.2.7 bfAzSl (beamforming azimuth sidelobe parameter)

Description: This parameter is the azimuth beamforming sidelobe suppression value in dB. The valid range of values is O-RU-specific and is conveyed from the O-RU to the O-DU as part of the initialization procedure via the M-Plane. The value of bfAzSl corresponds to a value of 10 dB for all zeros and increments in 5 dB steps (e.g. 001b corresponds to 15 dB, 010b corresponds to 20 dB, and so on). The value 111b corresponds to having a sidelobe suppression of 45 dB or more.

Value range: {all zeros – all ones}.

Type: unsigned integer.

Field length: 3 bits.

7.7.2.8 bfZeSl (beamforming zenith sidelobe parameter)

Description: This parameter is the zenith beamforming sidelobe suppression value in dB. The valid range of values is O-RU-specific and is conveyed from the O-RU to the O-DU as part of the initialization procedure via the M-Plane. The value of bfZeSl corresponds to a value of 10 dB for all zeros and increments in 5 dB steps (e.g. 001b corresponds to 15 dB, 010b corresponds to 20 dB, and so on). The value 111b corresponds to having a sidelobe suppression of 45 dB or more.

Value range: {all zeros – all ones}.

Type: unsigned integer.

Field length: 3 bits.

7.7.2.9 zero-padding

Description: This parameter is intended to pad out the data to the next 4-byte boundary. Because the preceding parameters are of varying bitwidth, in most cases padding will be needed to achieve the 4-byte alignment. Transmitter shall send value "0", while receiver shall ignore the value received.

Value range: {all zeros}.

Type: variable.

Field length: variable.

7.7.3 SE 3: DL precoding parameters

7.7.3.1 Overview

Section Extension 3 applies only to Category B O-RU and to Section Types 1 and 3. Section Extension 3 may be used only for precoding schemes Transmit Diversity and Spatial Multiplexing with Large Delay CDD; for other LTE transmission modes and for NR, if there is precoding in the DL processing chain, precoding shall be included in the beamforming operation (that is, encoded in the beamforming weights).

Table 7.7.3.1-1 shows the format of Section Extension 3 for the first data layer.

Table 7.7.3.1-1: Format of Section Extension 3 for the first data layer

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x03							1	Octet N
	extLen = 0x04 (4 words)							1	N+1
	codebookIndex							1	N+2
	layerId = 0000b or 1111b			numLayers[3:0]				1	N+3
	txScheme[3:0]			crsReMask[11:8]				1	N+4
	crsReMask[7:0]							1	N+5
crsShift	reserved			crsSymNum[3:0]				1	N+6
	reserved							3	N+7
reserved	beamIdAP1[14:8]							1	N+10
	beamIdAP1[7:0]							1	N+11
reserved	beamIdAP2[14:8]							1	N+12
	beamIdAP2[7:0]							1	N+13
reserved	beamIdAP3[14:8]							1	N+14
	beamIdAP3[7:0]							1	N+15

Table 7.7.3.1-2 shows the format of Section Extension 3 for the non-first data layer.

Table 7.7.3.1-2: Format of Section Extension 3 for non-first data layer

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x03							1	Octet N
	extLen = 0x01 (1 word)							1	N+1
	codebookIndex							1	N+2
	layerId ≠ 0000b or 1111b			numLayers				1	N+3

LTE precoding schemes Transmit Diversity and Spatial Multiplexing with Large Delay CDD support two or four antenna ports, hence two or four beamIds are needed (same beamId for user data and CRS REs). For Antenna Port 0,

the beamId is contained in the C-Plane data section header, while the Antenna Ports 1-3 beamIds are contained in this Section Extension. When there are two antenna ports, the Section Extension only contains the second Antenna Port beam ID ("beamIdAP1") and the Section Extension length is 3 words ("extLen" = 0x3). When there are four antenna ports, the Section Extension contains the second, third and fourth Antenna Port beam IDs ("beamIdAP1", "beamIdAP2", and "beamIdAP3") and the Section Extension length is 4 words ("extLen" = 0x4).

For the txScheme indicating Transmit Diversity, one Pcid is used for all the user data and one section instantiation is needed using the corresponding C-Plane Rtcid, providing all the beam IDs (up to 4) for the user data; a second section instantiation (same sectionId) with a different reMask may be used to provide the CRS RE beam IDs (also up to 4). In the Transmit Diversity case the layerId shall be set to 1111b ("TxD").

For the txSchemes indicating Spatial Multiplexing with Large Delay CDD, each layer shall have its own Pcid for the user data with a corresponding C-Plane Rtcid conveying the user data's beamId, with this Section Extension showing a different layer number for each layer. Only within the Rtcid for layer ID zero, the CRS REs will be provided with their beamIds (one beamId in the section header and the other beamIds in this Section Extension). For the non-zero layer number Rtcids, this Section Extension will still be provided to guide the precoding operation (provide the layer ID) but will not include beamIds (extLen = 0x1), and the beamId in the section header should be ignored by the O-RU and should be set to the default value by the O-DU.

Considering scheduling of multiple UEs within a slot, each using Spatial Multiplexing with Large Delay CDD, association of first data layer with respective non-first data layer while using Section Extension 3 at the O-RU is required. To create this association the O-DU shall use a unique eAxC_ID value for each data layer which shall not change within a slot. The value of eAxC_ID for a given layer shall be static at a per slot, per carrier level in C-Plane messages with Section Extension 3. The static value of eAxC_ID per layer shall be used by the O-RU to create eAxC_ID to layerId mapping. This mapping shall be same across all PRB allocations (across UEs) within a slot.

7.7.3.2 codebookIndex (precoder codebook used for transmission)

Description: This parameter is not used and shall be set to zero.

Value range: {0000 0000b - 1111 1111b}.

Type: unsigned integer.

Field length: 8 bits.

DefaultValue: 0000 0000b (used for invalid mode)

7.7.3.3 layerID (Layer ID for DL transmission)

Description: This parameter defines the layer ID that are used for DL transmission.

Value range: {0000b-1111b}. 0000b implies layer0, 0001b implies layer1, 0010b implies layer2, 0011b implies layer3. (for Transmit Diversity, set to all ones)

Type: unsigned integer.

Field length: 4 bits.

DefaultValue: 1111b (used for Transmit Diversity mode)

7.7.3.4 txScheme (transmission scheme)

Description: This parameter defines the precoding scheme used in this Section Type.

Value range: {0000b-1111b} Table 7.7.3.4-1 shows meaning of defined values.

Table 7.7.3.4-1: txScheme values

txScheme	TM scheme
0000b	Spatial Multiplexing with Large Delay CDD
0001b	Reserved
0010b	Transmit Diversity

0011b-111b	Reserved
------------	----------

Type: unsigned integer.

Field length: 4 bits.

7.7.3.5 numLayers (number of layers used for DL transmission)

Description: This parameter defines the number of layers that are used for DL transmission.

Value range: {0000b-1111b}. 0000b implies 1 layer, 0001 implies 2 layers, 0010 implies 3 layers, 0011b implies 4 layers.

Type: unsigned integer.

Field length: 4 bits.

7.7.3.6 crsReMask (CRS resource element mask)

Description: This parameter defines the CRS Resource Element (RE) mask within a PRB. Each bit in the crsReMask indicates if the section control is applicable to the RE sent in U-Plane messages (0=not applicable, 1=applicable). MSB corresponds to the RE of the lowest frequency in a PRB.

Value range: {0000 0000 0000b-1111 1111 1111b}

Type: unsigned integer (bit mask).

Field length: 12 bits.

7.7.3.7 crsSymNum (CRS symbol number indication)

Description: This parameter defines the CRS symbol number within a PRB. The value of the crsSymNum index indicates the symbol number to the RE sent in U-Plane messages.

Value range: {0000b-1111b}, value indicates symbol number.

0000b – 1101b: use symbol number 0 – 13 respectively;

1110b – 1111b: reserved

Type: unsigned integer.

Field length: 4 bits.

7.7.3.8 crsShift (crsShift used for DL transmission)

Description: This parameter indicates the shift pattern to pick up the right index for CRS positions for N Antennas (see Figure I.4-3)

Value range: 0 or 1, implying shift patterns that are layer-dependent according to the Table 7.7.3.8-1 below (see Table I.5-1, Table I.5-2 and Table I.5-3 in Annex I for vShift)

Table 7.7.3.8-1: crsShift values

	1 Layer	2 Layers	4 Layers
crsShift	=0 for $0 \leq \text{vshift} \leq 5$	0 for $0 \leq \text{vshift} \leq 2$ 1 for $3 \leq \text{vshift} \leq 5$	0 for $0 \leq \text{vshift} \leq 2$ 1 for $3 \leq \text{vshift} \leq 5$

Type: binary.

Field length: 1 bit.

7.7.3.9 beamIdAP1 (beam id to be used for antenna port 1)

Description: This parameter defines the beam pattern to be applied to the U-Plane data. beamId = 0 means no beamforming operation will be performed. No beamforming operation implies that the RU shall not apply any phase or amplitude weights to the U-Plane data and that the resulting RF signal will be applied to all antenna elements in the group equally.

Value range: {000 0000 0000 0001b-111 1111 1111 1111b}

Type: unsigned integer.

Field length: 15 bits.

Default Value: 000 0000 0000 0000b (no beamforming).

7.7.3.10 beamIdAP2 (beam id to be used for antenna port 2)

Description: This parameter defines the beam pattern to be applied to the U-Plane data. beamId = 0 means no beamforming operation will be performed. No beamforming operation implies that the RU shall not apply any phase or amplitude weights to the U-Plane data and that the resulting RF signal will be applied to all antenna elements in the group equally.

Value range: {000 0000 0000 0001b-111 1111 1111 1111b}

Type: unsigned integer.

Field length: 15 bits.

Default Value: 000 0000 0000 0000b (no beamforming).

7.7.3.11 beamIdAP3 (beam id to be used for antenna port 3)

Description: This parameter defines the beam pattern to be applied to the U-Plane data. beamId = 0 means no beamforming operation will be performed. No beamforming operation implies that the RU shall not apply any phase or amplitude weights to the U-Plane data and that the resulting RF signal will be applied to all antenna elements in the group equally.

Value range: {000 0000 0000 0001b-111 1111 1111 1111b}

Type: unsigned integer.

Field length: 15 bits.

Default Value: 000 0000 0000 0000b (no beamforming).

7.7.4 SE 4: Modulation compression parameters

7.7.4.1 Overview

Section Extension 4 applies only to Section Types 1, 3 and 5. Section Extension 4 enables the O-DU to convey to the O-RU one set of "csf and modCompScaler values" which is needed for modulation compression described in Annex A.5.

Table 7.7.4.1-1 shows the format of Section Extension 4.

Table 7.7.4.1-1: Format of Section Extension 4 (modulation compression parameters)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x04							1	Octet N
	extLen = 0x01 (1 word)							1	N+1
csf	modCompScaler[14:8]							1	N+2
	modCompScaler[7:0]							1	N+3

7.7.4.2 csf (constellation shift flag)

Description: This binary flag indicates whether to shift the constellation (csf=1) or not (csf=0). "Shift" means subtract from (during compression) or add to (during decompression) the I and Q values the value $2^{-\text{udIqWidth}}$ where "udIqWidth" is the number of I and Q bits in the U-Plane representation (see Table 7.7.4.2-1).

Table 7.7.4.2-1: Constellation shift definition

udIqWidth	Shift value
1	1/2
2	1/4
3	1/8
4	1/16
5	1/32

Value range: {0b-1b}

Type: binary.

Field length: 1 bit.

7.7.4.3 modCompScaler (modulation compression scaler value)

Description: This parameter conveys the scale factor O-RU shall apply to the unshifted constellation points during decompression. It is a fractional floating-point value having an unsigned but negative 4-bit exponent and an unsigned fractional 11-bit mantissa.

Value range: { 0 through $+(1-2^{-11})$ }.

Type: unsigned fractional floating-point value.

The value of the scale factor conveyed as modCompScaler field shall be calculated with:

$$\text{modCompScaler} = \text{mantissa} \times 2^{-\text{exponent}}$$

where exponent is the most significant 4 bits of the 15-bit modCompScaler field and mantissa is the least-significant 11 bits of the modCompScaler field. The values of exponent and mantissa shall be calculated with:

$$\text{mantissa} = \sum_{k=0}^{10} \text{modCompScaler}[k] \times 2^{k-11}$$

$$\text{exponent} = \sum_{k=11}^{14} \text{modCompScaler}[k] \times 2^{k-11}$$

where "modCompScaler[k]" is the kth bit of the modCompScaler field.

Field length: 15 bits.

7.7.5 SE 5: Modulation compression additional parameters

7.7.5.1 Overview

This Section Extension applies only to Section Types 1, 3 and 5. Section Extension 5 enables the O-DU to convey one or more set(s) of mcScaleReMask, csf and mcScaleOffset values to the O-RU which is needed for modulation compression described in Annex A.5. Table 7.7.5.1-1 and Table 7.7.5.1-2 shows the Section Extension format when one set and two sets of "mcScaleReMasks, csf and mcScaleOffset values" are conveyed. Please note that Section Extension 5 may be used to convey more than two sets of "mcScaleReMasks, csf and mcScaleOffset values" in which case the frame structure is extended in similar fashion, i.e., the zero padding bits are added at the end of the Section Extension to maintain 4-byte alignment.

Table 7.7.5.1-1: Format of Section Extension 5 with one scaler value, modulation compression parameters

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x05							1	Octet N
extLen = 0x2 (2 words)								1	N+1
mcScaleReMask[11:4]								1	N+2
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+3
mcScaleOffset [11:4]								1	N+4
mcScaleOffset [3:0]			zero padding					1	N+5
zero padding								1	N+6
zero padding								1	N+7

Table 7.7.5.1-2: Format of Section Extension 5 with two scaler values, modulation compression parameters

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x05							1	Octet N
extLen = 0x03 (3 words)								1	N+1
mcScaleReMask[11:4]								1	N+2
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+3
mcScaleOffset [11:4]								1	N+4
mcScaleOffset [3:0]			mcScaleReMask[11:8]					1	N+5
mcScaleReMask[7:0]								1	N+6
csf	mcScaleOffset [14:8]							1	N+7
mcScaleOffset [7:0]								1	N+8
zero padding								1	N+9
zero padding								1	N+10
zero padding								1	N+11

For a given extLen value, there are two possible cases for the number of sets of 'mcScaleReMask-csf-mcScaleOffset'. For example, when extLen equals to 4, the number of sets of 'mcScaleReMask-csf-mcScaleOffset' may be either 3 or 4; i.e., both cases will fit within 16 bytes (extLen=4). This happens when extLen minus 2 bytes equals to an integer multiple of 3.5 bytes (28bits). In such cases, if the last 28 bits (length of one set) of Section Extension 5 parameters are all set to 0, then O-RU shall consider the smaller number of parameter sets. Otherwise, if they are not all set to 0, O-RU shall consider the larger number of parameter sets. Table 7.7.5.1-3 shows example of the format when 3 or 4 sets of parameters are included in the Section Extension.

Table 7.7.5.1-3: Format of Section Extension 5 when extLen = 4 (three or four modulation compression parameters may be present)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x05							1	Octet N
extLen = 0x04 (4 words)								1	N+1
mcScaleReMask[11:4]								1	N+2
mcScaleReMask[3:0]				csf	mcScaleOffset [14:12]			1	N+3
mcScaleOffset [11:4]								1	N+4
mcScaleOffset [3:0]				mcScaleReMask[11:8]			1	N+5	
mcScaleReMask[7:0]								1	N+6
csf	mcScaleOffset [14:8]							1	N+7
mcScaleOffset [7:0]								1	N+8
mcScaleReMask[11:4]								1	N+9
mcScaleReMask[3:0]				csf	mcScaleOffset [14:12]			1	N+10
mcScaleOffset [11:4]								1	N+11
mcScaleOffset [3:0]				mcScaleReMask[11:8] (or zero padding)			1	N+12	
mcScaleReMask[7:0] (or zero padding)								1	N+13
csf (or zero padding)	mcScaleOffset [14:8] (or zero padding)							1	N+14
mcScaleOffset [7:0] (or zero padding)								1	N+15

7.7.5.2 mcScaleReMask (modulation compression power scale RE mask)

Description: This parameter defines the Resource Element (RE) mask to indicate the position of RE with same scaling and modulation type within a PRB. Each bit setting in the mcScaleReMask indicates if the mcScaleOffset and csf fields are applicable to the RE sent in U-Plane messages or not (0=not applicable; 1=applicable). Most significant bit of this parameter indicates the value for the RE of the lowest frequency in a PRB.

Different REs in a PRB may be indicated by different invocations of mcScaleReMask within the Section Extension 5. If any RE in a PRB is never pointed to by a mcScaleReMask (but other REs in that PRB are), the "missing" RE should be considered to represent not populated REs (e.g. no user data to transmit).

There is a relationship between the mcScaleReMask values and the section's reMask: no bit in any of the mcScaleReMasks shall be set (=1) in a position where the reMask has a zero, and every reMask bit that is set (=1) shall have exactly one bit =1 in one of the mcScaleReMasks.

Value range: {0000 0000 0000b – 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 12 bits.

Default Value: 1111 1111 1111b (all REs in the block applicable).

7.7.5.3 csf (constellation shift flag)

Description: refer to clause 7.7.4.2

7.7.5.4 mcScaleOffset (scaling value for modulation compression)

Description: This parameter is the scale factor to apply to the unshifted constellation points during decompression. It is a fractional floating-point value having an unsigned but negative 4-bit exponent and an unsigned fractional 11-bit mantissa.

Value range: {0 through $+(1-2^{-11})$ }.

Type: unsigned integer.

exponent is the most significant 4 bits of the 15-bit mcScaleOffset field and mantissa is the least-significant 11 bits of the mcScaleOffset field. mcScaleOffset[k] refers to the kth bit of the mcScaleOffset field. Therefore, the actual value of mcScaleOffset is:

$$\text{mcScaleOffset} = \text{mantissa} \times 2^{-\text{exponent}}$$

Field length: 15 bits.

7.7.6 SE 6: Non-contiguous PRB allocation in time and frequency domain

7.7.6.1 Overview

Section Extension 6 applies only to Section Types 1, 3 and 5.

Section Extension 6 enables allocation of non-contiguous sets of PRBs (Resource Block Groups, or RBGs) in frequency and time domain. This reduces the C-Plane overhead when users or channels are allocated with non-contiguous sets of PRBs in time and frequency.

If Section Extension 6 is present in a section description, then the following requirements apply:

- 1) Field startSymbolId in the message header and the fields rb, symInc, and numSymbol in the section description shall not be used for identification of symbols and PRBs referred by the section description.
- 2) O-DU may set symInc and numSymbol to any allowed value without restriction to values corresponding to symbols actually referred to by the section description; the values of these parameters may be used to affect range of symbols referred by subsequent section descriptions in the message.
- 3) O-DU shall set 'rb' bit to the value of zero. However, as an exception, if the O-RU supports the optional capability indicated by setting M-Plane feature flag "se6-rb-bit-supported", then the O-DU may use the "rb" bit. In this case, when the O-DU sets the "rb" bit to one, it shall also set the value of "rbgsize" to zero and the O-RU shall interpret the value of 'rb' bit as applicable to this data section.
- 4) O-DU shall set startSymbolId to the earliest symbol referred by any of section descriptions in the message, including, but not limited to, section descriptions with Section Extension 6. The earliest symbol referenced by a section description with Section Extension 6 is the earliest symbol selected by a set bit in the symbolMask. That means symbolMask's bit n, such that $n < \text{startSymbolId}$, shall be zero.

NOTE 1: Usage of Section Extension 6 does not affect the operation of user plane. A data section in the user plane is not required to contain a contiguous range of PRBs. By invoking the sectionId multiple times in the user plane, either within a single message or in different messages, it is possible to handle sets of non-contiguous PRBs.

NOTE 2: Utilization of the Section Extension 6 does not put any restriction on utilization of sections with contiguous PRB utilization except for the general rules of utilization of sections e.g. in the context of one eAxC, a resource element is referenced by at most one data section description in the control plane.

In case C-Plane and U-Plane coupling via sectionId is used then control plane section descriptions with the same sectionId shall allocate same set of PRBs and symbols.

When utilizing this Section Extension together with Section Type 3 where freqOffset parameter is present, then freqOffset affects the frequency span for the specific range of PRB numbers.

The O-DU shall comply to the limits described in clause 7.8.2.1.2 assuming the number of PRB ranges in the section description with the Section Extension 6 is the number of noncontiguous PRB ranges allocated by Section Extension 6 i.e., number of sequences of 1s in rbgMask separated by at least one 0 as specified in clause 7.7.6.3.

The structure of Section Extension 6 is presented in Table 7.7.6.1-1.

Table 7.7.6.1-1: Section Format for Section Extension 6 (non-contiguous PRB allocation in time and frequency)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x06						1	Octet N	
extLen = 0x02 (2 words)								1	N+1
repetition	rbgSize [2:0]		rbgMask [27:24]					1	N+2
rbgMask [23:16]								1	N+3
rbgMask [15:8]								1	N+4
rbgMask [7:0]								1	N+5
priority[1:0]		symbolMask[13:8]						1	N+6
symbolMask[7:0]								1	N+7

7.7.6.2 rbgSize (resource block group size)

Description: This parameter indicates the size in number of PRBs of the resource block groups allocated by the bit mask. The size of the resource block group to be used by the application shall be taken from Table 7.7.6.2-1 below. See clause 7.7.6.3 for special handling of exceptional resource block groups. If the M-Plane feature flag "se6-rb-bit-supported" is not set, then the value of rbgSize value of zero shall be considered as "reserved" and shall not be used.

Value range: {000b-111b}. Table 7.7.6.2-1 presents meaning of values.

Table 7.7.6.2-1: rbgSize

rbgSize	Number of PRBs in resource block group
000b	rb bit may be valid, see requirement #3, clause 7.7.6.1
001b	1
010b	2
011b	3
100b	4
101b	6
110b	8
111b	16

Type: unsigned integer.

Field length: 3 bits.

Default Value: 001b.

7.7.6.3 rbgMask (resource block group bit mask)

Description: This parameter is a bit mask where each bit indicates whether a corresponding resource block group is allocated. If rbgSize = 0 then O-RU shall ignore rbgMask and apply the section description to all PRBs selected by startPrbc, numPrbc, rb and other section extensions that may be present in the section description. Otherwise, the O-RU shall apply rbgMask as follows. If bit n in the mask is set, then the resource block group n is allocated where n can take values in range [0, lastRbgid] where

$$\text{lastRbgid} = \text{ceiling}((\text{numPrbc} + (\text{startPrbc} \bmod \text{rbgSize})) / \text{rbgSize}) - 1.$$

The range of PRBs included in a resource block group is given as follows:

Resource block group 0 include PRBs with id in the range [startPrbc, startPrbc+f(0)-1].

Resource block group n include PRBs with id in the range [startPrbc + f(0) + (n-1) × f(n), startPrbc + f(0) + n × f(n)-1] where 0 < n < lastRbgid.

Resource block group n = lastRbgid include PRBs with id in the range [startPrbc + numPrbc – f(lastRbgid), startPrbc+ numPrbc -1].

Where f(n) specifies the number of PRBs included in resource block group n and is given as follows:

$$f(0) = \text{rbgSize} - (\text{startPrbc} \bmod \text{rbgSize})$$

$$f(n) = \text{rbgSize}, \text{ where } 0 < n < \text{lastRbgid}.$$

$$\text{if } (\text{startPrbc} + \text{numPrbc}) \bmod \text{rbgsize} > 0 \text{ then } f(\text{lastRbgid}) = (\text{startPrbc} + \text{numPrbc}) \bmod \text{rbgSize}$$

$$\text{else } f(\text{lastRbgid}) = \text{rbgSize}$$

Only PRBs within the range defined by startPrbc and numPrbc may be allocated i.e., [startPrbc, startPrbc+numPrbc-1]. In case numPrbc is zero the PRB range is all PRBs as defined by the eAxC and startPrbc shall be set to zero. The O-DU shall ensure that rbgMask does not have non-zero bits outside the valid range [0, lastRbgid]; if receiver detects non-zero bits outside the valid range, those shall be ignored.

The O-DU shall not use combinations of startPrbc, numPrbc and rbgSize leading to a value of lastRbgid larger than 27 (number of bits available in the rbgMask field - 1). Such combinations are invalid and if detected by the receiver the section shall be ignored.

MSB of rbgMask indicates the highest frequency resource block group.

When Section Extension 6 is used with Section Extension 13, value of f(0) for each frequency hop shall be same across each frequency hop specified using one instance of Section Extension 13. Stated another way, the number of PRBs in the first resource block group shall be the same in every hop.

Value range: {0000 0000 0000 0000 0000 0000 0000b - 1111 1111 1111 1111 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 28 bits.

Default Value: 1111 1111 1111 1111 1111 1111 1111b (all resource block groups allocated).

The value of rbgSize used in the equations shall be the values in the column "Number of PRBs in resource block group" of Table 7.7.6.2-1.

EXAMPLE: Figure 7.7.6.3-1 shows an example of resource block grouping when startPrbc = 3, rbgSize = 2, and numPrbc = 14.




Figure 7.7.6.3-1 Resource block group definition example

7.7.6.4 symbolMask (symbol bit mask)

Description: This parameter is a bit mask where each bit indicates whether the rbgMask applies to a given symbol in the slot. If bit n is set then the rbgMask shall be applied to symbol n, where n has range [0..13]. If no bits are set then the rbgMask shall not be applied to RBGs in the slot, in other words there are no RE allocations in the slot in this section description. If the symbolMask values indicate allocations beyond a slot boundary, such allocations shall be ignored (e.g., when there are fewer than 14 symbols in a slot).

LSB of symbolMask indicates symbol zero (the first symbol in a slot).

Value range: {00 0000 0000 0000b-11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

Default Value: 11 1111 1111 1111b (section repeated in 14 symbols).

7.7.6.5 priority

Description: This parameter is used in conjunction with C-Plane and U-Plane coupling via frequency and time with priorities (see clause 7.8.1.3). If this coupling method is not used (e.g. O-RU or O-DU do not have corresponding capability) O-DU shall send value zero. This parameter is used by O-RU to find the section description with the highest priority among section descriptions referencing the same RE that are present in one C-Plane message. O-RU shall restrict the scope of the search to one C-Plane message to avoid beamforming configuration errors that would occur if search would span multiple messages and one of them was lost. If Section Extension 6 is absent, then data section description priority is zero (this corresponds to the default priority). Within one C-Plane message a set of priority values shall be restricted to up to two priority values including value zero i.e. {0, +1} or {0, -1}.

Value -2 is reserved and shall not be used by O-DU.

Value range: {-1, 0, +1}.

Type: signed integer, 2's complement.

Field length: 2 bits.

Default Value: 0 (normal priority).

7.7.6.6 repetition (repetition flag)

Description: This parameter is used in conjunction with C-Plane and U-Plane coupling via frequency and time with priorities (optimized) described in clause 7.8.1.4. If this coupling method is used and the section description is duplicated (see requirement 3 in clause 7.8.1.3) then O-DU shall set this flag to 1. Otherwise, O-DU shall set this flag to zero. O-RU may use this flag to identify duplicated section descriptions.

Value range: {0b=no repetition; 1b=repeated highest priority data section in the C-Plane message}.

Type: binary bit.

Field length: 1 bit

7.7.7 SE 7: eAxC mask

7.7.7.1 Overview

Section Extension 7 allows a specific C-Plane message to apply to multiple eAxC_ID values. The intended use case applies to Section Type 0, wherein the intention is to allow transmission blanking (designated by Section Type 0) to be indicated for multiple eAxC_ID values e.g., all component carriers. Table 7.7.7.1-1 shows the format of Section Extension 7.

Table 7.7.7.1-1: Format of Section Extension 7 (eAxC mask)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x07							1	Octet N
extLen = 0x01 (1 word)								1	N+1
eAxCmask[15:8]								1	N+2
eAxCmask[7:0]								1	N+3

When a C-Plane message with Section Extension 7 is received by an O-RU supporting Section Extension 7, the O-RU shall interpret it as if multiple C-Plane messages were received with all of the eAxC_ID values permitted by the mask. O-RU shall ignore invalid eAxC_ID values (that is, eAxC_ID values not configured over M-Plane) permitted by the mask.

7.7.7.2 eAxCmask (eAxC Mask)

Description: This binary mask indicates which eAxC_ID values the C-Plane message applies to. The C-Plane message shall be applied by O-RU to any eAxC_ID configured over M-Plane for the same direction (UL or DL) as the C-Plane message, such that for any non-zero bit of the mask the corresponding bits of the eAxC_ID configured over M-Plane have same values as the corresponding bits of eAxC_ID in the C-Plane message. In other words, the O-RU shall apply

the C-Plane message to configured eAxC_ID values matching the eAxC_ID value in the C-Plane message header with the matching operation checking only bits where the mask bits are non-zero.

NOTE 1: An all-zero eAxCmask means "use this message for all eAxC_ID values" while an all-ones eAxCmask means "use this message for only the designated eAxC_ID value". In the former case ("use this message for all eAxC_ID values"), the actual eAxC_ID value in the C-Plane message header could be any value because the message would apply to all eAxC_ID values regardless of the actual value in the C-Plane message. In the latter case ("use this message for only the designated eAxC_ID value") this Section Extension adds no value because that is the normal operation.

NOTE 2: The intended use of the mask is to indicate which sub-field of the eAxC_ID the message applies to. For example, if a single sector single band of the O-RU will be blanked (designated by Section Type 0) for all component carriers, the RU-Port-ID and CC-ID sub-field portions of the eAxCmask would be all zeros ("use this message for all layers/spatial streams and all component carriers") while the other subfield portions of the eAxCmask would be all ones ("use this message for the specific band, sector, and DU_Port_ID values provided in the eAxC_ID"). In this way a Section Type 0 message could indicate the blanking of all component-carriers in a band-sector in a single C-Plane message.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 16 bits.

Default Value: 1111 1111 1111 1111b (all bits of eAxC_ID are considered).

7.7.8 SE 8: Regularization factor

7.7.8.1 Overview

This Section Extension is used for conveying value of regularization factor used for generating multi-user beamforming weight for the scheduled UEs in a slot. Regularization factor is the noise variance used for MMSE (Minimum Mean Square Error) computation of multi-user beamforming weights. The intended use case applies to Section Type 5 for sending regularization factor instead of sending regularization factor over Section Type 6.

Support of this Section Extension is mandatory for O-RUs supporting Section Type 5 in the present document. If an O-RU reports a capability excluding support of this Section Extension (O-RUs supporting previous version of the present document or O-RUs not supporting Section Type 5), then the O-DU shall not employ it, and if the O-RU reports a capability including support of this Section Extension then O-DU supporting the present document shall employ this Section Extension instead Section Type 6 to convey regularization factor.

See Annex J for further description of the regularization factor's use in uplink and downlink.

The format of this Section Extension shall follow Table 7.7.8.1-1.

Table 7.7.8.1-1: Format of Section Extension 8 (Regularization factor)

0(msb)	1	2	3	4	5	6	7(lsb)	# of bytes	
ef	extType = 0x08							1	Octet N
	extLen = 0x01 (1 word)							1	N+1
	regularizationFactor[15:8]							1	N+2
	regularizationFactor[7:0]							1	N+3

7.7.8.2 regularizationFactor (regularization factor)

Description: This parameter provides a regularization factor to support MMSE operation when UE is scheduled, and applied to Section Type 5.

Value range: {0000 0000 0000 0000b – 1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.7.9 SE 9: Dynamic Spectrum Sharing parameters

7.7.9.1 Introduction

This Section Extension applies to all Section Types.

Table 7.7.9.1-1 shows the format of this Section Extension.

Table 7.7.9.1-1: Format for Section Extension 9 (DSS parameter)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x09							1	Octet N
extLen = 0x01 (1 word)								1	N+1
technology[7:0]								1	N+2
reserved								1	N+3

This Section Extension is used for Dynamic Spectrum Sharing (DSS) operation and to inform the O-RU of the information (LTE or NR) transmitted to the eCPRI interface. O-RU shall use technology field in SE 9 to interpret whether the section description in the C-Plane message received applies for LTE or NR. O-RU shall consider that the C-Plane message received is for NR, if Section Extension 9 is not present in the C-Plane message for a carrier configured as DSS via the M-Plane. If Section Extension 9 is received in a C-Plane message for a carrier configured as LTE or NR (not DSS), the O-RU shall ignore the Section Extension 9 and may log the error. More details for DSS support can be found in clause 7.2.6.

7.7.9.2 technology (interface name)

Description: This parameter indicates if the data transmitted through eCPRI interface is LTE or NR in real time. In this way, O-RU knows the data format transmitted from O-DU and can decode the associated user data.

Value range: {0000 0000b-1111 1111b}:

- 0000 0000b: LTE (i.e. the C-Plane section applies to LTE).
- 0000 0001b: NR (i.e. the C-Plane section applies to NR).
- 0000 0010b – 1111 1111b: reserved.

Type: unsigned integer.

Field length: 8 bits.

Default Value: 0000 0000b (LTE).

7.7.10 SE 10: Group configuration of multiple ports

7.7.10.1 Overview

This Section Extension applies to Section Types 1, 3 and 5. C-Plane section information for the multiple ports (i.e., layers or Tx/Rx paths) can be the same except beam IDs or UE IDs in most cases. When multiple ports share common section information within an O-RU, the C-Plane sections to be sent via corresponding ports are merged into one C-Plane section via a representative port using this Section Extension. M-Plane pre-configures representative port by grouping the ports to be merged to represent them. The beamId in the section of the representative port can be considered as an index which indicates a set of multiple beam vectors (beamIds) each of which applies to the corresponding port (layer or Tx/Rx path). The 2nd port (for example, 2nd port beamId or ueId as given in Table 7.7.10.1-2) for beamGroupType = 10b maps to the first member in the list of **member-tx-eaxc-id** and **member-rx-eaxc-id** specified in M-Plane. The **representative-tx-eaxc-id** and **representative-rx-eaxc-id** is the eAxC_ID in the transport header of the message carrying this Section Extension, and the beamId or ueId associated with the **representative-tx-eaxc-id** and **representative-rx-eaxc-id** is the beamId or ueId contained in the section header of the data section carrying this Section Extension. Section Extension 10 shall follow either Table 7.7.10.1-1 or Table 7.7.10.1-2 format depending on the value of the parameter beamGroupType.

If the O-RU reports the capability of 'non-scheduled-ueId-supported' = TRUE for an endpoint, and the O-DU configures 'non-scheduled-ueId-enabled' = TRUE, and SE 10 is applied to ST 5 and beamGroupType = 10b, if the O-DU set ueId of a port in SE 10 to 0x7FFF, the O-RU shall interpret that the PRBs in the section description are not scheduled for this port (refer to clause 7.9.13 for more details).

Table 7.7.10.1-1: Section Extension 10 for beamGroupType = 00b or 01b

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
ef	extType = 0x0A							1	Octet N	
extLen = 0x01 (1 word)								1	N+1	
beamGroupType		numPortc							1	N+2
reserved								1	N+3	

Table 7.7.10.1-2: Section Extension 10 for beamGroupType = 10b

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
ef	extType = 0x0A							1	Octet N	
extLen								1	N+1	
beamGroupType		numPortc							1	N+2
reserved	2 nd port beamId[14:8] (or ueId[14:8])							1	N+3	
2 nd port beamId[7:0] (or ueId[7:0])								1	N+4	
reserved	3 rd port beamId[14:8] (or ueId[14:8])							1	N+5	
3 rd port beamId[7:0] (or ueId[7:0])								1	N+6	
reserved	...							1	var	
...								1	var	
reserved	(numPortc+1)th port beamId[14:8] (or ueId[14:8])							1	var	
(numPortc+1)th port beamId[7:0] (or ueId[7:0])								1	var	
filler to ensure 4-byte boundary								var	var	

7.7.10.2 beamGroupType

Description: This parameter indicates the type of beam grouping.

Value range: {00b – 11b}.

- **00b** (common beam): the beamId in the section header is used as a common beamId for all the 'numPortc' ports which are grouped by M-Plane. This type is not used for Section Type 5, and extLen = 0x01.
- **01b** (beam matrix indication): the consecutive 'numPortc' beamIds subsequent to the beamId in the section header apply to the 'numPortc' ports. This type is not used for Section Type 5, and extLen = 0x01.
- **10b** (beam vector listing): the beamIds listed in the Section Extension apply to the 'numPortc' ports. The Section Extension shall include 'numPortc' beamIds or ueIds.
- **11b**: reserved.

Type: unsigned integer.

Field length: 2 bits.

Default Value: 00b (common beam).

7.7.10.3 numPortc

Description: This parameter indicates the number of eAxC ports indicated by the Section Extension. It can cover up to 64 ports.

Value range: {000000b – 111111b}. 00 0000b expresses 64 ports.

Type: unsigned integer.

Field length: 6 bits.

Default Value: 000000b.

7.7.10.4 Interaction with other Section Extensions

Section Extension 10 can be used in all situations when same information is present in section headers and in Section Extension headers across multiple eAxC_IDs. Interaction of Section Extension 10 with other Section Extensions is defined in Table.7.7.10.4-1.

Table 7.7.10.4-1: Section Extension 10 interactions with other Section Extensions

Section Extension	Title	Interaction with other Section Extensions
1	Beamforming Weights	Section Extension 1 can be used with Section Extension 10 only if the beamforming weights transferred using Section Extension 1 are the same for all streams (or layers). In general, the beamforming weight vector will be different per data layer (eAxC_ID) and Section Extension 10 cannot be used in conjunction. Once the beamforming weights are downloaded, then if the O-DU uses same beamforming vector by addressing beamId for a number of slots after the beamforming weight is downloaded, then in that case, O-DU can use Section Extension 10 to combine multiple C-Plane messages to one single C-Plane message using the representative eAxC_ID.
2	Beamforming Attributes	Section Extension 2 can be used with Section Extension 10 only if the beamforming attribute transferred using Section Extension 2 is same for all streams (or layers), otherwise, Section Extension 2 cannot be used with Section Extension 10. If O-DU uses beam attributes downloaded using Section Extension 2 later using beamId, then Section Extension 10 can be used to combine multiple C-Plane messages to one C-Plane message.
3	DL Precoding	This Section Extension has different format for different layers (first layer has one set of parameters and rest of the layers have different set of parameters). So Section Extension 3 cannot be used with the Section Extension 10.
4	Modulation Compression	Section Extension 4 can be used with Section Extension 10 (see clause 7.9.10).
5	Modulation Compression Additional Parameters	Section Extension 5 can be used with Section Extension 10 (see clause 7.9.10).
6	Non-contiguous PRB	Section Extension 6 can be used with the Section Extension 10, if the section information of non-contiguous PRB allocation has same pattern for each spatial stream or data layer (eAxC).
7	eAxC Mask	Section Extension 7 cannot be used with Section Extension 10 because the use of Section Extension 10 includes the functionality of Section Extension 7.
8	Regularization Factor	Section Extension 8 can be used with Section Extension 10 only if the regularization parameter transferred using Section Extension 8 is the same for all streams (or layers), otherwise Section Extension 8 cannot be used with Section Extension 10
9	DSS Parameters	Section Extension 9 can be used with Section Extension 10 because all the spatial streams or data layers (eAxCs) grouped by Section Extension 10 are used for the same technology.
11	Flexible Beamforming Weights	Section Extension 11 can be used with Section Extension 10 only if the beamforming weights transferred using Section Extension 11 are same for all streams (or layers). In general, the beamforming weight vector will be different per data layer (eAxC_ID) and Section Extension 11 cannot be used in conjunction. Once the beamforming weights are downloaded, then if the O-DU uses same beamforming vector by addressing beamId for a number of slots after the beamforming weights are downloaded, then in that case, the O-DU can use Section Extension 11 to combine multiple C-Plane messages to one single C-Plane message using the representative eAxC_ID.
14	Null-layer Info	Section Extension 14 can be used with Section Extension 10. In this case, extLen in the Section Extension 14 shall be set based on total number of nullLayerInd fields, i.e. equal to the total number of uelds conveyed using section description and Section Extension 10 (see Table 7.7.14.1-2).

7.7.11 SE 11: Flexible beamforming weights

7.7.11.1 Overview

This section applies to the flexible sending of beamforming weights from the O-DU to the O-RU. This enables the O-DU to provide different beamforming weights for different PRBs within one section to facilitate, e.g., zero-forcing precoding. The O-DU provides the numBundPrb parameter, which informs the O-RU how many PRBs are bundled together and share the same beamforming weights. The optional "little endian byte order" is applied to bfwI/bfwQ fields if chosen via M-plane. Refer to Annex D.2 for details of little endian byte order. This Section Extension applies only to Section Types 1 and 3. Section Extension 11 shall follow formatting defined in Table 7.7.11.1-1 when disableBFW = 0 and Table 7.7.11.1-2 when disableBFWs = 1.

Table 7.7.11.1-1: Section Extension 11 Data Format (when disableBFWs = 0)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
ef	extType = 0x0B							1	Octet N	
extLen[15:0]								2	Octet N+1	
disableBFWs	RAD	reserved (6 bits)							1	Octet N+3
numBundPrb[7:0]								1	Octet N+4	
bfwCompHdr[7:0]								1	Octet N+5	
bfwCompParam (for PRB bundle 0)								var		
contInd	beamId[14:8] (for PRB bundle 0)							1		
beamId[7:0] (for PRB bundle 0)								1		
bfwI (for TRX 0 and PRB bundle 0)								var		
bfwQ (for TRX0 and PRB bundle 0)								var		
remaining beamforming weights bfwI and bfwQ up to L TRXs and PRB bundle 0								var		
bfwCompParam (for PRB bundle 1)								var		
contInd	beamId[14:8] (for PRB bundle 1)							1		
beamId[7:0] (for PRB bundle 1)								1		
bfwI (for TRX 0 and PRB bundle 1)								var		
bfwQ (for TRX0 and PRB bundle 1)								var		
remaining beamforming weights bfwI and bfwQ up to L TRXs and PRB bundle 1								var		
...										
bfwCompParam (for last PRB bundle)								var		
contInd	beamId[14:8] (for last PRB bundle)							1		
beamId[7:0] (for last PRB bundle)								1		
bfwI (for TRX 0 and last PRB bundle)								var		
bfwQ (for TRX0 and last PRB bundle)								var		
remaining beamforming weights bfwI and bfwQ up to L TRXs and last PRB bundle								var		
zero pad to 4-byte boundary								var		

Table 7.7.11.1-2: Section Extension 11 Data Format (when disableBFWs = 1)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
ef	extType = 0x0B							1	Octet N	
extLen[15:0]								2	Octet N+1	
disableBFWs	RAD	reserved (6 bits)							1	Octet N+3
numBundPrb[7:0]								1	Octet N+4	
contInd	beamId[14:8] (for PRB bundle 0)							1		
beamId[7:0] (for PRB bundle 0)								1		
contInd	beamId[14:8] (for PRB bundle 1)							1		
beamId[7:0] (for PRB bundle 1)								1		
...										
contInd	beamId[14:8] (for last PRB bundle)							1		
beamId[7:0] (for last PRB bundle)								1		
zero pad to 4-byte boundary								var		

7.7.11.2 bfwCompHdr (beamforming weight compression header)

Description: refer to clause 7.7.1.2

This parameter defines the compression method and IQ bit width for the beamforming weights in the specific section in the C-Plane message. For the block compression methods, the block size is the vector of beamforming weights for a specific PRB bundle.

7.7.11.3 bfwCompParam for PRB bundle x (beamforming weight compression parameter)

Description: refer to clause 7.7.1.3

This parameter applies to the following vector of beamforming weights for a specific PRB bundle (i.e., bundle x).

7.7.11.4 numBundPrb (Number of bundled PRBs per beamforming weights)

Description: This parameter is the number of bundled PRBs per beamforming weight sets. The number of beamforming weight sets per TRX in Section Extension 11 should be equal to the total number of PRBs selected by section description in the C-Plane message (using startPrbc and numPrbc and other parameters present in other Section Extensions present in the section description) divided by numBundPrb. When the division outcome is fractional value, one additional beamforming weight for each respective TRX is included to cover the orphan PRBs (as defined in clause 7.9.4.2).

Value range: {all zeros – all ones}. The value zero is reserved.

Type: unsigned integer.

Field length: 8 bits.

7.7.11.5 bfwl (beamforming weight in-phase value)

Description: This parameter is the in-phase beamforming weight value. When optional "little endian byte order" is chosen via M-plane, refer to Annex D.2 for details on byte order. The total number of weights in the section depends on the number TRX at the O-RU and the number of bundled PRBs per beamforming weight.

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 1-16 bits.

7.7.11.6 bfwQ (beamforming weight quadrature value)

Description: This parameter is the quadrature beamforming weight value. When optional "little endian byte order" is chosen via M-plane, refer to Annex D.2 for detail on byte order. The total number of weights in the section depends on the number TRX at the O-RU and the number of bundled PRBs per beamforming weight

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 1 – 16 bits.

7.7.11.7 disableBFWs (disable beamforming weights)

Description: This parameter is used to indicate whether beamforming weights under Section Extension are disabled or not. The default is that the beamforming weights exist since this is the main objective of the Section Extension. However, these can be disabled by setting this parameter to convey only the beamIds. In this case, all bfwI, bfwQ, and bfwCompParam fields as well as the bfwCompHdr field are not included in the Section Extension and the formatting of Section Extension 11 shall follow Table 7.7.11.1-2.

Value range: {0b=beamforming weights included in Section Extension; 1b= beamforming weights excluded}.

Type: binary bit.

Field length: 1 bit.

Default Value: 0b (beamforming weights included in Section Extension).

7.7.11.8 RAD (Reset after PRB discontinuity)

Description: This parameter is used when Section Extension 11 is used in conjunction with Section Extension allowing non-contiguous frequency allocation (Section Extensions 6, 12 and 13). When the section parameters refer to a continuous set of PRBs, the default value should be used, which is RAD = 1. When Section Extension allowing non-contiguous frequency allocation is used with Section Extension 11, this parameter indicates whether the PRB bundle boundary is reset after discontinuity in the PRB allocation (RAD = 1), or the PRB bundle boundaries are mapped to the PRBs regardless of the discontinuity (see examples in clause 7.9.4.2). When this extension is used with Section Extension 13 the discontinuity shall be interpreted as discontinuity in set of frequency ranges from all hops combined over time (see examples in clause 7.9.4.4).

Value range: {0b = No reset; 1b = reset}.

Type: binary bit.

Field length: 1 bit.

Default Value: 1b (Reset after PRB discontinuity).

7.7.11.9 contInd (PRB region continuity flag)

Description: If the O-RU indicates support for this flag by M-plane parameter se11-continuity-flag-supported and O-DU enables this function by setting M-plane parameter se11-continuity-flag-enabled to TRUE then O-DU shall set the value of this flag to indicate the PRB region continuity between the current and next PRB bundle for the following cases:

- Current and next PRB bundles belong to the same "MU-MIMO user group" and hence belong to a continuous PRB region.
- Current and next PRB bundle belong to the same PRG allocated for a given UE and belong to a continuous PRB region.

The O-RU advertises its capability to support the handling of this flag by setting the M-Plane Boolean flag 'se11-continuity-flag-supported' to 'TRUE'. The O-DU shall indicate it will send continuity information by setting the M-Plane per endpoint flag 'se11-continuity-bit-enabled' to TRUE. If the O-DU does not support the flag 'se11-continuity-bit-enabled' or if the O-DU supports the M-Plane flag and sets 'se11-continuity-flag-enabled' = FALSE, then the O-DU shall set the contInd flag to value 0 (reserved). In this case, no information regarding PRB region continuity is sent to

the O-RU and the O-RU shall process the PRB bundles as it would have in the absence of the contInd flag. This behaviour shall be same as setting the field to reserved value (0b).

When Section Extension 11 is used in conjunction with Section Extension allowing non-contiguous frequency allocation (Section Extensions 6, 12 and 13), the O-DU may set the value of this flag to 1 in the presence of non-contiguous PRB allocation to indicate continuity across a gap described by the non-contiguous allocation.

The O-DU shall set the value of contInd flag to '0b' for the last PRB bundle.

The O-RU may use the continuity information to provide optimizations to MU-MIMO user group processing, as an example, for beamforming weight interpolation across REs of MU-MIMO user group boundaries.

Value range:

{0b = PRB region discontinuity between the current and next PRB bundle is indicated;

1b = PRB region continuity between current and next PRB bundle is indicated}.

Type: binary bit.

Field length: 1 bit.

Default Value: 0b

7.7.12 SE 12: Non-contiguous PRB allocation with frequency ranges

7.7.12.1 Overview

Section Extension 12 applies only to Section Types 1, 3 and 5. It cannot be used with Section Extension 6 in the same section description.

Section Extension 12 enables allocation of non-contiguous sets of PRBs (Resource Block Groups, or RBGs) in time domain and frequency domain. This extension reduces the C-Plane overhead when users or channels are allocated with non-contiguous sets of PRBs in time or frequency. This extension is more space-efficient than Section Extension 6 if the allocation is continuous in frequency or extends over a frequency span wider than the range defined by rbgMask size and rbgMaskSize value. Structure of Section Extension 12 is presented in Table 7.7.12.1-1.

Table 7.7.12.1-1: Format of Section Extension 12 (non-contiguous PRB allocation with frequency ranges)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType[6:0] = 0x0C							1	Octet N
extLen (variable)								1	N+1
priority[1:0]		symbolMask[13:8]						1	N+2
symbolMask[7:0]								1	N+3
offStartPrb(1)								1	variable
numPrb(1)								1	variable
offStartPrb(2)								1	variable
numPrb(2)								1	variable
...									
offStartPrb(R-2)								1	variable
numPrb(R-2)								1	variable
offStartPrb(R-1)								1	variable
numPrb(R-1)								1	variable

This extension consists of a fixed part (fields priority and symbolMask) and a variable size part (pairs of offStartPrb(r) and numPrb(r) for r = 1 ... R-1). The variable part may be empty allowing the use of priority and symbolMask without overhead.

If Section Extension 12 is present in a section description, then the section description shall be applied to symbols identified by symbolMask and one or more frequency ranges calculated from startPrbc and numPrbc from the section description and offStartPrb(r) and numPrb(r) pairs from this Section Extension.

If Section Extension 12 is present in a section description, then following requirements apply:

- 1) The field startSymbolId in the message header and the fields rb, symInc, and numSymbol in the section description shall not be used for identification of symbols and PRBs referred by the section description.
- 2) O-DU may set symInc and numSymbol to any allowed value without restriction to values corresponding to symbols actually referred by the section description; the values of these parameters may be used to affect range of symbols referred by subsequent section descriptions in the message.
- 3) O-DU shall set rb to the value of zero.
- 4) O-DU shall set startSymbolId to the earliest symbol referred by any of section descriptions in the message, including, but not limited to, section descriptions with Section Extension 12. The earliest symbol referenced by a section description with Section Extension 12 is the earliest symbol selected by a set bit in the symbolMask. That means symbolMask's bit n, such that $n < \text{startSymbolId}$, shall be zero.

If Section Extension 12 is present in a section description, then the section description shall be applied to one or more PRB ranges. Specifically, the section description shall be applied to R spectrum fragments identified by pairs startPrb(r) and numPrb(r) for $r = 0, \dots, R-1$ where startPrb(r) values are calculated by recurrence:

startPrb(0) and numPrb(0) are values of section description's parameter startPrbc and numPrbc; values of numPrb(1), numPrb(2), ... numPrb(R-1) are from the variable part of section description.

$\text{startPrb}(r) = \text{startPrb}(r-1) + \text{numPrb}(r-1) + \text{offStartPrb}(r)$ for $r = 1, \dots, R-1$

The section description applies to PRB #p if $\text{startPrb}(r) \leq p < \text{startPrb}(r) + \text{numPrb}(r)$ for any $r = 0, \dots, R-1$.

Number of frequency ranges R shall be derived from extLen: $R = (\text{extLen} - 1) \times 2 + 1$. Empty PRB ranges are allowed e.g., numPrb(r) = 0 may be used to shift the following PRB range beyond limit of 8-bit offset. Similarly, offStartPrb(r) = 0 may be used to concatenate two PRB ranges and effectively extend range width beyond the limit of 8-bit numPrb(r). If the last pair in the extension are present (due to extension size alignment to multiple of 4 bytes) but not used, then they shall have offStartPrb(R-1) = 0 and numPrb(R-1) = 0. The number of frequency ranges R per section description shall be less than or equal to the value of optional M-Plane parameter "max-prb-ranges-per-sec-ext-12" if this parameter exists. When an empty range (numPrb(r) = 0) is used for range extension the empty range shall not be counted against "max-prb-ranges-per-sec-ext-12", and when two PRB ranges are concatenated (offStartPrb(r)=0) the concatenated PRB ranges shall be counted as a single range when counted against "max-prb-ranges-per-sec-ext-12". Additionally, number of frequency ranges R per highest priority section description shall be less than or equal to the value of another optional M-Plane parameter "max-prb-ranges-per-hp-section-sec-ext-12" if this parameter exists and this shall be applicable when C-Plane message limits apply (as specified in clause 7.8.2.2).

The O-DU shall comply to the limits described in clause 7.8.2.1.2 assuming the number of PRB ranges in the section description with the Section Extension 12 is the number of noncontiguous, non-empty (i.e., with numPrb(r) > 0) frequency ranges in the Section Extension 12. Two concatenated frequency ranges (see case with offStartPrb(r) = 0 in the preceding paragraph) shall be considered contiguous.

NOTE 1: Usage of Section Extension 12 does not affect the operation of user plane. A data section in the user plane is not required to contain a contiguous range of PRBs. By invoking the sectionId multiple times in the user plane, either within a single message or in different messages, it is possible to handle sets of non-contiguous PRBs.

NOTE 2: Utilization of Section Extension 12 does not put any restriction on utilization of sections with contiguous PRB utilization except for the general rules of utilization of sections e.g. in context of one eAxC a resource element is referenced by at most one data section description in the control plane.

In case C-Plane and U-Plane coupling via sectionId is used then control plane section descriptions with the same sectionId shall allocate same set of PRBs and symbols.

When utilizing this Section Extension together with Section Type 3 where freqOffset parameter is present, then freqOffset affects the frequency span for the specific range of PRB numbers.

7.7.12.2 priority (priority of section description)

This parameter has same format and semantics as priority parameter in Section Extension 6 (see clause 7.7.6.5).

7.7.12.3 symbolMask (symbol bit mask)

This parameter has same format and semantics as symbolMask parameter in Section Extension 6 (see clause 7.7.6.4).

7.7.12.4 offStartPrb(r) (offset of PRB range start)

Description: This parameter indicates the offset to start of the r -th PRB range for $r=1, \dots, R-1$.

Value range: 0 ... 255.

Type: unsigned integer.

Field length: 8 bits.

7.7.12.5 numPrb(r) (number of PRBs in PRB range)

Description: This parameter indicates the number of PRBs in the r -th PRB range for $r=1, \dots, R-1$.

Value range: 0 ... 255.

Type: unsigned integer.

Field length: 8 bits.

7.7.13 SE 13: PRB allocation with frequency hopping

7.7.13.1 Overview

Section Extension 13 applies only to Section Types 1, 3 and 5.

Section Extension 13 allows to describe two or more PRB allocations starting at different symbols and different PRB. It is intended to be used for allocations with intra-slot frequency hopping. This Section Extension reduces the C-Plane overhead when users or channels are allocated with intra-slot frequency hopping. This extension is more space-efficient than the use of several section descriptions. Structure of Section Extension 13 is presented in Table 7.7.13.1-1.

Table 7.7.13.1-1: Format of Section Extension 13 (PRB allocation with Frequency hopping)

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
ef	extType[6:0] = 0x0D							1	Octet N	
extLen (variable)									1	N+1
reserved[1:0]		nextSymbolId(1)[3:0]			nextStartPrbc(1)[9:8]			1	N+2	
nextStartPrbc(1)[7:0]									1	N+3
...										
reserved[1:0]		nextSymbolId(R-2)[3:0]			nextStartPrbc(R-2)[9:8]			1	variable	
nextStartPrbc(R-2)[7:0]									1	variable
reserved[1:0]		nextSymbolId(R-1)[3:0]			nextStartPrbc(R-1)[9:8]			1	variable	
nextStartPrbc(R-1)[7:0]									1	variable

This extension is of variable size and conveys list of pairs $\text{nextSymbolId}(n)$ and $\text{nextStartPrbc}(n)$ for $n = 1, \dots, R-1$.

If this Section Extension is present in a section description, then the section description and all other Section Extensions included in it, shall be interpreted by O-RU as if startPrbc value was changed at specified symbols given in $\text{nextSymbolId}(n)$ to values of $\text{nextStartPrbc}(n)$ provided in the extension. The value of numPrbc from the section description applies to all frequency hops.

Specifically, the description and extensions shall be interpreted as if startPrbc value was as in $\text{nextStartPrbc}(n)$ for symbols s , where $\text{nextSymbolId}(n) \leq s < \text{nextSymbolId}(n+1)$ for $n=1$ to $R-1$ and value of $\text{nextSymbolId}(R)$ is a value greater by 1 than the number of the last symbol addressed by the description.

NOTE 1: $\text{nextSymbolId}(R)$ is a guard value and does not have to identify actual symbol. If the section description includes a Section Extension that provides symbolMask parameters (Section Extension 6, 12 and 19) then $\text{nextSymbolId}(R)$ is derived from symbolMask . Otherwise, $\text{nextSymbolId}(R)$ is derived from startSymbolId from the message header and symInc and numSymbol parameters in the section description and in section descriptions preceding it in the message.

startPrbc from the section description shall be applied to symbols $s_0 \leq s < \text{nextSymbolId}(1)$ where s_0 is the first symbol addressed by the section description.

NOTE 2: s_0 is determined by symbolMask if Section Extension 6, 12 or 19 is present; otherwise s_0 is determined by startSymbolId from the message header and symInc and numSymbol parameters in the section description and in section descriptions preceding it in the message.

The set of symbols where startPrbc or $\text{nextStartPrbc}(n)$ applies may be further restricted by symbolMask if Section Extension 6, 12 or 19 is also present.

Number of frequency hops R shall be derived from extLen : $R = \text{extLen} \times 2$. If the last pair in the extension are present (due to extension size alignment to multiple of 4 bytes) but not used, then they shall have values as in the preceding pair (e.g., if pair $R-1$ is not used then $\text{nextSymbolId}(R-1)$ and $\text{nextStartPrbc}(R-1)$ shall be set to $\text{nextSymbolId}(R-2)$ and $\text{nextStartPrbc}(R-2)$ respectively). The number of frequency hops R per section extension 13 shall be less than or equal to the value of optional M-Plane parameter "max-freq-hops-per-sec-ext-13" if this parameter exists.

7.7.13.2 $\text{nextSymbolId}(n)$ (next hop start symbol)

Description: This parameter indicates the symbol at which n -th frequency hop occurs for $n=1, \dots, R-1$. The value shall correspond to one of symbols addressed by the section description.

NOTE 1: If Section Extension 6, 12 and 19 are not present in the section description then set of symbols addressed by the description is determined by startSymbolId from the message header and symInc and numSymbol parameters in the section description and in section descriptions preceding it in the message

NOTE 2: If Section Extension 6, 12 or 19 is present in the section description then set of symbols addressed by the description is determined by symbolMask from Section Extension 6, 12 or 19.

O-DU shall ensure values are ordered in increasing order i.e. $\text{nextSymbolId}(n+1) \geq \text{nextSymbolId}(n)$.

O-DU shall set $\text{nextSymbolId}(n)$ to value of $\text{nextSymbolId}(n-1)$ and set $\text{nextStartPrbc}(n)$ to value of $\text{nextStartPrbc}(n-1)$ if pair ($\text{nextSymbolId}(n)$, $\text{nextStartPrbc}(n)$) is used for padding.

Value range: {0000b-1111b}.

Type: unsigned integer.

Field length: 4 bits.

7.7.13.3 $\text{nextStartPrbc}(n)$ (next hop startPrbc)

Description: This parameter indicates the value to be used instead of startPrbc for the n -th frequency hop for $n=1, \dots, R-1$.

O-DU shall set $\text{nextSymbolId}(n)$ to value of $\text{nextSymbolId}(n-1)$ and set $\text{nextStartPrbc}(n)$ to value of $\text{nextStartPrbc}(n-1)$ if pair ($\text{nextSymbolId}(n)$, $\text{nextStartPrbc}(n)$) is used for padding.

Value range: {00 0000 0000b-11 1111 1111b}.

Type: unsigned integer.

Field length: 10 bits.

7.7.14 SE 14: Nulling layer information for ueld based beamforming

7.7.14.1 Overview

This Section Extension applies to Section Type 5. This enables the O-DU to provide layer-by-layer indication, which denotes that the corresponding ueId is for nulling-layer indication.

Section Extension 14 shall follow the formatting defined in Table 7.7.14.1-1.

Section Extension 14 can be used together with Section Extension 10. In this case, SE14 has multiple nullLayerInd fields and shall follow the formatting defined in Table 7.7.14.1-2. SE 14 with more than one nullLayerInd field shall only be used when the O-RU supports both SE10 and SE14 and only when SE10 is used with the same section description.

Table 7.7.14.1-1: Section Extension 14 without Section Extension 10

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x0E							1	Octet N
	extLen = 0x01 (1 word)							1	N+1
	nullLayerInd							1	N+2
	reserved							1	N+3

O-RU generates beamforming weights considering all the ueIds, but once the beamforming weight is generated, O-RU shall set to zero weights for the layer(s) corresponding to those ueIds for which nullLayerInd is set to 0000 0001b.

Table 7.7.14.1-2: Section Extension 14 with Section Extension 10

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x0E							1	Octet N
	extLen = var [7:0]							1	N+1
	1 st nullLayerInd (for 1 st port ueld) [7:0]							1	N+2
	2 nd nullLayerInd (for 2 nd port ueld) [7:0]							1	N+3
	...								
	(numPortc+1) th nullLayerInd (for (numPortc+1) th port ueld) [7:0]							1	
	zero padding to ensure 4-byte boundary							var	

SE10 is used to convey (numPortc + 1) ueIds with one section description and when Section Extension 14 is used together with Section Extension 10, SE 14 shall carry (numPortc+1) nullLayerInd fields, and in this case, extLen of SE14 will be greater than 0x01. When extLen is greater than 1, the total number of nullLayerInd shall be determined based on the value of (numPortc+1) from SE10.

7.7.14.2 nullLayerInd (null layer indication)

Description: This parameter indicates whether the corresponding layer is nulling-layer or not.

Value range: {0000 0000b – 0000 0001b}.

The following mapping shall be used:

0000 0000b indicates the corresponding ueId is given for scheduled layer, no specific operation based on this field is performed by O-RU.

0000 0001b indicates the corresponding ueId is given for making beam-nulling dimension, so this layer shall be nulled. No user data is transmitted for the layer(s) corresponding to the ueId for which nullLayerInd is set to 0000 0001b, i.e., beamforming weights corresponding to the layer(s) shall be set to zero.

Type: unsigned integer.

Field length: 8 bits.

7.7.15 SE 15: Mixed numerology information for ueld based beamforming

7.7.15.1 Overview

This Section Extension applies to Section Types 5 and 6. When this Section Extension is applied to Section Type 6, the value of FFT type in frameStructure and cpLength can be set as 0.

Section Extension 15 shall follow the formatting defined in Table 7.7.15.1-1.

Table 7.7.15.1-1: Section Extension 15

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x0F							1	Octet N
extLen = 0x02 (2 words)								1	N+1
frameStructure								1	N+2
freqOffset								3	N+3
cpLength								2	N+6

7.7.15.2 frameStructure (frame structure)

See clause 7.5.2.13 for the description of this parameter.

7.7.15.3 freqOffset (frequency offset)

See clause 7.5.3.11 for the description of this parameter.

7.7.15.4 cpLength (cyclic prefix length)

See clause 7.5.2.14 for the description of this parameter.

7.7.16 SE 16: Antenna mapping in UE channel information based UL beamforming

7.7.16.1 Overview

This Section Extension applies to Section Type 5. The Section Extension includes bitmask per RX endpoint to indicate the antennas to be pre-combined into the RX endpoint (i.e., eAxC_ID) and shall follow formatting defined in Table 7.7.16.1-1. This Section Extension can be used together with Section Extension 10. In this case, it has a list of the bitmasks as many the RX endpoints as used for the Section Extension 10 and shall follow formatting defined in Table 7.7.16.1-2.

Table 7.7.16.1-1: Section Extension 16 without Section Extension 10

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x10							1	Octet N
extLen = 0x03 (3 words)								1	N+1
antMask [63:0]								8	N+2
zero padding to ensure 4-byte boundary								var	N+10

Table 7.7.16.1-2: Section Extension 16 with Section Extension 10

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x10							1	Octet N
extLen [7:0]								1	N+1
antMask (1st RX eAxC) [63:0]								8	N+2
antMask (2nd RX eAxC) [63:0]								8	N+10
...							
antMask (16th RX eAxC) [63:0]								8	N+122
zero padding to ensure 4-byte boundary								var	N+130

7.7.16.2 antMask

Description: This parameter indicates the indices of antennas to be pre-combined per RX endpoint.

Value range: {0x0000 0000 0000 0000 – 0xFFFF FFFF FFFF FFFF}.

Each bit indicates whether the antenna corresponding to the bit is pre-combined or not. The maximum number of antennas is 64.

Type: unsigned integer.

Field length: 64 bits.

Default Value: 0x0000000000000000.

7.7.17 SE 17: Indication of user port group

7.7.17.1 Overview

This Section Extension applies to Section Extension 10 for beamGroupType=10b within Section Type 5 and cannot be used in standalone manner. This Section Extension provides the number ueIDs of the users scheduled in preceding Section Type and Section Extension messages. A user may have more than one ueIDs (i.e., more than one channel information, for example if the UE supports Transmit Antenna Switching feature for SRS transmission; by which O-DU can obtain different channel information corresponding to each of the transmit antennas). When this Section Extension is used, two restrictions shall apply: first, ueIDs of each user are consecutive by exploiting the three reserved bits of ueId[2:0]. This means that the maximum number of ueIDs per user that can be supported is 8. Second, in Section Extension 10, ueIDs whose three reserved bits are all zero are configured repeatedly as many times as the number of layers scheduled for the corresponding user. Therefore, the preceding Section Type and extension messages implicitly provide the number of scheduled users (i.e., number of different ueIDs) and number of layers of each user (i.e., number of same ueIDs). Finally, the number of ueIDs associated with each user is provided in this Section Extension.

Section Extension 17 shall follow the formatting defined in Table.7.7.17.1-1.

Table 7.7.17.1-1: Section Extension 17

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x11							1	Octet N
extLen								1	N+1
numUeID of 1 st user				numUeID of 2 nd user				var	N+2
...									
..				numUeID of last user					
zero padding to ensure 4-byte boundary									

7.7.17.2 numUeID

Description: This parameter indicates the number of ueIDs per user.

Value range: {0001b – 1000b}. {1001b – 1111b} are reserved.

Type: unsigned integer.

Field length: 4 bits.

Default Value: 100b (4 uelds per user).

7.7.18 SE 18: Uplink transmission management

7.7.18.1 Introduction

This Section Extension applies to sending of transmission windows information for management of the uplink user data transmission from the O-RU. With the two parameters included in this Section Extension the O-RU shall transmit the corresponding user data within that time window. The data shall either be sent in normal transmission mode or be uniformly distributed in time depending on value of toT parameter. The format of Section Extension 18 is specified in Table 7.7.18.1-1

Table 7.7.18.1-1: Section Extension 18

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x12							1	Octet N
extLen								1	N+1
transmissionWindowOffset[15:8]								1	N+2
transmissionWindowOffset[7:0]								1	N+3
Reserved	transmissionWindowSize[13:8]							1	N+4
transmissionWindowSize[7:0]								1	N+5
Reserved						toT		1	N+6
zero pad to 4-byte boundary								1	N+7

7.7.18.2 transmissionWindowOffset

Description: This parameter indicates the start of the transmission window as an offset to when the transmission window would have been without this parameter, i.e. $(Ta3_max - Ta3_min)$. The resolution of the parameter is symbols, where the duration of a symbol is either as determined via M-Plane (when used with message Type 1 and 5) or according to information in parameter frameStructure (when used with message Type 3) and assuming normal cyclic prefix.

If numSymbol in the section header is greater than 1, then the number of different transmission windows is equal to numSymbol. The start of each transmission window follows the same procedure as described above for every symbol.

See clause 4.6.2 for more detailed information about overlapping transmission windows.

If message is used for an eAxC configured for delay-managed traffic, then transmissionWindowOffset value in combination with the value of transmissionWindowSize shall not result in calculated buffering load that would exceed the O-RU buffering capacity. See clause 4.6.4 for more details.

If message is used for an eAxC configured for non delay-managed traffic, then transmissionWindowOffset value can exceed the O-RU buffering capacity. See clause 4.6.4 for more details.

Value range: {0 – 65535}.

Type: unsigned integer.

Field length: 16 bits.

Default Value: 0

7.7.18.3 transmissionWindowSize

Description: This parameter indicates the size of the transmission window in resolution μ s. If window size is set to a value smaller than $(Ta3_max - Ta3_min)$ the O-RU will if possible, transmit the requested data reliably during that window size. If not possible due to for instance lack of transmission capability on link caused by other higher prioritized traffic, then the O-RU will use a default transmission window size of $(Ta3_max - Ta3_min)$.

Maximum window size is 10000 μ s = 10 ms.

See clause 4.6.2 for more detailed information about overlapping transmission windows.

If message is used for an eAxC configured for delay-managed traffic, then transmissionWindowSize value in combination with the value of transmissionWindowOffset shall not result in calculated buffering load that would exceed the O-RU buffering capacity.

If message is used for an eAxC configured for non delay-managed traffic, then transmissionWindowSize which is not affecting O-RU operation shall be set to 0 (zero). See clause 4.6.4 for more details.

Value range: {0x0 – 0x2710}

Type: unsigned integer.

Field length: 14 bits.

Default Value: 0x00.

7.7.18.4 toT (type of transmission)

Description: This parameter indicates to the O-RU if the associated user plane data is sent in normal transmission mode or transmitted uniformly in time over the transmission window. The O-RU shall indicate via the M-Plane if it supports the possibility to change the type of transmission. If the O-RU does not support this feature, the O-RU will ignore this parameter.

The value of Type of Transmission shall be set to the same value for all sections that have Section Extension 18 within one C-Plane message.

See clause 4.6.2 for more detailed information about uniformly distribution of data.

Value range: {00b – 11b}.

00b = normal transmission mode, data can be distributed in any way the O-RU is implemented to transmit data

01b = uniformly distributed over the transmission window.

10b = Reserved

11b = Reserved

Type: bits

Field length: 2 bits

Default Value: 00b

7.7.18.5 Interaction with other Section Extensions

The Section Extension 18 is used when the UL traffic needs to be managed e.g., avoid peaks in the traffic on shared links between an O-DU and multiple O-RUs.

Table 7.7.18.5-1 lists how Section Extension 18 is used in combination with other Section Extensions.

Table 7.7.18.5-1 Interaction table

Section Extension	Title	Interaction with section extension 18
1	Beamforming weights	No special handling needed
2	Beamforming attributes	No special handling needed
3	DL Precoding configuration parameters and indications	No special handling needed
4	Modulation compression parameters	No special handling needed
5	Modulation compression additional scaling parameters	No special handling needed
6	Non-contiguous PRB allocation	No special handling needed
7	Multiple-eAxC designation	The usage of this Section Extension has the effect that more than one eAxC is targeted with one C-Plane message. For UL this means that the same transmission time and window size will be used for all targeted eAxCs.
8	Regularization factor	No special handling needed
9	Dynamic spectrum sharing parameters	No special handling needed
10	Multiple ports grouping	The usage of this Section Extension has the effect that more than one eAxC is targeted with one C-Plane message. For UL this means that the same transmission time and window size will be used for all eAxCs.
11	Flexible BF weights	No special handling needed
12	Non-contiguous PRB allocation with frequency ranges	No special handling needed
13	Frequency hopping	No special handling needed
14	Null-layer information for UE-ID-based beamforming	No special handling needed
15	Mixed-numerology information for UE-ID-based beamforming	No special handling needed
16	Antenna mapping for UW channel information-based UL beamforming	No special handling needed
17	User port group indication	The usage of Section Extension 17 is always combined with Section Extension 10, same attention as for number 10 is thus applicable.

7.7.19 SE 19: Compact beamforming information for multiple port

7.7.19.1 Overview

This Section Extension applies to Section Types 1 and 3. This Section Extension shall be used for sending compact beamforming information for multiple antenna ports (the term 'port' used henceforth in context of this Section Extension refers to logical antenna port). The optional "little endian byte order" is applied to *bfwI/bfwQ* fields if chosen via M-plane. Refer to Annex D.2 for details of little endian byte order. CSI-RS channel will benefit the most from using this extension, considering large number of CSI-RS ports and multiple CSI resource sets.

This Section Extension is structured into a common Section Extension header specifying total number of ports, consolidated symbol bitmask and beamforming weight compression header applicable for all the ports. This shall be followed by fields describing per port information. Per port fields in the extension specify the location of port using a separate *portReMask* and *portSymbolMask*. The per port section further contains fields to specify per port beamforming information. When using this for CSI-RS one instance of this extension can be used to specify one CSI-RS resource set. For section description with this extension, *reMask* in section header is an aggregate of *portReMask* for all ports specified in SE 19. Also, if section description has SE 19, the *beamId* in section header shall be ignored.

The first instance of this extension shall be used with 'repetition=0' and shall contain beamweights/beamIds for all ports in the range of startPrb and numPrb. When used with highest priority sections, and the section is repeated, this extension can be sent with 'repetition=1', in which case, per port fields are excluded from the extension indicating the associated beamforming information has already been sent to the O-RU. Section Extension 19 invocations shall follow the format shown in Table 7.7.19.1-1. This Section extension can be used with any coupling method. Interaction with other extensions is presented in Table 7.7.19.14-1.

Table 7.7.19.1-1: Section Extension 19

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x13							1	Octet N
extLen[15:0]								2	N+1
disableBFWs	repetition	numPortc[5:0]						1	N+3
priority	symbolMask[13:8]							1	N+4
symbolMask[7:0]								1	N+5
bfwCompHdr								var	
reserved				portReMask[11:8] (for port 1)				1	
portReMask[7:0] (for port 1)								1	
reserved		portSymbolMask[13:8] (for port 1)						1	
portSymbolMask[7:0] (for port 1)								1	
reserved	beamId[14:8] (for port 1)							1	
beamId[7:0] (for port 1)								1	
bfwCompParam (for port 1)								var	
bfw (for port 1 and TRX 0)								var	
bfwQ (for port 1 and TRX 0)								var	
remaining beamforming weights bfw and bfwQ up to L TRXs								var	
...									
reserved				portReMask[11:8] (for last port)				1	
portReMask[7:0] (for last port)								1	
reserved		portSymbolMask[13:8] (for last port)						1	
portSymbolMask[7:0] (for last port)								1	
reserved	beamId[14:8] (for last port)							1	
beamId[7:0] (for last port)								1	
bfwCompParam (for last port)								var	
bfw (for last port and TRX 0)								var	
bfwQ (for last port and TRX 0)								var	
remaining beamforming weights bfw and bfwQ up to L TRXs								var	
Zero pad to 4-byte boundary								var	

7.7.19.2 disableBFWs (disable beamforming weights)

Description: refer to clause 7.7.11.7. This parameter is used to enable/disable sending of beamforming weights as part of this extension, same as in Section Extension 11.

7.7.19.3 repetition (repeat port info flag)

Description: This parameter shall be used to indicate repetition of port beamforming information within a C-Plane message with SE 19.(see clause 7.7.6.6).

Value range: {0b=per port info present in the extension; 1b=per port information not present in the extension}.

Type: binary bit.

Field length: 1 bit.

7.7.19.4 numPortc

Description: refer to clause 7.7.10.3.

For this extension, this parameter indicates the number of logical antenna ports for which associated RE information and beamforming information is contained in this extension. It can cover up to 64 ports.

7.7.19.5 priority (priority of section description)

Description: This parameter has same format and semantics as priority parameter in Section Extension 6 (see clause 7.7.6.5).

7.7.19.6 symbolMask (resource symbol bitmask)

Description: This parameter, similar in functionality to symbolMask in clause 7.7.6.4, is a consolidated symbol bitmask of all the ports described in this extension. This parameter applies to all PRBs specified by startPrbc and numPrbc in the section header containing this extension. If the symbolMask values indicate allocations beyond a slot boundary, such allocations shall be ignored (e.g. when there are fewer than 14 symbols in a slot). LSB of symbolMask indicates symbol zero (the first symbol to arrive in a slot).

This parameter in the common part of this extension along with reMask in section header helps the O-RU in optimized implementation; looping over per-port information in the extension is not required for consolidated port information (which is useful for puncturing low priority sections). Also, for the case where 'repetition = 1' in this extension, and per port information is not present in the extension, 'symbolMask' together with 'reMask (in section header)' is sufficient for carrying consolidated port information.

Value range: {00 0000 0000 0001b-11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

7.7.19.7 bwCompHdr

Description: refer to clause 7.7.1.2

If 'repetition = 1' or 'disableBFWs = 1' beamforming weights are not present in the extension, and hence 'bwCompHdr' parameter is also not present in the extension header.

7.7.19.8 portReMask (RE bitmask per port)

Description: This parameter defines the port Resource Element (RE) mask within a PRB. Each bit set in the portReMask indicates the RE associated with the port. MSB indicates the value for the RE of the lowest frequency in a PRB.

Value range: {0000 0000 0001b-1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 12 bits.

7.7.19.9 portSymbolMask (symbol bitmask per port)

Description: This bitmask specifies the symbols associated with a specific port. It is a subset of symbolMask specified in clause 7.7.19.6.

Value range: {00 0000 0000 0001b-11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

7.7.19.10 bfwCompParam (beamforming weight compression parameter)

Description: refer to clause 7.7.1.3.

This parameter provides beamforming compression parameters associated with a port.

7.7.19.11 beamId (beam identifier for a port)

Description: refer to clause 7.5.3.9.

This parameter provides the beamId associated with a port

7.7.19.12 bfwI (beamforming weight in-phase value)

Description: refer to clause 7.7.1.4.

7.7.19.13 bfwQ (beamforming weight quadrature-phase value)

Description: refer to clause 7.7.1.5.

7.7.19.14 Interaction with Other Section Extensions

Table 7.7.19.14-1: Section Extension 19 Interactions with other Section Extensions

Section Extension	Title	Interaction with existing Section Extensions
1	Beamforming Weights	Section Extension 1 shall not be used in the same section as Section Extension 19, because Section Extension 19 provides an alternate way of specifying beamforming weights.
2	Beamforming Attributes	Section Extension 2 shall not be used in the same section as Section Extension 19, because Section Extension 19 provides an alternate way of specifying beams.
3	DL Precoding	Section Extension 3 is not compatible with Section Extension 19 and shall not be used together in the same section.
4	Modulation Compression	Can be used with Section Extension 19.
5	Modulation Compression (Additional)	Can be used with Section Extension 19.
6	Non-Contiguous PRB	Section Extension 6 shall not be used, because Section Extension 19 provides an alternate way of specifying priorities and symbol Masks.
7	eAxC Mask	Can be used with Section Extension 19.
8	Regularization factor	Not applicable as Section Extension 19 is not intended to be used with Section Type 5.
9	DSS Parameters	Can be used with Section Extension 19.
10	Group Configuration for multiple ports	Section Extension 19 can be used with Section Extension 10 only if the Beamforming Weights transferred using Section Extension 19 are the same for all streams (or layers). In general, the BF weight vector will be different per data layer (eAxC) and Section Extension 19 cannot be used in this case. Once the Beamforming Weights are downloaded, then if the O-DU uses same BF vector by addressing beamId for a number of slots after the BF Weights are downloaded, then in that case, the O-DU can use Section Extension 19 to combine multiple C-Plane messages to one single C-Plane message using the representative eAxC_ID.
11	Flexible Beamforming Weights	Section Extension 11 can be used with Section Extension 19, for providing per port flexible beamforming weights. Section Extension 19 shall appear before Section Extension 11. There would be one instance of Section Extension 11 per port. beamIds specified in Section Extension 19 shall be ignored and disableBFWs in section Extension 19 shall be set to '1' to exclude beamweights in Section Extension 19.
12	Non-Contiguous PRB Allocation with Frequency Ranges	Section Extension 12 shall not be used, because Section Extension 19 provides an alternate way of specifying priorities and symbol Masks.
13	PRB Allocation with Frequency Hopping	Can be used with Section Extension 19.
14	Nulling-Layer Info	Not applicable because Section Extension 19 is not intended to be used with Section Type 5.
15	Mixed Numerology Info for ueld-based beamforming	Not applicable because Section Extension 19 is not intended to be used with Section Type 5.
16	Antenna Information in UE Channel Information based UL beamforming	Not applicable because Section Extension 19 is not intended to be used with Section Type 5.
17	Indication of User Port group	Not applicable because Section Extension 19 is not intended to be used with Section Type 5.
18	Uplink Transmission Management	Can be used with Section Extension 19.
20	Puncturing Extension	Can be used with Section Extension 19.

7.7.20 SE 20: Puncturing extension

7.7.20.1 Overview

This extension specifies a puncturing pattern to be applied to a section. This extension contains a common header with the number of puncturing patterns, followed by symbolMask, PRB ranges, puncReMask and optionally RBG mask fields for each puncturing pattern.

O-RU shall process the fields of the puncturing pattern and remove all the overlapping REs from the current section description. For coupling via time and frequency with priorities (or optimized), Section Extension 20 is restricted to be used with only the lower priority section to avoid any ambiguity in application of puncturing patterns. The number of puncturing patterns in Section Extension 20 and the number of highest priority sections in a C-Plane message together should be less than **max-highest-priority-sec-per-plane-message** limits defined in M-Plane.

The O-DU shall comply to the limits described in clause 7.8.2.1.2 assuming the number of PRB ranges in the section description with the Section Extension 20 is the number of noncontiguous PRB ranges allocated by Section Extension 20 with puncturing removing PRBs as specified in clause 7.7.6.3 if rbgMask is present in the Section Extension (i.e., rbgIncl=1 as described in clause 7.7.20.8).

Section Extension invocations shall follow the format in Table 7.7.20.1-1.

Table 7.7.20.1-1: Section Extension 20

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x14							1	Octet N
extLen[15:0]								2	N+1
numPuncPatterns[7:0]								1	N+3
symbolMask[13:6](1)								1	N+4
symbolMask[5:0](1)					startPuncPrb[9:8](1)			1	N+5
startPuncPrb[7:0](1)								1	N+6
numPuncPrb[7:0](1)								1	N+7
puncReMask[11:4](1)								1	N+8
puncReMask[3:0](1)			rb(1)	reserved(1)		rbgIncl(1)		1	N+9
reserved(1)	rbgSize[2:0](1)		rbgMask[27:24](1)					0-1	var
rbgMask[23:16](1)								0-1	var
rbgMask[15:8](1)								0-1	var
rbgMask[7:0](1)								0-1	var
...									
symbolMask[13:6](last)								0-1	
symbolMask[5:0](last)					startPuncPrb[9:8](last)			0-1	
startPuncPrb[7:0](last)								0-1	
numPuncPrb[7:0](last)								0-1	
puncReMask[11:4](last)								0-1	
puncReMask[3:0](last)			rb(last)	reserved(last)		rbgIncl(last)		0-1	
reserved(last)	rbgSize[2:0](last)		rbgMask[27:24](last)					0-1	
rbgMask[23:16](last)								0-1	
rbgMask[15:8](last)								0-1	
rbgMask[7:0](last)								0-1	
Zero pad to 4-byte boundary								var	

7.7.20.2 numPuncPatterns (number of puncturing pattern)

Description: This parameter is used to indicate total number of puncturing patterns contained within single instance of this extension.

Value range: {000000b-111111b}.

Type: unsigned integer.

Field length: 8 bits.

7.7.20.3 symbolMask (puncturing pattern symbol mask)

Description: This parameter (also see clause 7.7.6.4 and 7.7.19.9) is a bitmask where each bit indicates the symbols associated with the puncturing pattern. A value of '1' indicates that the symbol shall be considered for puncturing. A value of '0' indicates the symbol need not be considered for puncturing.

Value range: {00 0000 0000 0001b-11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

7.7.20.4 startPuncPrb (starting PRB to which one puncturing pattern applies)

Description: This parameter conveys the first PRB of the puncturing pattern.

Value range: {00 0000 000b-11 1111 1111b}.

Type: unsigned integer.

Field length: 10 bits.

7.7.20.5 numPuncPrb (number of contiguous PRBs to which one puncturing pattern applies)

Description: This parameter conveys the number of PRBs of the puncturing pattern

Value range: {0000 0000b-1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 8 bits.

7.7.20.6 puncReMask (puncturing pattern RE mask)

Description: This parameter defines the Resource Element (RE) mask of the puncturing pattern within a PRB. Each bit in the puncReMask indicates the presence/absence of a puncturing RE within a PRB. A value of '1' indicates presence of a puncturing RE, and a value of '0' indicates absence of a puncturing RE. MSB indicates the value for the RE of the lowest frequency in a PRB.

Value range: {0000 0000 0001b-1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 12 bits.

7.7.20.7 rb (resource block indicator)

Description: refer to clause 7.5.3.2

This shall not be applicable (value shall be ignored) if rbgIncl flag is set to '1' in this extension

7.7.20.8 rbgIncl (rbg included flag)

Description: This parameter is used to indicate presence/absence of resource block group for the case of non-contiguous PRB allocation. If this flag is set to 0, two fields following this parameter i.e., rbgSize and rbgMask are absent and vice versa.

Value range: {0b=resource block group not included; 1b=resource block group included}.

Type: binary bit.

Field length: 1 bit.

7.7.20.9 rbgSize (rbg size)

Description: refer to clause 7.7.6.2.

7.7.20.10 rbgMask (rbg bitmask)

Description: refer to clause 7.7.6.3.

7.7.20.11 Interaction with Other Extensions

Section Extension 20 is compatible with all other Section Extensions. Section Extension 20 is used to describe the puncturing information of *other channels and their associated sections* on top of the current section. Other Section Extensions are used to describe the Beamforming and Scheduling information of the current section.

7.7.21 SE 21: Variable PRB group size for channel information

7.7.21.1 Overview

This Section Extension only applies to Section Type 6. It enables the O-DU to send channel information with dynamic PRB group size that can match the channel variation in the frequency domain and reduce the fronthaul traffic.

Section Extension 21 shall follow the formatting defined in Table 7.7.21-1.

Table 7.7.21-1: Section Extension 21 for variable PRB group size for Section Type 6

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x15							1	Octet N
extLen = 1								1	N+1
ciPrbGroupSize								1	N+2
reserved								1	N+3

7.7.21.2 ciPrbGroupSize (channel information PRB group size)

Description: This parameter indicates the PRB group size of channel information. If this Section Extension exists, O-RU shall apply every pair of ciIsample/ciQsample for one PRB group received in Section Type 6 to ciPrbGroupSize number of continuous PRBs starting from index startPrbc + $i \times$ ciPrbGroupSize, where index i ranges from 0 to numPrbc -1. If this Section Extension exists, the 'numPrbc' in Section Type 6 indicates the number of PRB groups per antenna and is used to determine the length of the Section Type 6 message, and the actual number of PRBs to which Section Type 6 is applied is determined by the method described above. Also, numPrbc shall not be set to 0 when

ciPrbGroupSize is configured. If $\text{startPrbc} + i \times \text{ciPrbGroupSize}$ is larger than the maximum available PRBs configured via M-Plane (**number-of-prb**), the channel information shall only be applied to those PRBs within the maximum available PRBs. If both M-Plane configuration **ci-prb-group-size** and Section Extension 21 exist, O-RU shall use the ciPrbGroupSize received in Section Extension 21.

Value range: {2 – 254}.

Value 0, 1 and 255 are reserved for future use.

Type: unsigned integer.

Field length: 8 bits.

7.7.22 SE 22: ACK/NACK request

7.7.22.1 Overview

Section Extension 22 enables the O-DU to request ACK/NACK feedback using Section Type 8 C-Plane message for a section description. Refer to clause [7.2.8](#) for the usage of this Section Extension. Section Extension 22 shall follow the format defined in Table [7.7.22.1-1](#).

Table 7.7.22.1-1: Section Extension 22 for ACK/NACK request

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x16							1	Octet N
extLen = 0x01 (1 word)								1	N+1
ackNackReqId								2	N+3

7.7.22.2 ackNackReqId (ACK/NACK request Id)

Description: This parameter indicates the ACK/NACK request ID of a section description. The value of ackNackReqId shall be unique within a slot (the same value can be reused in different slots). The value of ackNackReqId shall be used for ackId and nackId in Section Type 8.

Value range: {0000 0000 0000 0000b-1111 1111 1111 1111b}.

Type: unsigned integer.

Field length: 16 bits.

7.7.23 SE 23: Arbitrary symbol pattern modulation compression parameters

7.7.23.1 Overview

This Section Extension enables specifying multiple sets of 'mcScaleReMask, csf and mcScaleOffset' values for one or more 'SymPrbPatterns'. The term 'SymPrbPattern' is used to specify set of PRBs that can span an entire PRB range (specified using prbPattern) and multiple symbols (specified using symMask). The proposed extension is motivated by the fact that in 5G NR reference signals like DM-RS, PT-RS, and data channel experience the same channel conditions (same beamId) but may use different MCS and hence different mcScaleOffset. This Section Extension applies to Section Types 1, 3 and 5.

This Section Extension has a nested structure comprising of two loops. The outermost loop which is bounded by the field 'numSymPrbPattern' shall specify multiple SymPrbPatterns. The innermost loop is bounded by the field "numMcScaleOffset" and shall specify multiple sets of 'mcScaleReMask, csf, and mcScaleOffset' per SymPrbPattern. Refer to Table 7.7.23.1 for details of the structure of SE 23.

If Section Extension 23 is present in a section description, then the following requirements shall apply:

- 1) Requirements 1, 2 and 3 as specified in clause 7.7.6.1 for SE 6.
- 2) SE 23 using a combination of symMask, prbPattern and mcScaleReMask shall specify mcScaleOffset values for all the symbols and REs whose scheduling information is specified in the section header (startSymbolId) and section description (numSymbols, reMask) or via the use of SE 6 or SE 12.

- 3) Specifically for SE 6 and SE 12 prbPattern shall apply to all allocated non-contiguous PRBs jumping over the un-allocated RBGs.
- 4) Any PRB on time-freq grid shall be addressed by only one SymPrbPattern in any instance of SE 23
- 5) Each section description shall specify only one instance of SE 23 per eAxC_ID. When SE 23 is used in combination with SE 10 refer to clause 7.9.10.
- 6) For every SymPrbPattern all REs in the PRBs as designated in the reMask in section header shall be assigned "mcScaleReMask, csf and mcScaleOffset" value. No bit in any of the mcScaleReMasks shall be set (=1) in a position where the reMask has a zero, and every reMask bit that is set (=1) shall have exactly one bit =1 in one of the mcScaleReMasks.

e.g. For section header reMask = 1111 1111 1111 b, union of mcScaleReMask-1 = 1010 1010 1010 b and mcScaleReMask-2 = 0101 0101 0101 b shall be equal to the reMask value.
- 7) When SE 23 is used in a section description, the number of sets of {mcScaleReMask, csf, mcScaleOffset} values per symPrbPattern shall be limited by the M-Plane O-RU capability parameter 'max-mcscleremask-per-prb'.

Table 7.7.23.1-1: Section Extension 23 for modulation compression for multiple symbols

0 (msb)	1	2	3	4	5	6	7(lsb)	# of bytes	
ef	extType = 0x16							1	Octet N
extLen[15:0]								1	N+1
numSymPrbPattern[3:0]				reserved				1	N+2
reserved								1	N+3
reserved		symMask[13:8] (1)						1	N+4
symMask[7:0] (1)								1	N+5
numMcScaleOffset[3:0] (1)				prbPattern[3:0] (1)				1	N+6
reserved								1	N+7
reserved				mcScaleReMask[11:8] (1.1)				1	N+8
mcScaleReMask[7:0] (1.1)								1	N+9
csf (1.1)	mcScaleOffset[14:8] (1.1)						1	N+10	
mcScaleOffset[7:0] (1.1)								1	N+11
reserved (1.2)				mcScaleReMask [11:8] (1.2)				1	N+12
mcScaleReMask [7:0] (1.2)								1	N+13
csf (1.2)	mcScaleOffset[14:8] (1.2)						1	N+14	
mcScaleOffset[7:0] (1.2)								1	N+15
...									
reserved		symMask[13:8] (n)							
symMask[7:0] (n)									
numMcScaleOffset[3:0] (n)				prbPattern[3:0] (n)					
reserved									
reserved				mcScaleReMask [11:8] (n.1)					
mcScaleReMask [7:0] (n.1)									
csf (n.1)	mcScaleOffset[14:8] (n.1)								
mcScaleOffset[7:0] (n.1)									

7.7.23.2 numSymPrbPattern (number of symbol and resource block patterns)

Description: This parameter specifies the number of SymPrbPatterns specified by SE 23 instance.

Value range: {0001b - 1111b} or {1 – 15} in decimal

Type: unsigned integer

Field length: 4 bits.

7.7.23.3 symMask (symbol mask part of symPrbPattern)

Description: This parameter is a bitmask for the symbols specified by SymPrbPattern

0: 'SymPrbPattern' does not apply to the associated symbol.

1: 'SymPrbPattern' applies to the associated symbol.

Value range: {00 0000 0000 0001b - 11 1111 1111 1111b}.

Type: unsigned integer (bit mask).

Field length: 14 bits.

7.7.23.4 prbPattern (resource block pattern part of symPrbPattern)

Description: This parameter is a 4-bit pattern mask for the PRBs specified by SymPrbPattern. This pattern repeats over all the allocated PRBs. When there are allocation discontinuities e.g. SE 6, SE 12, the pattern only applies to the allocated PRBs. If the prb range is not a multiple of 4 then the last prbPattern shall be truncated. In the specified mask LSB represents the lowest frequency PRB and MSB represents the highest frequency PRB in the prbPattern.

0: 'SymPrbPattern' does not apply to the associated PRB

1: 'SymPrbPattern' applies to the associated PRB.

Value range: {0000b - 1111b}.

Type: unsigned integer (bit mask).

Field length: 4 bits.

7.7.23.5 numMcScaleOffset (number of modulation compression scaling value per symPrbPattern)

Description: This parameter indicates the number of modulation compression parameter sets i.e., 'mcScaleReMask, csf and mcScaleOffset' values, present for each SymPrbPattern. Refer to requirement#6 in clause 7.7.23.1 for limits that apply to this parameter.

Value range: {0001b-1111b} or {1 – 15} in decimal

1 – 12: Valid range

0, 13, 14, 15: reserved

Type: unsigned integer.

Field length: 4 bits.

7.7.23.6 mcScaleReMask (modulation compression power scale RE mask)

Description: refer to clause 7.7.5.2 for description of this parameter. Usage of this parameter in SE 23 shall apply only to PRBs and symbols specified by SymPrbPattern.

7.7.23.7 csf (constellation shift flag)

Description: refer to clause refer to clause 7.7.4.2

7.7.23.8 Interaction with Other Section Extensions

Interaction of Section Extension 23 with other Section Extensions is defined in [Table 7.7.23.1-2](#).

Table 7.7.23.1-2: Section Extension 23 Interactions with other Section Extensions

Section Extension	Title	Interaction with existing Section Extensions
1	Beamforming Weights	This Section Extension is independent of SE 23
2	Beamforming Attributes	SE 2 can be used with SE 23 only if the Beamforming Attribute transferred using SE 2 is same for DL data and control channel (DM-RS and PT-RS)
3	DL Precoding	This Section Extension is independent of SE 23
4	Modulation Compression	SE 23 cannot co-exist with this Section Extension in the same data section
5	Modulation Compression (Additional)	SE 23 cannot coexist with this Section Extension in the same data section
6	Non-Contiguous PRB	SE 6 can be used with SE 23. SE 23 shall apply to PRB allocations with SE 6.
7	eAxC Mask	This Section Extension is independent of SE 23
8	Regularization factor	This Section Extension is independent of SE 23
9	DSS Parameters	This Section Extension is independent of SE 23
10	Group Configuration for multiple ports	No special handling needed . Refer to clause 7.9.10 for the interaction details.
11	Flexible Beamforming Weights	This Section Extension is independent of SE 23
12	Non-Contiguous PRB Allocation with Frequency Ranges	Interaction same as SE 6
13	PRB Allocation with Frequency Hopping	Interaction same as SE 6
14	Nulling-Layer Info	This Section Extension is independent of SE 23
15	Mixed Numerology Info for ueld-based beamforming	This Section Extension is independent of SE 23
16	Antenna Information in UE Channel Information based UL beamforming	This Section Extension is independent of SE 23
17	Indication of User Port group	This Section Extension is independent of SE 23
18	Uplink Transmission Management	This Section Extension is independent of SE 23
19	Compact multiple port beamforming information	SE=19 is used for specifying separate beamforming weights for data and reference signals (CSI-RS), usage of SE=23 with SE=19 hence is hence restricted
20	Puncturing Extension	This Section Extension is independent of SE 23
21	Variable PRB group size for channel information	This Section Extension is independent of SE 23

7.8 Aspects of C-Plane operation

7.8.1 Coupling of C-Plane and U-Plane

7.8.1.1 Overview

Selected types of C-Plane messages (Section Types 1, 3 and 5) carry data section descriptions that convey information applicable to data sections carried in U-Plane messages. This clause describes a method of coupling data section descriptions from C-Plane and data sections in U-Plane.

7.8.1.2 Coupling via sectionId value

This method of C-Plane and U-Plane coupling shall be used by default. The O-DU may configure the O-RU to use the other methods via the M-Plane. In this method, information from section description D in C-Plane message CM is

applied to resource element R in PRB P (PRB is understood as 12 resource elements of one OFDM symbol consecutive in frequency) in data section S in U-Plane message UM if:

- 1) CM corresponds to UM i.e. both have the same eAxC_ID and CM.**dataDirection** value matches direction of UM.
- 2) CM.**frameId** = UM.**frameId** and CM.**subframeId** = UM.**subframeId** and CM.**slotId** = UM.**slotId**.
- 3) symbol identified by UM.**symbolId** is described by the section description D.
This requires evaluation of CM.**startSymbolId**, **symInc** and **numSymbol** in CM or checking D.**symbolMask** in Section Extension 6, 12 or 19 if any of the extensions is present.
- 4) PRB P is described by the section description D. This requires checking that $D.startPrbc \leq P < (D.startPrbc + D.numPrbc \times (D.rb + 1))$, if $D.rb = 1$ checking that $D.startPrbc \text{ modulo } 2 = P \text{ modulo } 2$; it requires evaluating D.**rbgMask** in Section Extension 6 if this extension is present in D. Checking Section Extension 12, 13, 19 and 20 is required if any of them is present in D.
- 5) D.**reMask** has value 1 in a bit corresponding to RE R. Checking effect of Section Extension 20 is required if this extension is present in D.
- 6) D.**sectionId** = S.**sectionId**.

In addition, the following requirements and restrictions apply (in the context of one eAxC):

- 1) If a PRB is present in a U-Plane message, the PRB shall occur only once across all U-Plane messages related to the same OFDM symbol.
- 2) A resource element shall be "referenced" (see conditions 3, 4 and 5 above) by one section description.
- 3) All section descriptions that are applicable to a PRB (see conditions 3 and 4 above) shall have same value of sectionId.
- 4) Section descriptions that are not applicable to same PRB (see conditions 3 and 4 above) shall have different value of sectionId.
- 5) The O-DU shall send in sectionId a value that is unique within a slot.
- 6) Section descriptions that have the same value of sectionId shall have same value of **rb**, **startPrbc**, **numPrbc** and **numSymbol**.

7.8.1.3 Coupling via frequency and time

This coupling method shall be used if a corresponding capability is reported by the O-RU via the M-Plane and use of this coupling method for an eAxC is configured via the M-Plane. In this method, information from section description D in C-Plane message CM is applied to resource element R in PRB P (PRB is understood as 12 resource elements of one OFDM symbol consecutive in frequency) in data section S in U-Plane message UM if:

- 1) CM corresponds to UM i.e. both have the same eAxC_ID and CM.**dataDirection** value matches direction of UM.
- 2) CM.**frameId** = UM.**frameId** and CM.**subframeId** = UM.**subframeId** and CM.**slotId** = UM.**slotId**.
- 3) symbol identified by UM.**symbolId** is described by the section description D.
This requires evaluation of CM.**startSymbolId**, **symInc** and **numSymbol** in CM or checking D.**symbolMask** in Section Extension 6, 12 or 19 if any of the extensions is present.
- 4) PRB P is described by the section description D. This requires checking that $D.startPrbc \leq P < (D.startPrbc + D.numPrbc \times (D.rb + 1))$, if $D.rb = 1$ checking that $D.startPrbc \text{ modulo } 2 = P \text{ modulo } 2$; it requires evaluating D.**rbgMask** in Section Extension 6 if this extension is present in D. Checking Section Extension 12, 13, 19 and 20 is required if any of them is present in D.
- 5) D.**reMask** has value 1 in a bit corresponding to RE R. Checking effect of Section Extension 20 is required if this extension is present in D.

In addition, the following requirements and restrictions apply (in context of one eAxC):

- 1) If a PRB is present in a U-Plane message, the PRB shall occur only once across all U-Plane messages related to the same OFDM symbol.
- 2) A resource element shall be "referenced" (see conditions 3, 4 and 5 above) by one section description.
- 3) The O-DU shall send in sectionId a value 4095.
- 4) The O-DU shall order section descriptions such that any section descriptions referencing to any RE (non-zero value of reMask) are present before section descriptions that do not reference any RE (i.e. zero value in reMask). This allows O-RU to more quickly identify section descriptions that carry beamforming configuration for any RE and also makes them localized in memory which can improve message processing performance.

7.8.1.4 Coupling via frequency and time with priorities

This coupling method shall be used if a corresponding capability is reported by the O-RU via the M-Plane and use of this coupling method for an eAxC is configured via the M-Plane. In this method information from section description D in C-Plane message CM is applied to resource element R in PRB P (PRB is understood as 12 resource elements of one OFDM symbol consecutive in frequency) in data section S in U-Plane message UM if:

- 1) CM corresponds to UM i.e. both have the same eAxC_ID and CM.dataDirection value matches direction of UM.
- 2) CM.frameId = UM.frameId and CM.subframeId = UM.subframeId and CM.slotId = UM.slotId.
- 3) symbol identified by UM.symbolId is described by the section description D. This requires evaluation of CM.startSymbolId, symInc and numSymbol in CM or checking D.symbolMask in Section Extension 6, 12 or 19 if any of the extensions is present.
- 4) PRB P is described by the section description D. This requires checking that $D.startPrbc \leq P < (D.startPrbc + D.numPrbc \times (D.rb + 1))$, if $D.rb = 1$ checking that $D.startPrbc \text{ modulo } 2 = P \text{ modulo } 2$; it requires evaluating D.rbgMask in Section Extension 6 if this extension is present in D. Checking Section Extension 12, 13, 19 and 20 is required if any of them is present in D.
- 5) D.reMask has value 1 in a bit corresponding to RE R and D has highest priority among data section descriptions referencing the RE R in the message CM. Checking effect of Section Extension #20 is required if this extension is present in D. Priority of data section description is the value of priority field in Section Extension 6 (see clause 7.7.6.4), 12 (see clause 7.7.12.1) or 19 (see clause 7.7.12.1) if any of the extensions is present in D and zero if the extensions are absent. The scope of search for highest priority description is restricted to one C-Plane message to avoid beamforming configuration errors that would be unavoidable if scope would be covering multiple messages and one of them would be lost. See also note in restriction #3 below.

In addition, the following requirements and restrictions apply (in context of one eAxC):

- 1) If a PRB is present in a U-Plane message, the PRB shall occur only once across all U-Plane messages related to the same OFDM symbol.
- 2) Data section descriptions that refer to the same RE and are conveyed in the same C-Plane message shall have different priority in order to avoid ambiguity.
- 3) If a complete beamforming configuration does not fit into one C-Plane message, then sender shall duplicate highest priority data section descriptions, specifically each C-Plane message shall contain the highest priority section description referring to any RE that is referred in a message. This is required to ensure that O-RU will interpret configuration received in each message correctly. If such duplication is needed to avoid misinterpretation it does not violate restriction #4 below. Optionally, each C-Plane message can also contain highest priority data section description(s) referring to REs that are not referred to by non-highest priority data section description(s) in a message. As an alternate to duplicating highest priority section descriptions, Section Extension 20 can be used (refer to clause 7.7.20, and clause 7.9.12 for more details).

- 4) Sender shall avoid sending unnecessary data section descriptions e.g., descriptions that carry the same configuration (but refer to the requirement #3 above).
- 5) The O-DU shall send in sectionId a value 4095.
- 6) The O-DU shall order section descriptions such that any section descriptions referring to any RE (non-zero value of reMask) are present before section descriptions that do not reference any RE (i.e., zero value in reMask); section descriptions referring to any RE shall be ordered by effective priority (highest priority first). Ordering section descriptions with non-zero reMask first allows the O-RU to more quickly identify section descriptions that carry beamforming configuration for any RE and also makes them localized in memory which can improve message processing performance. Ordering section descriptions by highest priority first is intended to optimize O-RU C-Plane message processing i.e., RE beamforming configuration can be determined from the earliest section description occurring in the message.

7.8.1.5 Coupling via frequency and time with priorities (optimized)

All the requirements and restrictions remain the same as "coupling via frequency and time with priorities" as described in clause 7.8.1.4 except requirement #5, which is replaced with:

- 5) The O-DU shall send in sectionId value 4095 for section description with lower priority. For highest priority sections unique sectionIds shall be set in a defined range starting from "0" to "max-highest-priority-sections-per-slot" specified through M-Plane, across eAxC_ID values (limiting this range helps the O-RU reduce the search space for repeated highest priority section descriptions detection). While duplicating highest priority sections, all fields including sectionId shall be duplicated. In addition to this, for C-Plane message processing O-RUs, additional limits to restrict the number of highest priority sections per C-Plane message on top of eAxC limits shall be applied (clause 7.8.2.2).

In addition to associating unique sectionId to highest priority section as described above, the O-DU shall set repetition=1 for every repeated C-Plane message (see clause 7.7.6.6). The O-RU may ignore the flag based on its implementation.

7.8.2 O-RUs per endpoint and per C-Plane message limits

7.8.2.1 O-RU per endpoint processing limits

7.8.2.1.1 Overview

When O-RUs processing granularity is endpoint based i.e., processing resources in O-RU to handle C/U-Plane messages are allocated per endpoint, certain limits may be imposed by the O-RU per endpoint e.g., endpoint-section-capacity, endpoint-beam-capacity, endpoint-prb-capacity. When such limits are imposed by the O-RU, the O-DU shall comply with the limits, otherwise the stated O-RU capacity may be compromised.

7.8.2.1.2 Limits on Number of PRB ranges

One type of per endpoint limits is related to overall number of PRB ranges addressed in C-plane messages. If O-RU reports the capability in M-Plane parameter "max-prb-ranges-per-symbol" then O-DU shall ensure the number of PRB ranges in C-plane messages meets the following condition:

The total number of PRB ranges in all section descriptions in a symbol shall not exceed the value reported by the O-RU in the M-plane parameter "max-prb-ranges-per-symbol". The restriction shall be applied per an endpoint and per groups of endpoints sharing the capacity.

For evaluation of the above condition the following PRB range counting rules shall be applied:

- 1) Evaluation of PRB ranges per symbol shall be done for each symbol separately.
- 2) Number of PRB ranges in a symbol is the smallest number of sections without section extensions that could represent the beamforming configuration requested for the symbol.

7.8.2.2 O-RU C-Plane message limits

For O-RUs with per C-Plane message processing limits in addition to per endpoint processing limits, the O-RU may choose to advertise its limitations on a per C-Plane message basis. The O-DU may indicate that it adheres to the associated limitations. The defined limits are the maximum number of beams, maximum number of highest priority sections and maximum number of PRB ranges per highest priority section with Section Extension 12 within a C-Plane message. Refer to clause 15.8 of the M-Plane Specification [7] for details on usage of this feature for various scenarios.

7.9 C-Plane optimizations

7.9.1 C-Plane optimization using Section Extension 6

Section Extension 6 can be used for non-contiguous PRB allocation in both time and frequency domains. Section Extension 6 conveys two bitmasks: symbolMask and rbgMask. The first allows to select an arbitrary subset of symbols within a slot. The second allows to select arbitrary subset of blocks of subcarriers (each block has $12 \times \text{rbgSize}$ subcarriers) between startPrbc and startPrbc+numPrbc. The selected set of RE is a cross-section of symbols and subcarriers selected by both masks. This allows to describe a wide range of non-contiguous resource allocations with one section description.

In addition, O-RU may support optional feature allowing it to interpret a non-zero value priority carried in this Section Extension. This optimization is described in clause 7.9.5.

7.9.2 C-Plane Optimization using Section Extension 7

Void.

7.9.3 C-Plane optimization using Section Extension 10

7.9.3.1 Overview

In general, O-DU uses unique eAxC_IDs to address each layer or spatial stream when sending C-Plane and U-Plane messages to the O-RU. In many situations, information contained in C-Plane messages for the different spatial streams is the same or similar. For example, a single user MIMO allocation with 8 layers may have same values for startPrbc, NumPrbc, reMask, and numSymbol in the section header for all 8 C-Plane messages, one message per layer. In this case, Section Extension 10 can be used along with a 'representative eAxC_ID' (configured via M-Plane) to reduce C-Plane overhead of sending multiple messages by sending one single C-Plane message. This is an optional feature which can be taken advantage of by the O-DU if the O-RU capability indicates that the O-RU supports Section Extension 10. Based on this capability, the O-RU can be configured with a 'representative eAxC_ID' representing all 8 eAxC_IDs in above example. Once a 'representative eAxC_ID' is configured, the O-DU can send one single C-Plane message addressed to the 'representative eAxC_ID' along with the Section Extension 10 instead of 8 C-Plane messages in above example. Once the O-RU receives a C-Plane message addressed to the 'representative eAxC_ID' and with Section Extension 10, then the O-RU applies this message to all endpoints pointed to by the 'representative eAxC_ID' whereby performing same operation just like O-RU had received 8 different C-Plane messages. Figure 7.9.3.1-1 below illustrates the example.

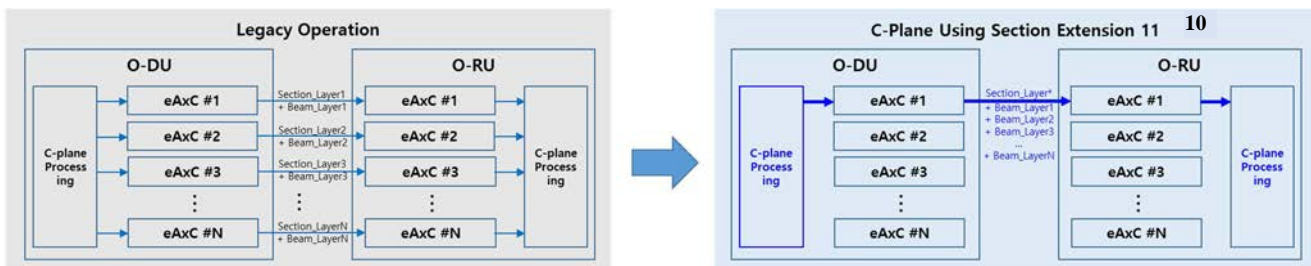


Figure 7.9.3.1-1: Operation of Section Extension 10

7.9.4 C-Plane optimization using Section Extension 11

7.9.4.1 Overview

Section Extension 11 refers to flexible beamforming weight transmission. One way to transfer per-PRB beamforming weights is using Section Extension 1 and creating one section for each PRB. This introduces overhead in the form of section headers and Section Extensions. The objective of Section Extension 11 is to eliminate this problem by allowing the inclusion of the beamforming weights for multiple PRBs in a single section.

Example use case include zero-forcing beamforming (or similar methods) by using per-PRB channel information to calculate per-PRB beamforming weights. Specifically, the O-RU sends uplink pilots (SRS) to O-DU, which in return calculates per-PRB beamforming weights in the O-DU, then transfer them to the O-RU. These weights can be updated every slot, as an example.

Section Extension 11 gives the flexibility of either sending the sets of beamforming weights along with the Section Extension beamIds, or only the Section Extension beamIds (after, loading the beamforming weights in a previous time).

When Section Extension 11 is used, O-RU shall ignore the beamId in the section description

7.9.4.2 Interaction between Section Extension 11 and Section Extension 6

7.9.4.2.1 Overview

Section Extension 6 is used for non-contiguous PRB allocation in both time and frequency domains. The main parameter in Section Extension 6 that is used to identify the PRB groups is `rbgSize`. On the other hand, `numBundPrb` is used in Section Extension 11 to identify the number of bundled PRBs that share the same beamforming weights vector for L TRXs.

7.9.4.2.2 Interaction when $RAD = 1$

When Section Extension 11 and Section Extension 6 is used together, following requirements shall apply when RAD is set to 1:

- 1) The value of the `numBundPrb` field may be different from `rbgSize` field.
- 2) For each contiguous set of PRBs as defined under Section Extension 6 (if used), the O-RU shall start counting for the number of bundled PRBs and apply the sets of beamforming weights to PRB bundles in an increasing order of PRBs.
 - a. If the last bundled PRB does not coincide with the last RBG in the contiguous set of PRBs, then the unassigned PRBs (i.e., PRBs with unassigned beamforming weights) will be considered orphan PRB(s).
- 3) Orphan PRBs are allocated the next set of beamforming weights (i.e., set after the last set that was assigned to a regular PRB bundle) although the number of PRBs is less than a PRB bundle. For any set of PRBs that are not being assigned using Section Extension 6, no beamforming weight sets shall be assigned.
- 4) For the set of contiguous PRBs allocated in the same symbol after the discontinuity, the O-RU shall apply the next set of beamforming weights (starting from the last set of beamforming weights allocated to the last PRB bundle in the previous contiguous PRB set).
- 5) The process continues until all the PRBs defined under Section Extension 6 are assigned beamforming weights.
- 6) O-RU shall apply the same sets of beamforming weights to the corresponding PRB bundles in all symbols selected via `symbolMask`.

EXAMPLE 1: $RAD = 1b$ (see Figure 7.9.4.2.2-1 below)

`symbolMask = 10 0001 0001 0000b`

`rbgSize = 3`

`numBundPrb = 2`

Total number of L beamforming weights set = 7

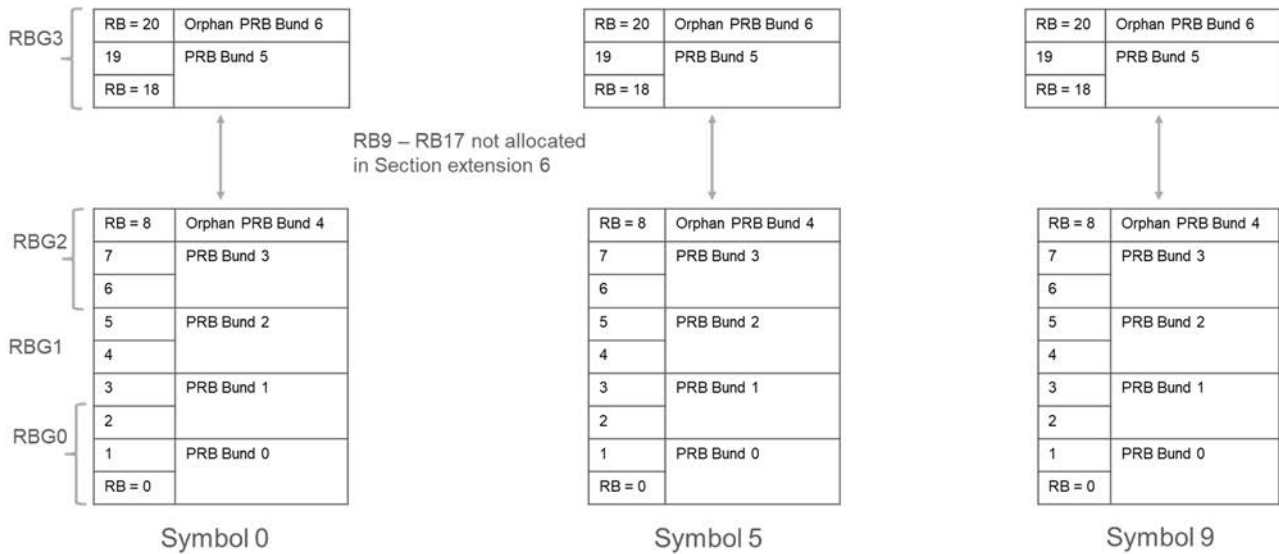


Figure 7.9.4.2.2-1: Example for how Section Extension 6 interacts with Section Extension =11: Case RAD=1

7.9.4.2.3 Interaction when RAD = 0

When Section Extension 11 and Section Extension 6 is used together, following requirements shall apply when RAD is set to 0. The set of beamforming weights for a given bundle shall be applied to the PRBs in increasing order regardless of the discontinuity in the PRBs allocation in Section Extension 6. O-DU shall send beamforming weights for a PRB bundle if and only if a PRB allocated using Section Extension 6 is part of the PRB bundle

EXAMPLE 1: RAD = 0b (see Figure 7.9.4.2.3-1 below)

symbolMask = 10 0001 0001 0000b

rbgSize = 3

numBundPrb = 10

Total number of L beamforming weights set = 3

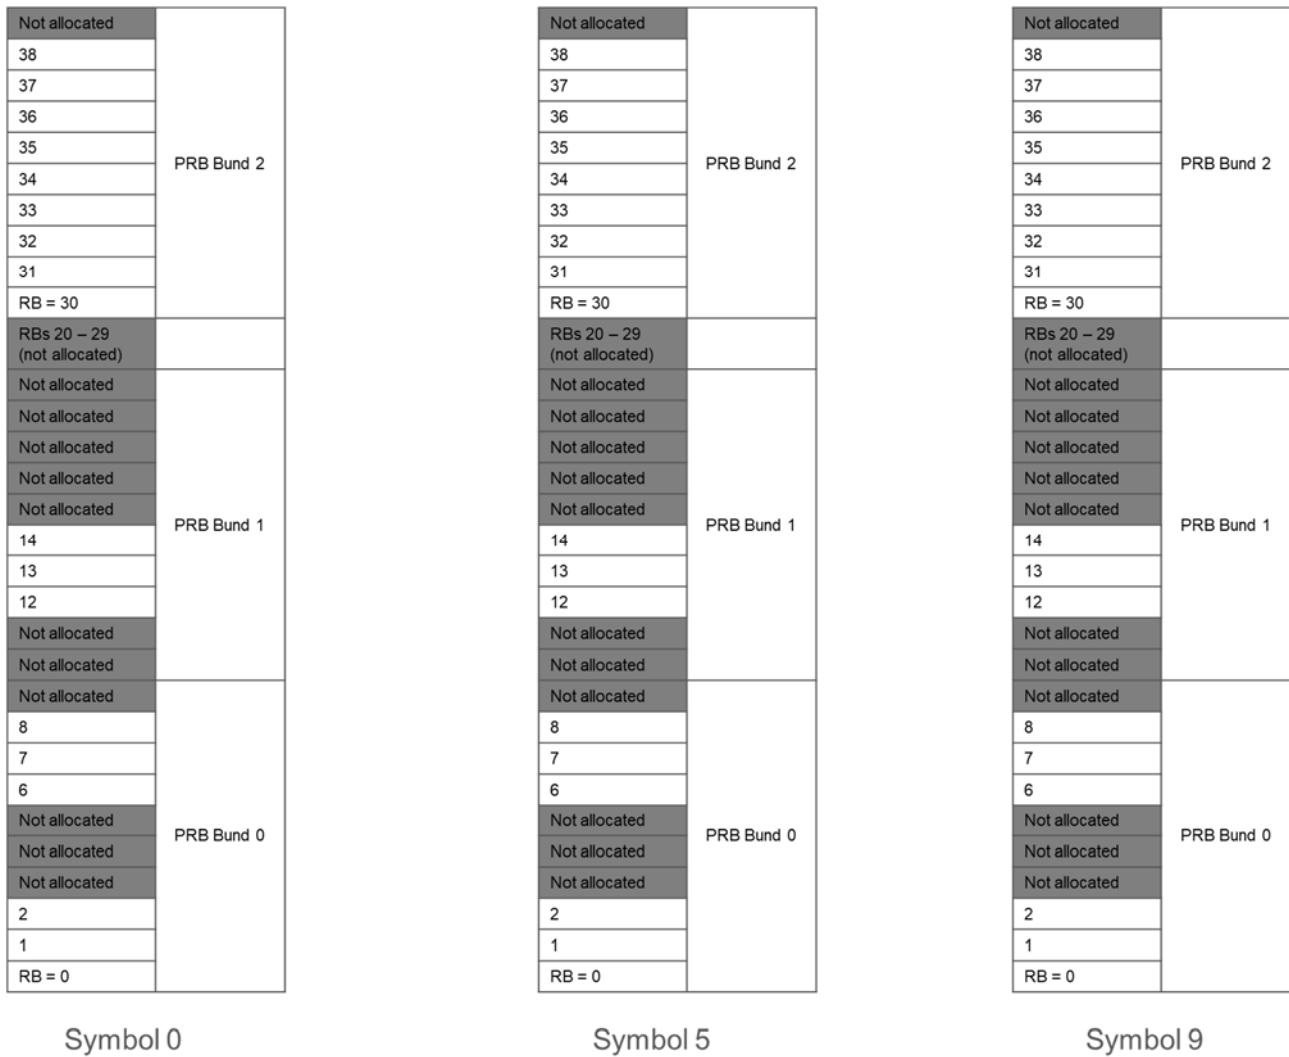


Figure 7.9.4.2.3-1: Example for how Section Extension =6 interacts with Section Extension 11: Case RAD = 0

7.9.4.3 Interaction between Section Extension 11 and Section Extension 12

Section Extension 12 is used for non-contiguous PRB allocation in both time and frequency domains. Section Extension 11 can be combined with Section Extension 12 in one section description in the same way as how Extension 11 may be combined with Section Extension 6. That is, description in clause 7.9.4.2 applies also to Section Extension 12. See also the examples below.

EXAMPLE 1: for RAD = 0 (see Figure 7.9.4.3-1 below)

symbolMask = 00 0000 0000 0101b

grey rectangles represent unallocated frequency ranges (a single PRB discontinuity starting at PRB 5)

numBundPrb = 3

Total number of L beamforming weights set = 4

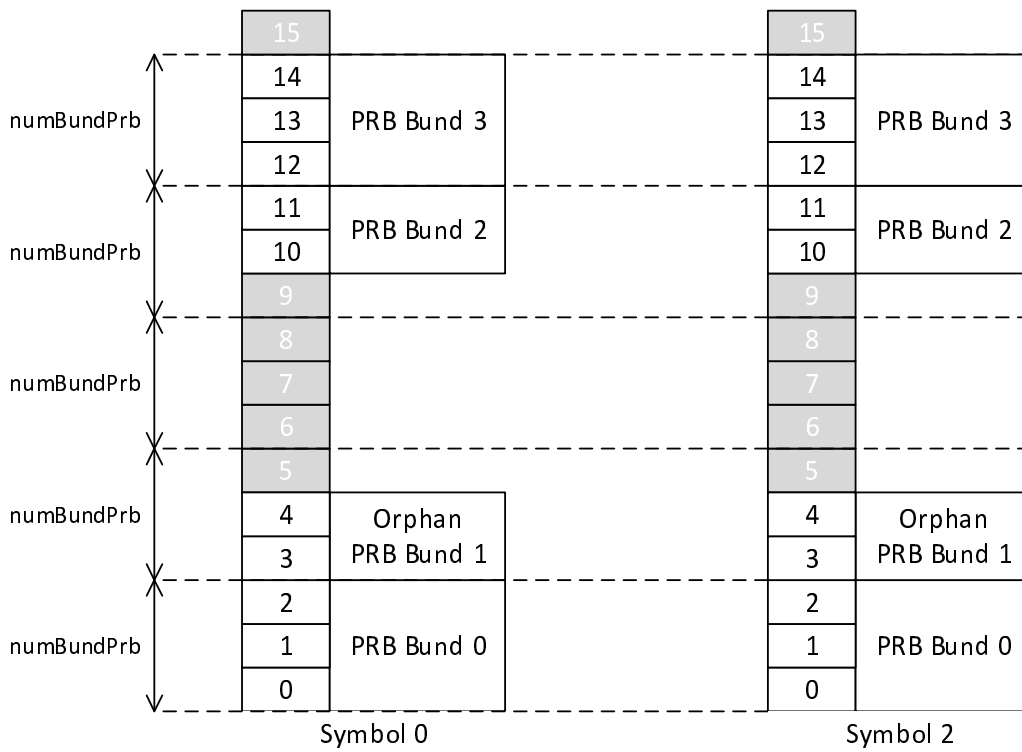


Figure 7.9.4.3-1: Example for how Section Extension 12 interacts with Section Extension 11: Case RAD=0

EXAMPLE 2: for RAD = 1 (see Figure 7.9.4.3-2 below)

symbolMask = 00 0000 0000 0101b

grey rectangles represent unallocated frequency ranges (a single PRB discontinuity starting at PRB 5)

numBundPrb = 3

Total number of L beamforming weights set = 4

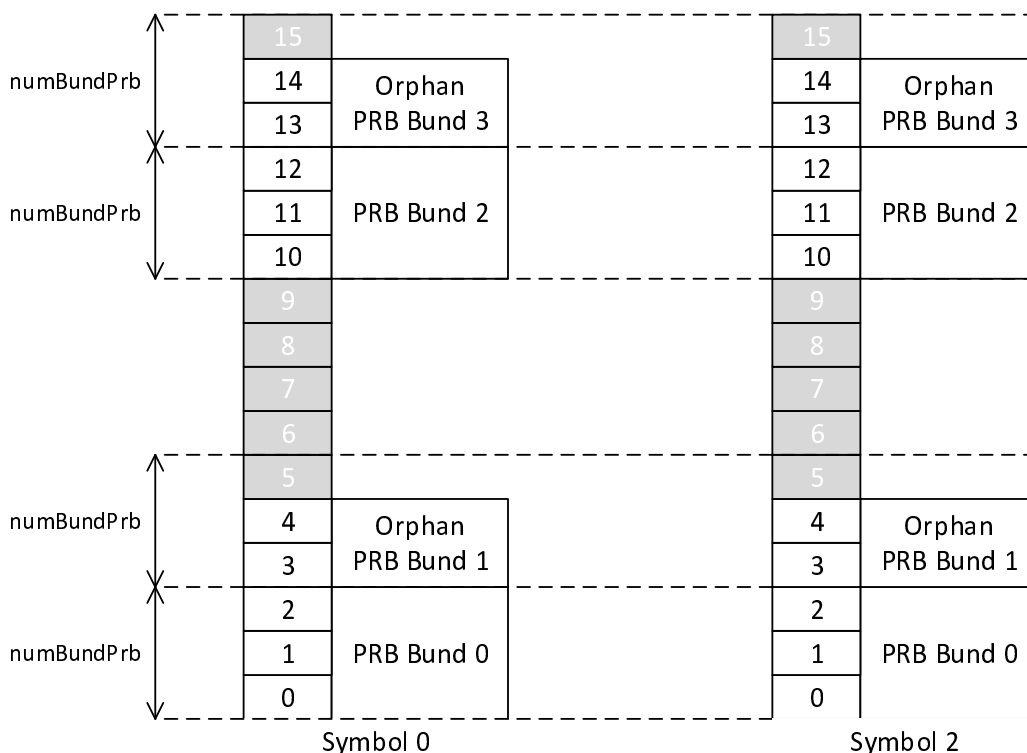


Figure 7.9.4.3-2: Example for how Section Extension 12 interacts with Section Extension 11: Case RAD = 1

7.9.4.4 Interaction between Section Extension 11 and Section Extension 13

Section Extension 13 is used for non-contiguous PRB allocation with frequency hopping. Section Extension 11 can be combined with Section Extension 13 in one section description in the same way as Section Extension 11 is combined with Section Extension 6 as described in clause 7.9.4.2 with following additional requirements:

- beamIds and beamforming weights carried by Section Extension 11 shall be provided for a combined PRB set that is result of merging PRB sets of all hops. This rule is general and applies to Section Extension 13 combined with Section Extensions allowing non-contiguous PRB allocation (Section Extensions 6 and 12).
- PRB discontinuities are discontinuities of the combined PRB set.

EXAMPLE 1: for RAD = 0 (see Figure 7.9.4.4-1 below)

startSymbolId=0 (and message includes only one section description)
 startPrbc = 10, numPrbc = 5, numSymbol=3

nextStartSymbolId(0) = 1, nextStartPrbc(0) = 0,

nextStartSymbolId(1) = 2, nextStartPrbc(1) = 7,

grey rectangles represent unallocated frequency ranges

note the combined set of PRB has one discontinuity starting at PRB 5

numBundPrb = 3

Total number of L beamforming weights set = 5

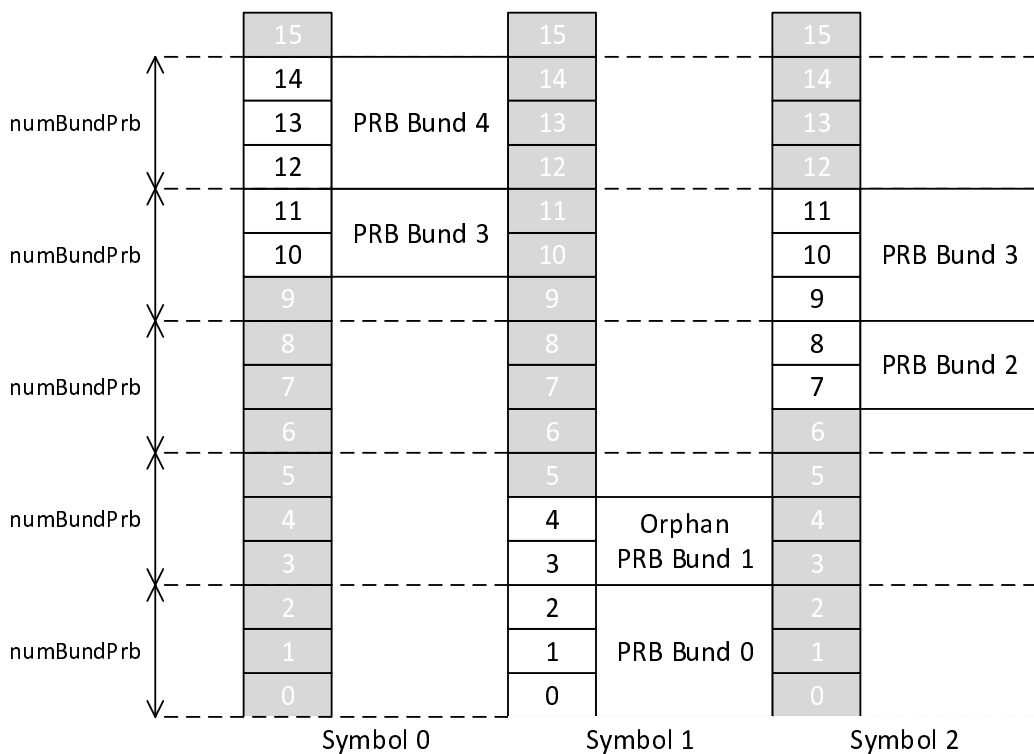


Figure 7.9.4.4-1: Example for how Section Extension 13 interacts with Section Extension 11: Case RAD = 0

EXAMPLE 2: for RAD = 1 (see Figure 7.9.4.4-2 below)

startSymbolId=0 (and message includes only one section description)
 startPrbc=10, numPrbc=5, numSymbol=3

nextStartSymbolId(1)=1, nextStartPrbc(1) = 0,

nextStartSymbolId(2)=2, nextStartPrbc(2) = 7,

gray rectangles represent unallocated frequency ranges

note the combined set of PRB has one discontinuity starting at PRB 5

numBundPrb = 3

Total number of L beamforming weights set = 5

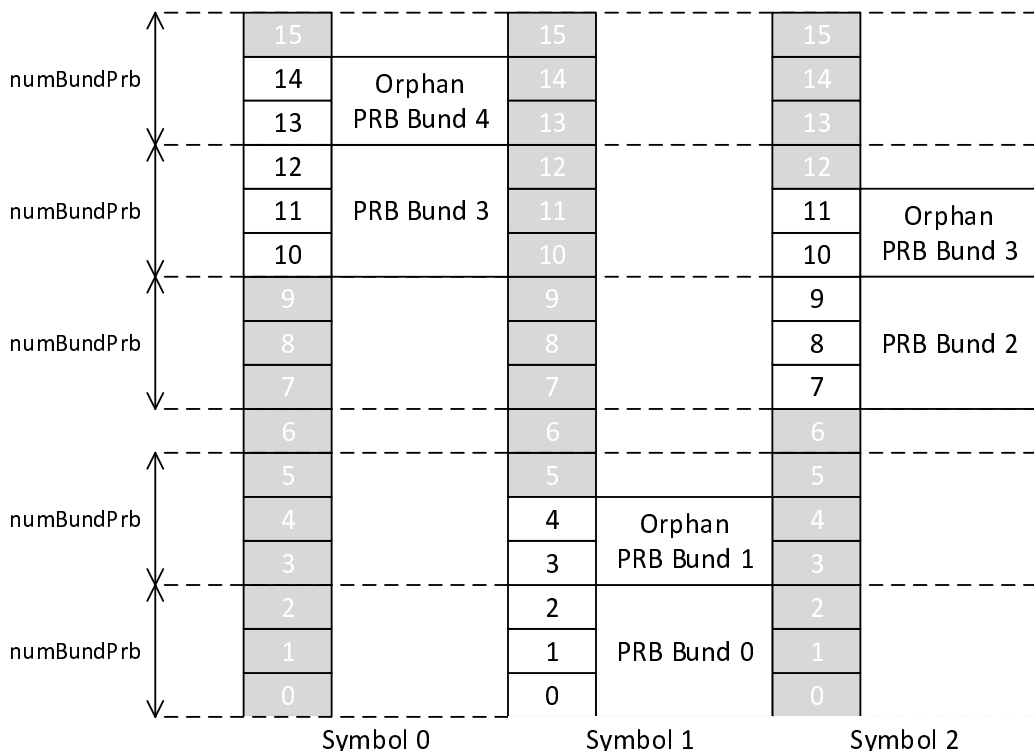


Figure 7.9.4.4-2: Example for how Section Extension 13 interacts with Section Extension 11: Case RAD = 1

EXAMPLE 3: for RAD = 0 with a non-contiguous PRB allocation using Section Extension 12 and 13 (see Figure 7.9.4.4-3)

startPrbc=8, numPrbc=2,

nextStartSymbolId(1)=2, nextStartPrbc(1) = 0,

symbolMask= 00 0000 0000 0101b

offStartPrb(1)=4, numPrb(1)=2

gray rectangles represent unallocated frequency ranges

note the combined set of PRB has two discontinuities starting at PRB 2 and PRB 10

numBundPrb = 3

Total number of L beamforming weights set = 5

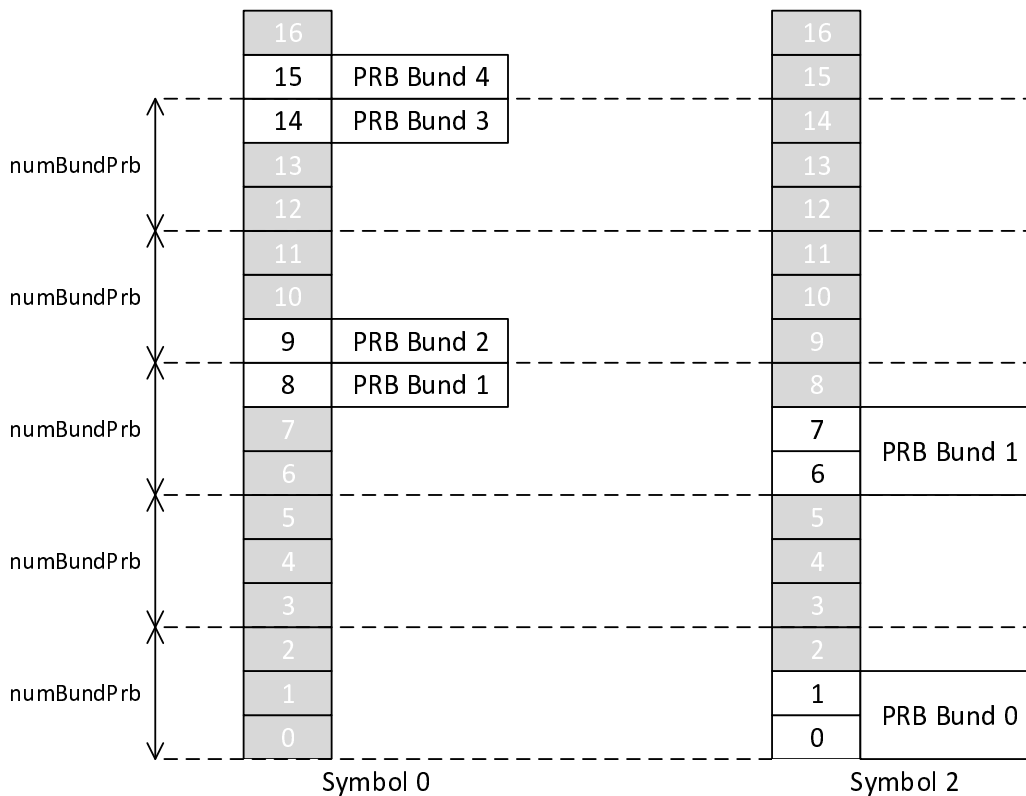


Figure 7.9.4.4-3: Example for how Section Extension 13 and 12 interact with Section Extension 11: Case RAD = 0

EXAMPLE 4: for RAD = 1 with a non-contiguous PRB allocation using Section Extension 12 and 13 (see Figure 7.9.4.4-4)

startPrbc=8, numPrbc=2,

nextStartSymbolId(1)=2, nextStartPrbc(1) = 0,

symbolMask= 00 0000 0000 0101b

offStartPrb(1)=4, numPrb(1)=2

gray rectangles represent unallocated frequency ranges

note the combined set of PRB has two discontinuities starting at PRB 2 and PRB 10

numBundPrb = 3

Total number of L beamforming weights set = 4

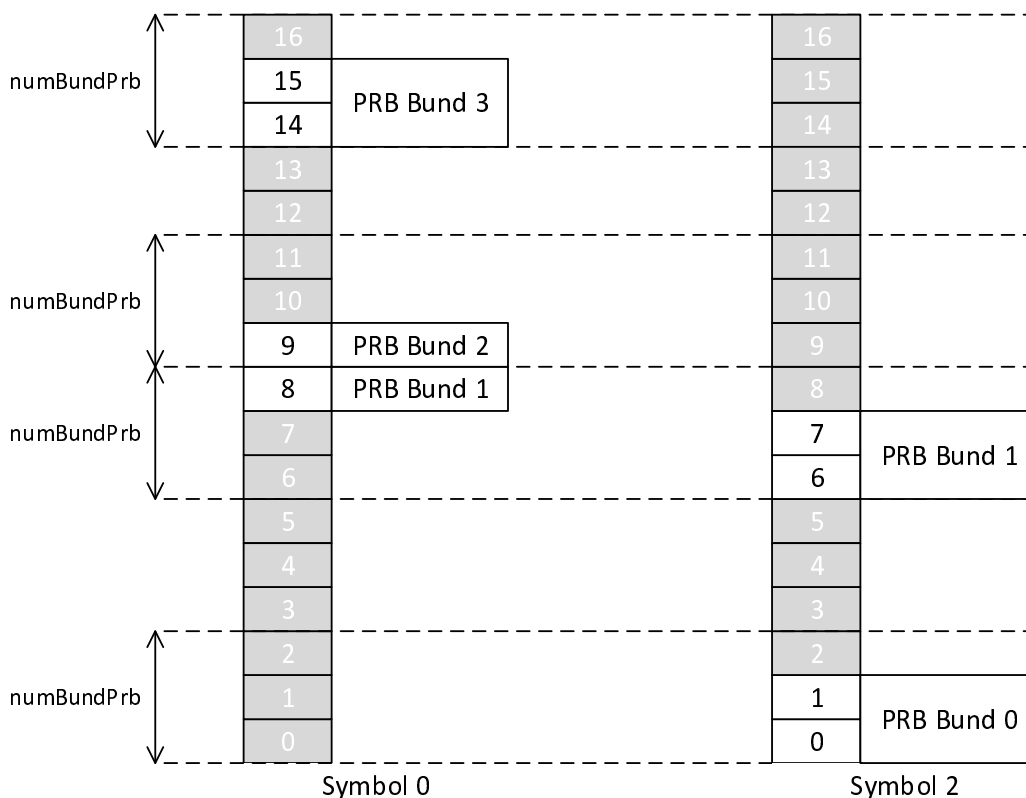


Figure 7.9.4.4-4: Example for how Section Extension 13 and 12 interact with Section Extension 11: Case RAD = 1

7.9.5 C-Plane optimization using Section Extension 12

The Section Extension 12 can be used for non-contiguous PRB allocation in time and frequency domains. The Section Extension conveys priority, symbolMask and a variable size list of frequency ranges. Section Extension 12 is more space efficient than Section Extension in scenarios where:

- 1) PRB allocation is contiguous in frequency domain but non-contiguous in time domain (in this case the list of frequency ranges can be empty).
- 2) PRB allocation consists of few contiguous fragments of spectrum that are spread over wide frequency span (in this case size of rbgMask is the factor limiting application of Section Extension 6).

This optimization cannot be combined with Section Extension 6 in one section description.

7.9.6 C-Plane optimization using Section Extension 13

The Section Extension 13 can be used for PRB allocation with intra-slot frequency hopping. The Section Extension 13 conveys a variable size list of pairs of parameters indicating a value to be used in place of startPrbc and a time instant (symbol number) starting from which the startPrbc shall be substituted.

When using intra-slot frequency hopping, Section Extension 13 is more size efficient than realizing the frequency hopping using only section descriptions. Section Extension 13 can be combined with other Section Extensions. In such case interpretation of frequency parameters conveyed by the section description and all Section Extensions shall be as if value of startPrbc was modified at time instances according to parameters in Section Extension 13.

7.9.7 Coupling via frequency and time with priorities

This optimization uses the coupling mechanism described in clause 7.8.1.4 and allows the reduction of the number of section descriptions. The coupling mechanism with priorities allows to unambiguously describe beamforming

configuration even if section descriptions overlap (i.e., refer to the same RE) provided they have different priorities. Message size reduction can be achieved by moving beamforming configuration of selected REs (e.g., REs allocated to reference signals) to a combined section description with a higher priority. With this, other section descriptions of lower priority can be simplified (e.g., details of RE allocation in PRB conveyed in reMask can be simplified by setting reMask to all ones) and may be combined. Combining section descriptions requires use of coupling mechanism that does not use sectionId. The meaning of "higher priority" and "lower priority" is relative e.g., priority 0 is higher than priority -1.

To enable O-RU to decide beamforming of a RE by processing a single C-Plane message, O-DU shall ensure that each message that has a section description referring to a RE includes the section description referring to the RE with highest priority. This may result in duplicating a subset of section description in more than one message but does not result in beamforming configuration ambiguity. The "highest priority" is relative e.g., if a RE is referenced only by section description with priority -1 then this is the section description with the highest priority.

The following example is simplified (e.g., symbolMask and reMask have less bits than actual fields in a message) to demonstrate the optimization mechanism and does not reflect a realistic resource allocation.

EXAMPLE: Initial set of section descriptions (assuming coupling via sectionId) is presented in Table 7.9.7-1 below.

Table 7.9.7-1: Initial set of section descriptions

#	sectionId	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	100	0110110	0	4	111111	100	UE1
2	101	1001000	0	4	110110	100	UE1
3	101	1001000	0	4	001000	1	Reference signal
4	101	1001000	0	4	000001	2	Reference signal
5	200	0110110	4	4	111111	200	UE2
6	201	1001000	4	4	110110	200	UE2
7	201	1001000	4	4	001000	1	Reference signal
8	201	1001000	4	4	000001	2	Reference signal
9	300	0110110	8	4	111111	300	UE3
10	301	1001000	8	4	110110	300	UE3
11	301	1001000	8	4	001000	1	Reference signal
12	301	1001000	8	4	000001	2	Reference signal
13	400	0000001	5	2	111111	3	Special channel

Optimized set of section descriptions (sectionId value is fixed and not significant) is presented in Table 7.9.7-2 below.

Table 7.9.7-2: Optimized set of section descriptions

#	priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	+1	1001000	0	12	001000	1	Reference signal
2	+1	1001000	0	12	000001	2	Reference signal
3	0	1111110	0	4	111111	100	UE1
4	0	1111110	4	4	111111	200	UE2
5	0	1111110	8	4	111111	300	UE3
6	0	0000001	5	2	111111	3	Special channel

In this example, in the optimized set of section descriptions, section descriptions 3, 4, 5 and 6 span over a continuous set of symbols and have priority 0 (default). They can be represented without Section Extension 6.

If optimized set of section descriptions does not fit in one message, then it is divided into several messages e.g., as presented in Table 7.9.7-3, Table 7.9.7-4 and Table 7.9.7-5 Table 7.9.7-3below. The split into messages is only to demonstrate duplication of highest priority section descriptions.

Table 7.9.7-3: Message #1

#	priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	+1	1001000	0	12	001000	1	Reference signal
2	+1	1001000	0	12	000001	2	Reference signal

3	0	1111110	0	4	111111	100	UE1
4	0	1111110	4	4	111111	200	UE2

Table 7.9.7-4: Message #2

#	priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	+1	1001000	0	12	001000	1	Reference signal
2	+1	1001000	0	12	000001	2	Reference signal
3	0	1111110	8	4	111111	300	UE3

In the message #2 presented in Table 7.9.7-4 section descriptions 1 and 2 are duplicated because they are highest priority section descriptions referring to REs referred by section descriptions in the message (here section description 3).

Table 7.9.7-5: Message #3

#	Priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	0	0000001	5	2	111111	3	Special channel

In the message #3 presented in Table 7.9.7-5 section description 1 is the section description with highest priority and there is no need to duplicate section descriptions with priority +1. In other words: for every RE referenced from the message the message contains the highest priority section description referring to the RE.

7.9.8 Coupling via frequency and time with priorities (optimized)

This optimization is an extension of optimization described in clause 7.9.7 that allows to identify duplication of highest priority section descriptions. When highest priority section descriptions are duplicated, O-RU processes the duplicated sections multiple times. To identify duplication, if 'coupling via frequency and time with priorities optimized' is enabled, O-DU shall set each highest priority section to have a unique sectionId (see clause 7.4.1.2.4). This enables O-RU to identify duplicate highest priority sections and take an optimized action.

EXAMPLE: Using the same example as in clause 7.9.7 the optimized set of section descriptions is presented in Table 7.9.8-1 with sectionId column added (sectionId value is fixed to 4095 for lower priority sections whereas highest priority section will have unique sectionId values in a defined range).

Table 7.9.8-1: Optimized set of section descriptions

#	sectionId	priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	0	+1	1001000	0	12	001000	1	Reference signal
2	1	+1	1001000	0	12	000001	2	Reference signal
3	4095	0	1111110	0	4	111111	100	UE1
4	4095	0	1111110	4	4	111111	200	UE2
5	4095	0	1111110	8	4	111111	300	UE3
6	4095	0	0000001	5	2	111111	3	Special channel

As described in in clause 7.9.7 when section description does not fit in one message, the message is divided across several message where the highest priority sections are duplicated.

Table 7.9.8-2: Message #1

#	sectionId	priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	0	+1	1001000	0	12	001000	1	Reference signal
2	1	+1	1001000	0	12	000001	2	Reference signal
3	4095	0	1111110	0	4	111111	100	UE1
4	4095	0	1111110	4	4	111111	200	UE2

In the message #1 presented in Table 7.9.8-2 highest priority sections are associated with unique sectionIds to be identified by O-RU when same highest priority sections are duplicated across C-Plane messages. Sections with priority 0 are still assigned 4095 as the sectionId since they are not duplicated.

Table 7.9.8-3: Message #2

#	sectionId	priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	0	+1	1001000	0	12	001000	1	Reference signal
2	1	+1	1001000	0	12	000001	2	Reference signal
3	4095	0	1111110	8	4	111111	300	UE3

Duplicated section descriptions across message #1 and message #2 presented in Table 7.9.8-3 can now be identified by O-RU using sectionId.

Table 7.9.8-4: Message #3

#	sectionID	Priority	symbolMask	startPrbc	numPrbc	reMask	beamId	Note
1	4095	0	0000001	5	2	111111	3	Special channel

The message #3 presented in Table 7.9.8-4 is the same as in clause 7.9.7.

7.9.9 U-Plane Operation Without C-Plane

O-RU may support U-Plane operation without C-Plane. In general, this function can be used for channels with resource allocation known (at least approximately, see below) during eAxC activation (e.g. PRACH and SRS). If this function is enabled via M-Plane for a given eAxC then for this eAxC, C-Plane is not used to provide control information. Instead, the O-DU shall configure the O-RU with a static control information (resource allocation details and beamforming configuration are set before activation) via M-Plane. For more details see clause 8.2.2.

This functionality can also be used for channels where only approximate resource allocation is known during activation (e.g. SRS channel reception is scheduled periodically at known interval even if no UE is present). In this case it is possible resource allocation conveyed via M-Plane exceeds actual resource allocation e.g. O-RU would be statically configured to receive channel related REs without reflecting actual, varying over time allocations. Such approximations generally cause additional UL U-Plane traffic related to unallocated (at given time) resource elements; allocation of DL U-Plane traffic is not expected.

This C-Plane optimization method is not compatible with any Section Extensions since C-Plane messages are not used for selected eAxC.

7.9.10 Modulation compression with Section Extension 10

Section Extension 10 is used for group configuration of multiple ports. Section Extension 4, Section Extension 5 and Section Extension 23 are used for modulation compression. Extension type 10 can be used together with extension type 4 or 5 or 23. When all parameter values of Section Extension 4 or 5 is same for all eAxC_IDs, one single Section Type 4 or 5 or 23 shall be appended after Section Extension 10. When parameter values are different for all eAxCs, Section Extension 4 or 5 or 23 for all eAxC_IDs in sequence based on eAxC_ID order shall be appended after Section Extension 10.

If O-RU receives only one Section Extension 4 or 5 or 23 with Section Extension 10, O-RU applies same parameters to all eAxC_IDs. If O-RU receives equal to the number of the eAxC_IDs grouped, O-RU applies the Section Extension 4 or 5 or 23 in the order of how eAxC_IDs are grouped. Any other number of Section Extension 4 or 5 or 23 is an error condition. O-RU shall use numPortc parameter in Section Extension 10 to identify number of Section Extension 4 or 5 or 23 that are present.

EXAMPLE 1: Example with 4 Layers (numPortc=3), beamGroupType=00b or 01b for Section Extension 10 and all eAxC_IDs share same modulation compression parameters is shown in Table 7.9.10-1. One Section Extension 5 is appended after Section Extension 10 in this case (Section Extension 5 with two scaler values, modulation compression parameters is used in this example).

Table 7.9.10-1: Section Extension 10 for beamGroupType = 00b or 01b with one Section Extension 5

0 (msb)	1	2	3	4	5	6	7(lsb)	# of bytes	
ef	extType = 0x0A							1	Octet N
extLen = 0x01 (1 word)								1	N+1
beamGroupType		numPortc=3						1	N+2
reserved								1	N+3
ef	extType = 0x05							1	N+4
extLen = 0x3 (3 words)								1	N+5
mcScaleReMask[11:4]								1	N+6
mcScaleReMask[3:0]				csf		mcScaleOffset [14:12]		1	N+7
mcScaleOffset [11:4]								1	N+8
mcScaleOffset [3:0]				mcScaleReMask[11:8]				1	N+9
mcScaleReMask[7:0]								1	N+10
csf	mcScaleOffset [14:8]							1	N+11
mcScaleOffset [7:0]								1	N+12
zero padding								1	N+13
zero padding								1	N+14
zero padding								1	N+15

EXAMPLE 2: Example with 4 Layer (numPortc=3), beamGroupType=00b or 01b for Section Extension 10 and eAxC_IDs use different modulation compression parameters is shown in Table 7.9.10-2. Four instances of Section Extension 5 are appended after Section Extension 10 in this case (Section Extension 5 with one scaler value, modulation compression parameters is used in this example).

Table 7.9.10-2: Section Extension 10 for beamGroupType = 00b or 01b with multiple Section Extension 5

0 (msb)	1	2	3	4	5	6	7(lsb)	# of bytes	
ef	extType = 0x0A							1	Octet N
extLen = 0x01 (1 word)								1	N+1
beamGroupType	numPortc=3							1	N+2
reserved								1	N+3
ef	extType = 0x05							1	N+4
extLen = 0x2 (2 words)								1	N+5
mcScaleReMask[11:4]								1	N+6
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+7
mcScaleOffset [11:4]								1	N+8
mcScaleOffset [3:0]			zero padding					1	N+9
zero padding								1	N+10
zero padding								1	N+11
ef	extType = 0x05							1	N+12
extLen = 0x2 (2 words)								1	N+13
mcScaleReMask[11:4]								1	N+14
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+15
mcScaleOffset [11:4]								1	N+16
mcScaleOffset [3:0]			zero padding					1	N+17
zero padding								1	N+18
zero padding								1	N+19
ef	extType = 0x05							1	N+20
extLen = 0x2 (2 words)								1	N+21
mcScaleReMask[11:4]								1	N+22
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+23
mcScaleOffset [11:4]								1	N+24
mcScaleOffset [3:0]			zero padding					1	N+25
zero padding								1	N+26
zero padding								1	N+27
ef	extType = 0x05							1	N+28
extLen = 0x2 (2 words)								1	N+29
mcScaleReMask[11:4]								1	N+30
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+31
mcScaleOffset [11:4]								1	N+32
mcScaleOffset [3:0]			zero padding					1	N+33
zero padding								1	N+34
zero padding								1	N+35

EXAMPLE 3: Example with 4 Layer (numPortc=3), beamGroupType=10b for Section Extension 10, and all eAxC_IDs share same modulation compression parameters is shown in Table 7.9.10-3. One Section Extension 5 is appended after Section Extension 10 in this case (Section Extension 5 with two scaler values, modulation compression parameters in this example).

Table 7.9.10-3: Section Extension 10 for beamGroupType = 10b with one single Section Extension 5

0 (msb)	1	2	3	4	5	6	7(lsb)	# of bytes	
ef	extType = 0x0A							1	Octet N
extLen = 0x03 (3words)								1	N+1
beamGroupType		numPortc=3						1	N+2
reserved	2 nd port beamId[14:8] (or ueld[14:8])							1	N+3
2 nd port beamId[7:0] (or ueld[7:0])								1	N+4
reserved	3 rd port beamId[14:8] (or ueld[14:8])							1	N+5
3 rd port beamId[7:0] (or ueld[7:0])								1	N+6
reserved	4 th port beamId[14:8] (or ueld[14:8])							1	N+7
4 th port beamId[7:0] (or ueld[7:0])								1	N+8
zero padding								1	N+9
zero padding								1	N+10
zero padding								1	N+11
ef	extType = 0x05							1	N+12
extLen = 0x3 (3 words)								1	N+13
mcScaleReMask[11:4]								1	N+14
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+15
mcScaleOffset [11:4]								1	N+16
mcScaleOffset [3:0]			mcScaleReMask[11:8]					1	N+17
mcScaleReMask[7:0]								1	N+18
csf	mcScaleOffset [14:8]							1	N+19
mcScaleOffset [7:0]								1	N+20
zero padding								1	N+21
zero padding								1	N+22
zero padding								1	N+23

EXAMPLE 4: Example with 4 Layer (numPortc=3), beamGroupType=10b for Section Extension 10, and eAxC_IDs use different modulation compression parameters is shown in Table 7.9.10-4. Four instances of Section Extension 5 are appended after Section Extension 10 in this case (Section Extension 5 with one scaler value, modulation compression parameters is used in this example).

Table 7.9.10-4: Section Extension 10 for beamGroupType = 10b with multiple Section Extension 5

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
ef	extType = 0x0A							1	Octet N
extLen = 0x03 (3words)								1	N+1
beamGroupType		numPortc=3						1	N+2
reserved	2 nd port beamId[14:8] (or ueld[14:8])							1	N+3
2 nd port beamId[7:0] (or ueld[7:0])								1	N+4
reserved	3 rd port beamId[14:8] (or ueld[14:8])							1	N+5
3 rd port beamId[7:0] (or ueld[7:0])								1	N+6
reserved	4 th port beamId[14:8] (or ueld[14:8])							1	N+7
4 th port beamId[7:0] (or ueld[7:0])								1	N+8
zero padding								1	N+9
zero padding								1	N+10
zero padding								1	N+11
ef	extType = 0x05							1	N+12
extLen = 0x2 (2 words)								1	N+13
mcScaleReMask[11:4]								1	N+14
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+15
mcScaleOffset [11:4]								1	N+16
mcScaleOffset [3:0]			zero padding					1	N+17
zero padding								1	N+18
zero padding								1	N+19
ef	extType = 0x05							1	N+20
extLen = 0x2 (2 words)								1	N+21
mcScaleReMask[11:4]								1	N+22
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+23
mcScaleOffset [11:4]								1	N+24
mcScaleOffset [3:0]			zero padding					1	N+25
zero padding								1	N+26
zero padding								1	N+27
ef	extType = 0x05							1	N+28
extLen = 0x2 (2 words)								1	N+29
mcScaleReMask[11:4]								1	N+30
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+31
mcScaleOffset [11:4]								1	N+32
mcScaleOffset [3:0]			zero padding					1	N+33
zero padding								1	N+34
zero padding								1	N+35
ef	extType = 0x05							1	N+36
extLen = 0x2 (2 words)								1	N+37
mcScaleReMask[11:4]								1	N+38
mcScaleReMask[3:0]			csf		mcScaleOffset [14:12]			1	N+39
mcScaleOffset [11:4]								1	N+40
mcScaleOffset [3:0]			zero padding					1	N+41
zero padding								1	N+42
zero padding								1	N+43

7.9.11 Optimization with Section Extension 19

The benefits of Section Extension 19 can be explained with a CSI-RS example. Consider a CSI-RS resource with 32 ports, FD-CDM2. Each port is associated with a separate beam. Up to 16 ports (belonging to one CDM group) can be sent in a single layer (eAxC). Without use of Section Extension 19 this would require 16 section invocations. Each section invocation is 8 bytes for a total of 128 bytes. Assuming 4 CSI-RS resources, this amounts to 64 sections and 512 bytes.

When used with Time and Frequency coupling with priorities, this example will result in 64 highest priority sections. Fragmentation of C-Plane packets would further result in duplicating these highest priority sections in every C-Plane message, adding further to the fronthaul load in addition to O-DU and O-RU processing.

With Section Extension 19 in the same example as above, as opposed to processing up to 16 section headers per CSI-RS resource, only one section header would need to be processed by the O-RU. Processing highest priority sections for removing overlapping reMask bits would become more efficient. Rather than accumulating reMasks and symbolMask across 16 high priority sections, the combined reMask and symbolMask of the entire CSI-RS resource would be provided in one shot via reMask in the section header & symbolMask in the Section Extension. This would make the O-DU and O-RU processing more efficient, especially when dealing with fragmentation of C-Plane packets.

Moreover, when repeating highest priority sections, the port information associated with Section extension 19 would not be repeated, leading to fronthaul savings.

7.9.12 Optimizations with Section Extension 20

Section Extension 20 (see clause 7.7.20) is like coupling via frequency and time with priorities, see clause 7.8.1.4, but more efficient and more widely applicable. Using Section Extension 19 rather than Section Extension 20, if there is C-Plane application layer packet fragmentation, each C-Plane message shall contain the highest priority section descriptions referring to any RE that is referred in a message (see clause 7.8.1.3 restriction item #3). This can lead to increased fronthaul and processing overhead.

Section Extension 20 optimizes the method of specifying the overlapping RE information (puncturing information) with reduced fronthaul and processing overhead. This extension contains only the puncturing information (specified via Resource Element masks) of the channels without any beamforming information. It further consolidates the puncturing information of several channels into a single extension. If there is C-Plane application layer packet fragmentation, beamforming information needs to be sent only once and only puncturing information (which is compact) needs to be repeated in every C-Plane message. Using Section Extension 20 simplifies the section description (e.g., details of RE allocation in PRB conveyed in reMask may be simplified by setting reMask to all ones). This can reduce the number of sections required.

This Section Extension can be used with all coupling methods as described in clause 7.8.1, thereby benefiting all coupling methods with reduced number of sections.

7.9.13 Optimizations with ueId indicating non-scheduled PRB range

In channel-information-based beamforming, the O-RU calculates the BF weight based on the channel information of the ueId(s) from Section Type 5 and/or Section Extension 10. In MU-MIMO scheduling, co-scheduled UEs/layers may be scheduled with different number of PRBs, as shown in Figure 7.9.13-1.

Layer 5	UE4 (ueId = 30)			
Layer 4	UE3 (ueId = 24)			
Layer 3	UE2 (ueId = 11)			
Layer 2	UE2 (ueId = 10)			
Layer 1	UE1 (ueId = 2)			
Layer 0	UE1 (ueId = 1)			
	RB0~RB18	RB19~RB33	RB34~RB46	RB47~RB99

Figure 7.9.13-1 Example of co-scheduled MU UEs with different PRB allocation

In the example shown in Figure 7.9.13-1, in a C-Plane message, Section Type 5 with Section Extension 10 can only be sent to indicate allocation for first 19 PRBs and rest of the allocation needs to be indicated using individual C-plane messages per port or eAxC id. Even if the O-DU tries to use SE 10 for rest of the frequency and layer allocation, it would require to use multiple C-plane messages.

This can be solved by indicating the 'non-scheduled' information using ueId = 0x7FFF in Section Type 5 and/or Section Extension 10 with beamGroupType = 10b to indicate the non-scheduled ueId in specific PRB ranges. When ueId =

0x7FFF is used, O-RU shall interpret that those PRBs in the specified port are not allocated and hence shall not be considered for BF weight generation. Table 7.9.13-1 and Figure 7.9.13-2 shows an example of describing the layer and PRB allocation in Figure 7.9.13-1 using four section descriptions (sectionId#A, B, C and D) using non-scheduled-ueld. SectionId#A describes RB0 to RB18 for Layer0 to Layer5 using SE10, sectionId#B describes RB19 to RB33 for Layer0 to Layer5 using SE10 and so on.

	sectionId#A	sectionId#B	sectionId#C	sectionId#D
Layer 5	UE4 (ueld = 30)	UE4 (ueld = 30)	UE4 (ueld = 30)	UE4 (ueld = 30)
Layer 4	UE3 (ueld = 24)	ueld = 0x7FFF	ueld = 0x7FFF	ueld = 0x7FFF
Layer 3	UE2 (ueld = 11)	UE2 (ueld = 11)	UE2 (ueld = 11)	ueld = 0x7FFF
Layer 2	UE2 (ueld = 10)	UE2 (ueld = 10)	UE2 (ueld = 10)	ueld = 0x7FFF
Layer 1	UE1 (ueld = 2)	UE1 (ueld = 2)	ueld = 0x7FFF	ueld = 0x7FFF
Layer 0	UE1 (ueld = 1)	UE1 (ueld = 1)	ueld = 0x7FFF	ueld = 0x7FFF
	RB0~RB18	RB19~RB33	RB34~RB46	RB47~RB99

Figure 7.9.13-2 Example of four ST 5 sections with SE 10 when non-scheduled-ueld is enabled

Table 7.9.13-1 Example of four ST 5 sections with SE 10, non-scheduled-ueld enabled

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes		
transport header, see clause 5.1.3								8	Octet 1	
Radio application header								8	Octet 9	
sectionId = A								1	Octet 17	
sectionId			rb		symInc	startPrbc		1	Octet 18	
startPrbc = 0								1	Octet 19	
NumPrbc = 19								1	Octet 20	
reMask[11:4]								1	Octet 21	
reMask[3:0]				numSymbol					1	Octet 22
ef = 1	ueld[14:8]								1	Octet 23
ueld[7:0] = 1									1	Octet 24
ef = 0	extType = 0x0A								1	Octet 25
extLen									1	Octet 26
beamGroupType=10b		NumPortc = 5							1	Octet 27
reserved	2nd port beamId[14:8] (or ueld[14:8])								1	Octet 28
2nd port beamId[7:0] (or ueld[7:0]) = 2									1	Octet 29
reserved	3rd port beamId[14:8] (or ueld[14:8])								1	Octet 30
3rd port beamId[7:0] (or ueld[7:0]) = 10									1	Octet 31
reserved	4th port beamId[14:8] (or ueld[14:8])								1	Octet 32
4th port beamId[7:0] (or ueld[7:0]) = 11									1	Octet 33
reserved	5th port beamId[14:8] (or ueld[14:8])								1	Octet 34
5th port beamId[7:0] (or ueld[7:0]) = 24									1	Octet 35
reserved	6th port beamId[14:8] (or ueld[14:8])								1	Octet 36
6th port beamId[7:0] (or ueld[7:0]) = 30									1	Octet 37
Padding to 4-byte boundary									3	Octet 38
sectionId = B									1	Octet 41
sectionId			rb		symInc	startPrbc		1	Octet 42	
startPrbc = 19									1	Octet 43
NumPrbc = 15									1	Octet 44
reMask[11:4]									1	Octet 45
reMask[3:0]				numSymbol					1	Octet 46
ef = 1	ueld[14:8]								1	Octet 47
ueld[7:0] = 1									1	Octet 48
ef = 0	extType = 0x0A								1	Octet 49
extLen									1	Octet 50

0 (msb)	1	2	3	4	5	6	7 (lsb)	# of bytes	
beamGroupType=10b		NumPortc = 5						1	Octet 51
reserved	2nd port beamId[14:8] (or ueld[14:8])						1	Octet 52	
2nd port beamIdbeamId[7:0] (or ueld[7:0]) = 2								1	Octet 53
reserved	3rd port beamId[14:8] (or ueld[14:8])						1	Octet 54	
3rd port beamId[7:0] (or ueld[7:0]) = 10								1	Octet 55
reserved	4th port beamId[14:8] (or ueld[14:8])						1	Octet 56	
4th port beamId[7:0] (or ueld[7:0]) = 11								1	Octet 57
reserved	5th port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 58	
5th port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 59
reserved	6th port beamId[14:8] (or ueld[14:8])						1	Octet 60	
6th port beamId[7:0] (or ueld[7:0]) = 30								1	Octet 61
Padding to 4-byte boundary								3	Octet 62
sectionId = C								1	Octet 65
sectionId				rb	symInc	startPrbc		1	Octet 66
startPrbc = 34								1	Octet 67
NumPrbc = 13								1	Octet 68
reMask[11:4]								1	Octet 69
reMask[3:0]				numSymbol				1	Octet 70
ef = 1	ueld[14:8] = 0x7F						1	Octet 71	
ueld[7:0] = 0xFF								1	Octet 72
ef = 0	extType = 0x0A						1	Octet 73	
extLen								1	Octet 74
beamGroupType=10b		NumPortc = 5						1	Octet 75
reserved	2nd port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 76	
2nd port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 77
reserved	3rd port beamId[14:8] (or ueld[14:8])						1	Octet 78	
3rd port beamId[7:0] (or ueld[7:0]) = 10								1	Octet 79
reserved	4th port beamId[14:8] (or ueld[14:8])						1	Octet 80	
4th port beamId[7:0] (or ueld[7:0]) = 11								1	Octet 81
reserved	5th port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 82	
5th port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 83
reserved	6th port beamId[14:8] (or ueld[14:8])						1	Octet 84	
6th port beamId[7:0] (or ueld[7:0]) = 30								1	Octet 85
Padding to 4-byte boundary								3	Octet 86
sectionId = D								1	Octet 89
sectionId				rb	symInc	startPrbc		1	Octet 90
startPrbc = 47								1	Octet 91
NumPrbc = 53								1	Octet 92
reMask[11:4]								1	Octet 93
reMask[3:0]				numSymbol				1	Octet 94
ef = 1	ueld[14:8] = 0x7F						1	Octet 95	
ueld[7:0] = 0xFF								1	Octet 96
ef = 0	extType = 0x0A						1	Octet 97	
extLen								1	Octet 98
beamGroupType=10b		NumPortc = 5						1	Octet 99
reserved	2nd port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 100	
2nd port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 101
reserved	3rd port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 102	
3rd port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 103
reserved	4th port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 104	
4th port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 105
reserved	5th port beamId[14:8] (or ueld[14:8]) = 0x7F						1	Octet 106	
5th port beamId[7:0] (or ueld[7:0]) = 0xFF								1	Octet 107
reserved	6th port beamId[14:8] (or ueld[14:8])						1	Octet 108	
6th port beamId[7:0] (or ueld[7:0]) = 30								1	Octet 109
Padding to 4-byte boundary								3	Octet 110

8 U-Plane Protocol

8.1 General

8.1.1 U-Plane transport

Either eCPRI or IEEE 1914.3 (see clause 5.1.3) shall be used as an encapsulation mechanism for the User Plane messages. Due to the nature of these messages (very strict delay constraints), it is assumed that message acknowledgements are not possible. Likewise, different data flows may be used for the User Plane and Control Plane messages.

8.1.2 U-Plane data compression

U-Plane IQ data, (both DL and UL) including user data, PRACH and control channels may be transmitted in compressed format. There are several envisioned compression methods including an "uncompressed" format. The method of compression may vary based on sectionId but shall be constant for every U-Plane data section referenced in a single C-Plane message. Likewise, the I and Q bitwidth of the compressed samples may vary based on sectionId but shall be constant for every U-Plane data section referenced in a single C-Plane message. The block compression methods shall be performed on a Physical Resource Block (PRB) basis (that is, block compression shall be executed on blocks of 12 Resource Elements in a single PRB). More specific details on the supported compression methods may be found in Annex A.

A static IQ format and compression method, in which case the M-Plane configures the static IQ format (bitwidth) and compression method, is also defined. In this case there are four fewer bytes in the U-Plane section header, at the cost of less flexibility. This allows the taking advantage of the static format while still allowing some compression flexibility e.g., allowing modulation compression on the DL (which only applies to DL) and some other compression method on the UL.

8.1.3 Digital power scaling

8.1.3.1 Definition of IQ power in dBFS

IQ power level in dBFS (dB full scale) is a logarithmic representation of the power level for an IQ sample carried over the digital interface. IQ power level in dBFS is proportional to logarithm of I^2+Q^2 :

$$\text{IQ power level [dBFS]} = 10 \cdot \log_{10}(I^2+Q^2) - 10 \cdot \log_{10}(FS) = 10 \cdot \log_{10}(I^2+Q^2) - 10 \cdot \log_{10}(FS_0 \times 2^{-FS_Offset})$$

Where:

I is the in-phase portion of a received constellation point;

Q is the quadrature portion of a received constellation point;

FS is the Full Scale (maximum) permitted value of I or Q based on their digital representations;

FS_Offset is an M-Plane parameter (value 0 is used if this parameter is not supported by O-RU or not set by O-DU);

$FS_0 = \max(I) = \max(Q) = \max(I^2+Q^2)$ with max over all IQ values that can be represented by IQ data format in U-Plane message.

The actual IQ values that may occur in a U-Plane message are restricted by

$$I^2 + Q^2 \leq FS = FS_0 \cdot 2^{-FS_Offset}.$$

For frequency domain IQ data, 0 dBFS shall be the maximum power level which can be carried by one subcarrier. The smallest non-zero IQ power level is defined by the interface resolution.

It can be expected that an O-RU will normalize any received DL value to its internal representation of full scale so that a 0 dBFS can be properly handled.

EXAMPLE 1:

$$FS_Offset = 0$$

$$I = \min I, Q = 0$$

With 9-bit mantissa 2's complement + 4-bit exponent compression: $\min I = -256 \times 2^{15} = -2^{23} \rightarrow$

$$FS_0 = 2^{46}, FS = FS_0 \times 2^{FS_Offset} = 2^{46}$$

$$0 \text{ dBFS} \Leftrightarrow \text{average}(P+Q^2) = 2^{46}$$

$$\text{Interface resolution} \Leftrightarrow \frac{1}{2} 2^{46} \Leftrightarrow -138.47 \text{ dBFS}$$

EXAMPLE 2:

$$FS_Offset=10$$

$$I=\min I, Q=0$$

With 14-bit mantissa 2's complement + 4-bit exponent compression: $\min I = -2^{13} \times 2^{15} = -2^{28} \rightarrow$

$$FS_0 = 2^{56}, FS = FS_0 \times 2^{-FS_Offset} = 2^{56-10} = 2^{46}$$

$$0 \text{ dBFS} \Leftrightarrow \text{average}(P+Q^2) = 2^{46}$$

$$\text{Interface resolution} \Leftrightarrow \frac{1}{2} 2^{46} \Leftrightarrow -138.47 \text{ dBFS}$$

8.1.3.2 Definition of Gain over Fronthaul Interface

8.1.3.2.1 DL and UL gain definition commonality

The gain of an array is the relation between the levels of a test signal seen at its input and output, also called digital power scaling. The gain of an array may be calculated from the gain of one element of the array while assuming all elements have same gain. Figure 8.1.3.2.1-1 depicts the gain relations between the digital interface and the RF reference plane to an assumed lossless antenna (i.e., antenna insertion losses are counted as part of the gain in both DL and UL direction).

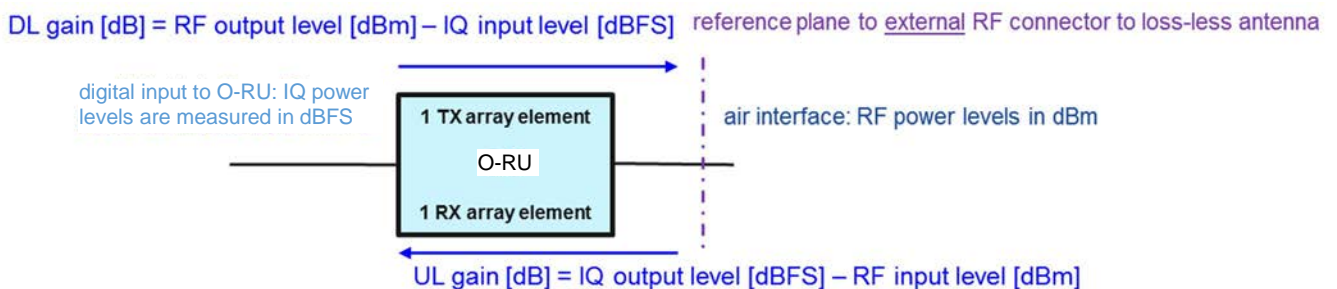


Figure 8.1.3.2.1-1: DL and UL gain over fronthaul interface

Both DL and UL gains in [dB] shall be described by the following relations

$$\text{DL gain [dB]} = \text{RF output level [dBm]} - \text{IQ input level [dBFS]}$$

$$\text{UL gain [dB]} = \text{IQ output level [dBFS]} - \text{RF input level [dBm]}$$

where:

'RF output level [dBm]' and 'RF input level [dBm]' are RF signal level in dBm at a reference plane to an assumed lossless antenna when the antenna is integrated inside the O-RU, or at the antenna connector if the antenna is not integrated inside the O-RU. The rms output power [dBm] is measured per array element as Total Radiated Power TRP (i.e., excluding any antenna directivity gain). The rms input power [dBm] is measured per array element after including the antenna directivity gain. Note that the antenna directivity gain results from the array element directivity + $10 \cdot \log_{10}(\text{number of elements})$. Hence, when evaluating the array element gain, the antenna directivity is reduced to the array element directivity.

'IQ input level [dBFS]' and 'IQ output level [dBFS]' are $10 \cdot \log_{10}(\text{average}(P+Q^2))$ normalized such that 0 dBFS is the maximum nominal (r.m.s.) power level which can be achieved with a constant IQ signal with arbitrary phase (i.e., for frequency-domain IQ signal for **one sub-carrier**, constant over time).

In O-RUs supporting beamforming, the actual DL gain and UL gain of the array element may be impacted by the gain level of beamforming weight used and which may change dynamically during operation. **Such dependency on the beam pattern shall be excluded when describing any configurable gain or reported gain (as O-RU capability).**

For an exact definition of DL and UL gain, the respective input test signals shall be declared by the O-RU vendor; this is not in scope of the present document. [For information only: This can be e.g., continuous-wave (CW) signal at carrier center or test signals/reference signals defined by 3GPP; for DL gain those being defined to test accuracy of the maximum output power (3GPP TS 36.141 [40], clause 6.2 or TS 38.141 [41], clause 6.2) and for UL gain those being defined to test absolute accuracy of Received Interference Power (3GPP TS 36.133 [38], clause 10.1.1 or TS 38.133 [37], clause 10.1.1)]. In Figure 8.1.3.2.1-2 the details of DL gain are described. For every tx-array, the O-RU shall report (as a capability) the maximal configurable DL gain of **one element of the array** (alternatively the nominal power per tx-array element mapped to 0 dBFS). In addition, the O-DU shall configure over the M-Plane the DL gain to be used per tx-array element for a certain carrier configuration.

The values for the max configurable DL gain and configured DL gain shall be configured assuming:

- No power loss/gain due to beamforming weights;
- all available DL power may be allocated to one RE in a single eAxC (if dynamic range optimization is used by configuring non-zero value in M-Plane parameter Reference_Level then single RE can have allocated available DL power less Reference_Level value; see clause 8.1.3.3);

$$\text{DL gain [dB]} = \text{RF output level [dBm]} - \text{IQ input level [dBFS]}$$

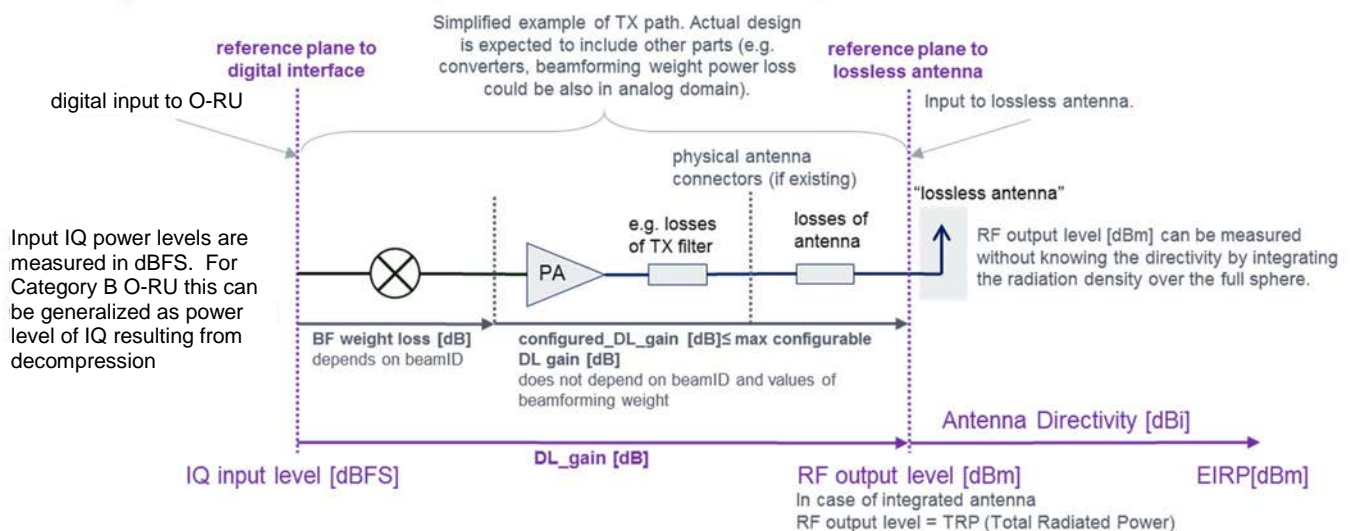


Figure 8.1.3.2.1-2: Details of DL gain

In Figure 8.1.3.2.1-3 the details of UL gain are described. For every rx-array, O-RU shall report (as a capability) an UL gain_correction_range in dB of **one element of the array (applicable to all the elements of the array)**. The gain_correction_range shall be signed, have a max and min value, and a step size. In addition, the O-DU shall configure over the M-Plane a gain_correction value to be applied to the signal received by each array-element to calculate the IQ data in each eAxC (for purpose of backward compatibility this shall be configured as a sum of a common gain correction applicable to all eAxCs of given array carrier and an individual gain correction of each eAxC). The O-RU may then configure its internal UL gain of the rx-array element for that carrier (rx-array carrier element) if the IQ compression method is configured as static. In case the compression method is dynamic, the internal gain of the O-RU may be dynamic depending on compression information received in realtime over the C-Plane.

The values for the UL gain shall be configured assuming no power loss/gain due to beamforming weights.

$$\text{UL gain [dB]} = \text{IQ output level [dBFS]} - \text{RF input level [dBm]}$$

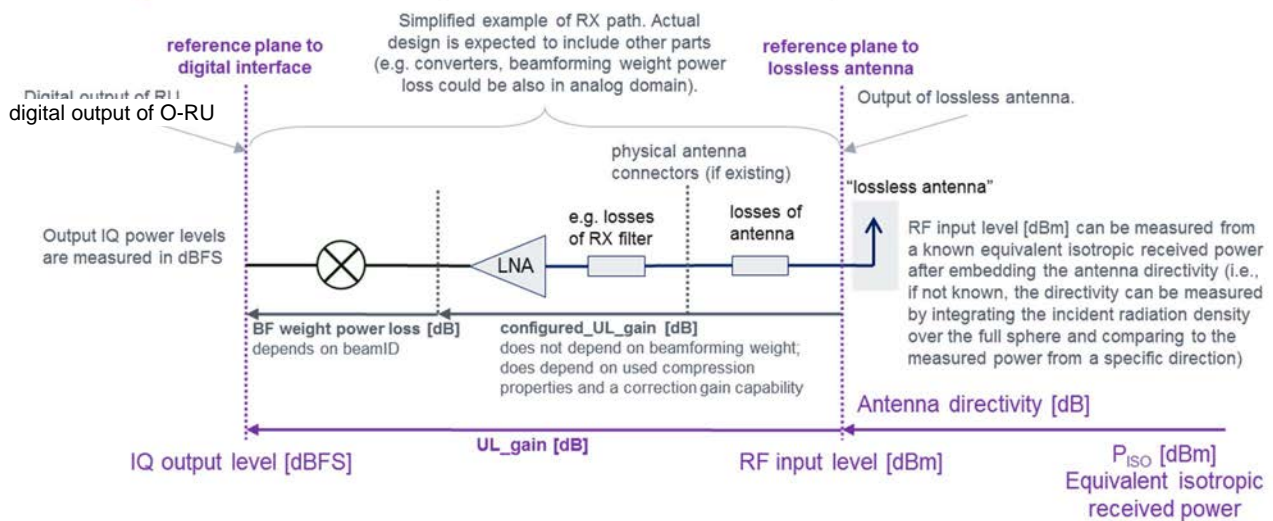


Figure 8.1.3.2.1-3: Details of UL gain

8.1.3.2.2 DL gain definition

The DL gain of a TX array carrier element shall be provided as part of the O-RU carrier set-up procedure. The configured (by the O-DU over the M-Plane) DL gain of a TX array carrier element (i.e., "TX array carrier element" refers to the TX array element serving the respective carrier configured on the respective TX array) shall be defined by mapping the Reference_Level dBFS (see clause 8.1.3.3) carried over any subcarrier (of the carrier) on the digital interface to the maximum RMS power the TX array carrier element may output at the assumed-lossless antenna connector port (i.e., including all insertion losses). In single-carrier operation, the configured DL gain shall apply to the carrier. In multiple-carrier operation, the configured DL gain should be lower accommodating the number of multiple carriers and their bandwidths. The O-DU shall scale down the level of the signals at the interface and/or set tx-array-carrier gain values to avoid saturation in the O-RU for single or multiple-carrier configurations.

$$\text{Configured DL_gain [in dB]} \leq \text{maximum TX power per array element [in dBm]} - \text{Reference_Level [in dBFS]},$$

which is valid for each individual spatial stream served under the TX array carrier element.

The following statements shall apply to the above equation:

- 1) The maximum configured DL gain of a TX array carrier is equivalent to mapping 0 dBFS carried over any subcarrier (of the carrier) on the digital interface to the maximum RMS power the TX array carrier element is supposed to output $+ 10 \times \log_{10}$ (the number of array elements) – Reference_Level [in dBFS].
- 2) For a category A O-RU, the tx-array output rms power is measured over 1 polarization
- 3) For a category B O-RU, the tx-array output rms power shall be measured over all polarizations addressed by the respective precoding operations
- 4) Configured DL_gain [in dB] = target TX carrier power per array element [in dBm] - Reference_Level [in dB] where target TX carrier power per array element \leq maximum TX power per array element.
- 5) If the O-RU reports a minimum DL gain capability, then the O-DU shall also ensure that Configured DL_gain [in dB] \geq minimum DL gain per array element [in dB] – Reference_Level [in dBFS] where the minimum DL gain per array element is the capability reported by O-RU.

8.1.3.2.3 UL gain definition

The UL gain or scaling of an rx-array carrier element (i.e., "rx-array carrier element" refers to the rx-array element serving the respective eAxC configured on the respective rx-array) shall be defined by mapping -152 dBm at the assumed-lossless antenna port to the smallest power level an IQ sample can carry over the digital interface (i.e., $\text{average}(I^2+Q^2) = 1$) while considering the configured gain_correction value and the IQ compression properties to be used; configured gain_correction is sum of RX carrier specific gain_correction and eAxC specific gain_correction. In addition, in order to avoid saturation over the interface when beamforming is used over the rx-array carrier, the largest power level that can be received at the assumed-lossless antenna port by the rx-array carrier element shall be equivalent to 0 dBFS - $10 \cdot \log_{10}(\text{number of array elements})$. The configured gain-correction shall allow the adjustment of the level of the smallest & largest receivable power.

The following principles shall apply:

- 1) The UL gain of an RX eAxC shall be equal to the UL gain of the rx-array carrier element, and contributions to IQ in the eAxC from all rx-array elements shall have the same gain
- 2) The minimum and maximum power levels that can be received by the RX eAxC shall be scaled by $+10 \cdot \log_{10}(\text{number of array elements})$ from the respective values of the rx-array carrier element (i.e. when signals are received with equal levels from all the elements of the rx array meaning no tapering used and all rx elements have same design properties). In case an rx-array has only one rx-array element, then the scale factor shall be unity. The reason behind this principle is to allow applying tapering (beam ID dependent) to an rx-array without impacting the quality of the IQ data being transferred over the interface (i.e., in a worst case scenario, the signal is received from only one rx-array element from the rx-array).
- 3) Assumed-lossless antenna port means antenna losses shall be part of the unit under test.
- 4) The -152 dBm shall be considered as the smallest level that can be measured by a narrow subcarrier of 1.25 kHz for a system with 3 dB of equivalent noise figure and without being impacted by the interface noise (i.e., 20 dB margin considered).
- 5) $\text{Configured_UL_gain [in dB]} = \text{Interface resolution [dBFS]} - (-152 \text{ dBm}) + \text{gain_correction [dB]}$; valid for each received individual spatial stream (i.e. configured_UL_gain is configured by O-RU and not by the O-DU).

The UL gain shall depend on the digital interface resolution (in dBFS) representing the smallest level that can be used. The interface resolution depends on the compression scheme which can be static or dynamic. For this reason, when multiple compression methods or IQ bitwidths are used for data streams received from an rx-array carrier element, the configured gain shall accommodate all the intended compression methods and IQ bitwidths.

For example, when using block floating point compression, the interface resolution shall be defined as

$$\text{Interface resolution [dBFS]} = -20 \times \log_{10} \left(2^{\text{Mantissa_bits} - 1} \times 2^{2^{\text{Exponent_bits} - 1}} \right)$$

where:

Mantissa_bits is the number of mantissa bits in the defined block floating point notation.

Exponent_bits is the number of exponent bits in the defined block floating point notation.

It is assumed that the gain_correction is 0 dB by default, unless a different value is needed due to special circumstances (details beyond the scope of the present document).

The gain_correction is configured as one value per rx array carrier and does not change once an array carrier is activated. The O-RU applies a configured value of gain_correction regardless of used compression.

EXAMPLE: With rx-array formed by 10 array elements and gain_correction of 0 dB:

9-bit mantissa 2's complement + 4-bit exponent → Interface resolution = -138.5 dBFS

Default UL gain [dB] = -138.5 dBFS + 152 dBm + 0 dB = 13.5 dB

The UL gain and power scaling are summarized in Table 8.1.3.2.3-1.

Table 8.1.3.2.3-1: Example of UL gain and power scaling for an rx array with 10 elements and for block floating point compression based on 9-bit mantissa and 4-bit exponent

Interface resolution -138.5 dBFS	rx-array carrier element	Rx-array carrier with 10 elements
UL Gain	13.5dB	13.5 dB
Lowest received power level without degradation due to interface (excluding losses due to beamforming weights)	-152 dBm → -138.5 dBFS	-142 dBm → -128.5 dBFS
Largest possible received power level	-23.5 dBm → -10 dBFS	-13.5 dBm → 0 dBFS

8.1.3.3 TX power budget for Category A and Category B O-RUs

This clause describes the handling of the power budget in a category A and category B O-RU. In general, care shall be taken to avoid exceeding the maximum rms power rating of a tx-array element.

In this clause, $m_{a,k}$ is the maximum rms power rating (in W) of tx-array element k of a tx-array a with K elements. For simplicity it is assumed all K elements of the array have the same maximum rms power rating (i.e., for every array a and every k and k' , $m_{a,k} = m_{a,k'}$). **For an O-RU with tx-array a , the maximum rms power rating of tx-array element shall be reported as read-only parameter. This will be a common value for all array elements of the tx-array a .**

NOTE 1: The maximum rms power rating can be reported as form of gain when mapped to 0dBFS.

The maximum rms power rating of array can be derived from $m_{a,k}$ by scaling linearly with the number of elements K (i.e., $m_a = K \cdot m_{a,k}$)

In addition, let $g_{c,a,k}$ be the gain (in dB) configured for tx-array element k of tx-array a for array carrier c . It is considered that the configured gain is same to every element (i.e., for every tx-array carrier c , every tx-array a and every element k and k' , $g_{c,a,k} = g_{c,a,k'}$). **Hence, the O-DU should configure the O-RU by providing a value for $g_{c,a,k}$ in parameter 'gain' of tx-array-carrier.**

The total gain of array carrier c served on array a can be derived from the configured gain $g_{c,a,k}$ as:

$$g_{c,a} [\text{dB}] = g_{c,a,k} [\text{dB}] + 10 \times \log_{10}(K).$$

where:

$g_{c,a}$ is the total gain of array carrier c served on array a

$g_{c,a,k}$ is the configured gain of array element k of array a for carrier c

Hence, the power $p_{c,a,k}$ of array carrier c served on tx-array element k of array a can be derived as

$$p_{c,a,k} [\text{dBm}] = g_{c,a,k} [\text{dB}] + 0 \text{ dBFS}.$$

As a result, the total power of $p_{c,a}$ of array carrier c served on tx-array a is derived as

$$p_{c,a} [\text{dB}] = p_{c,a,k} [\text{dB}] + 10 \times \log_{10}(K),$$

where:

K is the total number of array elements in array a

Finally, the total power $p_{a,k}$ used on tx-array element k of tx-array a serving all configured carriers is derived as

$$p_{a,k} = 10 \times \log_{10} \sum_c 10^{\frac{p_{c,a,k}}{10}}.$$

In scenarios where a tx-array element k is shared between multiple tx-arrays, the O-DU shall ensure that the configured gains $g_{c,a,k}$ are constrained by:

$$\sum_a \sum_c 10^{\frac{g_{c,a,k} + \text{Reference_Level}_c}{10}} \leq 1000 \times m_k$$

where:

$$m_k = \min_a (m_{a,k})$$

$Reference_Level_c$ (in dB) is the array-carrier specific IQ normalization level optionally configured via M-Plane per array carrier (the value 0 is used if this parameter is not supported by the O-RU or not configured by the O-DU).

The index a spans over every tx-array a that shares array element k and has array carrier configured. The summing over c includes every array carrier c that is configured for tx-array a .

NOTE 2: In general $m_{a,k}$ can be different for different values of tx-array a . However, tx-arrays not used in a given configuration do not contribute to the above constraint.

If O-RU reports via M-Plane the minimum gain of the tx-array element then the O-DU shall ensure that the configured gains $g_{c,a,k}$ are constrained by:

$$Min_Gain_{a,k} \leq g_{c,a,k} + Reference_Level_c$$

where $Min_Gain_{a,k}$ (in dB) is the tx-array element specific minimum gain optionally reported by the O-RU via the M-Plane parameter min-gain.

For every configured array carrier c over a tx-array a , the O-DU shall ensure that the input power levels of all resource elements used over all eAxCs x are constrained by:

$$\sum_x \sum_n 10^{\frac{RE_{n,x,c,a}^{dBFS}}{10}} \leq 10^{\frac{Reference_Level_c}{10}}$$

and

$$RE_{n,x,c,a}^{dBFS} \leq 0$$

where:

$RE_{n,x,c,a}^{dBFS}$ = input power level in dBFS of an RE n of an eAxC x of array carrier c configured over tx-array a .

$$RE_{n,x,c,a}^{dBFS} = 10 \cdot \log_{10}(I_{n,x,c,a}^2 + Q_{n,x,c,a}^2) - 10 \cdot \log_{10}(FS_0 \cdot 2^{-FS_Offset})$$

where:

$I_{n,x,c,a}$ and $Q_{n,x,c,a}$ are decompressed IQ sample value components, FS_0 and FS_Offset are defined in clause 8.1.3.1.

The summing over x includes every eAxC x of array carrier c that is used simultaneously in DL on the tx-array a . The summing over n includes every RE of the eAxC x that is used simultaneously in DL.

When beamforming is used in O-RU, an additional constraint applies to all the beamforming weights to be used for beamforming in order to ensure that the tx power per tx-array element after the beamforming has well-defined upper bound and does not exceed the limit of the maximum rms power rating of the tx-array element. To simplify the notation and cover all beamforming types, i.e. frequency-domain beamforming and time-domain beamforming, any individual beamforming weight (a complex number multiplier used by O-RU in beamforming operation) is denoted as w .

For any beamforming weight w to be used by O-RU to perform beamforming operations, the entity controlling the generation of the weight (i.e., O-DU or O-RU) shall ensure that

$$|w|^2 \leq 1.$$

This constraint is applicable to every beamforming weight, used in frequency or time domain. The constraint is applicable to every explicit or implied beamforming weight applied by O-RU regardless of the source of the beamforming weight (e.g. received from O-DU, predefined by O-RU, generated by O-RU from beam attributes or from channel information, etc.).

8.2 Elementary procedures

8.2.1 IQ data transfer procedure

This procedure (an alternative procedure IQ data transfer is described in clause 8.2.2) shall be used to transfer frequency domain IQ data samples between the O-DU and O-RU. Data shall be transmitted symbol by symbol as U-Plane messages. U-Plane data (UL and DL) shall always be preceded by control information (C-Plane) except for U-Plane data covered under M-Plane static configuration e.g., see clauses 8.3.5.3 and 8.3.5.4 pertaining to PRACH and SRS. U-Plane data shall be sent in a different data flow from the C-Plane with its own header and encapsulated payload. C-

Plane messages pertaining to U-Plane data may be sent for every slot (or for LTE, TTI). An overview of the logic for transmission of both the data-associated control information and IQ data is depicted in Figure 8.2.1-1.

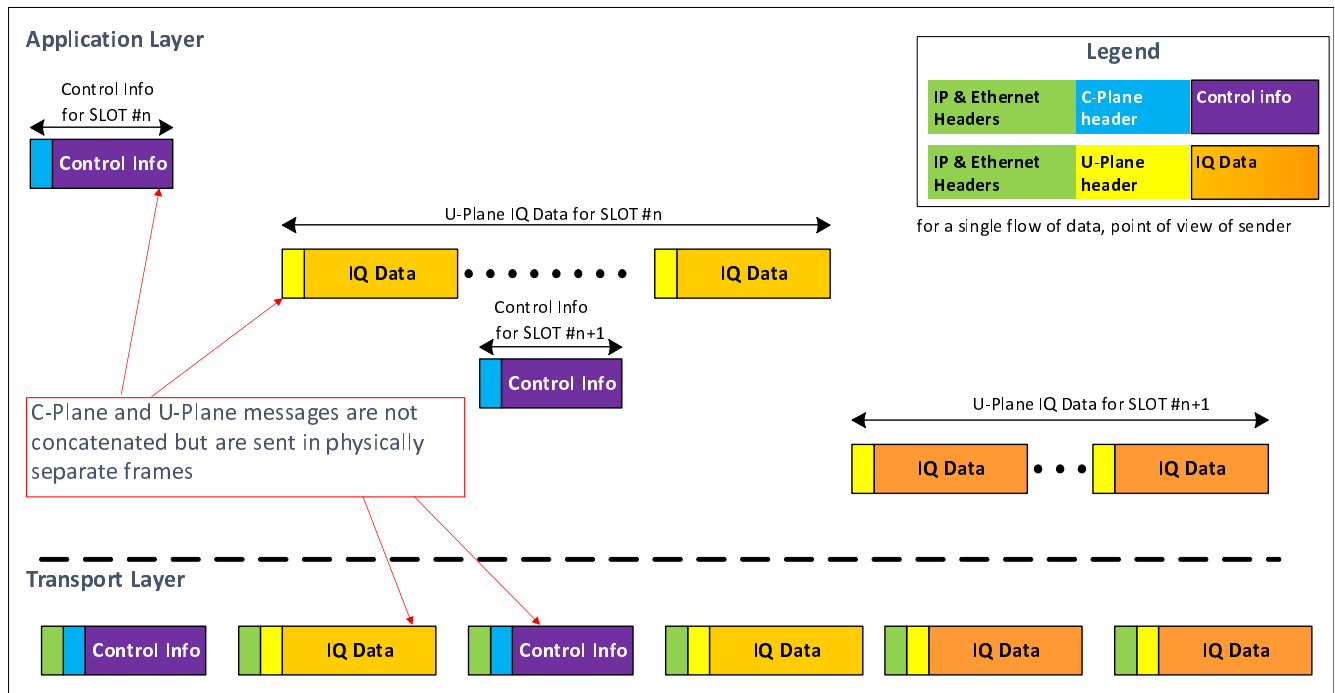


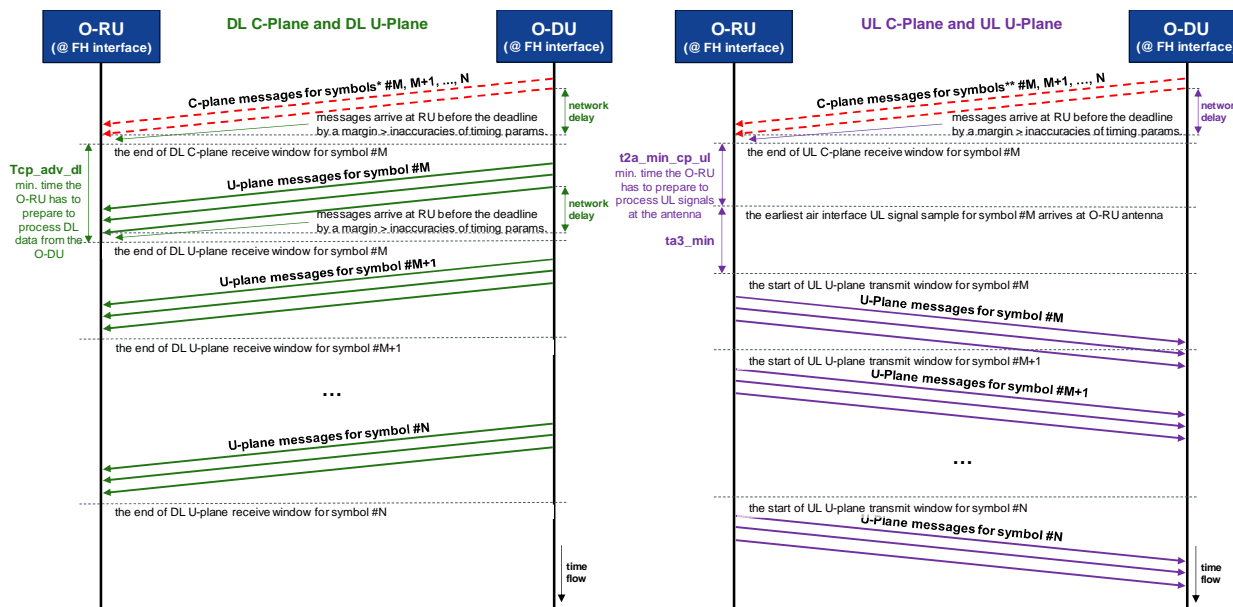
Figure 8.2.1-1: DL IQ data transfer overview

The IQ transfer procedure shall comprise of two types of messages:

- Message containing scheduling and beamforming commands information (i.e., data-associated control).
- Message containing frequency domain IQ samples.

The data-associated control information shall be bundled in messages within a single data flow (albeit possibly in multiple C-Plane messages). However, control information for UL and DL shall be sent separately (See Figure 8.2.1-2). As shown in Figure 8.2.1-1, the C-Plane messages for a given slot and given symbol(s) exchanged between the O-DU and O-RU are followed by the U-Plane data for that slot and symbol(s), one symbol at a time. U-Plane messages shall be sent by the O-DU and O-RU in order of the symbol for which they carry IQ data. C-Plane and DL U-Plane messages shall be sent by the O-DU in advance (see clause 4.5 for details) such that they arrive at O-RU within a time window that is early enough to leave the O-RU time to process them. Figure 8.2.1-2 shows the case where O-DU sends DL C-Plane message describing symbols #M, M+1, ..., N of slot S between "T1a_max_cp_dl" and "T1a_min_cp_dl" before the start of DL symbol #M (the earliest symbol described by the message). As shown in Figure 8.2.1-2, the end of receive time window for DL C-Plane messages describing symbol #M, M+1, ..., N of slot S (where M is the earliest symbol described by each of messages) is by "Tcp_adv_dl" earlier than the end of receive time window for DL U-Plane messages carrying IQ data for symbol #M. The O-DU shall send and the O-RU shall receive DL U-Plane message as per the transmission window and reception window specified in clause 4.5. Similarly, the O-DU shall send UL C-Plane messages describing symbols #M, M+1, ..., N of slot S between "T1a_max_cp_ul" and "T1a_min_cp_ul" before the start of the UL symbol #M (the earliest symbol described by the message). The end of receive time window for UL C-Plane messages describing symbol #M, M+1, ..., N (where M is the earliest symbol described by each of messages) shall be "T2a_min_cp_ul" earlier than the start of UL symbol #M. The O-RU shall send and the O-DU shall receive UL U-Plane message as per the transmission window and reception window specified in clause 4.5. For non-PRACH channels the O-RU shall send UL U-Plane message carrying IQ data for symbol #M between "Ta3_min" and "Ta3_max" after the reception of the earliest air interface sample of UL symbol #M (this corresponds to start of cyclic prefix). For PRACH channels the O-RU shall send UL U-Plane messages carrying IQ data for a PRACH Symbol between "Ta3_min" and "Ta3_max" after the reception of the earliest air interface sample that is necessary to generate the IQ data in the message (this corresponds to the start of the FFT sampling window).

Also shown in Figure 8.2.1-2 is the fact that there is a certain period of time between the O-RU's receipt of C-Plane messages for a symbol and the need for the O-RU to process U-Plane data for that symbol. In particular, in the DL there shall be a period of time, " Tcp_adv_dl " which provides for some number of microseconds for the O-RU to e.g. update beamforming weights prior to processing the DL data arriving from the O-DU. In the UL there shall be a period of time, " $t2a_min_cp_ul$ " between the O-RU's receiving the C-Plane messages governing the processing of UL data and the receipt of UL signals at the O-RU's antennas. These time intervals, when combined with network delays and other processing latencies, result in the RAN's HARQ loop to be closed allowing feedback in the air interface processing.



NOTE 1: A DL C-Plane message describing multiple symbols shall arrive at O-RU at least Tcp_adv_dl before the end of DL U-Plane receive window for symbol $startSymbolId$ (the earliest symbol described by the message).

NOTE 2: A UL C-Plane message describing multiple symbols shall arrive at O-RU at least $t2a_min_cp_ul$ before the earliest air interface UL signal sample of symbol $startSymbolId$ (the earliest symbol described by the message) arrives at O-RU antenna.

Figure 8.2.1-2: C-Plane and U-Plane message transfer procedure (DL & UL shown)

Regardless of the level of compression, user plane data may exceed the maximum packet size (e.g., MTU of 1500 Bytes for regular not Jumbo frames) even in case of one symbol of granularity. Thus, IQ data for each symbol may require packetization over multiple packets.

User plane messages are sent as resource blocks ("PRBs") and the data for each PRB shall start on a byte boundary. If due to the IQ bit-width being used the natural end of the data in the PRB does not fall on a byte boundary, then zero bits shall be appended until a byte boundary is reached.

PRACH data shall be handled similarly, such that the PRACH REs are packaged into 12-RE blocks analogous with data PRBs. Each 12-RE block shall be padded with zero bits to end on a byte boundary. In cases when there is not an even factor of 12 REs in the PRACH data (e.g. 839 PRACH REs), added zero-value REs shall be added to fill out the block to 12 REs, and then if needed zero-padding will then be appended to reach a byte boundary.

8.2.2 IQ data transfer procedure without C-Plane

This is an optional procedure used to transfer frequency domain IQ data samples between the O-DU and O-RU without C-Plane control. With respect to IQ data transfer, this procedure shall follow the rules described in clause 8.2.1 with an exception: data-associated control information for selected eAxCs shall not be transferred in C-Plane messages but configured via M-Plane. This method is applicable to channels with predictable scheduling (like PRACH or SRS) which do not require changes of data-associated control information in real-time.

In the present document this procedure is applicable to eAxCs assigned to receive a subset of uplink signals and channels – specifically PRACH and SRS (see clause 8.3.5.3 and clause 8.3.5.4):

- 1) O-DU shall provide via M-Plane data-associated control information (this consists of scheduling and beamforming configuration) to O-RU and shall activate the eAxC. C-Plane messages shall not be sent for this eAxC.
- 2) O-RU will periodically receive uplink radio signal following the configured schedule and shall send U-Plane messages containing frequency domain IQ data (as in clause 8.2.1). Values of fields in U-Plane message headers and data section headers shall be derived from parameters configured via M-Plane as per rules described in M-Plane specification (see [7]).

For more details on parameters that convey scheduling and beamforming configuration and rules for mapping them to U-Plane messages refer to M-Plane specification (see [7]).

In addition, IQ data transfer without C-plane may be applied to DL U-plane data. This is applicable to deployment use cases where the O-RU is a category A O-RU without beamforming.

When the feature is supported by the O-RU as indicated via the M-plane feature **UPLANE-ONLY-DL-MODE** and enabled via the M-plane parameter **uplane-only-dl-mode-enable** the O-DU may omit sending DL C-plane messages to the O-RU. For minimal impact on the O-DU DL processing the O-DU may continue to transmit DL C-plane messages despite enabling the U-plane only DL mode on the O-RU.

The following aspects shall apply for this mode of operation:

- The O-RU is a category A O-RU and supports U-plane only DL mode operation.
- For any received DL U-plane messages the O-RU shall assume all bits in reMask are set for the received PRBs.
- The O-RU shall ignore any received DL C-plane messages.
- The mode, when enabled, shall apply to all TX endpoints of the O-RU.
- The usage of real-time endpoint configuration (scs, cpLength, etc.) via DL section type 3 is not supported.
- Beamforming is not supported.

NOTE: U-plane only DL mode is intended to entirely disable processing of DL C-plane messages to provide the opportunity for processing and power savings on the O-RU and O-DU. However, the O-DU may be actively interfacing to multiple O-RUs and as such may not be able to completely disable generating and transmitting DL C-plane messages.

8.3 Elements for the U-Plane protocol

8.3.1 General

U-Plane messages shall be encapsulated using a two-layered header approach. The first layer shall consist of an eCPRI or IEEE 1914.3 common header, including fields used to indicate the message type, while the second layer shall be an application layer including necessary fields for control and synchronization. When the optional "little endian byte order" is configured via M-Plane, the iSample/qSample fields shall use little endian byte order to transmit the complex numbers. Annex D.2 shows little endian byte order format for various I/Q data length.

8.3.2 DL/UL data

A common frame format shall be used for U-Plane messages consisting of a transport layer and an application layer. The application layer shall be within the transport payload and shall consist of a common header for time reference, followed by information and parameters dependent and specific to the Section Type in use. Data from multiple sections of the same value may be lined up one after another within the payload. To minimize packet rate over the interface, transmitter should fill messages with as many subsequent sections as possible. Data from sections of different Section Types may be mixed within a single U-Plane message for a given eAxC when the Section Types do not have differences in format and values of user plane application header or format of user plane section header. Whenever necessary, padding bits shall be added at the end of a section after (possibly compressed) I and Q sample data to achieve

1-Byte alignment. Within a resource block, unless an sReSMask is included in the udCompParam field, 12 REs shall always be sent across the interface. If some REs are meant to be blank some data (perhaps zero value) shall be sent anyway because the U-Plane data parser will expect exactly 12 complex RE values per resource block.

The following bullet list of information elements ("fields") shall be used within U-Plane messages as applicable, use of those fields shall observe the byte formatting as described in table 8.3.2-1:

- Transport Layer – see clause 5.1.3.
- Application Layer.
- Section Type 0 fields (used for indicating idle or guard periods from O-DU to O-RU):
 - No U-Plane messages are associated with Section Type 0.
- Section Type 1 fields (used for most Downlink and Uplink physical radio channels):
 - Common header fields
 - **dataDirection** (data direction (gNB Tx/Rx)) field: 1 bit.
 - **payloadVersion** (payload version) field: 3 bits:
 - Value = "1" shall be set (1st protocol version for payload and time reference format).
 - **filterIndex** (filter index) field: 4 bits.
 - **frameId** (frame identifier) field: 8 bits.
 - **subframeId** (subframe identifier) field: 4 bits.
 - **slotID** (slot identifier) field: 6 bits.
 - **symbolId** (symbol identifier) field: 6 bits.
 - Section header fields:
 - **sectionID** (section identifier) field: 12 bits.
 - **rb** (resource block indicator) field: 1 bit.
 - **symInc** (symbol number increment command) field: 1 bit.
 - **startPrbu** (starting PRB of user plane section) field: 10 bits.
 - **numPrbu** (number of contiguous PRBs per data section) field: 8 bits.
 - **udCompHdr** (user data compression header) field, not always present: 8 bits.
 - **reserved** (reserved for future use) field, only present with udCompHdr: 1 byte.
 - **udCompLen** (PRB field length) field, not always present: 16 bits.
 - PRB fields:
 - **udCompParam** (user data compression parameter) field: 0, 8 or 16 bits.
 - **iSample** (in-phase sample) field: 1-16 bits.
 - **qSample** (quadrature sample) field: 1-16 bits.
- Section Type 3 fields (used for PRACH and mixed-numerology channels (see Note)):
 - Timing header, section header and PRB fields same as for Section Type 1.
- Section Type 5 fields (used for UE scheduling information):
 - Timing header, section header and PRB fields same as for Section Type 1.
- Section Type 6 fields (used for sending channel information for a specific UE ID):
 - No U-Plane messages are associated with Section Type 6.

NOTE: Regarding Section Type 3, it is also true that Section Type 1 messages can also be used for PRACH channel data. The UL U-Plane processing in this case is same as a non-PRACH channel in O-RU.

U-Plane messages shall conform to the information element format described in table 8.3.2-1.

Table 8.3.2-1: IQ data frame format

Section Type 1,3: DL/UL IQ data msgs								# of bytes	
0 (msb)	1	2	3	4	5	6	7 (lsb)		
transport header, see clause 5.1.3								8	Octet 1
dataDirection	payloadVersion			filterIndex				1	Octet 9
frameId								1	Octet 10
subframeId				slotId				1	Octet 11
slotId		symbolId						1	Octet 12
sectionId								1	Octet 13
sectionId				rb	symlnc	startPrbu		1	Octet 14
startPrbu								1	Octet 15
numPrbu								1	Octet 16
udCompHdr (not always present)								0/1	Octet 17
reserved (not always present)								0/1	Octet 18
udCompLen (not always present)								0/2	Octet 17/19
udCompParam (not always present)								0/1/2	Octet 17/19/21
iSample (1 st RE in the PRB)								1	K= 17/19/20/21/23
qSample (1 st RE in the PRB)								1	K+1
...									
iSample (12 th RE in the PRB)								1	K+22
qSample (12 st RE in the PRB)								1	K+23
udCompParam (not always present)								0/1/2	K+24
iSample (1 st RE in the PRB)								1	K+24/25/26
qSample (1 st RE in the PRB)								1	K+25/26/27
...									
iSample (12 th RE in the PRB)								1	K+46/47/48
qSample (12 st RE in the PRB)								1	K+47/48/49
...									
sectionId								1	Octet M
sectionId				rb	symlnc	startPrbu		1	M+1
startPrbu								1	M+2
numPrbu								1	M+3
udCompHdr (not always present)								0/1	M+4
reserved (not always present)								0/1	M+5
udCompLen (not always present)								0/2	M+4/6
udCompParam (not always present)								0/1/2	M+4/6/8
iSample (1 st RE in the PRB)								1	K=M+4/6/7/8/10
qSample (1 st RE in the PRB)								1	K+1
...									
iSample (12 th RE in the PRB)								1	K+22
qSample (12 st RE in the PRB)								1	K+23
udCompParam (not always present)								0/1/2	K+24
iSample (1 st RE in the PRB)								1	K+24/25/26
qSample (1 st RE in the PRB)								1	K+25/26/27
...									
iSample (12 th RE in the PRB)								1	K+46/47/48
qSample (12 st RE in the PRB)								1	K+47/48/49

NOTE 1: Shading: yellow is transport header, pink is radio application header, others are repeated sections.
NOTE 2: Octet count given when the iqWidth = 8 but other iqWidth values are possible; the octet count assumes all REs are present which may not be the case when field sReSMask is present in udCompParam.

8.3.3 UL/DL data coding of information elements

8.3.3.0 General

See clause 5.1.3 for transport header information element details.

8.3.3.1 dataDirection (data direction (gNB Tx/Rx))

See clause 7.5.2.1.

8.3.3.2 payloadVersion (payload version)

See clause 7.5.2.2.

8.3.3.3 filterIndex (filter index)

See clause 7.5.2.3.

8.3.3.4 frameId (frame identifier)

See clause 7.5.2.4.

frameId in U-Plane message shall be set to the frameId value signaled in the corresponding C-Plane message.

8.3.3.5 subframeId (subframe identifier)

See clause 7.5.2.5.

subframeId in U-Plane message shall be set to the subframeId value signaled in the corresponding C-Plane message.

8.3.3.6 slotId (slot identifier)

See clause 7.5.2.6.

slotId in U-Plane message shall be set to the slotId value signaled in the corresponding C-Plane message.

8.3.3.7 symbolId (symbol identifier)

Description: This parameter identifies a symbol number within a slot. When a C-Plane message describes a single symbol (or in case of PRACH, a single PRACH Symbol), symbolId in the U-Plane message for that symbol (or PRACH Symbol) shall be set to the startSymbolId value signaled in the C-Plane message. When a C-Plane message describes multiple symbols (or in case of PRACH/NPRACH, multiple PRACH Symbols), symbolId in the U-Plane message for the first symbol (or PRACH Symbol) shall be set to the startSymbolId value signaled in the C-Plane message. symbolId in the U-Plane message for the subsequent symbols (or PRACH Symbols) are incremented thereon.

Value range: {00 0000b-11 1111b}

Type: unsigned integer.

Field length: 6 bits.

8.3.3.8 sectionId (section identifier)

See clause 7.5.3.1.

8.3.3.9 rb (resource block indicator)

See clause 7.5.3.2.

8.3.3.10 symInc (symbol number increment command)

See clause 7.5.3.3.

8.3.3.11 startPrbu (starting PRB of user plane section)

Description: This parameter is the starting PRB (lowest frequency) of a user plane data section. Values of rb, startPrbu and numPrbu shall ensure that data sections shall never overlap: a single PRB (a block of 12 resource elements consecutive in frequency) shall only exist within one data section for a given eAxC. For one section description in a C-Plane message, there may be multiple U-Plane data sections associated with it. The freqOffset parameter conveyed in the corresponding section description affects the frequency span for specific range of PRB numbers (formulas in clause 7.5.3.6 apply with startPrbu replacing startPrbc). Therefore "shall never overlap" shall consider the value of freqOffset.

U-Plane data sections shall include only PRBs addressed by the corresponding C-Plane section descriptions, but the O-RU shall be prepared for possibility of lost C-Plane or U-Plane messages. Depending on the method of coupling of C-Plane and U-Plane more specific rules apply.

If the coupling of C-Plane and U-Plane via sectionId value is used then values of rb, startPrbu and numPrbu shall ensure only PRBs addressed by section descriptions with the same sectionId as the user plane data section are present in the user plane data section. If the coupling of C-Plane and U-Plane via frequency and time (with or without priorities) is used then values of rb, startPrbu and numPrbu shall ensure only PRBs addressed by section descriptions are present in the user plane data section.

Due to fragmentation (an application level fragmentation to meet restriction of message size or a fragmentation of section in frequency, e.g. as a result of presence of Section Extension 6 or 12 in a corresponding section description) there may be multiple user plane data sections for a C-Plane section description. Presence of rb = 1 (see clause 7.5.3.2) in a C-Plane section description shall require the presence of rb = 1 in the associated user plane data section and does not require multiple user plane data sections.

Value range: {00 0000 0000b-11 1111 1111b}.

Type: unsigned integer.

Field length: 10 bits.

8.3.3.12 numPrbu (number of PRBs per user plane section)

Description: This parameter defines the number of PRBs (blocks of 12 resource elements consecutive in frequency) in the user plane data section. If the parameter rb (see clause 8.3.3.9) is zero, then the PRBs in the user plane data section shall be consecutive in frequency. Otherwise the set of PRBs includes only every other PRB in frequency. Values of rb, startPrbu and numPrbu shall ensure that data sections shall never overlap: a single PRB may only exist within one data section for a given eAxC. The freqOffset parameter conveyed in the corresponding section description affects the frequency span for specific range of PRB numbers (formulas in clause 7.5.3.6 apply with startPrbu replacing startPrbc). Therefore "shall never overlap" shall consider the value of freqOffset.

U-Plane data section shall include only PRBs addressed by the corresponding C-Plane section descriptions, but the O-RU shall be prepared for possibility of lost C-Plane or U-Plane messages. Depending on the method of coupling of C-Plane and U-Plane more specific rules apply.

If the coupling of C-Plane and U-Plane via sectionId value is used then values of rb, startPrbu and numPrbu shall ensure only PRBs addressed by C-Plane section descriptions with the same sectionId as the user plane data section are present in the user plane data section. If the coupling of C-Plane and U-Plane via frequency and time (with or without priorities) is used then values of rb, startPrbu and numPrbu shall ensure only PRBs addressed by C-Plane section descriptions are present in the user plane data section.

Due to fragmentation (an application level fragmentation to meet restriction of message size or a fragmentation of section in frequency, e.g. as a result of presence of Section Extension 6 or 12 in a corresponding section description) there may be multiple user plane data sections for a C-Plane section description. Presence of rb = 1 (see clause 7.5.3.2) in a C-Plane section description shall require the presence of rb = 1 in the associated user plane data section and does not require multiple user plane data sections.

Value range: {0000 0001b-1111 1111b, 0000 0000b = all PRBs in the specified SCS and carrier bandwidth }.

Value 0000 0000b shall be reserved for NR cases wherein the total number of PRBs may be more than 255; for other cases a non-zero value of numPrbc shall be used. In this case, the sending node shall set the startPrbu value to zero and the receiving node shall ignore whatever startPrbu value is received and assume it is zero.

Type: unsigned integer.

Field length: 8 bits.

8.3.3.13 udCompHdr (user data compression header)

Description: This parameter defines the compression method and IQ bit width for the user data in a data section. This means that, in the DL at least, each data section may have a different udCompHdr value. In the UL, the O-RU shall copy the received udCompHdr value in the C-Plane message to the udCompHdr field in the UL U-Plane message. This C-Plane instruction provides a single udCompHdr value for all data sections defined in the C-Plane message. This field shall be absent from U-Plane messages when the static IQ format and compression method is configured via the M-Plane. In this way a single compression method and IQ bit width is provided (per UL and DL, per LTE and NR) without adding more overhead to U-Plane messages. Table 8.3.3.13-1 provides the udIqWidth and udCompMeth bit allocation, and Table 8.3.3.13-2 and 8.3.3.13-3 provide udIqWidth and udCompMeth definitions respectively, all of which shall be followed in the construction of U-Plane messages.

Value range: {0000 0000b-1111 1111b}.

Table 8.3.3.13-1: Bit allocations

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
udIqWidth				udCompMeth				1	Octet 1

Table 8.3.3.13-2: udIqWidth definition

udIqWidth	Bit width of each I and each Q
0000-1111b	16 for udIqWidth=0, otherwise equals udIqWidth e.g. udIqWidth = 0000b means I and Q are each 16 bits wide; e.g. udIqWidth = 0001b means I and Q are each 1 bit wide; e.g. udIqWidth = 1111b means I and Q are each 15 bits wide

Table 8.3.3.13-3: udCompMeth definition

udCompMeth	compression method	udIqWidth meaning
0000b	no compression	bitwidth of each uncompressed I and Q value
0001b	block floating point (BFP)	bitwidth of each I and Q mantissa value
0010b	block scaling	bitwidth of each I and Q scaled value
0011b	μ -law	bitwidth of each compressed I and Q value
0100b	modulation compression	bitwidth of each compressed I and Q value
0101b	BFP + selective RE sending	bitwidth of each compressed I and Q value
0110b	mod-compr + selective RE sending	bitwidth of each compressed I and Q value
0111b – 1111b	reserved for future methods	depends on the specific compression method

Type: unsigned integer (concatenated bit fields).

Field length: 8 bits.

8.3.3.14 reserved (reserved for future use)

Description: This parameter provides 1 byte for future definition, should be set to all zeros by the sender and ignored by the receiver. This field is only present when udCompHdr is present, and is absent when the static IQ format and compression method is configured via the M-Plane.

Value range: {0000 0000b-1111 1111b}, but shall be set to all zeros.

Type: unsigned integer.

Field length: 8 bits.

8.3.3.15 udCompParam (user data compression parameter)

Description: This parameter applies to whatever compression method is used for the PRB (configured statically via M-Plane or specified in udCompHdr of the data section containing the PRB). U-Plane messages shall conform to the information element formatting as described in Table 8.3.3.15-1.

Value range: {0000 0000b-1111 1111b}.

Table 8.3.3.15-1: Bit allocations

udCompMeth	0 (msb)	1	2	3	4	5	6	7 (lsb)	compParam size
0000b = no compression	absent								0 octets
0001b = block floating point (BFP)	reserved (set to all zeros)				exponent (unsigned)				1 octet
0010b = block scaling	sblockScaler (unsigned, 1 integer bit, 7 fractional bits)								1 octet
0011b = μ -law	compBitWidth				compShift				1 octet
0100b = modulation compr.	absent								0 octets
0101 = BFP + selective RE sending	sReSMask[11:8]				exponent (unsigned)				1 octet
0110 = mod-compr + selective RE sending	sReSMask[11:8]				reserved				1 octet
0111b – 1111b	reserved (set to all zeros)								? octets

Type: variable.

Field length: zero for udCompMeth values 0000b and 0100b, 8 bits for udCompMeth values 0001b, 0010b and 0011b, 16 bits for udCompMeth values 0101b and 0110b; other udCompMeth values may imply other lengths but will always be an integer number of bytes.

8.3.3.16 iSample (in-phase sample)

Description: This parameter is the In-phase sample value. When the optional "little endian byte order" is configured via M-Plane, refer to Annex D.2 for detail byte order, otherwise see Annex D.1 for example sample formatting.

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 1-16 bits.

8.3.3.17 qSample (quadrature sample)

Description: This parameter is the Quadrature sample value. When the optional "little endian byte order" is configured via M-Plane, refer to Annex D.2 for detail byte order, otherwise see Annex D.1 for example sample formatting.

Value range: {all zeros – all ones}.

Type: signed integer.

Field length: 1-16 bits.

8.3.3.18 sReSMask (selective RE sending mask)

Description: This parameter defines the IQ usage mask when using the compression methods block floating point + Selective RE sending or modulation compression + Selective RE sending. IQ-samples in corresponding PRB that are included in the U-Plane message shall be indicated with bit-value 1b in the mask. IQ-samples that are not included in U-Plane message are indicated with the bit-value 0b. The notation of the IQ-samples 1st, 2nd ... is according to the same order as shown in Table 8.3.2-1. The most significant bit in this mask is defined as the bit for the highest frequency RE in the PRB. The U-Plane message using sReSMask shall use the formatting shown in Table 8.3.3.18-1.

Value range: { 0000 0000 0000b – 1111 1111 1111b}. See Table 8.3.3.18-1 and also Table 8.3.3.15-1.

Bit allocations:

Table 8.3.3.18-1: sReSMask bit allocations

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
12 th	11 th	10 th	9 th	not part of sReSMask					Octet 1
8 th	7 th	6 th	5 th	4 th	3 rd	2 nd	1 st	1	Octet 2

Type: unsigned integer.

Field length: 12 bits.

8.3.3.19 udCompLen (PRB field length)

Description: This parameter specifies the total number of octets in the current section. This includes bit-padding (performed per PRB after the last part of the last IQ sample of the PRB) in all PRBs in the current section. The maximum supported PRB field length is $2^{16}-1$, but the actual size may be further limited by the maximum payload size of the underlying transport network. This field shall only be present for the following compression methods:

udCompMeth 0101b = BFP + selective RE sending.

udCompMeth 0110b = Modulation compression + selective RE sending.

If an O-RU declares support of the field then udCompLen shall be included in the U-Plane message (DL and UL). If the O-RU does not declare support of this field then this field shall not be included in the U-Plane message. This ensures backward-compatibility with v03.00 of the present document.

Value range: {0000 0000 0000 0010b - 1111 1111 1111 1111b},
{0000 0000 0000 0000b, 0000 0000 0000 0001b} Reserved.

Type: unsigned integer.

Field length: 16 bits.

8.3.4 DL data precoding

When Section Extension 3 is used for C-Plane and there is present an associated sectionId for U-Plane, then:

- O-RU should read 12 REs which have CRS reference signals in that PRB.
- O-RU shall understand the crsShift and crsReMask field to map appropriately the CRS REs to each antenna port.

Table 8.3.4-1 shows an example of DL data precoding showing several information element values.

Table 8.3.4-1: DL data precoding example

C-Plane example DL Precoding configuration parameters and indications (Section Extension 3)	O-RU Outcome									
<ul style="list-style-type: none"> • Section ext = 3 • txScheme='txD' • codeBookIndex= '00000000' (invalid) • numLayer = 4 • layerId = '0000' (TxD) • crsReMask= 001001001001 (Assuming MSB..LSB) • crsSymNum=0000 • crsShift=0 • beamIdAP1 	<ul style="list-style-type: none"> • Based on numLayers, crsShift, crsSymNum and crsReMask bit positions, CRS ports are mapped to the appropriate RE position and rest are left blank. <table border="1"> <thead> <tr> <th>CRS (4 Layer)</th> <th>If crsSymNum</th> <th>crsReMask Bit Position</th> </tr> </thead> <tbody> <tr> <td>(crs_Ant0)</td> <td>0</td> <td>{0,6} + vShift</td> </tr> <tr> <td>(crs_Ant1)</td> <td>0</td> <td>{3,9} + vShift</td> </tr> </tbody> </table>	CRS (4 Layer)	If crsSymNum	crsReMask Bit Position	(crs_Ant0)	0	{0,6} + vShift	(crs_Ant1)	0	{3,9} + vShift
CRS (4 Layer)	If crsSymNum	crsReMask Bit Position								
(crs_Ant0)	0	{0,6} + vShift								
(crs_Ant1)	0	{3,9} + vShift								

<ul style="list-style-type: none"> • beamIdAP2 • beamIdAP3 	
--	--

8.3.5 Data transfer for special cases

8.3.5.1 Data message mapping and packetization

See clause 8.3.2.

8.3.5.2 Uplink data transfer

Uplink IQ data transfer is covered in clause 8.3.2. In particular, PRACH and other common channels as well as SRS and other reference signal channels may use the same frequency domain IQ data packetization as with user data channels (PDSCH, PUSCH) following the transfer procedure described in clause 8.2.1. Optionally, PRACH and SRS related IQ data may be transferred as described in clause 8.3.5.3 and clause 8.3.5.4.

8.3.5.3 PRACH data transfer without C-Plane

In addition to general uplink data transfer (see clause 8.3.5.2), PRACH related IQ data may be transferred using the IQ data transfer procedure without C-Plane (see clause 8.2.2). In this case parameters controlling signal reception and sending U-Plane messages shall be provided via M-Plane. Refer to "Static PRACH" in M-Plane specification (see [7]) for description of M-Plane parameters and rules for sending U-Plane messages and interpretation of fields in U-Plane messages.

8.3.5.4 SRS data transfer without C-Plane

In addition to the general uplink data transfer (see clause 8.3.5.2), SRS related IQ data may be transferred with IQ data transfer procedure without C-Plane (see clause 8.2.2). In this case parameters controlling signal reception and sending U-Plane messages shall be provided via M-Plane. Refer to "Static SRS" in M-Plane specification (see [7]) for description of M-Plane parameters and rules for sending U-Plane messages.

8.4 U-Plane optimizations

8.4.1 Coupling via frequency and time

Coupling between C-Plane section descriptions and U-Plane data sections via frequency and time (for more details see clause 7.8.1.3) is a mode of operation enabled per eAxC via the M-Plane. It allows to optimize U-Plane message size by combining data sections that are continuous in frequency and are within the same symbol. Combined data sections shall share the same data section header (sectionId, rb, symInc, startPrbu, numPrbu and optionally udCompHdr) resulting in reduced message size. Value of sectionId shall be the same for all sections (predefined value) and does not prevent combining. If udCompHdr is present, then data sections can be combined only if value of udCompHdr is the same. Details of resource allocation (e.g., discontinuity in frequency, differences in compression) can restrict combining but resulting U-Plane message size shall always be less than or equal to the size without combining.

This U-Plane optimization method is compatible with all Section Types and Section Extensions with an exception for Section Type 3. With this coupling method, when multiple Section Type 3 messages are sent each with two (or more) different "freqOffset" values and/or "frameStructures", but same startPrbc and numPrbc, each C-Plane message can be uniquely identified by "freqOffset" and/or "frameStructure" fields, however associated U-Plane message with sectionId as 4095 and with no "freqOffset" or "frameStructure" cannot be uniquely identified. Hence, Section Type 3 shall not be used in conjunction with Time-Frequency Coupling. A frequency discontinuity of beamforming configuration in C-Plane does not imply discontinuity of data sections in U-Plane.

EXAMPLE 1: Before combining in symbol s there are two data sections: A and B such that $A.udCompHdr = B.udCompHdr$ and $A.startPrbu < B.startPrbu$ and $A.rb = B.rb = 0$. If $B.startPrbu = A.startPrbu + A.numPrbu$ then A and B can be combined into data section C such that $C.rb = 0$, $C.startPrbu = A.startPrbu$, $C.numPrbu = A.numPrbu + B.numPrbu$.

EXAMPLE 2: Before combining in symbol s there are two data sections: A and B such that $A.\text{udCompHdr} = B.\text{udCompHdr}$ and $A.\text{startPrbu} < B.\text{startPrbu}$ and $A.\text{rb} = B.\text{rb} = 1$. If $B.\text{startPrbu} = A.\text{startPrbu} + 2 \times A.\text{numPrbu}$ then A and B can be combined into data section C such that $C.\text{rb} = 1$, $C.\text{startPrbu} = A.\text{startPrbu}$, $C.\text{numPrbu} = A.\text{numPrbu} + B.\text{numPrbu}$.

EXAMPLE 3: Before combining in symbol s there are two data sections: A and B such that $A.\text{udCompHdr} = B.\text{udCompHdr}$ and $A.\text{startPrbu} < B.\text{startPrbu}$ and $A.\text{rb} = B.\text{rb} = 1$. If $B.\text{startPrbu} = A.\text{startPrbu} + 1$ and $A.\text{numPrbu} = B.\text{numPrbu}$ then A and B can be combined into data section C such that $C.\text{rb} = 0$, $C.\text{startPrbu} = A.\text{startPrbu}$, $C.\text{numPrbu} = A.\text{numPrbu} + B.\text{numPrbu}$.

The above examples are non-exhaustive.

8.4.2 Data blanking

8.4.2.1 Overview

Data blanking refers to the process of not sending PRBs on the fronthaul interface for some time and frequency resources in DL and/or UL directions. The objectives of data blanking include reduction of the fronthaul interface bitrate, power saving, and to provide an interval for antenna calibration. Data blanking may be performed by the O-DU in case of DL. In the UL case, the O-RU shall follow the O-DU instructions received in the C-Plane message(s).

Data blanking can take one of two forms: PRB blanking and Transmission blanking.

8.4.2.2 PRB blanking

PRB blanking is an inherent feature of the frequency-domain-based fronthaul split 7-2x, where unallocated PRBs are not sent across the fronthaul interface to reduce the fronthaul throughput. An unallocated PRB is a PRB with all REs unallocated.

In the case of DL, if one or more sets of PRBs are not allocated within one or more symbols of a specific slot (based on the scheduler), the O-DU should not refer to those unallocated PRBs within the C-Plane message(s). Consequently, the O-DU should not send zeros in the corresponding time/frequency REs related to those unallocated PRBs in the U-Plane messages.

Similarly, for UL, only allocated PRBs should be indicated by the O-DU in the UL C-Plane message(s). The O-RU shall send U-Plane messages only for the indicated PRBs.

8.4.2.3 Transmission blanking

Transmission blanking is a feature directly related to Section Type 0, where the O-DU indicates to the O-RU that certain PRBs or symbols in DL or UL will not be used (idle periods, guard periods). Likewise, there are no associated U-Plane messages containing IQ data for Section Type 0. The purpose is to inform the O-RU that transmissions may be halted during the specified idle interval for example to save power or to provide an interval for antenna calibration.

8.5 Aspects of U-Plane operation

8.5.1 O-RU U-Plane limits

8.5.1.1 O-RU U-Plane message limits

For the O-RUs with per U-Plane message processing limits, the O-RU may advertise its limitations on a per U-Plane message basis. The O-DU may indicate that it adheres to the associated limitations. The defined limit is the maximum number of section headers per downlink U-plane message. If the O-DU sets **uplane-message-section-header-limit-enabled** to 'true' then the O-DU shall comply to the U-Plane message processing limits. If the O-DU does not set **uplane-message-section-header-limit-enabled** to 'true' then the O-RU shall process the U-Plane message even if the maximum number of section headers per message exceeds the limit indicated by **max-section-headers-per-uplane-message**, however the stated O-RU capacity may be compromised. Refer to clause 15.10 of the M-Plane Specification [7] for details on usage of this feature.

When O-DU adheres to the per U-Plane message limit `max-section-headers-per-uplane-message`, the O-DU cannot rely on transport layer fragmentation to ensure that the limitation is not exceeded. The O-DU may use application layer fragmentation to ensure the limit is not exceeded.

9 Counters and KPIs

9.1 Counters

This clause defines recommended O-RAN C/U-Plane specific performance counters for the fronthaul interface. Table 9.1-1 defines the set of mandatory and optional O-RAN performance counters to be implemented within an O-RU or O-DU. These counters are defined from the perspective of the O-RU's or O-DU's Ethernet interfaces. These counters increment on eCPRI/IEEE1914.3 (see clause 5.1.3) message-based events. A message is defined as a complete eCPRI, or IEEE1914.3 message including header and payload. Unless otherwise noted in Table 9.1-1, these counters may be implemented to count events with a granularity of per O-RU, per Transport Flow, or per eAxC as indicated in the fronthaul M-Plane specification [7].

NOTE: Higher granularity counters provide more detailed information and can simplify managing and troubleshooting the front haul network.

The counters defined in Table 9.1-1 should only increment if the message is contained in a valid (error free) layer 2 or layer 3 packet. For example, if the packet has no layer 2 CRC or no layer 3 checksum errors it may be counted. All counters defined in the present document are wrap-around counters. Wrap-around counters are counters that automatically go from their maximum/final value to zero and continue to operate. These are unsigned counters. These counters shall provide explicit means to return them to their minimum/zero state, i.e., reset during normal operation. Due to their nature, wrap-around counters should be read frequently enough to avoid loss of information.

Only the transport related counters are listed herein; other O-RU counters may also be relayed across the fronthaul interface and these are described in Annex B of the M-Plane Specification [7].

Table 9.1-1: O-RU and O-DU Performance Counter Definitions

Counter name	Size (bits)	Mandatory or Optional	Counter definition
RX_TOTAL	64b	Mandatory	The total number of control/user plane eCPRI or IEEE 1914.3 messages received. This counter is the sum of all valid and errored messages received. (previously O-RAN-WG4.CUS.0-v02.00 named this Total_msgs_rcvd)
RX_ON_TIME	64b	Mandatory	The number of inbound user plane (ecpri type 0) messages that arrived within the specified time window. Some "on time" messages may have sequence number errors or corruption errors but as long as they arrived within specified window time, this counter should include them. If the received message has been transport-fragmented, the full message shall be reassembled before checking its arrival window. <ul style="list-style-type: none"> • Implementation of a single global counter is required. • Implementation of per transport counters are optional. • Implementation of per eAxC counters are optional.
RX_EARLY	64b	Mandatory	The number of inbound user plane messages which were detected to have arrived before the start of their designated receive window time. This counter increments whether the message is subsequently processed or dropped. <ul style="list-style-type: none"> • Implementation of a single global counter is required. • Implementation of per transport counters are optional. • Implementation of per eAxC counters are optional.
RX_LATE	64b	Mandatory	The number of inbound user plane messages which were detected to have arrived after the end of their designated receive window time. This counter increments whether the message is subsequently processed or dropped. <ul style="list-style-type: none"> • Implementation of a single global counter is required. • Implementation of per transport counters are optional. • Implementation of per eAxC counters are optional.

Counter name	Size (bits)	Mandatory or Optional	Counter definition
RX_SEQID_ERR	64b	Optional	<p>The number of inbound on time user plane messages in which a sequence identifier number error is detected. When this counter is implemented with an eAxC granularity, separate uplink and downlink counters shall be implemented.</p> <p>This error occurs when the ecpriSeqId field does not increment as specified in clause 5.1.3. Both the Sequence ID, and Subsequence ID fields shall be checked if transport fragmentation is supported otherwise only the Sequence ID field may be checked.</p> <p>In addition to identifying a sending equipment sequencing error, this counter can increment when packets are dropped prior to reception by the O-RU, or when packets reordered by the network exceed the receiving device's capabilities.</p>
RX_ON_TIME_C	64b	Mandatory	<p>The number of valid inbound control plane (ecpri type 2) messages that arrived within the specified time window. Some "on time" messages may have sequence number errors or corruption errors but as long as they arrived within specified window time, this counter should count them.</p> <ul style="list-style-type: none"> • Implementation of a single global counter is required. • Implementation of per eAxC counters are optional. • Implementation of per transport counters are optional. • Implementation of counters for each combination of eAxC and dataDirection fields are optional.
RX_EARLY_C	64b	Mandatory	<p>The number of inbound control plane messages which were detected to have arrived before the start of their designated receive window time. This counter increments whether the message is subsequently processed or dropped.</p> <ul style="list-style-type: none"> • Implementation of a single global counter is required. • Implementation of per eAxC counters are optional. • Implementation of per transport counters are optional. • Implementation of counters for each combination of eAxC and dataDirection fields are optional.
RX_LATE_C	64b	Mandatory	<p>The number of inbound control plane messages which were detected to have arrived after the end of their designated receive window time. This counter increments whether the message is subsequently processed or dropped.</p> <ul style="list-style-type: none"> • Implementation of a single global counter is required. • Implementation of per eAxC counters are optional. • Implementation of per transport counters are optional. • Implementation of counters for each combination of eAxC and dataDirection fields are optional.
RX_SEQID_ERR_C	64b	Optional	<p>The number of inbound on time control messages in which a sequence identifier number error is detected. This counter increments under the same conditions as the rx_seqid_num_err except for control plane messages.</p> <ul style="list-style-type: none"> • Implementation of a global counter is optional. • Implementation of eAxC counters is optional. • Implementation of per transport counters are optional. • Implementation of counters for each combination of eAxC and dataDirection fields are optional.

Counter name	Size (bits)	Mandatory or Optional	Counter definition
RX_CORRUPT	64b	Optional	<p>The number of inbound on time messages with a correct ecpriSeqId (no sequence number error) which are dropped by the terminating entity due to the message containing one or more eCPRI/IEEE 1914.3 or O-RAN protocol errors. This counter's granularity is limited to per O-RU or per transport flow implementations.</p> <p>Protocol errors are defined as when eCPRI/IEEE 1914.3/O-RAN defined fields contain invalid values or indicate unsupported capabilities. Some examples of this are:</p> <ol style="list-style-type: none"> 1) Pcid or section Id number which has not been configured. 2) Unexpected use of C bit, 3) Unconfigured or Unsupported udCompHdr setting. 4) Unsupported Section Extension. 5) Wrong ecpriVersion and/or payloadVersion information 6) ecpriMessage field does not contain 0, 2, or 5.
RX_ERR_DROP		Optional	The total number of inbound messages which are discarded by the receiving O-RAN entity for any reason.
RX_PKT_DUPL	64b	Optional	Duplicated packet. This counter is deprecated.
TX_TOTAL	64b	Mandatory	The number of valid outbound control/user plane messages.
TX_TOTAL_C	64b	Mandatory	The number of valid outbound control plane messages. This counter is required on an O-RU only if O-RU supports LAA/LBT capabilities.

10 Specification mandatory and optional capabilities

10.1 General

Table 10.1-1 lists the general system capabilities supported by the present document.

Table 10.1-1: Supported LTE/NR channels and system capabilities

Feature	System capability support
LTE DL Channels	PDSCH, PBCH, PCFICH, PDCCH, ePDCCH, MPDCCH, PHICH, CRS, MBSFN RS, UE-RS, DMRS for ePDCCH/MPDCCH, PRS, CSI-RS, PSS, SSS, Discovery RS
LTE UL Channels	PUSCH, PUCCH, DMRS-PUSCH, DMRS-PUCCH, SRS, PRACH (incl. eMTC)
Narrow band IoT DL Channels	NB-DMRS, NB-PDSCH, NB-PBCH, NB-PDCCH, NB-RS, NB-PRS, NB-PSS, NB-SSS
Narrow band IoT UL Channels	NB-PUSCH, NB-PRACH
NR DL Channels	PDSCH, PDCCH, DMRS-PDSCH, PTRS-PDSCH, DMRS-PDCCH, DMRS-PBCH, CSI-RS, PSS, SSS, SS Block/PBCH
NR UL Channels	PUSCH, PUCCH, PRACH, DMRS-PUSCH, PTRS-PUSCH, DMRS-PUCCH, SRS
Technologies	LTE TDD, FDD (normal and extended CP) NR TDD, FDD
Channel Bandwidth	LTE: 1.4, 3, 5, 10, 15, 20 MHz
	NR: up to 400MHz
Subcarrier Spacing	LTE: 15 kHz, 7.5 kHz, 1.25 kHz LTE PRACH: 1.25 kHz, 7.5 kHz NB-IoT PRACH: 3.75 kHz
	NR: 15, 30, 60, 120, 240 kHz
	NR Multi Numerology
	NR PRACH: 1.25, 5, 15, 30, 60, 120 kHz
LTE Specific Features	DL Transmission Modes: TM1 - TM10
	UL Transmission Modes: TM1, TM2
	Carrier Aggregation
	eMBMS
	TTI-Bundling
	Semi-Persistent Scheduling (SPS)
	MIMO (SU/MU-MIMO)
	UE TAS (Tx Antenna Selection)
	FeICIC (ABS)
	CoMP (DL/UL), Joint Transmission
	Short TTI
	eMTC
	NB-IOT (in band/guard band/standalone)
	License Assisted Access (LAA)
	Sidelink (Proximity Services)
	Dynamic TDD (eIMTA)
Mission Critical PS-LTE Features (MCPTT)	
Positioning (PRS, OTDOA etc)	
V2X	

Feature	System capability support
	Distributed Antenna System Support
	CBRS Support
NR Specific Features	EN-DC
	SSBlock
	BW Part
	Supplementary UL
	Mini-slot
	LTE-NR Co-existence
Beamforming	Analog Beamforming
	Digital Beamforming
	Hybrid Beamforming
	O-RU Support for 64 TRX
Transport	L2: Ethernet
	L3: IPv4, IPv6
	QoS over Fronthaul

10.2 Mandatory and optional capabilities

This clause provides details regarding which capabilities within the present document are mandatory and which are optional. The capability requirements of O-DU can be different from the O-RU because in many cases, the O-DU needs to implement multiple options as mandatory to ensure interoperability with O-RUs that have optional capabilities. For example, the ability to support many compression methods may be mandatory in the O-DU while in O-RUs there may be only a single mandatory compression method to allow simplicity in O-RU design (while vendors may enhance their O-RU product offering by implementing some of the optional compression methods).

Table 10.2-1 describes the capabilities required of O-DU and O-RU units. There are three choices:

Mandatory: The unit shall support the described capability to be O-RAN compliant.

Conditional Mandatory: The unit shall support the described capability to be O-RAN compliant, but the additional information column describes the conditions under which the capability is mandatory.

Optional: The unit need not support the capability and still be O-RAN compliant, but if the unit does support the described capability it shall support it in the way described within the present document.

Table 10.2-1: O-RAN mandatory and optional features

Category	Feature	O-DU support	O-RU support	Additional information
O-RU Category Support	Support for Category-A O-RU (up to 8 spatial streams)	Mandatory	NA	The O-DU may only support fewer than 8 spatial streams; that number of spatial streams shall however be supported for Category-A O-RUs.
	Support for Category-A O-RU (> 8 spatial streams)	Optional	NA	
	Support for Category-B O-RU (precoding in O-RU)	Mandatory	NA	
Security	IEEE 802.1X supplicant functionality [51]	Optional	Mandatory	

Category	Feature	O-DU support	O-RU support	Additional information
	Grand Master Clock Redundancy	Optional	Optional	Principles for Grand Master Clock redundancy are provided in Annex P.2.3.1 of IEEE 1588 [33]. Guidelines are provided in Annex G of O-RAN Synchronization Architecture and Solution Specification [54].
Beam-forming	Beam index based	Mandatory	Conditional Mandatory	Condition applies to UE-specific BF for any O-RU capable of BF; a non-BF O-RU shall be supplied a zero beamId if a C-Plane message containing a beamId is sent at all.
	Real-time BF Weights	Conditional Mandatory	Conditional Mandatory	Condition for O-DU: Mandatory only for O-DUs designed to support any kind of BF that involves the operation of updating BF weights in real-time; Condition for O-RU: for any O-RU internally using BF weights, the ability to update the weights in real-time via C-Plane messages (using "Real-time BF Weights") shall be mandatory.
	Real-time beamforming attributes	Optional	Optional	This is considered to not internally use BF weights.
	Real-time UE channel Info	Optional	Optional	This is considered to internally use BF weights.
	Predefined beam tilt for beam index based beamforming	Optional	Optional	
	Antenna calibration support	Optional	Optional	
	Null-layer Info. for ueld-based beamforming using Section Extension 14	Optional	Optional	
User port group indication for beamforming based on UE channel (Section Extension 17)	Optional	Optional		
Bandwidth Saving	IQ Data Formats			
	Fixed point (no compression)	Mandatory	Mandatory	16-bit mandatory (others optional).
	Block floating point compression	Conditional Mandatory	Conditional Mandatory	9, 12 & 14-bit mantissa mandatory if this compression method is supported (others optional). Condition: if an O-DU or O-RU supports any IQ compression it shall support this one.
	Block scaling compression	Optional	Optional	9 & 14-bit scaler mandatory if this compression method is supported (others optional).
	μ -law compression	Optional	Optional	9 & 14-bit width mandatory if this compression method is supported (others optional).
	Modulation compression	Optional	Optional	4-bit width mandatory if this compression method is supported (others optional).
	Block floating point compression + selective RE sending	Optional	Optional	9, 12 & 14-bit mantissa mandatory if this compression method is supported (others optional).
Modulation compression + selective RE sending	Optional	Optional	4-bit width mandatory if this compression method is supported (others optional).	

Category	Feature	O-DU support	O-RU support	Additional information
	Presence of udCompLen	Conditional Mandatory	Optional	If the O-RU declares support of udCompLen then it shall be used, otherwise it shall be omitted (in DL and UL U-Plane messages); this is only relevant for block floating point with selective RE sending or modulation compression, with selective RE sending.
	Real-time variable bit-width	Optional	Optional	IQ data format is determined by value of udCompHdr. This implies presence of udCompHdr in U-Plane messages.
	Real-time variable bit-width per Channel (per data section)	Optional	Optional	IQ data format is determined by value of udCompHdr. This implies presence of udCompHdr in U-Plane messages. Values of udCompHdr can be different in different sections in the U-Plane messages.
	Static configuration of U-Plane IQ format and compression header	Conditional Mandatory	Conditional Mandatory	IQ data format is determined by M-Plane configuration. This implies absence of udCompHdr in U-Plane messages. Mandatory for supported IQ formats listed in first 7 rows under "IQ Data Formats" in this table.
	Beamspace compression	Optional	Optional	This compression algorithm is specific to beamforming weights.
	Channel information compression	Optional	Optional	
	<ul style="list-style-type: none"> No compression / Fixed point 	Conditional Mandatory	Conditional Mandatory	16-bit width for cilsample and ciQsample mandatory if the channel information based BF feature is supported (other bitwidth optional).
	<ul style="list-style-type: none"> Block Floating Point 	Optional	Optional	
	<ul style="list-style-type: none"> Block Scaling 	Optional	Optional	
	<ul style="list-style-type: none"> μ-law 	Optional	Optional	
	Use of "symInc" flag to allow multiple symbols in a C-Plane section	Optional	Optional	
	Coupling via sectionId Value	Mandatory	Mandatory	
	Coupling via Frequency and Time	Conditional Mandatory	Conditional Mandatory	If O-RU or O-DU supports "Coupling via Frequency and Time with Priorities" it shall also support this one.
	Coupling via Frequency and Time with Priorities	Optional	Optional	
	Coupling via Frequency and Time with Priorities (Optimized)	Optional	Optional	Refer to clause 7.8.1.5 and clause 7.9.8
	PRACH data transfer without C-Plane	Optional	Optional	
	SRS data transfer without C-Plane	Optional	Optional	

Category	Feature	O-DU support	O-RU support	Additional information
Energy Savings	Transmission blanking	Optional	Optional	
O-DU – O-RU Timing	Defined Transport Method	Mandatory	Mandatory	
	Measured Transport Method (eCPRI Msg 5)	Optional	Optional	If O-RU supports Measured Transport Method it shall support both 1-Step and 2-step version of T12 measurement and at least one of 1-Step or 2-Step version of T34 measurement. If the O-DU supports T34 measurement it shall support both 1-Step and 2-Step version. See clause 4.4.4.4 for more detailed information.
	External Antenna Delay handling using Minimal O-DU Impact Method	Optional	Optional	Using Tda and Tau parameters as defined in clause 4.7.2.
Synchroni- zation	G.8275.1 [31]	Conditional Mandatory	Conditional Mandatory	When the G.8275.1 [31] profile is used: In LLS-C1/C2/C3, the O-RU shall be synchronized from a PTP source using this profile and may optionally use PLFS assistance as well In LLS-C1/C2/C3/C4, the O-DU may optionally be synchronized from a PTP source using this profile and may optionally use PLFS assistance as well. In LLS-C1, the O-DU shall synchronize the O-RU using this PTP profile on the fronthaul interface. The O-DU shall synchronize the O-RU using PLFS as well only if the O-RU requires it (otherwise it is optional). In LLS-C2, the O-DU shall synchronize the fronthaul network elements using this PTP profile and PLFS on the fronthaul interface. In LLS-C3/C4, the O-DU does not transmit synchronization signals to the O-RU. In LLS-C2 and LLS-C3 topologies supporting Shared cell, the network elements in the Fronthaul synchronization chain (FHM or cascaded O-Rus) that are on the path to other network elements in the synchronization chain shall synchronize them using this PTP profile and PLFS on the fronthaul interface.
	G.8275.2 [32]	Optional	Optional	When the G.8275.2 [32] profile is used: O-DU (in LLS-C1/C2/C3/C4) or O-RU (in LLS-C1/C2/C3) may optionally be synchronized from a PTP source using this PTP profile. In LLS-C1/C2, O-DU may optionally synchronize the O-RU using this PTP profile on the fronthaul interface.

Category	Feature	O-DU support	O-RU support	Additional information
	Local PRTC	Optional	Optional	O-DU (in LLS-C1/C2/C3/C4) or O-RU (in LLS-C4) may optionally be synchronized from a local time source, for example GNSS-based.
Transport Features	L2: Ethernet	Mandatory	Mandatory	
	L3: IPv4, IPv6 (CUS Plane)	Optional	Optional	
	QoS over Fronthaul	Mandatory	Mandatory	
	Prioritization of different U-Plane traffic types	Optional	Optional	
	Support of Jumbo Ethernet frames	Optional	Optional	
	eCPRI	Mandatory	Mandatory	
	support of eCPRI concatenation	Optional	Optional	
	IEEE 1914.3 transport header	Optional	Optional	See clause 5.1.3.3.
	Application layer fragmentation	Mandatory	Mandatory	C-Plane and U-Plane.
Radio Transport layer fragmentation	Optional	Optional	U-Plane (see clause 5.1.3.2.8).	

Category	Feature	O-DU support	O-RU support	Additional information
Section Types and Section Extensions	Section Type 0	Optional	Mandatory	O-RU may ignore message if blanking or other Section Type 0 utility is not supported.
	Section Type 1	Mandatory	Mandatory	
	Section Type 3	Mandatory	Mandatory	
	Section Type 4	Optional	Optional	Specific configuration Control
	Section Type 5	Optional	Optional	Specific to Channel-Info beamforming.
	Section Type 6	Optional	Optional	Specific to Channel-Info beamforming.
	Section Type 7	Optional	Optional	Specific to LAA which is an optional capability.
	Section Type 8	Conditional Mandatory	Conditional Mandatory	Condition: Mandatory if SE 22 for ACK/NACK request is supported. Specific to ACK/NACK reporting for section descriptions in C-plane messages
	Beamforming weight transfer using Section Extension 1	Conditional Mandatory	Conditional Mandatory	Condition for O-DU: Mandatory only for O-DUs designed to support any kind of BF that involves the operation of updating BF weights in real-time; Condition for O-RU: for any O-RU internally using BF weights, the ability to update the weights in real-time via C-Plane messages (using "Real-time BF Weights") shall be mandatory.
	Beamforming attribute transfer using Section Extension 2	Optional	Optional	Attribute-based beamforming is optional.
	DL precoding configuration using Section Extension 3	Optional	Optional	While Category B is mandatory, it is possible to precode using beamforming so use of this extension is optional.
	<ul style="list-style-type: none"> First Data Layer and Non-first data layer association 	Conditional Mandatory	Conditional Mandatory	Condition: DL precoding configuration using Section Extension 3 is supported by O-DU and O-RU.
	Modulation compr. Parameters using Section Extension 4	Conditional Mandatory	Conditional Mandatory	Condition for O-DU: If modulation compression is supported then it is mandatory for O-DU to support Section Extension 4 Condition for O-RU: If modulation compression is supported and the O-RU does not support Section Extension 5 then it is mandatory for O-RU to support Section Extension 4
	Modulation compr. Parameters using Section Extension 5	Conditional Mandatory	Conditional Mandatory	Condition for O-DU: If modulation compression is supported then it is mandatory for O-DU to support Section Extension 5 Condition for O-RU: If modulation compression is supported and the O-RU does not support Section Extension 4 then it is mandatory for O-RU to support Section Extension 5
Non-contiguous PRB allocation using Section Extension 6	Optional	Optional	Use of non-contiguous PRBs is optional.	
eAxC masking using Section Extension 7	Optional	Optional	Use of eAxC masking is optional.	

Category	Feature	O-DU support	O-RU support	Additional information
	Provide MMSE parameters using Section Extension 8	Conditional Mandatory	Conditional Mandatory	Specific to Channel-Info beamforming; O-DU condition: if O-RU and O-DU both support Channel-Info BF, then the O-DU shall use Section Extension 8 to convey MMSE parameters; O-RU condition: if the O-RU supports Channel-Info BF, then the O-RU shall accept Section Extension 8 MMSE parameters from the O-DU.
	LTE/NR DSS using Section Extension 9	Optional	Optional	DSS using overlapping carriers is possible; DSS using this Section Extension is optional.
	Group configuring of multiple ports using Section Extension 10	Optional	Optional	Use of multiple port (multiple eAxC) grouping is optional.
	Flexible Beamforming Weights using Section Extension 11	Optional	Optional	Use of flexible beamforming weights Section Extension is optional.
	contInd flag in SE 11	Optional	Optional	See clause 7.7.11.9
	Non-contiguous PRB allocation with frequency ranges using Section Extension 12	Optional	Optional	
	PRB allocation with frequency hopping using Section Extension 13	Optional	Optional	
	Nulling-layer Info. for ueld-based beamforming using Section Extension 14	Optional	Optional	
	Mixed-numerology Info. for ueld-based beamforming using Section Extension 15	Optional	Optional	
	Antenna mapping in UE channel information based UL beamforming using Section Extension 16	Optional	Optional	Use of antenna mapping in UE channel information based UL beamforming is optional.
	Indication of user port group using Section Extension 17	Optional	Optional	
	Uplink traffic management using Section Extension 18	Conditional Mandatory	Conditional Mandatory	Mandatory if uplink traffic management using C-Plane is supported. Not permitted if uplink traffic management using C-Plane is not supported.
	Compact beamforming information for multiple port using Section Extension 19	Optional	Optional	See clause 7.7.19 and 7.9.11.
	Dedicated puncturing Section using Section Extension 20	Optional	Optional	See clause 7.7.20 and 7.9.12.
	Variable PRB group size for Channel Information using Section Extension 21	Optional	Optional	
	ACK/NACK request using Section Extension 22	Optional	Optional	See clause 7.7.22 and 7.2.8 for more details.
	Multiple symbol mcScaleOffset using Section Extension=23	Optional	Optional	Multiple symbol mcScaleOffset using Section Extension=23
Section Type 4 commands	Time-domain beamforming configuration using st4CmdType(1) TIME_DOMAIN_BEAM_CONFIG	Optional	Optional	See clause 7.5.3.38 and clause 7.2.9.2.1 for more details
	TDD pattern configuration using st4CmdType(2) TDD_PATTERN_CONFIG	Optional	Optional	Refer to clause 7.5.3.38 and clause 7.2.9.2.2 for more details
Other features	LAA LBT O-DU congestion window mgmt	Conditional Mandatory	Conditional Mandatory	Mandatory only for O-DUs and O-RUS supporting LAA.
	LAA LBT O-RU congestion window mgmt	Optional	Optional	
	UL gain correction per eAxC	Optional	Optional	See clause 8.1.3.2.3.

Category	Feature	O-DU support	O-RU support	Additional information
	DL reference level adjustment	Optional	Optional	See Reference_Level in clause 8.1.3.3.
	FS adjustment	Optional	Optional	See FS_Offset in clause 8.1.3.
	Ordered transmission	Optional	Optional	See clause 4.6.3.
	Uplink traffic management using M-Plane	Optional	Optional	See clause 4.6.4.
	Uplink traffic management using C-Plane	Optional	Optional	See clause 4.6.4.
	Uniformly distributed transmission	Optional	Optional	In accordance with clause 4.6.2. Requires support of uplink traffic management (using M-Plane or C-Plane).
	Independent U-Plane transmission window control	Optional	Optional	According to clause 4.6.4. Requires support of uplink traffic management (using M-Plane or C-Plane).
	C-Plane Message processing O-RU limits	Optional	Optional	As specified in clause 7.8.2 and M-Plane specification [7], clause 15.8.
	O-RU U-Plane message limits	Optional	Optional	As specified in clause 8.5.1.1 and M-Plane specification [7], clause 15.10
	beam-update-contention-control	Optional	Optional	See clause 12.4.3
	Provision of beam-context-gap-period	N/A	Optional	See clause 12.4.3
	UPLANE-ONLY-DL-MODE	Optional	Optional	See clause 8.2.2
	Shared Cell	Optional	Optional	For O-DU: support for configuring an FHM and/or cascaded O-RUs is optional. For O-RU: support for the shared-cell copy/combine function is optional.
NB-IoT	NB-IoT support	Optional	Optional	See clause 15
	Extended PRB grid method for NB-IoT in DL	Conditional Mandatory	Conditional Mandatory	Condition: O-DU / O-RU supports guard-band NB-IoT in DL. See clause 15.2.1 for the description of the method
	Extended PRB grid method for NB-IoT UL	Optional	Optional	See clause 15.4.3 for the description of the method
	SSSC method for NB-IoT in UL using ST 3	Conditional Mandatory	Conditional Mandatory	Condition: O-DU / O-RU supports guard-band or in-band NB-IoT. See clause 15.4.1 for the description of the SSSC method
	SSMC method for NB-IoT in UL using ST 3	Optional	Optional	See clause 15.4.2 for the description of the SSMC method
	Extended PRACH configuration for PRACH data transfer without C-Plane for NPRACH	Optional	Optional	See feature EXTENDED-PRACH-CONFIGURATION in M-plane.
NOTE 1: When a capability that is "per endpoint" is cited as mandatory in the O-RU, it means at least one endpoint shall support it (a minimum number meaningful for the functionality of the O-RU), not that ALL endpoints shall support it.				
NOTE 2: For some mandatory capabilities, only a subset of the full set of parameter values is mandatory, this is marked in the "Additional information" column. If an optional capability is supported, the "additional information" column lists the parameter values that shall be supported; support for other parameter values of the optional capability remains optional.				

11 S-Plane Protocol

11.1 General

11.1.1 Overview

Time and frequency synchronization can be distributed to the O-DU and O-RU from the same source using different methods. However, synchronization accuracy is mostly impacted by implementation (e.g., timestamping near the interfaces, number of hops) than by the technology itself. The following synchronization options are available over an Ethernet network:

- Frequency synchronization where clocks are aligned in frequency
- Phase synchronization where clocks are aligned in phase
- Time synchronization where clocks are aligned to a common base time

Together the above parameters define a profile for the network, requiring a set of features and option selections for bridges and end stations operation. Further, the profile also states the conformance requirements for supporting equipment and user applications.

The present document considers frequency, phase and time synchronization of all the network elements (O-DUs, intermediate switches and O-RUs) for TDD and FDD features requiring specific TAE(Time Alignment Error). Frequency-only configurations (like LTE FDD or 5G FDD) are currently outside the scope of the present document. The present document considers a configuration where the O-DU and the O-RU are separated by a fronthaul link or network, and not configurations that have both O-DU and O-RU functions in the same piece of equipment.

11.2 Synchronization baseline

11.2.1 List of reference documents

See clause 2.

11.2.2 Clock model and synchronization topology

11.2.2.1 Introduction

Different synchronization topologies are necessary to address different deployment market needs. The following four topology configurations are considered by as compliant topologies for supporting the O-RU synchronization needs. A configuration label is used for easier reference through the present document:

- Configuration LLS-C1: with this topology, the O-DU is part of the synchronization chain towards the O-RU. Network timing is distributed from O-DU to O-RU via direct connection between O-DU site and O-RU site.
- Configuration LLS-C2: with this topology, the O-DU is part of the synchronization chain towards the O-RU. Network timing is distributed from O-DU to O-RU between O-DU sites and O-RU sites. One or more Ethernet switches are allowed in the fronthaul network. Interconnection among switches and fabric topology (for example mesh, ring, tree, spur etc.) are deployment decisions which are out of the scope of the present document.
- Configuration LLS-C3: with this topology, the O-DU is not part of the synchronization chain towards the O-RU. Network timing is distributed from PRTC/T-GM to O-RU typically between central sites (or aggregation sites) and O-RU sites. One or more Ethernet switches are allowed in the fronthaul network. Interconnection among switches and fabric topology (for example mesh, ring, tree, spur etc.) are deployment decisions which are out of the scope of the present document.
- Configuration LLS-C4: This topology concerns the case where the synchronization reference is provided to the O-RU with no involvement of the transport network (typically with a local GNSS receiver).

NOTE 1: For all LLS-C1 to LLS-C4 configurations, with exception of the O-DU and O-RU, the only clock types and classes (PRTC/T-GM, ePRTC/T-GM T-BC, T-TC) as defined in ITU-T G.8272 [25], ITU-T G.8272.1 [26], ITU-T G.8273.2 [28] and ITU-T G.8273.3 [29] respectively, and considered for the network limits in Recommendation ITU-T G.8271.1 [23] are covered by the present document.

NOTE 2: Also applying to all LLS configurations, how an O-DU is synchronized is not in the scope of this classification of the synchronization topologies. O-DU may be synchronized from either a local or remote PRTC.

11.2.2.2 Topology configuration LLS-C1 and LLS-C2 synchronization

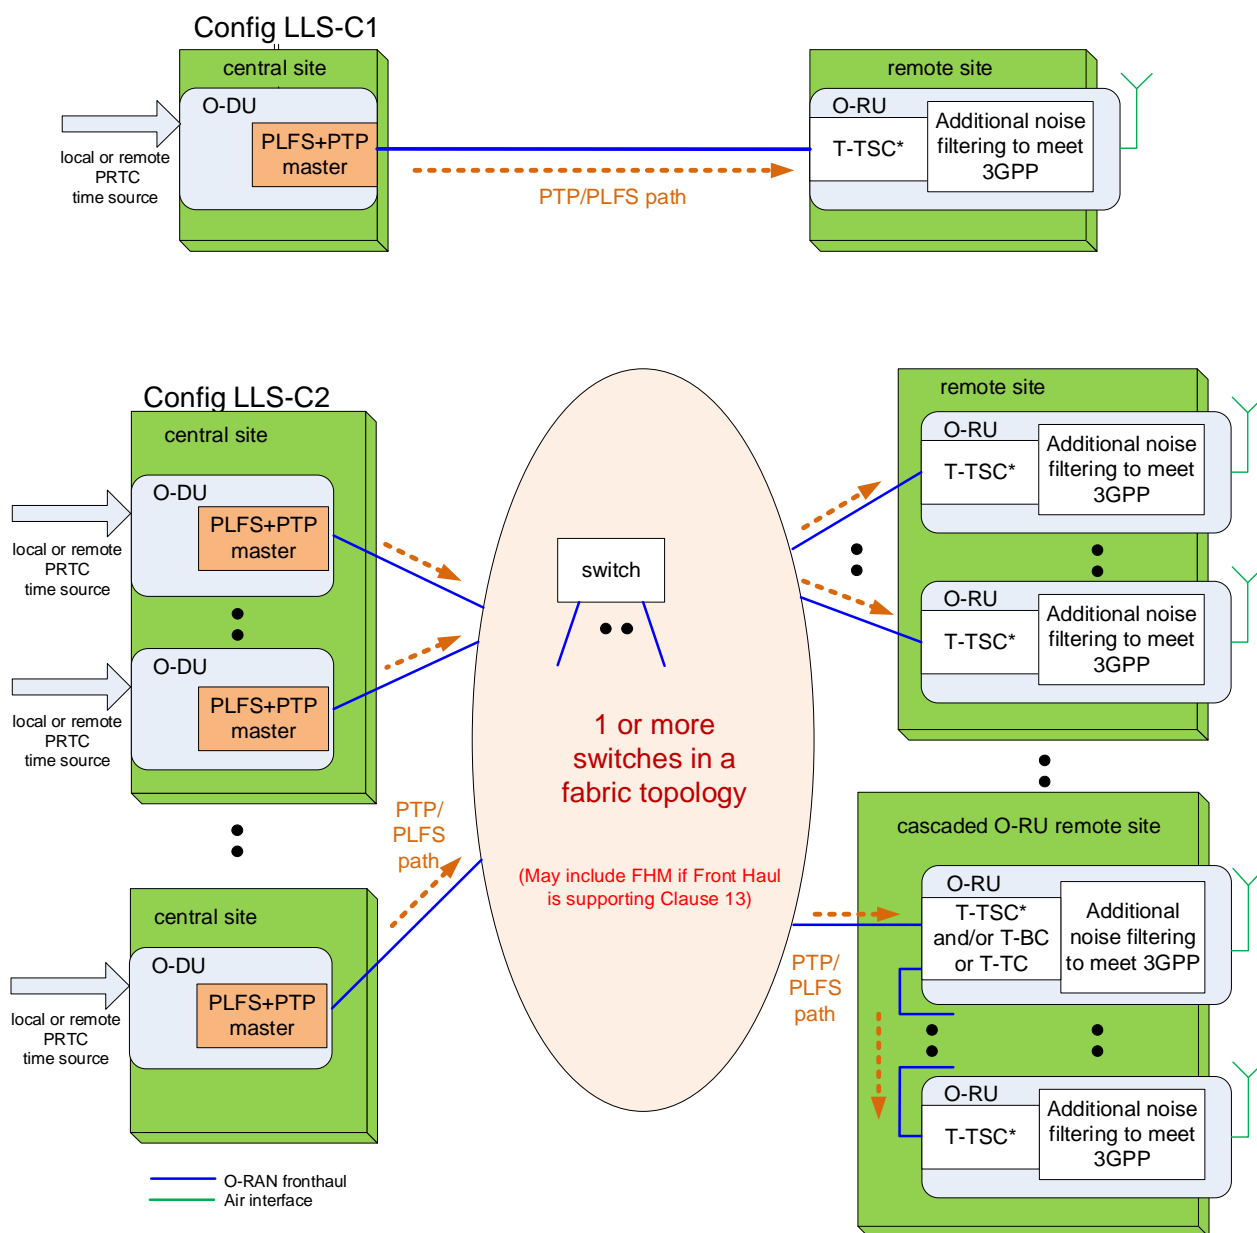
Configuration LLS-C1 is based on point-to-point connection between O-DU and O-RU using network timing option. As shown in the top portion of Figure 11.2.2.2-1 below, it is basically the simplest topology for network timing option, where O-DU directly synchronizes O-RU.

Configuration LLS-C2 is similar to LLS-C1 with O-DU acting as master to distribute network timing toward O-RU. One or more Ethernet switches are allowed between the central site (hosting O-DUs) and the remote sites (hosting O-RUs). The allowed number of switches in the synchronization path is limited by frequency and time error contributions by all clocks in the chain.

With Full Timing Support, the allowed network noise limit in the budget can be met by a certain number (refer to Annex H for more details) of Class B or class C T-BC switches as shown in Figure 11.2.2.2-1 below. Additional T-BC switches may be allowed if total noise limit can be met. The synchronization master shall be located at the O-DU. Further, all Ethernet switches in the fronthaul function as T-BC as specified by Recommendation ITU-T G.8273.2 [28], T-TC as specified by Recommendation ITU-T G.8273.3 [29] shall be also allowed as T-BC replacement with the same expectation based on Recommendation ITU-T G.8271.1 [23].

With Partial Timing Support, which network limits are specified in G.8271.2 [24], non-T-BC switches may also be deployed. Further investigation is required to specify appropriate frequency and timing budgets and network configuration to ensure 4G and NR TAE requirements as described in Table 11.3.2.1-1 and 3GPP frequency accuracy requirements can be met.

Interconnection among switches and fabric topology (for example mesh, ring, tree, spur etc.) are deployment decisions which are out of the scope of the present document.



NOTE: This is for illustrative purposes and does not provide specific deployment guidance (for example the number of switches in a clock chain and fabric topology).

Figure 11.2.2.2-1: Configurations LLS-C1 and LLS-C2 synchronization (see considerations below for additional information on O-DU, O-RU, switch functionality).

Considerations for O-DU: O-DU acts as PLFS + PTP synchronization master towards the fronthaul interface, but there are different possible sub-configurations based on the O-DU sync source. Two main cases are possible: O-DU is itself the master at the top of the synchronization chain (e.g., PRTC directly connected to the O-DU) or it gets synchronization from the network. In particular:

- 1) If O-DU synchronization source is from a local PRTC (typically a GNSS receiver), it may act as a T-GM, or as a specific PLL (Phase Locked Loop) with higher jitter and wander filtering capability.
- 2) If O-DU synchronization source is from a remote PRTC through a network (typically PTP, with or without PLFS, which grandmaster can be located anywhere in the network), O-DU shall act as the receiver of synchronization toward the upstream network. Two sub-configurations are possible:
 - If the PTP profile used in the upstream network is different (typically Recommendation ITU-T G.8275.2 [32]) from the fronthaul one (which is typically Recommendation ITU-T G.8275.1 [31], annex A), then the O-DU shall act as an Inter-Working Function (IFW) clock to bridge between the profiles, as

described in Recommendation ITU-T G.8271.2 [24], Appendix I (but no clock specification exists for such IWF). Informative guidelines are provided in Recommendation ITU-T G.8275 [30] Appendix III.

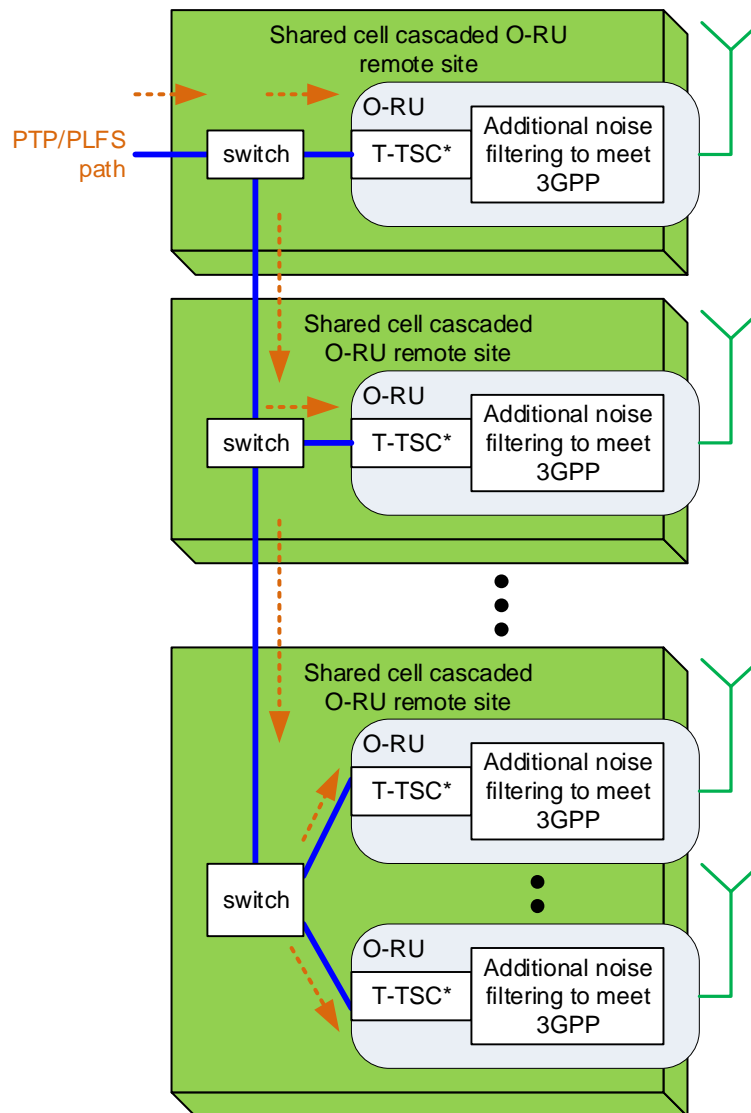
- If the PTP profile used in the upstream network is same as the fronthaul one (typically Recommendation ITU-T G.8275.1 [31], annex A), then the O-DU may act as a combined receiver and transmitter of PTP synchronization with higher jitter and wander filtering capability. NOTE: O-DU acting as a Recommendation ITU-T G.8273.2 [28] T-BC clock does not provide enough wander cleaning to guarantee the 15 ppb limit and is therefore outside the scope of LLS-C1/LLS-C2.
- 3) When multiple O-DUs are directly connected to the central site aggregation switch supporting T-BC in configuration LLS-C2, only one active master will be in the clock chain as determined by BMCA. One O-DU is on the synchronization path towards the O-RUs, including the ones controlled by other O-DUs over the M-Plane. If master redundancy is present, the synchronization path via another O-DU will serve as a backup synchronization reference.

Considerations for Switches:

- 1) ITU-T "Full Timing Support" Clock: Recommendation ITU-T G.8273.2 [28] T-BC (or Recommendation ITU-T G.8273.3 [29] T-TC).
- 2) In the case of FHM or cascaded O-RU used in Shared cell topologies described in clause 13, this equipment also behaves as an ITU-T "Full Timing Support" Clock from its synchronization timing receiver-to master port from S-Plane point of view.
- 3) Partial timing support using non-T-BC switches may be used. The PTP profile defined in Recommendation ITU-T G.8275.2 annex A shall apply for partial timing support. Performance aspects and budgets associated with this mode require further investigation.

Considerations for O-RU:

- 1) As an end application, the O-RU includes an application-specific PTP T-TSC* clock that may be compliant to Recommendation ITU-T G.8273.2 [28], T-TSC. The combination of the T-TSC* and the local RF end application clock needs to fulfill additional requirements to ensure 3GPP air interface compliance on the RF interface. T-TSC shall be per IEEE 802.1CM [11, 12] interface condition Case 1 in clause 6.
- 2) In the case of Shared cell using Cascade mode with direct chaining of the O-RUs, the O-RU also implements either Recommendation ITU-T G.8273.2 [28], T-BC or Recommendation ITU-T G.8273.3 [29], T-TC on the path between the S-Plane synchronization timing receiver-and S-Plane master ports (as per above Considerations for Switches).
If the O-RU is not on the synchronization path to other O-RUs, the O-RU need not support T-BC or T-TC functions. These may be supported by a collocated switch. Figure 11.2.2.2-2 presents an example not applicable to LLS-C1 where the O-RUs from the U-Plane perspective can be cascaded, but from the S-Plane perspective the synchronization chain is arranged over the external switches implementing T-BC or T-TC. Other examples are possible where one switch is directly connected to the O-RUs implemented in other remote sites. This would further minimize the relative Time Error between O-RUs implemented in different sites. The topology discovery procedure (Management Plane Specification, clause 17.6.2) in this case would need to take into consideration that intermediate switches are deployed between O-RUs.
- 3) An end application connected to an external T-TSC per IEEE 802.1CM [11, 12] clause 6.4 interface condition Case 2. This case need not be considered since the O-RU generally does not provide a separate PPS/ToD interface for external T-TSC connection, and there can be performance concerns about the 1pps distribution interface compared to PTP over Ethernet.



NOTE: The lines between the switches and the O-RUs are logical connections, and for the User Plane they would involve a pair of physical links, implementing the northbound and southbound connections in the cascaded O-RUs. For some O-RU implementations this set up might not be feasible.

Figure 11.2.2.2-2: Configuration LLS-C2 synchronization: example on Shared Cell Cascade O-RU remote site configuration alternative options

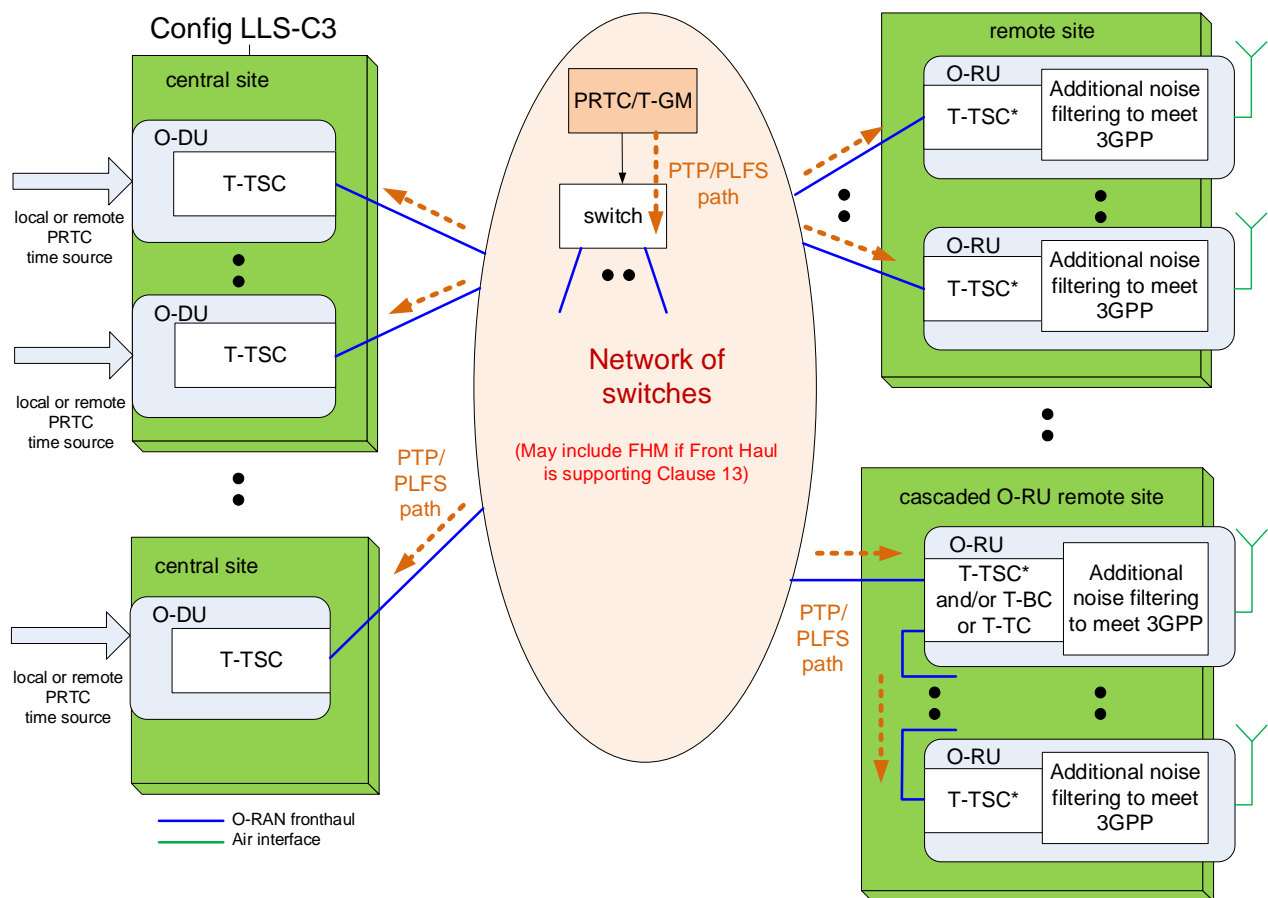
11.2.2.3 Topology configuration LLS-C3 synchronization

Configuration LLS-C3 is similar to LLS-C2 except frequency and time distribution is made by the fronthaul network itself (not by the O-DU). That means that one or more PRTC/T-GM are implemented in the fronthaul network to distribute network timing toward O-DU and O-RU. One or more Ethernet switches are allowed between the central site (hosting O-DUs) and the remote sites (hosting O-RUs). The permitted number switches in the synchronization path is limited by frequency and time error contributions by all clocks in the chain. Figure 11.2.2.3-1 shows a diagram of an LLS-C3 topology.

With Full Timing Support, all Ethernet switches in the fronthaul function as either T-BC (specified by Recommendation ITU-T G.8273.2 [28]) or a T-TC (specified by Recommendation ITU-T G.8273.3 [29]). The synchronization performance expectation is the same as described in Recommendation ITU-T G.8271.1 [23].

Partial Timing Support using non-T-BC switches may also be allowed. The PTP profile specified in Recommendation ITU-T G.8275.2 [32] shall be required for Partial Timing Support. Performance aspects and budgets associated with this mode requires further investigation.

Interconnection among switches and fabric topology (for example mesh, ring, tree, spur etc.) are deployment decisions which are out of the scope of the present document.



NOTE: This is for illustrative purpose and does not provide deployment guidance (for example the number of switches in network)

Figure 11.2.2.3-1: Configuration LLS-C3 synchronization (see below for additional information on O-DU, O-RU, switch functionality).

Considerations for O-DU: Unlike LLS-C1 & LLS-C2, O-DU does not act as synchronization master towards the fronthaul interface. It can select its own synchronization from local or remote PRTC like in LLS-C1/LLS-C2 but can also select the same synchronization master from the fronthaul as the O-RU and act as a Recommendation ITU-T G.8273.2 [28] T-TSC.

- 1) One possible LLS-C3 implementation consists of having one of the Recommendation ITU-T G.8273.2-compatible [28] T-BC of the chain being embedded in or co-located with the O-DU.

NOTE: An O-DU acting as a Recommendation ITU-T G.8273.2 [28] T-BC clock does not provide enough wander cleaning to guarantee the 15 ppb limit and thus is instead considered part of an LLS-C3 configuration.

- 2) Whatever synchronization configuration is used, the O-DU shall meet the time error limit for latency management as specified in clause 4.4.

Considerations for Switches: These are the same as for LLS-C2 in clause 11.2.2.2.

Considerations for O-RU: These are the same as for LLS-C1 and LLS-C2 in clause 11.2.2.2.

Considerations for PRTC/T-GM: Depending on the operator choice in deployment are the following:

- 1) The PRTC (typically a GNSS receiver) can be embedded either in an external T-GM connected to any T-BC, or an embedded function inside any T-BC in the network.

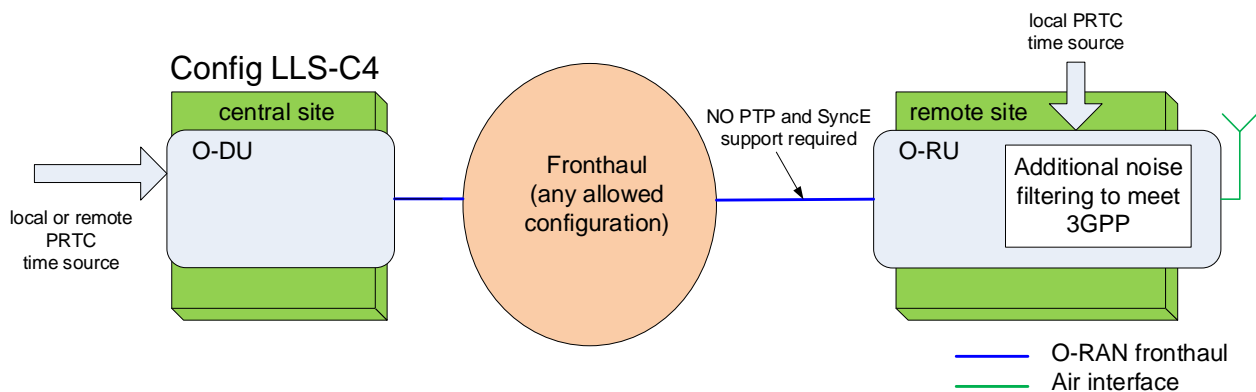
- 2) Multiple PRTC/T-GMs can offer redundancy. Recommendation ITU-T G.8275.1 BMCA shall be used to optimize the time distribution through the clock chain.

11.2.2.4 Topology configuration LLS-C4 synchronization

The network timing distribution is the preferred approach within the fronthaul network, however, there can be some deployment use cases that prevent the fronthaul network (or only a section of the network) from being upgraded to Recommendation ITU-T G.8271.1 [23] compliance and meeting the target performance at the O-RU. To cover these use cases, the synchronization strategy shall allow a local PRTC (typically GNSS, and typically meeting Recommendation ITU-T G.8272 clause 6 [25]) timing option at the O-RU. Figure 11.2.2.4-1 shows a diagram of an LLS-C4 topology.

Providing this type of timing support at O-RU requires extra timing interface or embedded timing function. Examples of implementation include:

- a local PRTC that provides time synchronization to the O-RU via a 1PPS/ToD interface.
- a local PRTC/T-GM that provides synchronization to the O-RU via a local PTP interface.
- a local PRTC embedded in the O-RU.



NOTE: For the O-RU shown above additional noise filtering is needed to filter sync source dynamic noise to meet 3GPP frequency accuracy requirement and satisfy target maxTAE.

Figure 11.2.2.4-1: Configuration LLS-C4 synchronization

Considerations for O-DU: Unlike within LLS-C1 & LLS-C2 configurations, O-DU does not act as synchronization master towards the fronthaul interface. It can select its own synchronization from local or remote PRTC like in LLS-C1/LLS-C2 but can also select the same synchronization master from the fronthaul.

The O-DU shall meet the time error limit for latency management as specified in clause 4.4.

11.2.3 Clock Synchronization

The following requirements are provided for clock synchronization:

- For Full Timing Support networks (either upstream O-DU or fronthaul between O-DU and O-RU), PLFS (typically SyncE) shall be used within the fronthaul network distribution per Recommendation ITU-T G.8271.1 [23], G.8275.1 [31] and G.8273.2 [28]. In an LLS-C1 configuration, the O-DU shall drive the PLFS regardless of its own selected time source. A derogation applies under the operator's responsibility in the specific case where O-RU does not make use of PLFS, in which case PLFS emission shall be optional for the O-DU. In an LLS-C2 configuration, the O-DU shall drive the PLFS regardless of its own selected time source. In an LLS-C3 configuration, the network shall deliver the PLFS to all switches and O-RUs. However, an O-DU or O-RU end application is not required to use a PLFS to achieve clock (frequency) synchronization. They can use PTP alone to achieve both frequency and phase/time synchronization. The ITU recommendations listed above refer to any PLFS and not just SyncE (for example SDH). However only SyncE has been fully

studied for the T-BC so far (Recommendation ITU-T G.8273.2 [28]), and SyncE is the most common implementation as it is in-band over Ethernet.

- For Partial Timing Support networks (either upstream of the O-DU or via the fronthaul between the O-DU and O-RU), PLFS may be used, as per Recommendation ITU-T G.8271.2 [24] and G.8275.2 [32]
- When used for clock synchronization, SyncE shall comply with the following ITU-T specifications:
 - Definitions: Recommendation ITU-T G.8260 [17].
 - Architecture: Recommendation ITU-T G.8261 [18].
 - SSM transport channel and format: Recommendation ITU-T G.8264 [21].
 - Clock specifications: Recommendation ITU-T G.8262 [19] (EEC) for the original, non-enhanced Synchronous Ethernet.
 - Clock specifications: Recommendation ITU-T G.8262.1 [21] (eEEEC) for the enhanced Synchronous Ethernet.
 - Functional model and SSM processing: Recommendation ITU-T G.781 [15].
- PLFS implementations other than SyncE are currently outside the scope of the present document. All network elements in the fronthaul network should use the same PLFS method for optimal interoperability and performance.
- Recommendation ITU-T G.8271.1 [23] shall be consulted for guidance on balancing the need to meet a target air interface performance at the O-RU with the maximum number of hops allowed. The network operator shall plan for specific fronthaul network deployment based on this guidance.
- The use of eEECs is recommended as they generate less noise (Time error, FFO) in normal operation and during network rearrangements than EECs, allowing a larger number of hops while meeting a given time error budget.

11.2.4 Profiles

11.2.4.1 Physical layer frequency signals (PLFS)

An implementation providing SyncE shall:

- Support Recommendation ITU-T G.781 [15] Option 1 Quality Level (per clause 5.4 Quality Code of Recommendation ITU-T G.781 [15])
 - Support Recommendation ITU-T G.781 [15] other options, are currently out of the scope of the present document. Recommendation ITU-T G.8271.1 [23] analysis has only been done with a Synchronous Ethernet network based on option 1 EECs or eEECs.
- Support Recommendation ITU-T G.8264 [21] message types, format, transmission, and reception (per Table 11.3 of Recommendation ITU-T G.8264 [21]) if SyncE is used.
 - Support of extended SSM TLV (per of Recommendation ITU-T G.8264 [21]) is optional.

Alternate PLFS implementations are out of scope of the present document.

NOTE: The accepted Quality Level may be limited to PRC or better value for optimal operation, while other values may be used under at the operator's discretion considering the HOLDOVER capability of the SyncE source and the way the end applications use SyncE.

11.2.4.2 PTP

11.2.4.2.1 Full timing support

Use of IEEE 1588 [3], [33] or PTP for time/phase synchronization shall be according to its clauses referred by Recommendation ITU-T G.8275.1 [31] (Full Timing Support).

- The T-TSC inside the O-RU and O-DU are considered as T-TSC inside 3GPP end application modules. Such T-TSC are not required to provide a 1PPS measurement interface, and Recommendation ITU-T G.8273.2 [28] Appendix IV applies: the combined performance within each module may not behave as a stand-alone T-TSC described in the normative section of the recommendation.
- PLFS and PTP distribution shall be supported (with derogations described in the "Additional information" column of the Synchronization Category in Table 10.2-1), either using congruent or non-congruent topologies, as per Recommendation ITU-T G.8271.1 [23].
- O-RUs shall support L2 (Ethernet) and multicast communication when acting as a Recommendation ITU-T G.8275.1 [31] PTP synchronization/timing receiver in a fronthaul network.
- O-RUs and O-DUs supporting LLS-C1, LLS-C2, and LLS-C3 shall follow Recommendation ITU-T G.8275.1 [31] clause 6, including requirements related to one-step and two-step clock mode.
- For configurations LLS-C1 and LLS-C2, the O-DU shall support L2 (Ethernet) and multicast communication when acting as a Recommendation ITU-T G.8275.1 [31] PTP master in a fronthaul network to synchronize the O-RUs.
- For all configurations, O-DU shall support L2 (Ethernet) and multicast communication when acting as a Recommendation ITU-T G.8275.1 [31] synchronization timing receiver in a fronthaul or other network is optional.

11.2.4.2.2 Partial timing support

Support of Partial Timing Support (PTS) using Recommendation ITU-T G.8275.2 [32] Telecom Profile is permissible but requires additional considerations:

- Partial Timing Support allows switches with no T-BC or T-TC, hence there is no guarantee of synchronization performance based on ITU-T standard specifications such as Recommendation ITU-T G.8273.2 [28]. As a result, the system operator shall ensure the network components will have adequate performance to meet frequency and phase error budgets to allow an accurate detection of frequency accuracy and phase for proper network operation.
- When implementing PTS, the relevant specifications are Recommendation ITU-T G.8273.4 for PTP aware clocks and Recommendation ITU-T G.8271.2 for network limits.
- O-RUs may support L3 (UDP/IP) and unicast communication when acting as a Recommendation ITU-T G.8275.2 [32] PTP synchronization timing receiver in a fronthaul network. This is considered "optional" in the present document.
- O-RUs and O-DUs supporting LLS-C1, LLS-C2, and LLS-C3 shall follow Recommendation ITU-T G.8275.2 [32] clause 6, including requirements related to one-step and two-step clock mode.
- For configurations LLS-C1 and LLS-C2, the O-DU may support L3 (UDP/IP) and unicast communication when acting as a Recommendation ITU-T G.8275.2 [32] PTP master in a fronthaul network to synchronize the O-RUs. This is considered "optional" in the present document.
- For all configurations, the O-DU may support L3 (UDP/IP) and unicast communication when acting as a Recommendation ITU-T G.8275.2 [32] PTP synchronization/timing receiver. This is considered "optional" in the present document.

Partial Timing Support has currently been defined only for relatively coarse timing accuracy (1.5 μ s). Using Partial Timing Support for S-Plane fronthaul distribution to the O-RU to satisfy tighter requirements than levels of accuracy 4 and 4A (as per Recommendation ITU-T G.8271 [22] Table 1) is currently out of scope of the present document.

11.2.5 Synchronization accuracy

11.2.5.1 Introduction

The parameters Time Error and other derived parameters are used in the subsequent clauses of this document. The definition of Time Error function, $TE(t)$ is given in clause 3.1. For a synchronized clock or timing signal, the Time Error function is composed of several different error components which contribute to the total, and these individual components of time error have limits which are specified in subsequent clauses. To clarify the specification tables later in this clause, definitions of these additional time error parameters are provided here.

- **$|TE(t)|$** : Absolute value of the Time Error with respect to an ideal time reference (recognized time standard). The limit of the time error, $\max|TE|$ (i.e. the furthest point away from zero, either positive or negative) is defined in Recommendation ITU-T G.8260 [17], clause 3.1.20.
- **Constant Time Error, cTE** : The mean value of the time error function, $TE(t)$, over a measurement period. See Recommendation ITU-T G.8260 [17], clause 3.1.20.
- **Dynamic Time Error, $dTE(t)$** : The change in the time error function, $TE(t)$, over a measurement period. See Recommendation ITU-T G.8260 [17], clause 3.1.20.
The limits on $dTE(t)$ in the present document are based on peak-to-peak values.
 $dTE(t)$ is by definition an unfiltered quantity.
When setting the limits, a filtered version of this signal may be considered. In particular $dTE(t)$ may also be divided into high and low frequency components by filtering:
 - **$dTE_L(t)$** : The dynamic time error after low-pass filtering.
For packet timing signals (e.g. PTP), the measurement filter bandwidth is typically 0.1 Hz, and all limits on $dTE_L(t)$ in the present document are defined in terms of MTIE (Maximum Time Interval Error, which is specified in ITU-T G.810 [16] and based on peak-to-peak values).
 - **$dTE_H(t)$** : The dynamic time error after high-pass filtering.
For packet timing signals (e.g. PTP), the measurement bandwidth is typically 0.1 Hz; all limits on $dTE_H(t)$ in the present document are defined in terms of peak-to-peak value.
- **$TE_L(t)$** : The slow changes in time error after low-pass filtering.
For packet timing signals (e.g. PTP), the measurement bandwidth is typically 0.1 Hz. The maximum absolute value of the low pass filtered time error function is denoted $\max|TE_L|$.

Further information about the accumulation of time error through a synchronization network can be found in Appendix IV of Recommendation ITU-T G.8271.1 [23].

The above defined time error parameters and their limit can be applied between different reference points.

These two parameters listed below define time error with respect to an ideal time reference.

- **$|TE|_{\text{antenna}}$** : Absolute value of the Time Error as measured at the antenna of the O-RU with respect to an ideal time reference (recognized time standard).
- **$|TE|_{\text{PRTC}}$** : Absolute value of the Time Error contributed by a PRTC as measured with respect to an ideal time reference (recognized time standard).

These three parameters listed below define the time error between two different branches of the front haul network.

- **$|TE_R|$** : Absolute value of the Relative Time Error between two network elements or between two master ports of a network element.
- **$|TE_{RL}|$** : Absolute value of the first order low pass filtered Relative Time Error between two network elements or between two master ports of a network element.

- **TAE:** Time Alignment Error defined as largest timing difference between any two different antenna connectors.

These four parameters listed below define time errors contributed from input to output of equipment or network.

- $|TE|_{O-RU}$: Absolute value of the Time Error contributed by the O-RU (measured from its network input to its antenna).
- $|TE_L|_{network}$: Absolute value of the first order low pass filtered Time Error generated by the network (measured from its input to its output).
- $|TE|_{network}$: Absolute value of the Time Error contributed by the front haul network (measured from its input to its output).
- $|TE_L|_{O-DU}$: Absolute value of the $|TE_L|$ contributed by the O-DU (measured from its network input to output).

11.2.5.2 Jitter

Within the O-RAN fronthaul network, all network equipment (NE) supporting SyncE transport across the network shall comply with input and output jitter requirements specified in Recommendation ITU-T G.8262 [19] clause 8.3.1 (for EEC) or Recommendation ITU-T G.8262.1 [20] clause 9.3.1 (for eEEEC).

Alternate PLFS implementations are currently out of scope of the present document.

11.2.5.3 Wander

Within the O-RAN fronthaul network, all network equipment (NE) supporting SyncE transport across the network shall comply with input and output wander requirements specified in Recommendation ITU-T G.8262 [19] clauses 8.1, 9.1, 9.2 (for EEC) or Recommendation ITU-T G.8262.1 [20] clauses 8.1, 9.1, 9.2 (for eEEEC).

Alternate PLFS implementations are out of scope of the present document.

11.2.5.4 Air interface frequency error

The O-RAN fronthaul network shall ensure O-RU meeting the 3GPP air interface frequency error requirement (i.e., 3GPP TS 36.104 clause 6.5.1 [9] for LTE and TS 38.104 clauses 6.5.1 and 9.6.1 [8] for 5G) specifying the short-term average frequency error in 1 ms duration.

NOTE: The present document derives the subsequent limits specified in clause 11 and in Annex H assuming the ± 50 ppb strictest limit from 3GPP (i.e., ± 50 ppb Wide Area Base Station limit).

Refer to clause 11.3.2 for more detailed information.

11.2.5.5 Air interface maximum time error

The O-RAN fronthaul network shall ensure O-RU meets the following air interface time alignment error ($|TAE|$ absolute or relative) requirements based on different features in LTE and 5G technologies. For features covered by 3GPP, they are specified in TS 36.104 clause 6.5.3 [9] for LTE and TS 38.104 clauses 6.5.3 and 9.6.3 [8] for 5G.

Figure 11.2.5.5-1 shows the reference points to define the network time error $|TE|$ vs air interface time alignment error (TAE) and the concept of relative vs absolute.

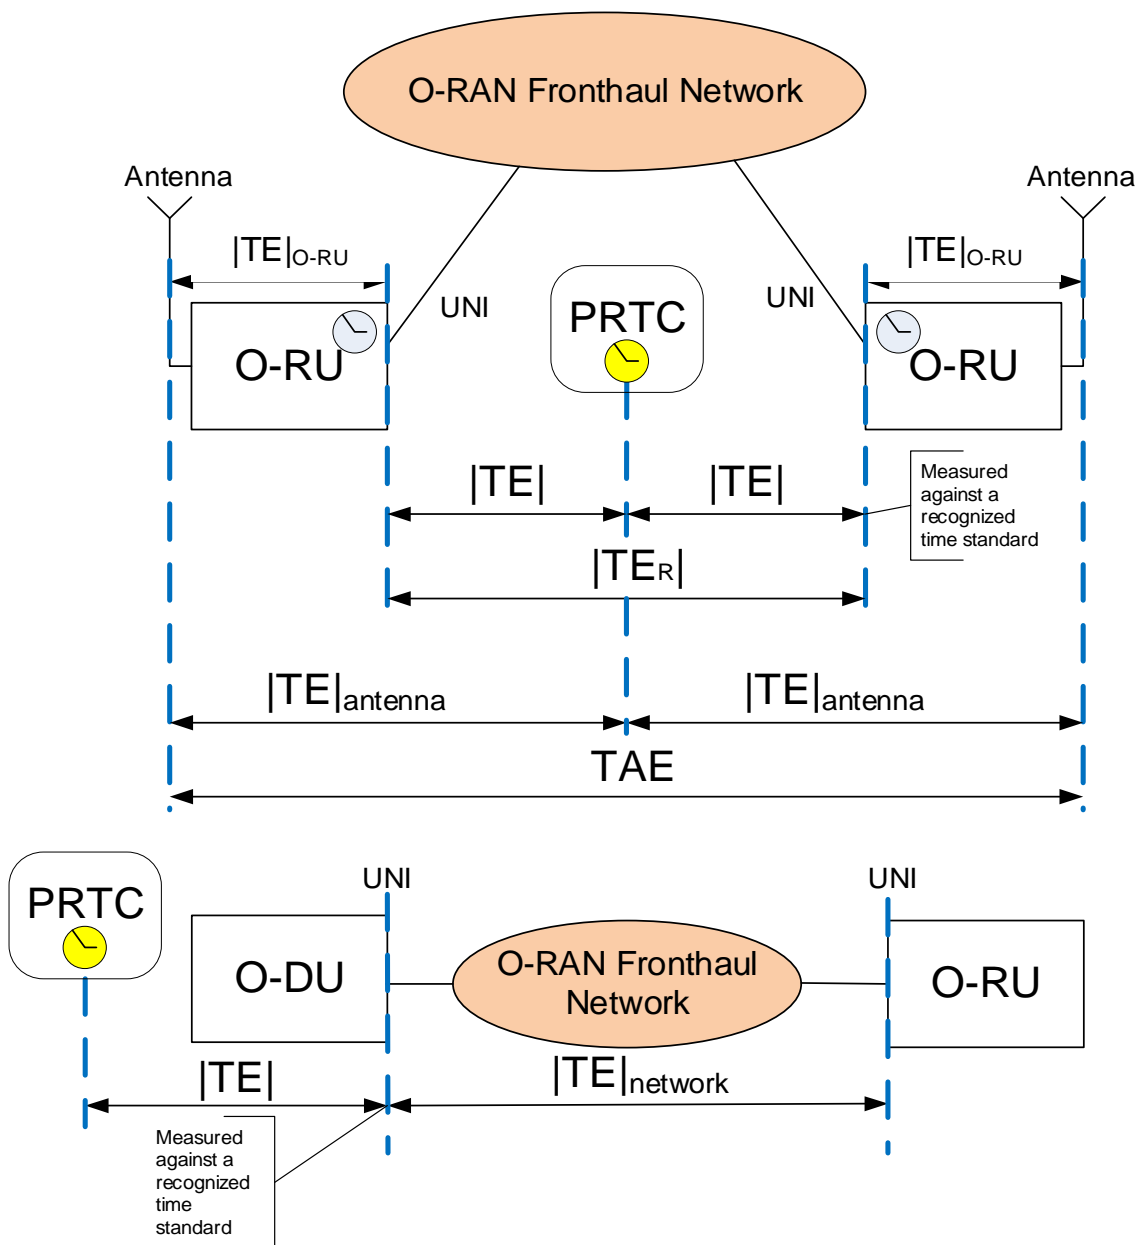


Figure 11.2.5.5-1: Definition of |TE| and TAE and UNI

Definitions for the timing parameters shown in Figure 11.2.5.5-1, are described in clause 11.2.5.1.

Table 11.2.5.5-1 shows the TAE requirements for LTE features:

Table 11.2.5.5-1: LTE features with time alignment error requirement at the air interface

LTE Features	TAE	Corresponding Timing Category in 802.1CM [11,12] and eCPRI [2] standards (Informative)
TDD	maxTAE $\leq 3 \mu\text{s}$ (cell radius < 3 km) TAE $\leq 10 \mu\text{s}$ (cell radius > 3 km) (NOTE 1) 3GPP TS 36.133 [38] clause 7.4	Timing Category C
Dual connectivity	maxTAE $\leq 3 \mu\text{s}$ (NOTE 2) 3GPP TS 36-133 [38] clause 7.4,7.13.2 and 7.15	Timing Category C
MIMO or Tx Diversity	maxTAE $\leq 65 \text{ ns}$ (NOTE 3, NOTE 5) 3GPP TS 36.104 [9] clause 6.5.3	
CA (intra-band contiguous)	maxTAE $\leq 130 \text{ ns}$ (NOTE 3, NOTE 6) 3GPP TS 36.104 [9] clause 6.5.3	Timing Category A (NOTE 4)
CA (inter-band or intra-band non-contiguous)	maxTAE $\leq 260 \text{ ns}$ (NOTE 3) 3GPP TS 36.104 [9] clause 6.5.3	Timing Category B (NOTE 4)
OTDOA	maxTE _{antenna} at O-RU antenna $\ll 1.5\mu\text{s}$, $\sim 100\text{-}200 \text{ ns}$ (not defined by 3GPP)	Not covered since it is not defined by 3GPP
<p>NOTE 1: For TDD the maxTAE requirement is applied to any pair of O-RUs that have overlapping coverage areas including when operated by different network operators. Hence maxTAE specification indirectly leads to a per O-RU requirement of maxTE_{antenna} $\leq 1.5 \mu\text{s}$.</p> <p>NOTE 2: In the case of Dual connectivity 3GPP has specified requirements in terms of MRTD at the input of the UEs. When specifying these requirements, it has been assumed a limit of 3 μs for maxTAE.</p> <p>NOTE 3: When these features are supported within 1 O-RU, TAE is impacted by O-RU internal TE only.</p> <p>NOTE 4: When these features are supported by multiple cooperating O-RUs, TAE is also impacted by TE_R _{network} where 802.1CM [11,12] clauses 6.4 Timing Category is applicable to limit the network error contribution.</p> <p>NOTE 5: In the present document, it is assumed that "MIMO or Tx Diversity" is applicable within a single O-RU. MIMO or Tx Diversity with multiple cooperating O-RUs is currently out of scope of the present document.</p> <p>NOTE 6: Applicable when the aggregated carriers are generated by co-located O-RUs (i.e., where the distance between the O-RUs does not create significant differential delay between the O-RUs and the UE, when compared with the timing requirement).</p>		

Table 11.2.5.5-2 shows the TAE requirements for 5G features:

Table 11.2.5.5-2: 5G features with time alignment error requirement

5G Features	TAE	Corresponding Timing Category in 802.1CM [11,12] and eCPRI [2] standards (Informative)
TDD	maxTAE $\leq 3 \mu\text{s}$ (NOTE 1) 3GPP TS 38.133 [37] clause 7.4	Timing Category C
Dual connectivity	maxTAE $\leq 3 \mu\text{s}$ (NOTE 2) 3GPP TS 38.133 [37] clause 7.6	Timing Category C
MIMO and Tx Diversity	maxTAE $\leq 65 \text{ ns}$ (NOTE 3, NOTE 5) 3GPP TS 38.104 [8] clause 6.5.3	

5G Features	TAE	Corresponding Timing Category in 802.1CM [11,12] and eCPRI [2] standards (Informative)
CA (intra-band contiguous in FR2)	$\max\text{TAE} \leq 130 \text{ ns}$ (NOTE 3, NOTE 6) 3GPP TS 38.104 [8] clause 9.6.3.	Timing Category A (NOTE 4)
CA (intra-band contiguous in FR1)	$\max\text{TAE} \leq 260 \text{ ns}$ (NOTE 3, NOTE 6) 3GPP TS 38.104 [8] clause 6.5.3, 9.6.3	Timing Category B (NOTE 4)
CA (intra-band non-contiguous in FR2)	$\max\text{TAE} \leq 260 \text{ ns}$ (NOTE 3, NOTE 6) 3GPP TS 38.104 [8] clause 9.6.3	Timing Category B (NOTE 4)
CA (intra-band non-contiguous in FR1)	$\max\text{TAE} \leq 3 \mu\text{s}$ (NOTE 3) 3GPP TS 38.104 [8] clause 6.5.3, 9.6.3	Timing Category C (NOTE 4)
CA (inter-band)	$\max\text{TAE} \leq 3 \mu\text{s}$ (NOTE 3) 3GPP TS 38.104 [8] clause 6.5.3, 9.6.3	Timing Category C (NOTE 4)
OTDOA	$\max \text{TE} _{\text{antenna}} \ll 1.5 \mu\text{s}$ (not defined by 3GPP)	Not covered since it is not defined by 3GPP
<p>NOTE 1: For TDD the maxTAE requirement is applied to any pair of O-RUs that have overlapping coverage areas including when operated by different network operators. Hence maxTAE specification indirectly leads to a per O-RU requirement of $\max \text{TE} _{\text{antenna}} = 1.5 \mu\text{s}$.</p> <p>NOTE 2: In the case of Dual connectivity 3GPP has specified requirements in terms of MRTD at the input of the UEs. When specifying these requirements, it has been assumed a limit of 3 us for maxTAE.</p> <p>NOTE 3: When these features are supported within 1 O-RU, maxTAE is impacted by O-RU internal TE only.</p> <p>NOTE 4: When these features are supported by multiple cooperating O-RUs, maxTAE is also impacted by $\text{TE}_R _{\text{network}}$ where 802.1CM [11, 12] clause 6.4 Timing Category is applicable to limit the network error contribution.</p> <p>NOTE 5: In the present document, it is assumed that "MIMO or Tx Diversity" is applicable within a single O-RU. MIMO or Tx Diversity with multiple cooperating O-RUs is out of scope of the present document.</p> <p>NOTE 6: Applicable when the aggregated carriers are generated by co-located O-RUs (i.e, where the distance between the O-RUs does not create significant differential delay between the O-RUs and the UE, when compared with the timing requirement).</p>		

11.3 Time and frequency synchronization requirements

11.3.1 Allowed PTP and PLFS clock types and clock classes

11.3.1.1 Overview

A network element (NE) may use the following clock types and classes to support PTP and PLFS and can be used among other such NEs to build an O-RAN-compliant fronthaul network meeting end-to-end frequency synchronization requirements as well as time synchronization requirements at the air interface.

- EEC (per Recommendation ITU-T G.8262 [19]).
- eEEEC (per Recommendation ITU-T G.8262.1 [20]).
- PRC (per Recommendation ITU-T G.811 [34]).
- PRTC (per Recommendation ITU-T G.8272 [25] and G.8272.1 [26]). NOTE 1
- T-BC (per Recommendation ITU-T G.8273.2 [28]).
- T-TSC (per Recommendation ITU-T G.8273.2 [28]).
- T-TC (per Recommendation ITU-T G.8273.3 [29]). NOTE 2

NOTE 1: Recommendation ITU-T G.8272 [25] specifies 2 types of PRTC: PRTC-A and PRTC-B. The time output of a PRTC-B is more accurate than that of a PRTC-A. PRTC-B is suitable for locations where it is possible to guarantee optimized environmental conditions (e.g., controlled temperature variation in indoor deployments). Typical examples are central location and large aggregation sites. This means that PRTC-B may be challenging and impractical for LLS-C4 deployments (e.g., due to generally outdoor deployments). Recommendation ITU-T G.8272.1 [26] specifies 2 types of ePRTC, ePRTC-A and ePRTC-B. The main difference between the PRTCs and the ePRTCs is the input from an external autonomous primary reference clock (e.g., cesium clock) required in case of ePRTCs. This means that ePRTCs are only suitable for deployment in centralized location (therefore not applicable as a solution for LLS-C4).

NOTE 2: According to Recommendation ITU-T G.8271.1 [23] and IEEE 802.1CM [11, 12], T-TC is allowed as T-BC replacement and T-TC/T-BC accuracy performance is equivalent. T-TC as accepted clock type is allowed. However, using T-TC has potential issue with broadcast storms in a bridging network as stated in Recommendation ITU-T G.8271.1 [23] and hence guidance by Appendix I in Recommendation ITU-T G.8275.1 [31] is applicable.

The time error analysis in Annex H only covers T-BC deployments, and not the T-TC deployments.

O-DU and O-RU are considered end points in an O-RAN-compliant fronthaul network. O-DU and O-RU can support specific PTP clock and specific classes based on O-DU and O-RU synchronization capability options.

Alternate PLFS implementations are out of scope of the present document.

11.3.1.2 O-DU synchronization capability options:

For configurations LLS-C1 and LLS-C2: PLFS and PTP Master.

- In case of local PRC/PRTC (typically a local GNSS receiver) used as frequency and time source: O-DU acts a PLFS and PTP master. The performance requirements are specified in clause 11.3.2.
- In case of a remote PRC/PRTC used as frequency and time source via a Full Timing Support packet network (defined by Recommendation ITU-T G.8271.1 [23]): O-DU shall act as an embedded end application with better wander and jitter filtering capability, as per Appendix IV of Recommendation ITU-T G.8273.2 [28]. Acting as a true Recommendation ITU-T G.8273.2 [28] T-BC does not guarantee the frequency accuracy required in clause 11.3.2 and should be considered analogous to the LLS-C3 configuration.
- In the case of a remote PRC/PRTC used as frequency and time source via a Partial Timing Support packet network (defined by Recommendation ITU-T G.8271.2 [24]): O-DU may act as an IWF P-F InterWorking Function to drive Full Timing Support (defined by Recommendation ITU-T G.8271.1 [23]) to the Fronthaul.

For both Full Timing Support and Partial Timing support, the remote PRC/PRTC used as frequency and time source can be located anywhere in the network.

For configuration LLS-C3 and LLS-C4:

- O-DU can be synchronized from any possible time source (either local or remote PRC/PRTC-traceable one, like for LLS-C1 and LLS-C2).
- O-DU does not need to meet the 3GPP frequency and TAE target specification as required in the O-RU. However, a more relaxed phase alignment between the O-DU and the O-RU timing should be kept to avoid data buffer overflow/underflow (impact to delay management topic and irrespective of the O-RU synchronization configuration). The O-DU shall meet the time error limit for latency management as specified in clause 4.4.

11.3.1.3 O-RU synchronization capability options:

- For configuration LLS-C1, LLS-C2 and LLS-C3 with full timing support, the T-TSC may be either embedded or external, as per Recommendation ITU-T G.8271.1 [23] and IEEE 802.1 CM [11,12]. In the case where an

O-RU includes a T-TSC, this one is "embedded in end application" as specified in Recommendation ITU-T G.8273.2 [28] Appendix IV and shall include additional necessary filtering function to ensure 3GPP air interface compliance. The requirements for Partial Timing Support are currently out of scope for the present document because there is not yet any available ITU recommendation.

- For configuration LLS-C4, the O-RU is simply synchronized from the local PRC/PRTC which may be using proprietary timing interface.

O-RU shall support network timing as mandatory synchronization capability to cover configuration LLS-C1, LLS-C2 and LLS-C3. Local PRC/PRTC (typically a GNSS receiver) may be a synchronization capability to cover configuration LLS-C4.

11.3.2 Frequency and time synchronization requirements across fronthaul network elements

11.3.2.1 Configurations LLS-C1 and LLS-C2

Based on IEEE 802.1CM [11, 12] and Recommendation ITU-T G.8271.1 [23] guidance, the following table summarizes the frequency and time error budgets across different elements of an O-RAN-compliant fronthaul network.

- O-DU: shall not exceed allocated frequency error budget and time error budget (for chosen air interface target)
- O-RU: shall not exceed allocated frequency error budget and time error budget (for chosen air interface target). The two O-RU classes described by IEEE 802.1CM [11, 12] clause 6 are deployment examples. The requirement shall be met only for O-RUs compliant with the limits of these classes.
- O-RAN fronthaul network: shall not exceed network limit to satisfy both frequency error budget and time error budget (for chosen air interface target). Allowed number of switches in a deployment can be derived based on allowed network limit vs chosen switch specification. Annex H shows the analysis of number of switches based on T-BC Class B and C switches.

Table 11.3.2.1-1 shows the budget allocation for configuration LLS-C1 and LLS-C2 (Refer to Figure 11.2.5.5-1 for reference point definition). The required specification limits in Table 11.3.2.1-1 are in **BOLD text**.

Table 11.3.2.1-1: Frequency and time error budget allocation (for topology configuration LLS-C1 and LLS-C2)

Frequency error budget allocation			
Timing Reference	O-RAN fronthaul network contribution limit	O-RU	Air interface target
O-DU PTP/PLFS master class A, with max freq error ≤ 15 ppb <ul style="list-style-type: none"> • @ O-DU UNI (see NOTES 1, 8, 9) 	max(dTE) ≤ 90ns <ul style="list-style-type: none"> • Between O-DU UNI and O-RU UNI • For LLS-C2: Allowed number of hops (see NOTE 2) • For LLS-C1: single hop by definition 	O-RU max freq error ≤ 35 ppb including both <ul style="list-style-type: none"> • FFO after O-RU filtering of dTE @ O-RU UNI • O-RU internal additive frequency error (see NOTE 3) 	max frequency error ≤ 50 ppb (NOTE 11)
O-DU PTP/PLFS master class B, with max freq error ≤ 5 ppb <ul style="list-style-type: none"> • @ O-DU UNI (see NOTES 1, 8, 9) 	max(dTE) ≤ 114ns <ul style="list-style-type: none"> • Between O-DU UNI and O-RU UNI • For LLS-C2: Allowed number of hops (see NOTE 2) • For LLS-C1: single hop by definition 	O-RU max freq error ≤ 45 ppb including both <ul style="list-style-type: none"> • FFO after O-RU filtering of dTE @ O-RU UNI • O-RU internal additive frequency error (see NOTE 3) 	max frequency error ≤ 50 ppb (NOTE 11)

Time error budget allocation			
Timing Reference	O-RAN fronthaul network contribution limit	O-RU (All allowed class options are shown) (see NOTE 6)	Air interface target
No $\max TE_{RL} $ contribution by O-DU since O-DU is common PTP and PLFS master to all co-operated O-RU (NOTE 1)	$\max TE_{RL} \leq 60 \text{ ns}$ <ul style="list-style-type: none"> Between 2 O-RUs UNI per IEEE 802.1CM [11] For LLS-C2: Allowed number of hops (see NOTE 4) For LLS-C1: single hop by definition 	$\max TE _{O-RU}$ includes as per IEEE 802.1CM [11] (see NOTE 5): <ul style="list-style-type: none"> (enhanced O-RU can be used) $\max TE _{O-RU} \leq 35 \text{ ns}$	$\max TAE \leq 130 \text{ ns}$ Category A (NOTE 12)
	$\max TE_{RL} \leq 100 \text{ ns}$ (using regular O-RU) $\max TE_{RL} \leq 190 \text{ ns}$ (using enhanced O-RU) <ul style="list-style-type: none"> Between 2 O-RUs UNI per IEEE 802.1CM [11] For LLS-C2: Allowed number of hops (see NOTE 4) For LLS-C1: single hop by definition 	$\max TE _{O-RU}$ includes as per IEEE 802.1CM [11] (see NOTE 5): <ul style="list-style-type: none"> Either regular O-RU $\max TE _{O-RU} \leq 80 \text{ ns}$ Or enhanced O-RU $\max TE _{O-RU} \leq 35 \text{ ns}$ 	$\max TAE \leq 260 \text{ ns}$ Category B (NOTE 12)
For LLS-C2: $\max TE_L \leq 1.325 \mu\text{s}$ For LLS-C1: $\max TE_L \leq 1.420 \mu\text{s}$ <ul style="list-style-type: none"> @ O-DU UNI includes HOLDOVER budget see NOTES 8, 9, 10 	$\max TE_L _{\text{network}} \leq 95 \text{ ns}$ (using regular O-RU) $\max TE_L _{\text{network}} \leq 140 \text{ ns}$ (using enhanced O-RU) <ul style="list-style-type: none"> Between O-DU UNI and O-RU UNI For LLS-C2: Allowed number of hops (see NOTE 4) For LLS-C1: single hop by definition 	$\max TE _{O-RU}$ includes as per IEEE 802.1CM [11] (see NOTE 5 and 7): <ul style="list-style-type: none"> Either regular $\max TE _{O-RU} \leq 80 \text{ ns}$ Or enhanced $\max TE _{O-RU} \leq 35 \text{ ns}$ 	$\max TE _{\text{antenna}} \leq 1.5 \mu\text{s}$ Category C (NOTE 12)

- NOTE 1: O-DU implements PTP and PLFS master function in this use case.
Two frequency error limits are defined in the present document, conservative class A with 15 ppb, and more advanced class B with 5 ppb.
Both frequency error limits and $\max|\text{TE}_L|$ time error limits are measured on the O-DU UNI after applying a first-order measurement low-pass filter bandwidth of 0.1 Hz to the time samples.
Measurement condition shall be applicable when O-DU Master Clock is either in locked state or HOLDOVER but excluding rare and temporary transients:
- Resynchronization to recovered source after HOLDOVER.
 - Resynchronization to newly selected source after failure of the previous one, in case of redundancy.
- It is considered that all master ports of the O-DU are fully synchronized together, and there is no port-to-port time error (see Annex H).
- NOTE 2: $d\text{TE}$ is the accumulated dynamic time error of a T-BC clock chain (excluding O-DU contribution) based on Recommendation ITU-T G.8271.1 [23] Appendix IV calculation method and considering the noise to be centered. For the lower frequency range, which depends on O-RU vendor-specific low pass filtering characteristics, a higher limit may be allowed. As per clause 11.2.5.1 of the present document this is peak-to-peak value. Refer to Annex H for detailed analysis for maximum number of T-BC Class B and C switches.
- NOTE 3: The O-RU requirement is not specified in IEEE 802.1CM [11,12] or Recommendation ITU-T G.8271.1 [23]. Only the total O-RU frequency error budget is specified to allow design flexibility by different solution vendors. The O-RU solution vendor can make a tradeoff between FFO (Fractional Frequency Offset after applying O-RU filtering) and internal additive frequency error as long as the total frequency error budget with network total $\max(d\text{TE})$ limit is met. Refer to Annex H for analysis.
- NOTE 4: Refer to Annex H for analysis of the number of switches to satisfy the allowed network limit. For the Time Error budget (both network and relative time errors) it is considered that the O-RU attenuates the $d\text{TE}_H$ part of the noise present at its input, and therefore the budget uses $\max|\text{TE}_L|$ using a 0.1 Hz low-pass filter as described in Recommendation ITU-T G.8271.1 [23].
- NOTE 5: The T-TSC are considered T-TSC embedded in end application as specified in Recommendation ITU-T G.8273.2 [28], Appendix IV and the performance may not behave as a standalone T-TSC described in the normative part of the Recommendation ITU-T G.8273.2 [28].
However, to ensure interoperability among O-RAN O-RU vendors, the agreed performance shall be based for O-RUs compliant with the classes defined in eCPRI and 802.1CM [11, 12] recommendation: under eCPRI Transport Network Requirements Specification [39] clause 4 and IEEE 802.1CM [11] clause 6, a regular O-RU with $\max|\text{TE}| = 80$ ns and an enhanced O-RU with $\max|\text{TE}| = 35$ ns. The O-RU time error accumulates linearly with the other contributors in the network.
- NOTE 6: $\max|\text{TE}|_{\text{O-RU}}$ is the total budget proposed in IEEE 802.1CM [11] clause 6 for two examples of O-RU classes. O-RU internal split as shown in the table is recommended based on eCPRI and 802.1CM guidance. It is also acceptable for RAN solution vendor to use a different O-RU internal split to meet the same $\max|\text{TE}|_{\text{O-RU}}$ total budget.
- NOTE 7: Meeting the specified budget split between O-DU and $\max|\text{TE}|_{\text{O-RU}}$ across all categories is recommended based on eCPRI [2] clause 6.2 and 802.1CM [11] clause 6 guidance. A different budget split than what is recommended may meet the overall radio interface requirements but shall be considered under the operator's responsibility. Typically, configuration LLS-C1 and LLS-C2 with category C target may allow more flexibility in $\max|\text{TE}|_{\text{O-DU}}$, $\max|\text{TE}_L|_{\text{network}}$ and $\max|\text{TE}|_{\text{O-RU}}$ budget split as long as the network limit ($\max|\text{TE}_L|_{\text{network}}$ for time error and $d\text{TE}$ for frequency error) is still respected. For example, if an operator can guarantee the $\max|\text{TE}_L|_{\text{network}} < 75$ ns which is below the 95 ns limit set in the present document, 20 ns (95 ns-75 ns) can be moved into the $\max|\text{TE}|_{\text{O-DU}}$ and/or $\max|\text{TE}|_{\text{O-RU}}$ budget. Inter-operability can be guaranteed by the network operator when all participating O-DU and O-RU vendors design to meet the operator-chosen budget split.
- NOTE 8: In case the O-DU is synchronized using a local PRTC (typically a GNSS receiver or local T-GM), the time error limit at the input of the O-DU is specified as being Reference point A (or B) defined by Recommendation ITU-T G.8271 [22], Recommendation ITU-T G.8271.1 [23] and G.8271.2 [24]. The O-DU acting as a T-GM in normal operating mode may meet Time Error limits slightly larger than what is specified in Recommendation ITU-T G.8272 [25] and is expected to be significantly lower than 1420 ns or 1325 ns.
- NOTE 9: In case the O-DU is synchronized using network distribution (typically PTP and PLFS), the network time error limit at the input of the O-DU is specified as being Reference point C defined by Recommendation ITU-T G.8271 [22], Recommendation ITU-T G.8271.1 [23] for Full Timing Support or Recommendation ITU-T G.8271.2 [24] for Partial Timing Support. Similar to NOTE 4, the relevant limit of the O-DU is the 0.1 Hz low-pass filtered $\max|\text{TE}_L|$, which implies that the O-DU low-pass filtering bandwidth is lower than 0.1 Hz (per clause 11.2.2.2 of the present document).
- NOTE 10: The maximum relative time error of the S-Plane measurement signals between any two O-DUs shall be twice the maximum absolute time error requirement for an O-DU at O-DU UNI provided they share same timescale (i.e., PTP) to meet the requirements of clause 11.7 of the present document.
- NOTE 11: Observed over 1 ms as per 3GPP Spec TS 36.104 [9] clause 6.5.1, and 38.104 [8] clauses 6.5.1 and 9.6.1, Wide Area BS class
- NOTE 12: As per IEEE 802.1CM [11] clause 6.4.

The synchronization accuracy for the OTDOA feature is not defined by 3GPP. Hence $\max|\text{TE}|_{\text{antenna}}$ and the corresponding $\max|\text{TE}_L|_{\text{network}}$ are out of scope of this edition of the present document and are not covered in this analysis. The two O-RU $\max|\text{TE}|$ classes indicated earlier may be considered also for this use case. In a deployed network, an operator can choose a target $\max|\text{TAE}|_{\text{antenna}}$ and then derive the corresponding $\max|\text{TE}_L|_{\text{network}}$ and the

allowed number of hops and required types of T-BC(s)/T-TC(s) based on Recommendation ITU-T G.8271.1 [23] guidance.

With Partial Timing Support, when using non-T-BC switches, network contribution limit requires further investigation.

11.3.2.2 Configuration LLS-C3

Based on IEEE 802.1CM [11, 12] and Recommendation ITU-T G.8271.1 [23] guidance, the following table summarizes the frequency and time error budgets across different elements of an O-RAN-compliant fronthaul network.

- PRTC/T-GM: shall not exceed allocated frequency error budget and time error budget (for chosen air interface target).
- O-RU: shall not exceed allocated frequency error budget and time error budget (for chosen air interface target). The two O-RU classes described by IEEE 802.1CM [11, 12] clause 6.4 are deployment examples. The requirement shall be met only for O-RUs compliant with the limits of these classes.
- O-RAN fronthaul network: It shall not exceed network limit to satisfy both frequency error budget and time error budget (for chosen air interface target). The network limits for time error, applicable at the input of the O-RU, are those defined in clause 7.3 and 7.5 of Recommendation ITU-T G.8271.1 [23] (Network limits at reference point C and for deployment cases with the PRTC deployed in the access network; only one of these would be applicable depending on the specific deployment), where the O-RU corresponds to the End Application with integrated PTP clock (Deployment Case 1, see Figure 7-1 in Recommendation ITU-T G.8271.1 [23]). The limits in Recommendation ITU-T G.8271.1 [23] are expressed in terms of MTIE, $\max|TE_L|$ and peak-to-peak $\max(dTE_H)$. Applicability of Recommendation ITU-T G.8271.1 [23] peak-to-peak $\max(dTE_H)$ limit is not defined in the present document. Further details and references are provided in Annex H. The requirements provided in this clause are aligned with IEEE 802.1CM [11].
- O-DU output measurement signal (1PPS): There is no such allocated frequency error budget as for LLS-C1/C2. Only $\max|TE|$ shall be within 1500 ns limits as specified in clause 4.3 of the present document. O-DU and O-RU can be traceable to a different PRTCs, as long as they share the same timescale. In case the O-DU is synchronized using a local PRTC (typically a GNSS receiver or local T-GM), the time error limit at the input of the O-DU $\max|TE_L|_{O-DU}$ shall be specified as being the limit at Reference point A (or B) defined by Recommendation ITU-T G.8271.1 [23] clause 7 and Recommendation ITU-T G.8271.2 [24] clause 7. In case the O-DU is synchronized using network distribution (typically PTP and PLFS), the network time error limit $\max|TE_L|_{network}$ at the input of the O-DU shall be specified as being the limit at Reference point C defined by Recommendation ITU-T G.8271.1 [23] clause 7 for Full Timing Support or Recommendation ITU-T G.8271.2 [24] clause 7 for Partial Timing Support.
- O-DU: shall not exceed allocated time error budget required by latency management as specified in clause 4.3 of the present document.

Table 11.3.2.2-1 shows the budget allocation for configuration LLS-C3 (Refer to Figure 11.2.5.5-1 for reference point definition). The required specification limits in Table 11.3.2.2-1 are in **BOLD text**.

Table 11.3.2.2-1: Frequency and time error budget allocation (for topology configuration LLS-C3)

Frequency error budget allocation			
Timing Reference	O-RAN fronthaul network contribution limit	O-RU	Air interface target
PRTC/T-GM $\max freq\ error \leq 2\ ppb$ • @ PRTC/T-GM UNI (see NOTE 1)	$\max(dTE) \leq 126ns$ (see NOTE 4) • Between PRTC output and O-RU UNI • Include PRTC/T-GM MTIE contribution • Allowed number of hops (see NOTE 2)	O-RU $\max freq\ error \leq 48\ ppb$ including both • FFO after O-RU filtering of dTE @ O-RU UNI • O-RU internal additive frequency error (see NOTE 3)	$\max frequency\ error \leq 50\ ppb$ (NOTE 12)

Time error budget allocation			
Timing Reference	O-RAN fronthaul network contribution limit	O-RU (All allowed class options are shown) (see NOTE 6) O-DU (see NOTE 8)	Air interface target
No $ TE_{RL} $ contribution by PRTC/T-GM as long as PRTC/T-GM is common PTP and PLFS master to all co-operating O-RU (see NOTE 1)	$\max TE_{RL} \leq 60$ ns <ul style="list-style-type: none"> Between 2 O-RUs UNI per IEEE 802.1CM [11] clause 6.4 Allowed number of hops (see NOTE 5) 	$\max TE _{O-RU}$ includes as per IEEE 802.1CM [11] clause 6.4: (Enhanced O-RU can be used) <ul style="list-style-type: none"> $\max TE _{O-RU} \leq 35$ ns 	$\max TAE \leq 130$ ns Category A (NOTE 13)
	$\max TE_{RL} \leq 100$ ns (for regular O-RU) $\max TE_{RL} \leq 190$ ns (for enhanced O-RU) <ul style="list-style-type: none"> Between 2 O-RUs UNI per IEEE 802.1CM [11] clause 6.4 Allowed number of hops (see NOTE 5) 	$\max TE _{O-RU}$ includes as per IEEE 802.1CM [11] clause 6.4 (see NOTE 6): <ul style="list-style-type: none"> Either regular O-RU: $\max TE _{O-RU} \leq 80$ ns Or enhanced O-RU: $\max TE _{O-RU} \leq 35$ ns 	$\max TAE \leq 260$ ns Category B (NOTE 13)
$\max TE _{PRTC} \leq 100$ ns or 40 ns (see NOTE 9) <ul style="list-style-type: none"> @ PRTC/T-GM UNI PRTC/T-GM spec per Recommendation ITU-T G.8272 (see NOTE 1) Not including HOLDOVER budget 	<ul style="list-style-type: none"> $\max TE _{network} \leq 1100$ ns at the O-RU UNI (see NOTE 7, NOTE 10) Short Clock Chain $\max TE _{network} \leq 100$ ns at the O-RU UNI (see NOTE 10, NOTE 11, NOTE 14) Allowed number of hops (see NOTE 5) 	$\max TE _{O-RU}$ includes as per IEEE 802.1CM [11] clause 6.4 (see NOTE 6): <ul style="list-style-type: none"> Either regular O-RU: $\max TE _{O-RU} \leq 80$ ns Or enhanced O-RU: $\max TE _{O-RU} \leq 35$ ns 	$\max TE _{antenna} \leq 1.5$ μ s Category C (NOTE 13)
$\max TE _{PRTC} \leq 100$ ns or 40 ns (see NOTE 9) <ul style="list-style-type: none"> @ PRTC/T-GM UNI PRTC/T-GM spec per Recommendation ITU-T G.8272 (see NOTE 1) Not including holdover budget	$\max TE _{network}$ is out of scope <ul style="list-style-type: none"> O-DU may get its synchronization from either local or remote PRTC 	$\max TE _{O-DU}$ at O-DU UNI < 1.5 μs (see NOTE 8)	$\max TE _{antenna} \leq 1.5$ μ s Category C (NOTE 13)

- NOTE 1: PRTC/T-GM follows Recommendation ITU-T G.8272 [25] specification for PRTC-A and PRTC-B, or Recommendation ITU-T G.8272.1 [26] for ePRTC. MTIE specification considers dynamic error during lock condition and its contribution is covered under dTE in NOTE 2. However, there is possible semi-static frequency error which is not part of MTIE specification. Therefore, it is included here. Both Recommendation ITU-T G.8272 [25] and G.8272.1 [26] specify that for the combined PRTC and T-GM function, time error samples are measured through a moving-average low-pass filter of at least 100 consecutive time error samples. Assuming a message rate of 16 PTP messages per second as described in Recommendation ITU-T G.8275.1, the low-pass filtering bandwidth is 0.05 Hz, which is not same as a 0.1 Hz low-pass-filtered $\max|TE_L|$. (The frequency response of a moving average filter is approximately similar to a first-order linear filter of bandwidth $1/(\pi T)$, where T is the width of the moving average window). In case the PRTC/T-GM is a multiple-port device, it is considered that all master ports are fully synchronized together, and there is no port-to-port time error. Refer to Annex H for detailed analysis.
- NOTE 2: dTE is the accumulated dynamic time error of a T-BC clock chain, including PRTC/T-GM contribution, to O-RU UNI (or O-DU UNI) based on Recommendation ITU-T G.8271.1 [23] Appendix IV calculation method, and considering the noise to be centered. For the lower frequency range, which depends on O-RU vendor-specific low pass filtering characteristics, a higher limit may be allowed (still within the limits specified by Recommendation ITU-T G.8271.1 [23] clause 7.3 as worst case). As per clause 11.2.5.1 of the present document this is peak-to-peak value. Refer to Annex H for analysis for maximum number of T-BC Class B and C switches.
- NOTE 3: The O-RU requirement is not specified in IEEE 802.1CM [11] or Recommendation ITU-T G.8271.1 [23]. Only the total O-RU frequency error budget is specified to allow design flexibility by different solution vendors. The O-RU solution vendor can make a tradeoff between FFO (Fractional Frequency Offset after applying O-RU filtering) and O-RU internal additive frequency error as long as the total frequency error budget with network total $\max(dTE)$ limit shall be met. Refer to Annex H for analysis.
- NOTE 4: The O-RAN fronthaul network contribution limits are stricter and therefore not consistent with peak-to-peak $\max(dTE_H)$ limits defined in Recommendation ITU-T G.8271.1 [23] for long synchronization chain (see Recommendation ITU-T G.8271.1 [23], as specified in clause 7.3). In order to tolerate the higher Recommendation ITU-T G.8271.1 [23] limits (applicable to a long network) the O-RU would be required to provide higher filtering capabilities. This is an option that may be considered but not further described in the present document. This requirement assumes shorter clock chains.
- NOTE 5: Refer to Annex H for analysis of the number of switches to satisfy the allowed network limit.
- NOTE 6: The T-TSC are considered T-TSC embedded in end application as specified in Recommendation ITU-T G.8273.2 [28], Appendix IV and the performance may not behave as a standalone T-TSC described in the normative section of the Recommendation ITU-T G.8273.2 [28]. However, to ensure interoperability among O-RAN O-RU vendors, the agreed performance shall be based for O-RUs compliant with the classes defined in eCPRI [2] and 802.1CM [11, 12] recommendation: under eCPRI Transport Network Requirements Specification [39] and IEEE 802.1CM [11], a regular O-RU with $\max|TE| = 80$ ns and an enhanced O-RU with $\max|TE| = 35$ ns. The O-RU time error accumulates linearly with the other contributors in the network.
- NOTE 7: $\max|TE_L|_{\text{network}} \leq 1100$ ns is specified by Recommendation ITU-T G.8271.1 [23] in clause 7.3 for the $\max|TE_L|$. An O-RU may consume only 35 ns (enhanced O-RU) or 80 ns (regular O-RU). $\max|TE_L| \leq 1100$ ns may leave extra budget that could be allocated for instance to holdover that could be especially relevant in case of TDD applications. This additional budget would be added on top of the 35 ns (enhanced O-RU) or 80 ns (regular O-RU), This implies that during sync failure conditions, O-RU TAE requirements for Category B may exceed. It is a network operator decision whether this extra budget can be allocated or not.
- NOTE 8: $\max|TE_R|$ between O-DU and O-RU limit (as specified in clause 4.4.1) translates into $\max|TE|$ for both O-DU and O-RU with half the budget. This shall be measured at any available O-DU output signal (either 1PPS or O-DU UNI).
- NOTE 9: PRTC-B ($\max|TE| = 40$ ns) is assumed in case of short clock chains by Recommendation ITU-T G.8271.1 [23].
- NOTE 10: The contribution from the network includes the noise generated by the PRTC/T-GM.
- NOTE 11: The network limits for short clock chain networks are specified in clause 7.5 of Recommendation ITU-T G.8271.1 [23].
- NOTE 12: Observed over 1 ms as per 3GPP Spec TS 36.104 [9] clause 6.5.1, and 38.104 [8] clauses 6.5.1 and 9.6.1, Wide Area BS class
- NOTE 13: As per IEEE 802.1CM [11] clause 6.4.
- NOTE 14: $\max|TE_L|_{\text{network}} \leq 1100$ ns at the input of the O-RU is based on assumptions made by Recommendation ITU-T G.8271.1 [23] and IEEE802.1CM [11] in terms of clock types and maximum number of cascaded clocks. This limit and the dimensioning rules are common to the networks delivering time to O-DU in all LLS configurations and to the O-RU in LLS-C3. As a comparison, in LLS-C2, due to the additional time error budget allocated to the O-DU and the fronthaul network clocks between the O-DU and O-RU, the network limit at the input of the O-RU is specified as $\max|TE_L|_{\text{network}} \leq 1420$ ns as indicated in Table 11.3.2.1-1 of clause 11.3.2.1. In principle even in LLS-C3 configuration, a network delivering $\max|TE_L|_{\text{network}} > 1100$ ns and up to 1420 ns at the O-RU input could still meet Category C and Category B requirements (assuming no extra budget to address failure conditions is required as part of the O-RU). This option is outside the scope of this specification.

The synchronization accuracy for the OTDOA feature is not defined by 3GPP. Hence $\max|\text{TE}|$ at O-RU antenna and the corresponding $\max|\text{TE}_L|_{\text{network}}$ are out of scope of the present document and therefore not covered in this analysis. The two $\max|\text{TE}|_{\text{O-RU}}$ classes indicated earlier may be considered also for this use case. In a deployed network, operator can choose a target $\max|\text{TE}|$ and then derive the corresponding $\max|\text{TE}_L|_{\text{network}}$ and the allowed number of hops and required types of T-BC(s)/T-TC(s) based on Recommendation ITU-T G.8271.1 [23] guidance.

With Partial Timing Support, when using non-T-BC switches, network contribution limit requires further investigation.

11.3.2.3 Configuration LLS-C4

Table 11.3.2.3-1 shows the frequency and time error budgets across different elements for LLS-C4.

- O-DU output measurement signal (1PPS): there is no such allocated frequency error budget as there is for LLS-C1/C2. Only time error shall be within ± 1500 ns limits as specified in clause 4.3.
In case the O-DU is synchronized using a local PRTC (typically a GNSS receiver or local T-GM), the time error limit $\max|\text{TE}_L|$ at the input of the O-DU shall be specified as being the limit at Reference point A (or B) defined by Recommendation ITU-T G.8271.1 [23] clause 7 and G.8271.2 [24] clause 7.
In case the O-DU is synchronized using network distribution (typically PTP and PLFS), the network time error limit $\max|\text{TE}_L|_{\text{network}}$ at the input of the O-DU shall be specified as being Reference point C defined by Recommendation ITU-T G.8271.1 [23] for Full Timing Support or Recommendation ITU-T G.8271.2 [24] for Partial Timing Support.
- O-RU: shall not exceed allocated frequency error budget and time error budget (for chosen air interface target). The two O-RU classes described by IEEE 802.1CM [11, 12] clause 6.4 are deployment examples. The requirement shall be met only for O-RUs compliant with the limits of these classes.

Table 11.3.2.3-1 shows the budget allocation for configuration LLS-C4. The required specification limits in Table 11.3.2.3-1 are in **BOLD text**.

Table 11.3.2.3-1: Frequency and time error budget allocation (for topology configuration LLS-C4)

Frequency error budget allocation			
Timing Reference in network or O-DU	O-RAN fronthaul network	O-RU	Air interface target
No dependency	No dependency	O-RU max freq error ≤ 50 ppb including both <ul style="list-style-type: none"> O-RU filtering on local PRTC noise O-RU internal additive frequency error (NOTE 3, 4) 	max frequency error ≤ 50 ppb (NOTE 6)
Time error budget allocation			
Timing Reference in network or O-DU	O-RAN fronthaul network	O-RU (See NOTE 1,2, 3, 4) O-DU (see NOTE 5)	Air interface target
No dependency	No dependency	Local time source: max TE ≤ 100 ns (PRTC-A), max TE ≤ 40 ns (PRTC-B) max TE _{O-RU} ≤ 30 ns	maxTAE ≤ 260ns Category B (NOTE 7)
		local time source: max TE ≤ 100 ns max TE _{O-RU} ≤ 1.4 μs including any holdover budget	max TE _{antenna} ≤ 1.5 μs Category C(NOTE 7)
max TE _{PRTC} ≤ 100 ns or 40 ns (see NOTE 1) <ul style="list-style-type: none"> @ PRTC/T-GM UNI PRTC/T-GM spec per Recommendation ITU-T G.8272 Not including holdover budget	max TE _{network} is out of scope <ul style="list-style-type: none"> O-DU may get its synchronization from either local or remote PRTC 	max TE _{O-DU} at O-DU UNI < 1.5 μs (NOTE 5)	max TE _{antenna} ≤ 1.5 μs Category C(NOTE 7)
NOTE 1: Local PRTC (typically GNSS receiver) to O-RU. Therefore, budget shall be based on Recommendation ITU-T G.8272 [25] (PRTC class A, with max TE = 100 ns, or PRTC class B, with max TE = 40 ns) on each GNSS receiver. Because there is no requirement in Recommendation ITU-T G.8272 [25] on the relative time error between two neighboring local PRTCs (GNSS receivers), it is not possible to remove any common time error in the budget at the O-RUs side.			
NOTE 2: When there is a need for low TAE, LLS-C4 configuration shall include 2x local PRTC-A max TE = 200 ns total, or 2x local PRTC-B max TE = 80 ns total. This disadvantage automatically prevents LLS-C4 configuration from meeting certain demanding maxTAE feature (as target by 802.1CM [11] clause 6.4 Category A). Hence, only 802.1CM [11] clause 6.4 Category B/C are shown.			
NOTE 3: max TE _{O-RU} is not governed by eCPRI [2], 802.1CM [11] or ITU-T standards since LLS-C4 configuration is not based on network timing solution. max TE _{O-RU} budget is basically the remaining budget to satisfy target feature TAE after excluding the local max TE _{PRTC} contribution. max TE _{O-RU} includes the O-RU clock recovery (i.e. deriving a clean clock from local time source) error and any O-RU internal error.			
NOTE 4: This O-RU option requires extra interface and extra hardware support including local PRTC (typically GNSS receiver and antenna) and likely a more expensive oscillator for noise filtering. Standard O-RU with network timing support (target for configuration LLS-C1, LLS-C2 and LLS-C3) cannot offer this option in general. A specific O-RU design is needed.			
NOTE 5: max TE _R between O-DU and O-RU limit (as recommended in clause 4.4.1) translates into max TE for both O-DU and O-RU with half the budget. This is measured at any available O-DU output signal (either 1PPS or O-DU UNI).			
NOTE 6: Observed over 1 ms as per 3GPP Spec TS 36.104 [9] clause 6.5.1, and 38.104 [8] clauses 6.5.1 and 9.6.1, Wide Area BS class			
NOTE 7: As per IEEE 802.1CM [11] clause 6.4.			

The synchronization accuracy for the OTDOA feature is not defined by 3GPP. Hence max|TE|_{antenna} at O-RU antenna and the corresponding max|TE|_{O-RU} are out of scope in the present document and is therefore not covered in this analysis. In a deployed network, operator can choose a target max|TE|_{antenna} and then derive the corresponding max|TE|_{O-RU}.

11.4 Node synchronization state behavior

11.4.1 Configurations LLS-C1 and LLS-C2

11.4.1.1 Overview

Clause 11.4.1 defines the synchronization state interactions between O-DU, O-RU, and SMO for O-RAN topology configurations LLS-C1 and LLS-C2 where the O-DU acts as PLFS and PTP master.

The operation of O-DU and O-RU during HOLDOVER and other related states is described in Table 11.4.1.1-1. O-RU HOLDOVER and O-DU HOLDOVER are independent events. Likewise, O-RU HOLDOVER behavior is optional (not mandatory to be supported by HW or SW).

In addition to synchronization state, the O-DU also considers estimated synchronization accuracy, because synchronization state alone does not necessarily reflect synchronization status – a node in the LOCKED or HOLDOVER mode may have synchronization accuracy outside of a required limit.

Table 11.4.1.1-1: Node behavior during HOLDOVER and FREERUN

O-DU State	Synch accuracy	O-RU State	M-Plane State	Action
-	-	-	Disconnected	Clause 11.4.1.2
-	N/A	FREERUN	Connected	Clause 11.4.1.3
FREERUN	-	-	Connected	Clause 11.4.1.4
LOCKED/ HOLDOVER	In a limit	LOCKED/ HOLDOVER	Connected	Clause 11.4.1.5
LOCKED/ HOLDOVER	Out of a limit	LOCKED/ HOLDOVER	Connected	Clause 11.4.1.5

11.4.1.2 M-Plane disconnected

O-RAN allows hybrid M-Plane model with M-Plane communication between

- O-RU and O-DU.
- O-RU and Service Management and Orchestration (SMO) function.

As a result, the following M-Plane disconnected events shall be considered:

- O-DU detects loss communication to O-RU.
- SMO detects loss communication to O-RU.
- O-RU detects loss communication to O-DU.
- O-RU detects loss communication to SMO.
- O-RU detects loss communication to both O-DU and SMO.

NOTE: The following behavior is an assumption and is described in M-Plane specification clause 15.3.3 [7].

- O-DU behavior:
 - If the O-DU detects a loss of M-Plane communication to an O-RU, the O-DU shall stop sending any IQ data towards the O-RU. The O-DU shall also send an explicit command to the O-RU to disable RF transmission.
 - The O-DU shall keep RF transmission on the O-RU off, and shall not turn it on, if M-Plane communication to the O-RU is broken.

- The requirement for the O-DU to send an explicit command is intended to prevent unsupervised radio operation, if M-Plane communication broken in one direction only, and the fault stays undetected on the O-RU. On the other hand, if the O-RU detects the fault, it disables RF transmission autonomously (see below), and the O-DU shall stop IQ transmission accordingly.
- O-RU behavior:
 - If the O-RU detects a loss of M-Plane communication to the O-DU or both O-DU and SMO, the O-RU stops RF transmission. The O-RU shall keep RF transmission off, and shall not turn it on, if M-Plane communication to the O-DU or SMO is broken.
 - If the O-RU detects a loss of M-Plane communication to the SMO only, the O-RU action shall continue RF transmission. It is expected SMO can eventually detect a loss of M-Plane communication to the O-RU based on some round-trip sanity check mechanism.
- SMO behavior:
 - If the SMO detects a loss of M-Plane communication to an O-DU or an O-RU, the SMO action is out of scope of the present document.

11.4.1.3 O-RU in the FREERUN state

- O-DU behavior:
 - If synchronization state on a connected O-RU transits to the FREERUN state, the O-DU shall stop sending C-Plane and U-Plane related data to the O-RU unless otherwise specified. When O-RU transits to LOCKED state, O-DU shall request carriers to be switched to ACTIVE to reenale transmission.
 - The O-DU receives a notification that the O-RU switched to the FREERUN state and, as consequence, all configured carriers on the O-RU were disabled. The O-DU shall disable carriers to be aligned on the carrier configuration. After O-RU reestablish synchronization, the O-RU notifies the O-DU that the O-RU switched to the LOCKED state and then if the O-DU intends to reenale transmission, the O-DU shall set carriers to active.
 - For carrier activation procedure and notifications related to carriers and synchronization state, refer to the M-Plane specification [7], clause 15.3.3.
- O-RU behavior:
 - If synchronization state on an O-RU transits to the FREERUN state, the O-RU shall autonomously stop RF transmission and switch all carriers to INACTIVE state. As specified in the O-RAN M-Plane specification [7] clause 10, and 15.3.3, when configured to do so The O-RU shall send a notification to the O-DU about synchronization and carriers state change. The O-RU shall re-enable RF transmission only when the O-DU requests it.
 - The O-RU is obliged to stop RF transmission as soon as the accuracy of the signal can no longer be guaranteed. The O-RU shall send notification to the O-DU about any changes in its states. The O-RU shall not activate carriers by itself as it is not guaranteed that the O-DU is providing valid U-Plane and C-Plane related data to the O-RU.

11.4.1.4 O-DU in the FREERUN state

- O-DU behavior:
 - If an O-DU transits to the FREERUN state, the O-DU shall disable RF transmission on all connected O-RUs, and keep it turned off until synchronization is reacquired.
 - The O-DU shall support configuration option that allows O-DU to operate outside of the required synchronization limits, or without any synchronization at all.

- O-RU behavior:
 - The O-RU shall only react on a change of Quality Level, received in PLFS SSMs, and Clock Class, received in PTP Announce messages:
 - If the received Quality Level and Clock Class are acceptable the O-RU shall continue using the reference signal.
 - If the received Quality Level or Clock Class becomes unacceptable the O-RU shall stop using the reference and transit to either the FREERUN state or a HOLDOVER state.
 - Otherwise, the O-RU shall rely on O-DU to take care of the changed synchronization state. The O-RU need not react on the FREERUN state at the O-DU in any distinct way.

11.4.1.5 Operation in LOCKED/HOLDOVER state

Whether in LOCKED or HOLDOVER state, the O-DU shall continuously monitor the LOCKED/HOLDOVER status of the O-RUs under its management.

- O-DU behavior:
 - In configuration LLS-C1 and LLS-C2: by collecting the O-RUs' LOCKED or HOLDOVER state, as well as the received PLFS and PTP quality status, O-DU in LOCKED or "HOLDOVER state is able to detect any self-estimated frequency and/or time accuracy degradation by the O-RUs.
 - All O-RUs in the LOCKED state and directly connected to the same master clock (typically the O-DU in LLS-C1, and the nearest T-BC in LLS-C2) preserve optimal relative time error between them, regardless the O-DU's status. As a result, intra-site features may remain enabled using different criteria compared to inter-site ones.
 - The port-to-port constant and dynamic time errors between two master ports of the same module (either the O-DU in configuration LLS-C1, and the T-BC ones in LLS-C2) may also contribute to the total error.
- O-RU behavior:
 - The O-RU shall only react on a change of Quality Level received in PLFS SSMs, and Clock Class received in PTP Announce messages:
 - If received Quality Level and Clock Class are acceptable, the O-RU shall keep on using the reference signal.
 - If received Quality Level or Clock Class become unacceptable, the O-RU shall stop using the reference and transit to either the FREERUN state or a HOLDOVER state.
 - Recommendation ITU-T G.8275.1 [31]-clause 6.4 define the notion of "within HOLDOVER specification" (clock class values 7 and 135), and "exceeding the HOLDOVER specification" (clock class values 140, 150, 160, 165).
 - Recommendation ITU-TG.8271.1 [23] Appendix V describes an example of HOLDOVER budget as 400 ns in the context of category C (as per IEEE 802.1CM [11] clause 6.4.1.4). This value is however too high and useless for categories A/B, highlighting that the notion of "acceptable" is dependent on the category of each feature.
 - Consequently, it shall be mandatory that each O-RU reports to O-DU the received Quality Level and Clock Class, allowing O-DU to enable and disable accordingly the various RF features.
 - Besides, the value for the HOLDOVER specification may be configurable, thus allowing each network operator to tune it to its own needs.

11.4.2 Configurations LLS-C3

11.4.2.1 Overview

Clause 11.4.2 defines the synchronization state interactions between O-DU, O-RU, and SMO for O-RAN topology configuration LLS-C3 based on IEEE 802.1CM [11] clause 7. PRTC/GM is provided by the fronthaul network.

The operation of the Fronthaul network elements, O-DU and O-RU during HOLDOVER and other related states is described in Table 11.4.2.1-1. O-RU HOLDOVER and O-DU HOLDOVER are independent events. Likewise, O-RU HOLDOVER behavior is optional (not mandatory to be supported by HW or SW).

More than one PRTC/GM may be considered as a deployment option to improve redundancy. Should a PRTC/GM fail, then another should be available as a backup time source and the PTP network tree would automatically re-arrange. Only a short HOLDOVER (NOTE: The duration needs to be defined from Recommendation ITU-T G.8271.1 [23]) shall be supported inside the various network elements (as well as O-DU and O-RU) to provide a safe operation during this rearrangement scenario.

NOTE: The O-DU can also be configured to provide backup PLFS+PTP like in LLS-C2.

In addition to synchronization state, the O-DU also considers estimated synchronization accuracy because the synchronization state alone does not necessarily reflect synchronization status; a node in the LOCKED or HOLDOVER mode may have synchronization accuracy outside of a required limit.

Table 11.4.2.1-1: Node behavior during HOLDOVER and FREERUN

Network State	Sync Accuracy	O-DU and O-RU State	M-Plane State	Action
-	-	-	Disconnected	Clause 11.4.2.2
-	N/A	FREERUN	Connected	Clause 11.4.2.3 and 11.4.2.4
LOCKED/ HOLDOVER	In limit	LOCKED/ HOLDOVER	Connected	Clause 11. 4.2.5
LOCKED/ HOLDOVER	Out of limit	LOCKED/ HOLDOVER	Connected	Clause 11. 4.2.5

11.4.2.2 M-Plane disconnected

This clause is same as 11.4.1.2

11.4.2.3 O-RU in the FREERUN State

This clause is same as 11.4.1.3

11.4.2.4 O-DU in the FREERUN state

If an O-DU and O-RU are synchronized from the same fronthaul network and are connected to neighbor nodes in this network, it is most probable that the event leading to O-RU transiting to the FREERUN state will also lead to the same transition at the O-DU.

If O-DU has backup frequency and time source, such as local or remote PRTC, it can become a backup Master like in configuration LLS-C2.

O-DU

- O-DU Behavior:

- If an O-DU transits to the FREERUN state because the synchronizing network delivers unacceptable synchronization quality, the O-DU shall disable RF transmission on all connected O-RUs, and keep it turned off until synchronization is reacquired again.
- The O-DU may support a configuration option that allows O-DU to operate outside of the required synchronization limits, or without any synchronization at all.

O-RU

- O-RU Behavior:
 - The O-RU is not synchronized from the O-DU. It may have no indication of the O-DU synchronization status, and therefore shall rely on the O-DU to take care of the changed synchronization state. The O-RU need not react to the FREERUN state at the O-DU in any distinct way.

11.4.2.5 Operation in LOCKED/HOLDOVER state

Whether in LOCKED or HOLDOVER state, the O-DU shall monitor the LOCKED or HOLDOVER state, as well as the received PLFS and PTP quality status of the O-RU. The behaviors of the O-DU and O-RU shall be the same as LLS-C1 and LLS-C2 configurations described in clause 11.4.1.5.

11.4.3 Configurations LLS-C4

11.4.3.1 Overview

Clause 11.4.3 defines the synchronization state interactions between O-DU, O-RU, and SMO for O-RAN topology configurations LLS-C4 where the O-RU is synchronized by local PRTC (typically a GNSS receiver).

The operation of O-DU and O-RU during HOLDOVER and other related states is described in Table 11.4.3.1-1. O-RU HOLDOVER and O-DU HOLDOVER are independent events. Likewise, O-RU HOLDOVER behavior is optional (not mandatory to be supported by HW or SW).

In addition to the synchronization state, the O-DU also considers estimated synchronization accuracy because the synchronization state alone does not necessarily reflect synchronization status; a node in the LOCKED or HOLDOVER mode may have synchronization accuracy outside of a required limit.

Table 11.4.3.1-1: Node behavior during HOLDOVER and FREERUN

O-DU State	Sync Accuracy	O-RU State	M-Plane State	Action
-	-	-	Disconnected	Clause 11.4.3.2
-	N/A	FREERUN	Connected	Clause 11.4.3.3
FREERUN	-	-	Connected	Clause 11.4.3.4
LOCKED/ HOLDOVER	In limit	LOCKED/ HOLDOVER	Connected	Clause 11.4.3.5
LOCKED/ HOLDOVER	Out of limit	LOCKED/ HOLDOVER	Connected	Clause 11.4.3.5

11.4.3.2 M-Plane disconnected

This clause is same as 11.4.1.2

11.4.3.3 O-RU in the FREERUN state

This clause is same as 11.4.1.3

11.4.3.4 O-DU in the FREERUN state

- O-DU behavior:
 - If an O-DU transits to the FREERUN state, the O-DU shall disable RF transmission on all connected O-RUs and keep it turned off until synchronization is reacquired again.
 - The O-DU may support a configuration option that allows O-DU to operate outside of the required synchronization limits, or without any synchronization at all.
- O-RU behavior:
 - If the O-RU is not synchronized from the O-DU, it may have no indication of the O-DU synchronization status, and therefore shall only rely on O-DU to take care of the changed synchronization state. The O-RU need not react to the FREERUN state at the O-DU in any distinct way.

11.4.3.5 Operation in LOCKED/HOLDOVER state

Whether in LOCKED or HOLDOVER state, the O-DU shall monitor the LOCKED or HOLDOVER state. The behaviors of the O-DU and O-RU shall be the same as for LLS-C1 and LLS-C2 configurations described in clause 11.4.1.5.

11.5 S-Plane handling in multiple link scenarios

Behavior of S-Plane in scenarios with multiple links shall be based on the following principles:

O-DU – Grand Master/synchronization transmitter (configurations LLS-C1 & LLS-C2): There shall be an input sync reference signal on at least one link to an O-RU. Likewise, it is permitted to have input reference signal on multiple or all links to a given O-RU.

O-RU (all configurations) – synchronization receiver: O-RU shall be able to adapt at startup to the reception of a sync reference signal on any port from their master. How the ingress signal is detected (usually SSM for SyncE and Announce for PTP), and how the active port is selected (e.g., via round-robin, Recommendation ITU-T G.8275.1 [31] BMCA, etc.) are implementation-specific.

If the input reference is present on multiple links, the O-RU may implement redundancy for the input reference signal if only capacity links are present on the module.

11.6 Announce messages

O-RU shall check the following advertised parameters against a list of acceptable values based on its own design. These parameters are M-Plane configurable.

- Domain Number: Default: 24 (for Full Timing Support per Recommendation ITU-T G.8275.1 [31] clause 6.2.1) or 44 (for Partial Timing Support per Recommendation ITU-T G.8275.2 [32] clause 6.2.1).
- PTP Acceptable Clock Classes:
 - Default: 6, 7.
 - Operator configurable: 6, 7, 135, 140, 150, 160, 248.

11.7 Elementary procedures

11.7.1 PTP time synchronization procedure

All procedures used to exchange time related information between a time synchronization master and subordinate shall be compliant to the Recommendation ITU-T G.8275.1 [31] or G.8275.2 [32] telecom profile, which provides necessary details on utilization of the IEEE 1588 [3], [33] protocol in telecom applications.

11.7.2 System frame number calculation from GPS time

The general framework for System Frame Number (SFN) calculation from GPS time is based on the following premises:

- PTP time on the fronthaul interface shall use PTP timescale.
- The PTP epoch is 1 January 1970 00:00:00 TAI, which is 31 December 1969 23:59:51.999918 UTC.
- PTP time on the fronthaul interface shall be traceable to a PRTC if a network wide synchronization of O-RUs at the air interface is required (as in TDD 5G)
- From PTP time, the GPS seconds elapsed since GPS epoch (midnight January 6th, 1980) can be calculated, since the difference between PTP and GPS epoch is a constant. The GPS seconds are expressed as a real number.
- GPS seconds shall be used to calculate the frame number according to:

$$\text{FrameNumber} = \text{floor}\left(\frac{\text{GPSseconds} - \beta \times 0.01 - \frac{\alpha}{1.2288 \times 10^9}}{\text{framePeriodinSeconds}}\right) \bmod (\text{maxFrameNumber} + 1)$$

where:

- FrameNumber expressed in seconds.
- framePeriodinSeconds = 0.01 seconds.
- maxFrameNumber = 1023.
- α and β are defined as shown in Table 11.7.2-1:

Table 11.7.2-1 : α and β parameter exchange

Direction	Parameter	Range	Resolution
O-DU to O-RU	Offset α (NR) of radio frame timing (NOTE 1)	α : 0~1.2288e7 α * (1/1.2288 ns) : 0,00 s~0.01 s	1/1.2288 ns
O-DU to O-RU	Offset β (NR) of radio frame timing (NOTE 1)	β : -32768~32767 β * (10 ms) : -327680 ms~327670 ms	10 ms
NOTE: Parameter data types and values for α and β are provided by the O-RAN M-Plane specification [7]. Epoch for α and β (i.e., SFN=0) is set to 1980.1.6 00:00 (UTC)			

Figure 11.7.2-1 below shows examples when α and β are set to different values. The frame number at the antenna reference point is shown in relation to the GPS_{seconds} time. For simplicity the GPS-time around 10 seconds is shown.


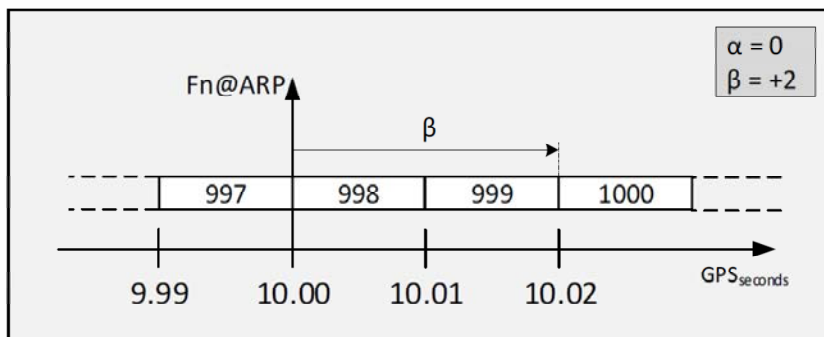
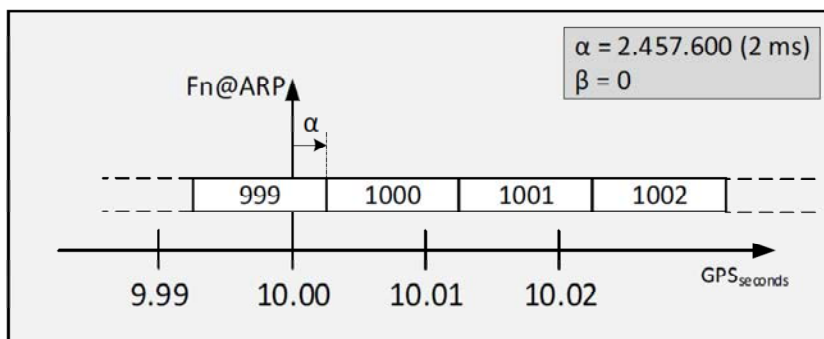




Figure 11.7.2-1: α and β parameter example

Table 11.7.2-2: Void

NOTE: Refer to clause 4.8 for the timing relation between UL radio frame and DL radio frame.

12 Beamforming functionality

12.1 General

The following clauses describe terminologies, rules, properties and uses cases related to beamforming and its functionalities. It is the baseline to follow by the O-DU, O-RU and modelling in M-Plane.

12.2 Hierarchy of radiation structure in O-RU

The hierarchy of radiation structure in O-RU is depicted in Figure 12.2-1 and described below:




Figure 12.2-1. Hierarchy of Radiation Structure

- **O-RU:** each O-RU can have 1 or several panels
- **panel:** each panel can have 1 or several tx-array and/or rx-array
- **tx-array and rx-array:**
 - tx-array and rx-array are logical constructs used for data routing; there is a relationship to physical antennas as defined by the O-RU construction, and eAxC_ID values may be assigned to tx-arrays and rx-arrays:
 - Multiple eAxC_ID values may be assigned to a single tx-array or rx-array
 - An eAxC_ID value for a given direction (DL or UL) shall not be assigned to more than one tx-array or rx-array
 - Each tx-array and rx-array has only one polarization for Category A O-RUs.
 - Each tx-array and rx-array may include both polarizations for Category B O-RUs.
 - Category B O-RU shall provide one or more single-polarization rx-arrays for each polarization.
 - Each tx-array and rx-array may be configured to terminate one or several RU_port_IDs.
 - Each tx-array and rx-array can have one or several array elements (one use case being antenna technologies wherein a single "array element" can support a high degree of beamforming, another use case being the need to address a single element in the array).
 - Each physical array element may be a member of one or multiple tx-arrays, multiple rx-arrays or both tx-arrays and rx-arrays.
 - Amplitude and phase of each array element inside one tx-array and rx-array can be controlled via the weights pointed to by the beamIDs which are received via the terminated RU_port_ID(s).

Alternatively, the beamIds may indicate the "state" of an antenna array not explicitly comprising separate radiating elements.

- It is assumed that when executing beamforming, array elements within an array are combined in phase to form directed energy, therefore the array elements belonging to a tx-array and rx-array are calibrated together as a group (although such a calibration operation can be hidden from the interface e.g., executed once upon O-RU manufacturing).
- **array element:** each array element can have 1 or several radiators (where amplitude and phase relation between the radiators cannot be changed dynamically during real-time)
- **radiator:** radiating element – see clause 2.306 in [50] for the definition of "radiating element"

12.3 Calibration

Calibration is the functionality of eliminating or minimizing relative amplitude and phase differences over frequency domain or time delay over the time domain between the array elements belonging to same tx-array or rx-array (including effect of front-end analog filters).

Calibration can also be applied between two or many tx-arrays or rx-arrays if those tx-arrays or rx-arrays belong to the same calibration group which is part of the O-RU capabilities.

12.4 beamId use for various beamforming methods

12.4.1 Overview

There are two main domains in which beamforming is executed, frequency-domain and time-domain; it is also possible to combine both (called "hybrid beamforming"). Frequency-domain beamforming is done between the RE mapping and FFT/iFFT processing stages (in UL and DL respectively) so is inherently a digital operation. Time-domain beamforming may be executed digitally or in the analog domain.

A characteristic of frequency-domain beamforming when used with OFDM is that different users may use the same time slot yet use different beams. In contrast, with time-domain beamforming all the users and signals in a time slot use the same beam. Hybrid beamforming allows different users in the same time slot to use different beams (the frequency-domain part) at the same time as all the users using a shared time-domain beam. An example is the case where the time-domain beam provides directivity in the elevation plane (so all users use the same elevation beam) while the frequency-domain beams provide directivity in the azimuth plane (so different users may use different azimuth beams).

The following beamforming methods (see also Annex J for more details) are supported:

- a) **Predefined-beam beamforming:** In this case, an index called "beamId" indicates the specific beam pre-defined in the O-RU to use. The beamId can indicate a frequency-domain beam or a time-domain beam or a combination of both ("hybrid" beam) and the O-DU needs to know it to ensure the beamId is properly applied e.g. the O-DU can not apply different time-domain beams to the different PRBs in the same OFDM symbol. The method the O-RU uses to generate the beam is otherwise not relevant, it could use the application of gain and phase controls on separate antenna elements, or use multiple shaped-energy antennas, or any other technology. The O-RU should convey to the O-DU via the M-Plane on startup beam characteristics but the O-DU may remain ignorant regarding how the beam is actually created by the O-RU.
- b) **Weight-based dynamic beamforming (based on real-time-updated weights):** Here the O-DU is meant to generate weights that create the beam, so the O-DU needs to know the specific antenna characteristics of the O-RU including how many antenna elements are present in the vertical and horizontal directions and the antenna element spacing, among other properties. The weight vector associated with each beam has a beamId value and the interpretation of this beamId value is addressed in clause 12.4.3.
- c) **Attribute-based dynamic beamforming (based on real-time-updated beam attributes):** Like predefined-beam beamforming, attribute-based beamforming allows the O-DU to tell the O-RU to use a specific beamId but in this case that beamId is associated with certain beam attributes as described in clause 7.7.2. How the O-RU achieves the implementation of the beams is not specified, however the O-DU needs to know whether the beam identified by the beamId is generated as a frequency-domain beam or a time-domain beam to ensure the beamId is properly applied e.g. the O-DU can not apply different time-domain beams to the different PRBs in

the same OFDM symbol. When the attribute-based beamforming method is used for at least one layer (or spatial stream), then O-DU shall use same beamforming method for all layers (or spatial streams) of the specific time-frequency resource element(s).

- d) Channel-information-based beamforming: In this case the O-DU provides channel information per UE periodically (generally less often than every slot) and then on a slot-by-slot basis the O-DU provides scheduling information which the O-RU uses along with the channel information to calculate the proper beamforming weights for the specific slot with its co-scheduled UEs. Here there is no beamId value associated with the beamforming, instead the ueId is associated with each data section. Therefore, this clause regarding beamId usage is not relevant for this beamforming method. When the channel-information-based beamforming method is used for at least one layer (or spatial stream), then O-DU shall use same beamforming method for all layers (or spatial streams) of the specific time-frequency resource element(s).

12.4.2 Predefined-beam beamforming

12.4.2.1 Overview

When implementing predefined-beam beamforming, the O-RU shall convey to the O-DU whether the beamforming type is frequency-domain, time-domain, or a mixture of the two ("hybrid beamforming"). In the case of frequency-domain-only or time-domain-only, the beamId is simply an index to the desired beamforming weight vector or other beamforming method. In the case of hybrid beamforming, there are present in the O-RU pre-loaded frequency-domain weight vectors and time-domain weight vectors (these are applied separately). The beamId points to a single combined frequency-domain and time-domain weight vector. However, in reality there will be the application of a frequency-domain beamforming weight vector and the separate application of multiple time-domain beamforming weight vectors, one per frequency-domain weight value.

For this case, O-RU shall convey via the M-Plane to the O-DU upon start-up as part of the O-RU self-description:


- the list of available beamIds and their characteristics, see clause 12.4.2.2 below.

NOTE: There is no requirement that beamIds be in sequential order or that there are no gaps between beamId values.

The beamId value of zero is reserved for no beamforming, so shall not be used for any other frequency-domain, time-domain or hybrid beam.

12.4.2.2 Beam characteristics

In order to use predefined-beam beamforming in a standardized way, the present document considers beamforming to be defined such that energy (in the DL) or sensitivity (in the UL) is focused into either a "coarse" or "fine" granularity with possible overlaps. In this way "broadcast" beams may be used to cover a wider area with less power or sensitivity, while higher-power or higher-sensitivity beams may be used in e.g. a per-UE fashion. Figure 12.4.2.2-1 shows an example of the assumed beam arrangement.



Note: ovals represent beams, numbers in ovals represent beamId

Figure 12.4.2.2-1: Example of updatable-weight frequency-domain plus fixed time-domain beamforming

Figure 12.4.2.2-1 introduces the concept of "beam-group". A "beam-group" is the set of beamIds that may be used simultaneously for some reason. In Figure 12.4.2.2-1 an assumption can be made that the elevation direction is time-domain while the azimuth direction is frequency-domain; the frequency-domain beams in the same time-domain group may be used at the same time but different time-domain groups shall be separated in time.

Figure 12.4.2.2-1 also shows the need to describe beam overlap and beam adjacency. It is necessary, for example, to convey that beamId=1 is a coarse beam overlapping with fine beamIds={9,10,17,18} and that beamId=18 is a fine beam overlapping with coarse beamId=1. It is also necessary to convey that beamId=1 has as neighbors beamId={2,5,11,19,25,26} and that beamId=18 has as neighbors beamIds={2,5,10,17,19,26}. Overlapping beams shall not be scheduled together to avoid interference, and neighboring beams should not be scheduled together to avoid interference where possible.

The specific beam characteristics are therefore as follows, all per beamId:

- 1) **beam-type** for the subject beamId enumerated as COARSE-BEAM or FINE-BEAM;
 - there is no attempt to quantify "coarse" and fine" in terms of beamwidth, this is just a relative relationship
- 2) **beam-group-id** for the subject beamId as an integer;
- 3) **coarse-fine-beam-relation** as a list: if the subject beamId is coarse, this is a list of the associated fine beams, and if the subject beamId is fine, this is a list of the associated coarse beams;
- 4) **neighbor-beam** as a list: a list of all beams that may interfere with the subject beamId

An O-RU which supports predefined beamforming may also support shifting the coverage area upon M-Plane command, referred to as the 'predefined-beam-tilt-offset' feature. The O-RU shall convey its capability and associated parameters for the 'predefined-beam-tilt-offset' feature if supported. This feature allows an operator to adjust the coverage of the O-RU in elevation and/or azimuth angle. Figure 12.4.2.2-2 shows an example of the predefined beam-tilt-offset operation. As a default, a predefined beam has its own steering angle denoted by (elevation angle: "theta", azimuth angle: "phi"). If the O-RU controller configures the elevation-tilt-offset-angle as "+A" degrees and azimuth-tilt-offset-angle as "+B" degrees, then the O-RU shall regenerate the beams to change its steering angles to be elevation angle value: "theta + A", azimuth angle value: "phi + B". An M-Plane command delivers the elevation and azimuth predefined beam-tilt-offset angles. To regenerate beams, O-RU can load one of the pre-stored multiple beamforming weight vectors corresponding to all supported predefined-beam tilt values or can regenerate beamforming weights upon receiving the M-Plane command; this is up to the O-RU implementation.

The elevation-tilt-offset-angle values smaller than 0 represents an up-shift of the default service area towards the zenith (i.e., corresponding to a decrease in zenith angle) and values larger than 0 represent a down-shift of the default service area away from the zenith (i.e., corresponding to an increase in zenith angle).

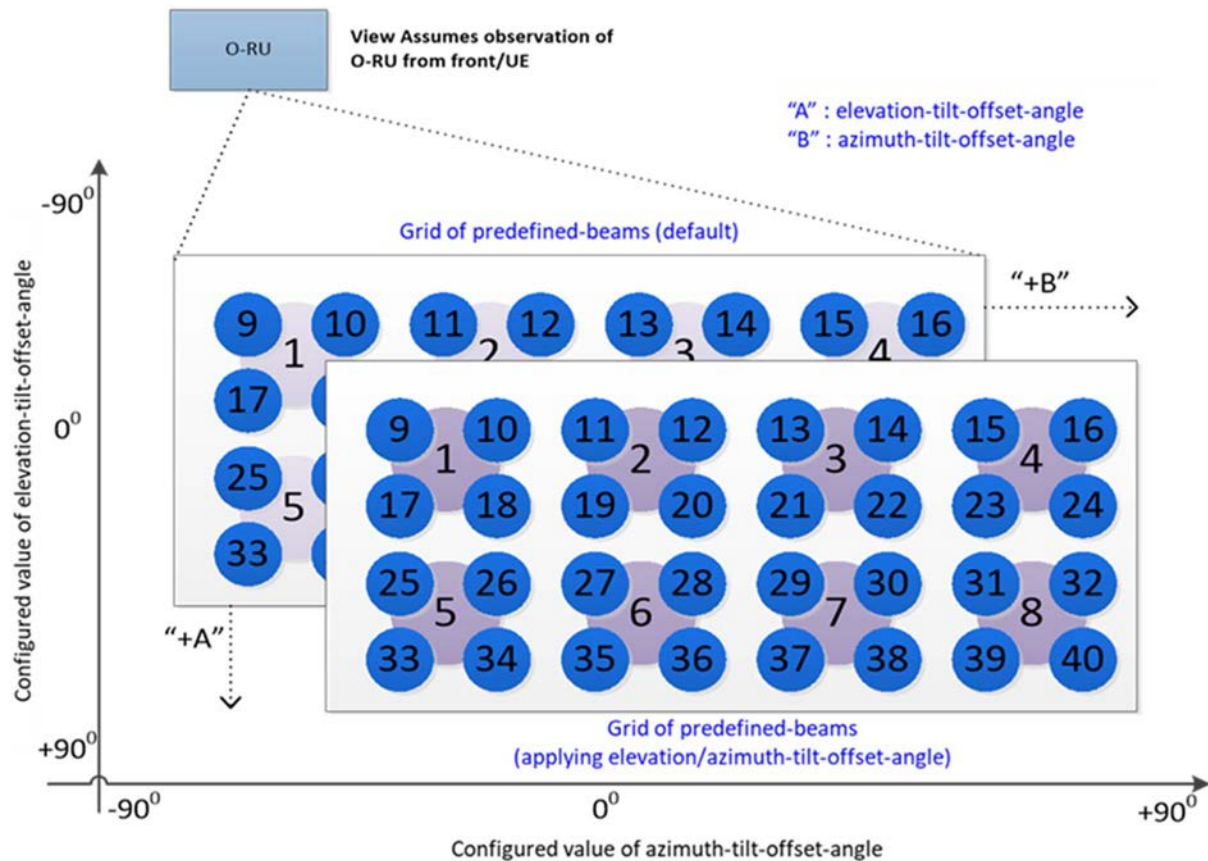


Figure 12.4.2.2-2: Example of shifting a grid of predefined beams with tilt-offset

12.4.3 Weight-based dynamic beamforming

12.4.3.1 Overview

Weight-based dynamic beamforming operates the same as predefined-beam beamforming, except that the O-DU shall convey beamforming weights to the O-RU. The following rules apply for the provision of new beamforming weights associated with a beamId, these apply for frequency-domain, time-domain, and hybrid beamforming cases; these are consistent with the principle that weights associated with a beamId when applied to applicable symbols in a slot shall be the same:

- 1) If beamforming weights are to be updated for a slot, then the O-DU shall provide the new weights for the first occurrence of the given beamId in the slot.
- 2) Weight updates (sending the same weights for a beamId is not considered a "weight update") shall not be provided more than once for a beamId in the slot. Where new beamforming weights are required in a single slot, a new beamId value shall be employed instead.
- 3) For PRACH a "slot duration" as considered here is considered to be a PRACH sequence duration (see Figure 7.2.3.4-1).

12.4.3.2 Optional restriction on sending of beamforming weights

To allow an optimization of O-RU processing, an optional restriction regarding the sending of beamforming weights is provided: for messages having the same startSymbolId within the same "context" (meaning with same numerology and array-carrier and direction (UL/DL), see also 3 paragraphs further down) within a slot, if beamforming weights are sent associated with a beamId in one C-Plane message, then any invocation of that same beamId in a subsequent C-Plane message having the same startSymbolId within the same slot shall also include the set of beamforming weights. For sections within the same C-Plane message, when beamIds are used in multiple data section descriptions the first such section shall contain the beamforming weights and the subsequent sections using the same beamId may omit the

beamforming weights. Because this restriction applies to the sending of beamforming weights, this restriction applies only to Section Types 1 and 3.

An exception exists: for high-priority data sections with a repetition bit, only the data section with repetition bit = 0 shall contain the set of beamforming weights; where the repetition bit = 1 for the same high-priority data section the set of beamforming weights shall not be sent, just the beamId value shall be sent.

This optional restriction is intended to enable the proper association of beamforming weights with beamId even if C-Plane messages within a given reception window are received out of order, which for some O-RU optimized designs may result in ambiguous results (i.e., uncertainty regarding when new beamforming weights will be applied). To support this optional restriction, O-RU shall report the M-Plane O-RU capability "beam-update-contention-control-limits-required" (an O-RU feature called "BEAM-UPDATE-CONTENTION-CONTROL"). The O-DU may enable or disable this feature at an endpoint granularity using boolean flag "beam-update-contention-control-enabled"; if set to FALSE (default) the above restriction shall not apply but if TRUE then the above restriction shall apply. If the "BEAM-UPDATE-CONTENTION-CONTROL" feature is supported by the O-RU, then the O-RU shall report "max-beams-per-symbol-with-beam-contention-control" and "max-beam-updates-per-symbol-with-beam-contention" which are limits that apply when the O-DU will comply with the restriction that O-RU requests (presumably higher limits when the restrictions will be obeyed). The O-DU may consider the relevant limits as being without "with-beam-contention" when the restriction will not be obeyed, and the O-DU may consider the relevant limits as being with "with-beam-contention" when the restriction will be obeyed. Though O-DU may set this flag "per-endpoint" the O-DU shall ensure all endpoints have the same "beam-update-contention-control-enabled" value, either TRUE or FALSE for all applicable endpoints (endpoints in use that may use the same beamId). In the case the O-DU chooses not to comply with the restriction when it has configured the value "TRUE", it is expected that either old or new beamforming weights will be applied to the subsequent C-Plane message.

A complication arises when updating beamforming weights for one data stream (eAxC) then applying those updated weights for a different data stream (eAxC). When the contexts of the two data streams are different there can be uncertainty when the new weights uploaded for the first data stream may be applied to second data stream in a different context. The uncertainty arises due to differences in the reception windows for C-Plane messages in different "contexts". The relevant contexts are:

- 1) DL vs UL: there can be several symbols' worth of time difference between these.
- 2) eAxC flows of different SCS: symbols are of different time-length hence reception windows likely are also offset between them.
- 3) Different array-carriers: reception windows are likely to be different length, and YANG model allows different reception windows to be defined.

To prevent uncertainty in the application of updated beamforming weights between contexts, an optional "beam-context-gap-period" (M-Plane parameter reported by the O-RU to the O-DU) defining the time between when beamforming weights are provided in a C-Plane message and the time the same beamId may be cited in a subsequent C-Plane message without weights. More precisely, after a beamId is updated with new weights by a specific C-Plane message "A", any subsequent C-Plane message "B" that cites that beamId without weights and whose reception window ends after the beam-context-gap-period starting from the end of the reception window for message "A" will use the updated weights. If the beam-context-gap-period is not honored then old or new weights (uncertain which ones) may be used by the second C-Plane message "B". See Figure 12.4.3.2-1 and Figure 12.4.3.2-2 for examples of how the beam-context-gap-period is applied for two C-Plane messages "A" and "B" in different contexts.

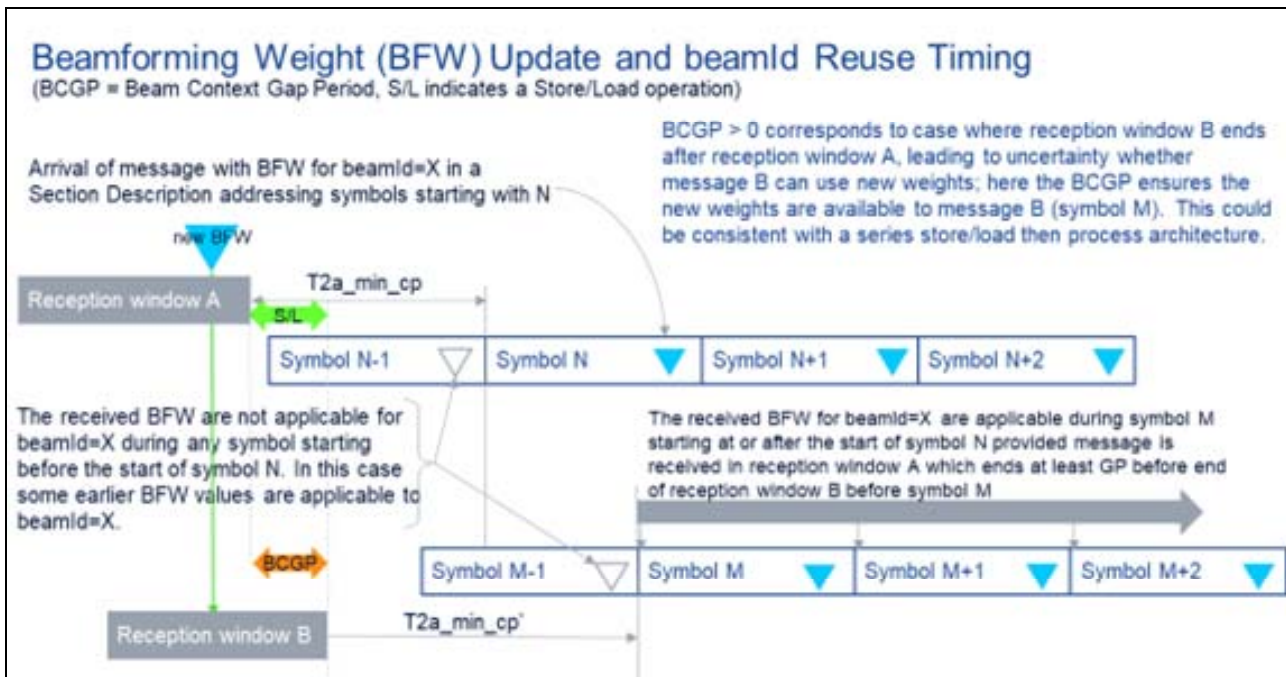


Figure 12.4.3.2-1: BFW Update Example #1 using beam-context-gap-period

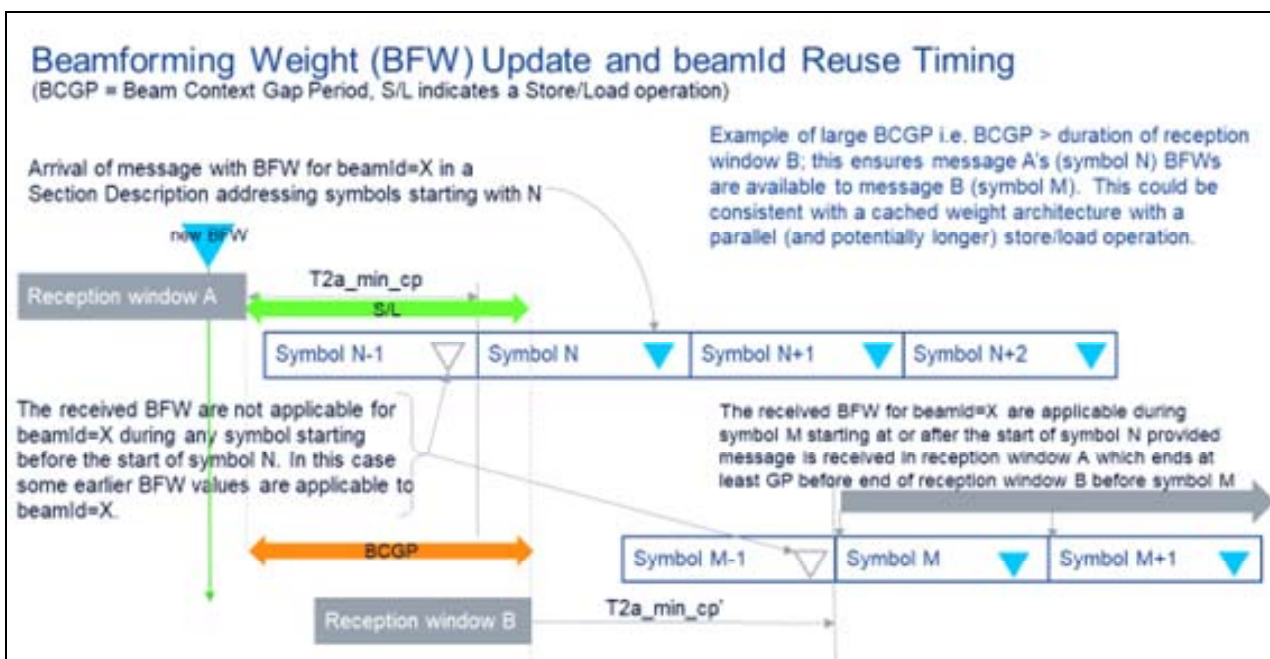


Figure 12.4.3.2-2: BFW Update Example #2 using beam-context-gap-period

If beam-context-gap-period is reported by the O-RU, the O-RU shall provide the beam-context-gap-period in microseconds as a 16-bit number, with the special value 0xFFFF meaning "infinite". Infinite means that for that O-RU, a beamId may not be reused at all in a different context (applicable only to weight-based dynamic beamforming). A zero value indicates the end of the message "A" and message "B" reception windows are perfectly aligned which should allow the second C-Plane message to use the new beamforming weights with a zero microsecond beam-context-gap-period. In order to ensure consistent operation across all SCS and carrier bandwidths, the beam-context-gap-period shall be set by the O-RU to be the same value across the entire O-RU. The value of beam-context-gap-period shall be provided by the O-RU for use by the O-DU to avoid a beamforming weight update race condition; this is independent of the "BEAM-UPDATE-CONTENTION-CONTROL" feature.

An O-DU not recognizing the beam-context-gap-period function (possible for an O-DU conforming to an earlier version of the present document), or an O-RU not providing a beam-context-gap-period (possible for an O-RU conforming to an earlier version of the present document), will operate in the same manner as prior to the introduction of the beam-context-gap-period, with an unpredictable result if a beamId has new weights loaded in one context and is reused in a different context.

12.4.3.3 Weight-based dynamic frequency-domain or time-domain beamforming (not hybrid)

In the case of either frequency-domain or time-domain beamforming wherein the beamforming weights can be updated in real-time and have a beamId value associated with the weights, the beamId is treated the same: it points to a set of weights that control the array elements' gain and phase and the number of weights equals the number of array elements. In many cases, the magnitude of each complex weight value will equal unity but this is not required; in particular "tapering" may require less-than-unity weight magnitudes for some array elements. The weight values prior to any compression will be fractional hence no I or Q value may exceed positive or negative unity.

The following list describes the information that the M-Plane shall carry from the O-RU to the O-DU upon start-up as part of the O-RU self-description, the information listed is per array (so per tx-array or per rx-array):

- 1) Beamforming type, possible choices are "frequency", "time" or "hybrid" , – here will be "frequency" or "time"
- 2) Maximum number of weight-based beamId values supported (could be zero): "numBeams"
 - O-RUs may have memory limitations that mean the number of beams is limited; zero means no weight-based beamforming is supported by this tx-array or rx-array
- 3) Initial value of weight-based beamId supported: "initBeamId"
 - Different ranges of beamId may support weight-based beamforming versus e.g. predefined beamforming
- 4) Frequency granularity of time-domain beamforming, enumerated as "per component carrier" or "per band".
 - Value is only present for time-domain beamforming
- 5) Time granularity of time-domain beamforming, enumerated as "per-OFDM-symbol" or "per-slot".
 - Value is only present for time-domain beamforming

Because the beams are to be generated by the O-DU the O-RU can not know the beam characteristics, so they are not reported.

The actual number of weights K in the frequency-domain or time-domain weight vectors will be clear from the O-RU antenna model, see the clause 12.5 on that topic.

12.4.3.4 Weight-based dynamic hybrid beamforming

12.4.3.4.1 Overview

Two sub-cases of weight-based dynamic hybrid beamforming are considered, wherein for one sub-case both the frequency-domain and time-domain weights may be updated in real-time (see clause 12.4.3.4.2), and for the second sub-case the frequency-domain weights may be updated in real-time but the time-domain beams are fixed (see clause 12.4.3.4.3).

12.4.3.4.2 Hybrid beamforming with updatable frequency-domain and time-domain weights

For this sub-case the beamforming weight vector is a composite of the frequency-domain weights and the time-domain weights so can be considered as simply a longer weight vector. Where a block-based beamforming weight compression is employed (block floating point, block scaling or μ -law compression), the block size is a single beamforming weight vector (both frequency-domain and time-domain parts). The actual number of weights in the composite frequency-domain plus time-domain weight vectors ($K' + K$) will be clear from the O-RU antenna model, see the Clause 12.5 on that topic.

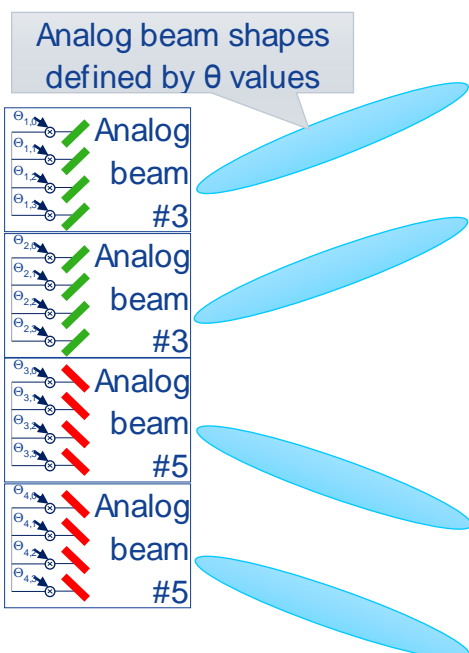
The following list describes the information that the M-Plane shall carry from the O-RU to the O-DU upon start-up as part of the O-RU self-description, the information listed is per array (so per tx-array or per rx-array):

- 1) Beamforming type, possible choices are "frequency", "time" or "hybrid" – here will be "hybrid"
- 2) Maximum number of weight-based beamId values supported (could be zero): "numBeams"

- O-RUs may have memory limitations that mean the number of beams is limited; zero means no weight-based beamforming is supported by this tx-array or rx-array
- 3) Initial value of weight-based beamId supported: "initBeamId"
 - Different ranges of beamId may support weight-based beamforming versus e.g. predefined beamforming
 - 4) Frequency granularity of time-domain beamforming, enumerated as "per component carrier" or "per band".
 - 5) Time granularity of time-domain beamforming, enumerated as "per-OFDM-symbol" or "per-slot".

The number of time-domain beamforming weights associated with a given beamId is the same as the number of array elements which is K , but the number of frequency-domain weights is less, being K' . p' represents the dimensionality of the time-domain beamforming operation, so that $K = K' * p'$. The total length of the beamforming weight vector, including both the K' frequency-domain weights and the K time-domain weights, is $K' + K = K' + (K' * p') = K' * (p'+1)$. Figure 12.4.3.4.2-1 shows an example where $K = 16$, $K' = 4$ and $p' = 4$, and the length of the beamforming weight vector (frequency-domain and time-domain combined) is 20 complex weights.

When different hybrid beams are used in the same symbol, here using beamId values 0x13 and 0x25, the time-domain weights shall be the same (here indicated by θ values) and the number of time-domain weights equals the number of array elements K (here, 16). The frequency-domain weights may differ (shown as two different sets of Φ values) and there are fewer of those, specifically there are K' (here, 4) frequency-domain weight values.



beam id	frequency-domain beamforming weights				time-domain beamforming weights															
0x13	$\Phi_{1,0}$	$\Phi_{1,1}$	$\Phi_{1,2}$	$\Phi_{1,3}$	$\theta_{1,0}$	$\theta_{1,1}$	$\theta_{1,2}$	$\theta_{1,3}$	$\theta_{2,0}$	$\theta_{2,1}$	$\theta_{2,2}$	$\theta_{2,3}$	$\theta_{3,0}$	$\theta_{3,1}$	$\theta_{3,2}$	$\theta_{3,3}$	$\theta_{4,0}$	$\theta_{4,1}$	$\theta_{4,2}$	$\theta_{4,3}$
0x25	$\Phi_{2,0}$	$\Phi_{2,1}$	$\Phi_{2,2}$	$\Phi_{2,3}$	$\theta_{1,0}$	$\theta_{1,1}$	$\theta_{1,2}$	$\theta_{1,3}$	$\theta_{2,0}$	$\theta_{2,1}$	$\theta_{2,2}$	$\theta_{2,3}$	$\theta_{3,0}$	$\theta_{3,1}$	$\theta_{3,2}$	$\theta_{3,3}$	$\theta_{4,0}$	$\theta_{4,1}$	$\theta_{4,2}$	$\theta_{4,3}$

Figure 12.4.3.4.2-1: Example of updatable-weight frequency-domain and time-domain beamforming

12.4.3.4.3 Hybrid beamforming with updatable frequency-domain weights and fixed time-domain beams

For this sub-case the beamforming weight vector is a composite of the frequency-domain weights and the time-domain beam numbers with the frequency-domain weights in the first half of the vector and the time-domain beam numbers in the second half of the vector. This vector shall not be considered as simply a longer weight vector because the frequency-domain weights may be compressed but the time-domain beam numbers shall not be compressed. Where a block-based beamforming weight compression is employed (block floating point, block scaling or μ -law compression), the block size is a single beamforming weight vector but only that half of the weight vector containing the frequency-domain weights. The remaining half of the vector containing the time-domain beam numbers may not be compressed and contains the integer time-domain beam numbers.

The actual number of weights in the composite frequency-domain weights plus time-domain beam-number vectors (K' and K) will be clear from the O-RU antenna model (see the Clause 12.5 on that topic) with the number of frequency-domain weights K' indicating which elements in the vector are subject to compression (the first K' complex values in the vector).

The following list describes the information that the M-Plane shall carry from the O-RU to the O-DU upon start-up as part of the O-RU self-description, the information listed is per array (so per tx-array or per rx-array):

- 1) Beamforming type, possible choices are "frequency", "time" or "hybrid" – here will be "hybrid"
- 2) Maximum number of weight-based beamId values supported (could be zero): "numBeams"
 - a. O-RUs may have memory limitations that mean the number of beams is limited; zero means no weight-based beamforming is supported by this tx-array or rx-array
- 3) Initial value of weight-based beamId supported: "initBeamId"
 - a. Different ranges of beamId may support weight-based beamforming versus e.g. predefined beamforming
- 4) Frequency granularity of time-domain beamforming, enumerated as "per component carrier" or "per band".
 - a. Value is only present for time-domain beamforming
- 5) Time granularity of time-domain beamforming, enumerated as "per-OFDM-symbol" or "per-slot".
 - a. Value is only present for time-domain beamforming
- 6) For each of the time-domain beams (the number is known from the O-RU antenna model), beam attributes (see 12.4.2.2)

The number of time-domain beam numbers associated with a given beamId (K') will be the same as the number of frequency-domain beamforming weights for that beamId; this is because each frequency-domain beamforming weight is applied to a data stream that is subsequently time-domain beamformed using a specific beam number, so if there are e.g. $K'=4$ frequency-domain weights associated with a given beamId there will be four time-domain beam numbers also associated with that same beamId. See Figure 12.4.3.4.3-1 for an example wherein four frequency-domain weights $\Phi_{x,y}$ are applied with four time-domain beam numbers (#3 and #5). Here the number of frequency-domain weights K' indicated by the antenna model would be "four" so the first four values in each vector would be complex fractional values and would be compressed, while the second four values would be real integers and not compressed.

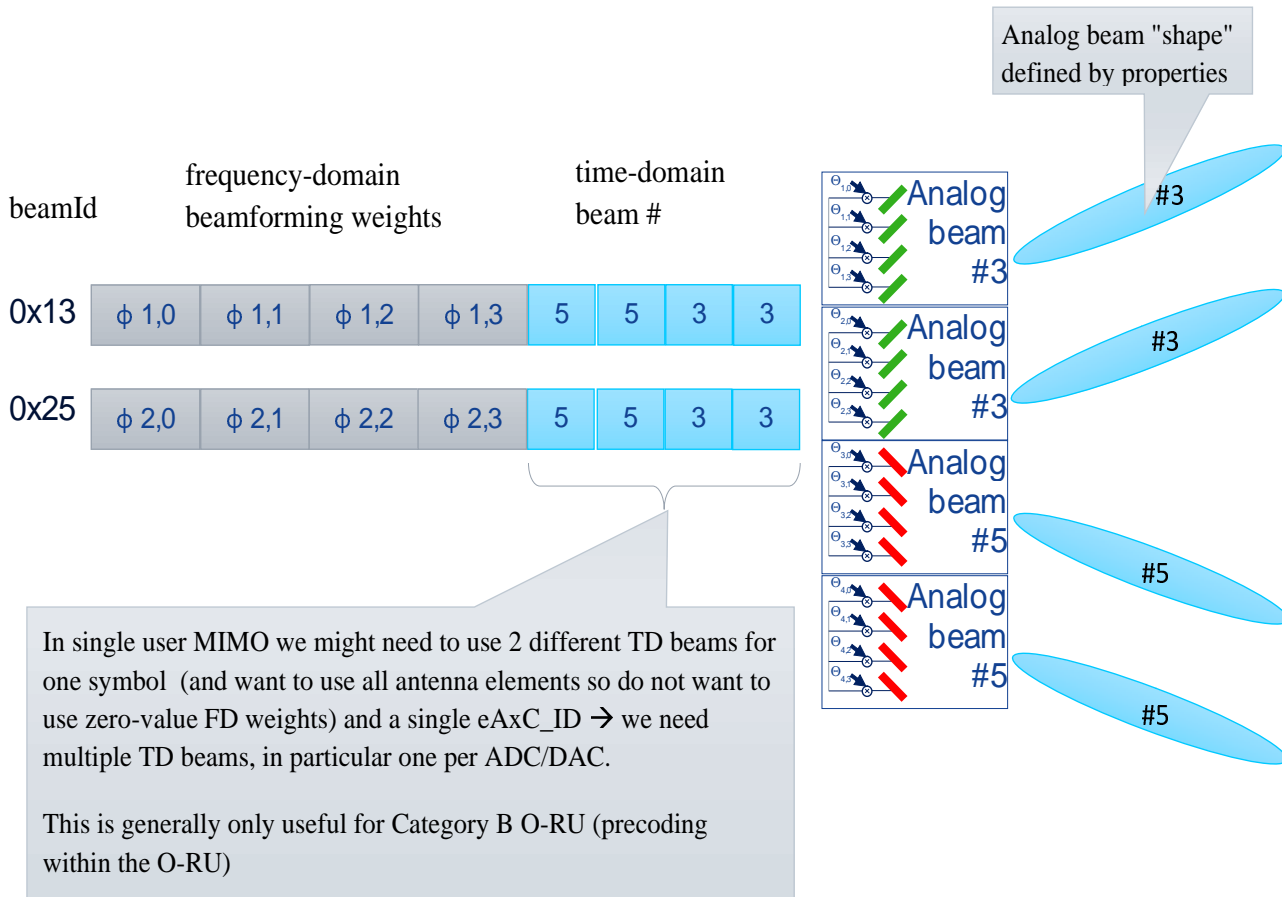


Figure 12.4.3.4.3-1: Example of updatable-weight frequency-domain plus fixed time-domain beamforming

12.4.3.4.4 Hybrid beamforming with Section Type 4 'TIME_DOMAIN_BEAM_CONFIG'

As described in clause 12.4.3.4.2 and clause 12.4.3.4.3, use of a common hybrid beamId to update two sets of beam information that have different update frequency (i.e. frequency-domain beamforming weights are expected to change per PRB and per spatial layer, however, the time-domain beamforming weights are expected to have a common value for all PRBs within a carrier and same for all spatial layers) results in repeating duplicate time-domain beamforming information per frequency-domain beamforming update resulting in redundant fronthaul usage. Also, analog beamforming information applies to the entire carrier or all carriers in a given band and hence is not subject to single endpoint processing.

Figure 12.4.3.4.4-1 and Figure 12.4.3.4.4-2 emphasize on the redundancy in each case as described in clause 12.4.3.4.2 and clause 12.4.3.4.3:

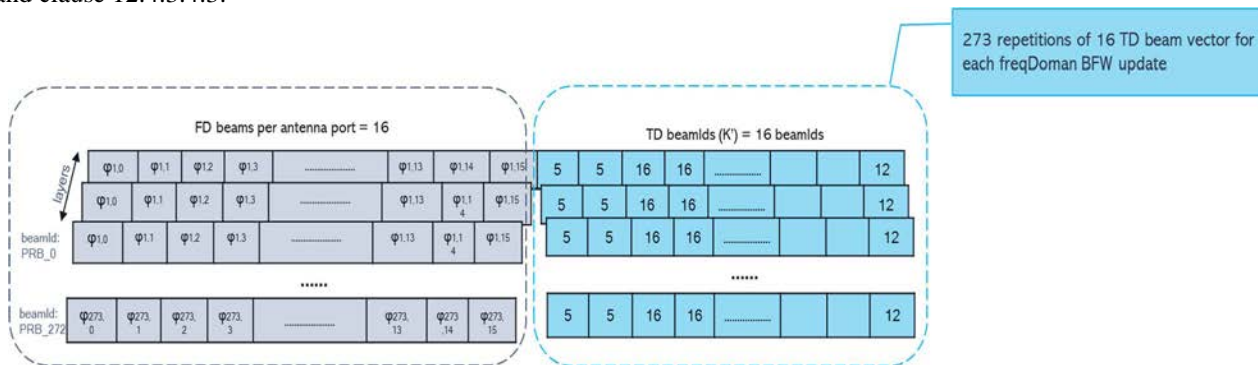


Figure 12.4.3.4.4-1: Example of repetition of time-domain beam number vector of size K' for every frequency-domain beamforming update




Figure 12.4.3.4.4-2: Example of repetition of time-domain beamforming weight vector of size $K' \cdot \text{numAntElements}$ for every frequency-domain beamforming update (per prb per layer)

The redundancy of repeating time-domain beamforming information as described in both the examples above may be avoided by decoupling time-domain and frequency-domain beamforming information using Section Type 4 with TIME_DOMAIN_BEAM_CONFIG command, see Figure 12.4.3.4.4-3. Using this Section Type 4 command, time-domain beamforming information may be sent separately.




Figure 12.4.3.4.4-3: Example of decoupled frequency-domain and time-domain beamforming information using Section Type 4 'TIME_DOMAIN_BEAM_CONFIG' command

When the O-DU configures the O-RU using the M-Plane parameter 'st4-for-time-domain-beamforming-weights-enabled', to use Section Type 4 to specify time-domain beamforming information for time-domain beamforming and/or hybrid beamforming, all other extensions i.e. SE 1, 3, 10, 11 used for specifying beamforming information shall contain frequency-domain beamforming information only. Also as defined in M-Plane, hybrid beamId space, time-domain beam-id space, and frequency domain beamId space are independently defined in M-Plane. Below requirements shall apply for specifying hybrid beams using Section Type 4 'TIME_DOMAIN_BEAM_CONFIG' command:

- M-Plane defined frequency-domain beam-id space shall be used to specify the frequency-domain portion of hybrid beam-id.
- M-Plane defined time-domain beamId space shall be used for each time-domain beam number specified in the 'TIME_DOMAIN_BEAM_CONFIG' command.
- Each time-domain beam number 'tdBeamNum' vector is identified by a 'tdBeamGrp'. This association is done by the O-RU as it receives a unique vector of 'tdBeamNum' associated with the unique 'tdBeamGrp' value. The total number of time-domain beam numbers per 'tdBeamGrp' shall be equal to the number of frequency-domain weights K' .
- Size of each time-domain beamforming vector is the number of antenna elements per frequency-domain beamforming port which is equal to p' . Value of p' shall be advertised by the O-RU as an M-Plane parameter 'p_dash'.

Note: $K'=K$ is a special case when $p'=1$.

12.4.4 Attribute-based dynamic beamforming

Attribute-based dynamic beamforming operates similarly to weight-based dynamic beamforming except that it is inherently a time-domain beamforming operation (both are "dynamic" meaning the definition of a beam as indicated by a beamId value may be changed via a C-Plane message). Also, instead of beamforming weights associated with a

specific beamId being sent from O-DU to O-RU in Section Extension 1, beam attributes associated with a specific beamId are sent instead from the O-DU to the O-RU in Section Extension 2.

The following list describes the information that the O-RU shall report via the M-Plane to the O-DU upon start-up as part of the O-RU self-description:

- 1) Beamforming type, possible choices are "frequency", "time" or "hybrid"– here will be "time".
- 2) Maximum number of beamId values supported (could be zero): "numBeams"
 - a. O-RUs may have memory limitations that mean the number of beams is limited; zero means no weight-based dynamic beamforming is supported by this tx-array or rx-array
- 3) Initial value of beamId supported: "initBeamId"
 - a. Different ranges of beamId may support generated beam beamforming versus e.g., predefined beamforming
- 4) Valid range of bfAzPt (see clause 7.7.2.3)
- 5) Valid range of bfZePt (see clause 7.7.2.4)
- 6) Valid range of bfAz3dd (see clause 7.7.2.5)
- 7) Valid range of bfZe3dd (see clause 7.7.2.6)
- 8) Valid range of bfAzSl (see clause 7.7.2.6)
- 9) Valid range of bfZeSl (see clause 7.7.2.7)

12.4.5 Channel-information-based beamforming

As stated in clause 12.4.1, beamId is irrelevant and unused in the case of channel-information-based beamforming

12.5 O-RU antenna model

12.5.1 Overview

Knowledge of O-RU antenna model is critical for certain types of beamforming. The following model is applicable for O-RU with one or more antennas, where each antenna has array of elements that are

- uniform (all elements have same properties) and
- organized into rectangular array (with rows and columns) that is planar (flat).

O-RU shall expose via M-Plane logical model of O-RU consisting of one or more arrays composed of one or more array elements. Array element represents independently controllable entity including one or more radiating elements and related RF processing elements (here, RF processing element is an entity that processes RF signal and is not related to processing element defined in M-Plane). RX and TX are in general independently controllable for that in the model tx-arrays and rx-arrays are described as separate entities. If O-RU supports beamforming, then beamforming is realized within each array separately i.e., beamforming weight vector is applicable within one array. One or more arrays can occupy same physical location e.g., rx-array and tx-array that use same set of radiators.

Beamforming methods that use dynamic beamforming with beamforming weights conveyed in C-Plane message (in contrast to predefined beams) require the O-DU to know antenna properties. Different beamforming methods require knowledge of different subsets of antenna properties.

12.5.2 Coordinate systems

Some of parameters describing model of antenna are related to coordinate system that defines three axes and three angles.

There are two coordinate systems defined:

- array coordinate system

- O-RU coordinate system

The **array coordinate system** is centered on the centre of the leftmost, bottom element of an array as presented on Figure 12.5.2-1. The array coordinate system is applicable within one array.




Figure 12.5.2-1: Array Coordinate System

The diagram presents view from the front of array-panel (from UE). Arrows indicate increasing values of coordinates.

x points towards broad-side

y increases to right, with antenna-array's columns

z points towards zenith

φ (phi) is azimuth angle, counter-clockwise rotation around z-axis, 0° points to broad-side, 90° points to y-axis

θ (theta) is zenith angle, counter-clockwise rotation around y-axis, 0° points to zenith, 90° points to horizon

ψ (psi) is angle, counter-clockwise rotation around x-axis, 0° points to horizon, 90° points to zenith

The **O-RU coordinate system** is a selected array coordinate system of an O-RU. The selection shall be fixed by O-RU design. The O-RU coordinate system is applicable within one O-RU.

12.5.3 O-RU antenna model parameters

The O-RU antenna model is described with the following parameters:

- **K** – number of array elements in array ($K = M \times N \times P \times Q$)
- **M** – number of rows of array elements in array. $M > 0$; value 0 is reserved for future use.
- **N** – number of columns of array elements in array. $N > 0$; value 0 is reserved for future use.
- **P** – number of polarizations in array. $P > 0$, value 0 is reserved for future use.
- **Q** – number of overlapping array elements (array-layers) in array. Each array-layer has M rows, N columns and P polarizations. $Q > 0$; value 0 is reserved for future use. See Annex K for more information on array-layers.
- **x, y, z** – position of centre of the leftmost, bottom element of array in O-RU coordinate system
- **φ, θ** direction of normal vector perpendicular to array's surface in O-RU coordinate system (array's normal vector corresponds to x axis in array's coordinate system)
- **dy** – mean distance between centres of nearby array elements in y direction in array coordinate system (distance between columns); value 0 is reserved for future use.

- **dz** – mean distance between centres of nearby array elements in z direction in array coordinate system (distance between rows); value 0 is reserved for future use.
- **list of polarizations in array** (this list has P elements, each representing p-th polarization); values ordered in ascending order of angle. Example: cross-polarized array having elements of one of two linear polarizations can be described by list: (-45°, +45°) indicating that array element with polarization index p=0 has linear polarization -45°, and array element with polarization index p=1 has linear polarization +45°.
- **independent power budget per layer** - in case of an array with multiple layers, corresponding elements (located in same row and column and same polarization) of different layers may have a shared power budget or have independent power budgets.

For an array supporting hybrid beamforming (see clause 12.4.2.2), additional parameters apply:

- **K'** – number of frequency domain beamforming weights $\phi_{k'}$ that can be applied within the array. $0 < K' \leq K$; value 0 is reserved for array not supporting hybrid beamforming.
- **h(k)** – mapping of array element k to frequency domain beamforming weight $\phi_{k'}$ where $k'=h(k)$. The mapping is represented as a list of lists: for every $0 < k' < K'$ a list of K/K' numbers identifying array elements where frequency domain beamforming weight $\phi_{k'}$ is applied. $k'=h(k)$ if number k is in the list corresponding to k' .
Clause 12.5.4 describes how numbers are assigned to array elements.

The model assumes the number of array elements corresponding to frequency domain beamforming weight $\phi_{k'}$ is the same for every k' ($0 < k' < K'$) and the elements corresponding to beamforming weights form a rectangular shape without overlapping i.e., every array element is linked with exactly one frequency domain beamforming weight $\phi_{k'}$.

In addition, the O-RU antenna model provides parameters describing key capabilities of array elements. The model assumes the array is uniform and all elements have the same properties. Each single value is applicable to all elements within the array.

Parameters describing array elements applicable to tx-array and rx-array:

- horizontal plane half power (-3 dB) beam width of array element's radiation pattern
- vertical plane half power (-3 dB) beam width of array element's radiation pattern
- horizontal plane quarter power (-6 dB) beam width of array element's radiation pattern
- vertical plane quarter power (-6 dB) beam width of array element's radiation pattern

Beam widths above are angles (expressed in degrees) between half-power (-3 dB) points or quarter-power (-6 dB) points respectively of the main lobe with reference to peak radiated power of main lobe. Horizontal and vertical plane correspond to the xy-plane and xz-plane respectively of the array in the array coordinate system.

The parameter describing array elements specific for tx-array:

- **$m_{a,k}$** - max rms power rating of array element of the array. Usage of max rms power rating is described in clause 8.1.3.3.

12.5.4 Identification and ordering of array elements

In many applications there is a need to assign to array element a number k such that $0 \leq k < K$ (see K in clause 12.5.3). One example is mapping position (represented by k, such that $0 \leq k < K$) of a beamforming weight in beamforming vector to an array element. Other example is identification of an array elements in the antenna model.

For purpose of identification and ordering a number k is assigned to each element of array by the function $f(m,n,p,q)$:

$$f(m,n,p,q) = q \cdot P \cdot M \cdot N + p \cdot M \cdot N + m \cdot N + n$$

where:

m - row (bottom to top), $0 \leq m < M$

n - column (left to right, view from the front of array), $0 \leq n < N$

p - polarization index, $0 \leq p < P$; polarization value of polarization index p is ψ_p (see ψ in clause 12.5.2)

q - array-layer, $0 \leq q < Q$

N, M, P, Q are defined in clause 12.5.3

NOTE: For a rectangular array, the function $f(m,n,p,q)$ can be inverted allowing to determine a "tuple" (m,n,p,q) of k -th element.

12.5.5 Relations between array elements

Beamforming methods that use dynamic beamforming with beamforming weights conveyed in C-Plane messages (in contrast to predefined beams) require the O-DU to know that specific elements of one array is co-located with elements of another array e.g., rx-array and tx-array that use same set of radiators. In addition, one or more tx-arrays may share elements and parts of RF processing paths (e.g., a power amplifier) resulting in a shared power budget described by a maximum rms power rating.

If element k_a of array A and element k_b of array B are in same position (same physical row and column) then k_a and k_b are co-located.

If element k_a of array A and element k_b of array B are co-located and share a power budget, then k_a and k_b are shared.

Relation of co-location is symmetric: k_a and k_b are co-located if and only if k_b and k_a are co-located.

Relation of co-location is transitive: if k_a and k_b are co-located and k_b and k_c are co-located then k_a and k_c are co-located.

Relation of sharing implies co-location: if k_a and k_b are shared then k_a and k_b are co-located.

O-RU shall report via M-Plane relations between array elements. O-RU shall avoid reporting redundant relations that can be derived from other relations by symmetric and transitive properties and implication of co-location relation by sharing relation. In addition, the O-RU shall provide a concise representation of the common case of two arrays that have all elements in relation (e.g., rx-array of -45° polarization and corresponding rx-array of $+45^\circ$ polarization).

12.5.6 Model usage

The O-RU antenna model reported by the O-RU consists of rx-arrays and tx-arrays. rx-arrays and tx-arrays represent a capability for transmitting/receiving RF signal related to an eAxC and - if beamforming is supported by O-RU on given array - beamforming capability. In this clause examples are presented: red and green bars represent array elements of different polarizations, grey box represents physical device, white rectangles represent arrays reported by O-RU.




Figure 12.5.6-1: Examples of Model Usage – TX as two single-polarization arrays or one cross-polarized array

As an example Figure 12.5.6-1 presents two O-RU designs: an O-RU with two tx-arrays each of one polarization and an O-RU with one tx-array of two polarizations. An O-RU that combines both above designs is possible as presented in Figure 12.5.6-2.

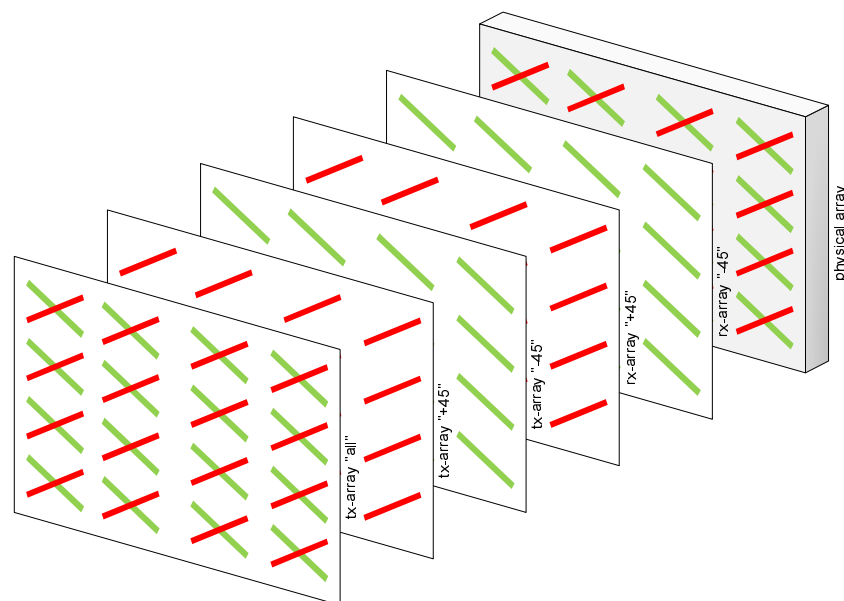


Figure 12.5.6-2: Examples of Model Usage – TX as two single-polarization arrays and one cross-polarized array

Similarly, in RX – if an O-RU does not support the collecting raw SRS by use of beams designed to mute signal from all elements except one then that O-RU – in addition to full rx-array – may expose rx-arrays with single elements as presented in Figure 12.5.6-3.




Figure 12.5.6-3: Example of Model Usage – RX with multi-element and multiple single-element arrays

An O-RU that does not support beamforming can be represented with arrays having one element each. Note that a category B O-RU could be represented with an array with two elements to accommodate two polarizations while a category A O-RU would generally be represented with a tx-array with only one polarization ($P=1$). Figure 12.5.6-4 presents the two design examples of a non-beamforming O-RU: an O-RU with two tx-arrays each of one polarization and an O-RU with one tx-array of two polarizations. An O-RU that combines both designs is also possible.

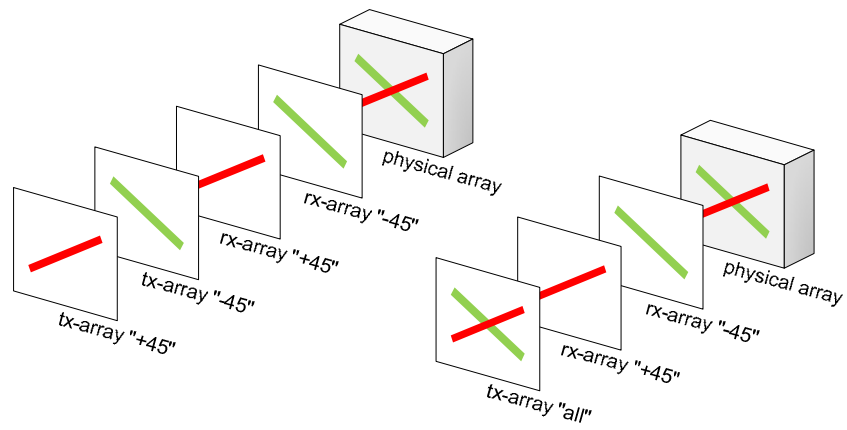


Figure 12.5.6-4: Examples of Model Usage – Non-beamforming O-RU with TX as two single-polarization arrays or one cross-polarized array

13 Support of Shared Cell

13.1 General

"Shared Cell" is defined as the operation for the same cell (the cell can have one or multiple component carrier(s)) by several O-RUs.

There are two approaches for realizing Shared Cell as shown in Figure 13.1-1.

- FHM mode:** Shared Cell is realized by FHM and several O-RUs. In this case, FHM is placed between the O-DU and multiple O-RUs. FHM is modelled as an O-RU with LLS Fronthaul support (same as normal O-RU) and copy and combine function (additional to normal O-RU), but without radio transmission/reception capability. For the copy and combine function, FHM may support selective transmission and reception function. In the present document, each O-RU under one FHM can be used for either operating the same cell (same cell scenario in Figure 13.1-1) or different cells (multiple cells scenario in Figure 13.1-1) by M-Plane configuration.
- Cascade mode:** Shared Cell is realized by several O-RUs cascaded in chain. In this case, one or more O-RU(s) are inserted between the O-DU and the O-RU. The O-RUs in the cascaded chain except for the last O-RU shall support Copy and Combine function. The O-RUs which support the Copy and Combine function are named "Cascade O-RU". Note that the last O-RU may also support Copy and Combine function although it is not used (i.e. not only normal O-RU but also Cascade O-RU may work as the last O-RU). In the present document, each Shared Cell O-RU in a physical chain can be used for either operating the same cell or different cells by M-Plane configuration; each cell is served by its own (logical) cascade chain.

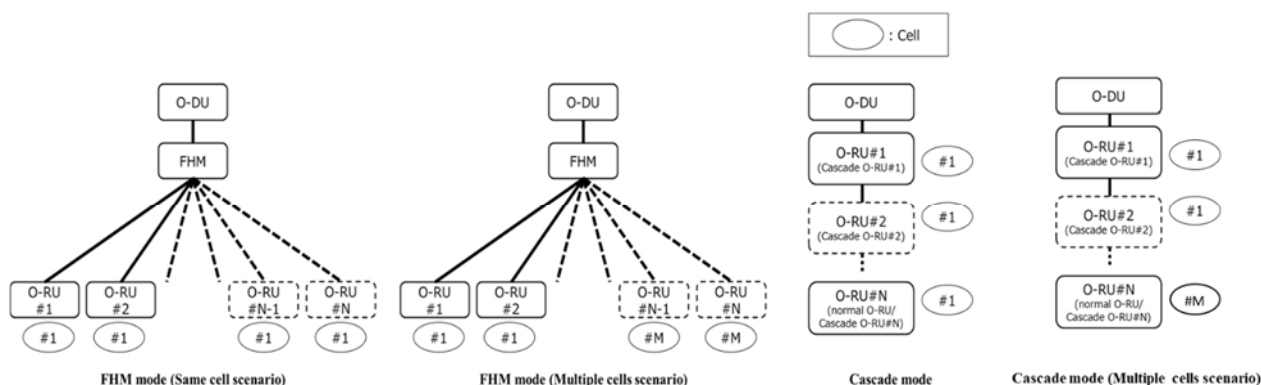


Figure 13.1-1: Shared Cell Concept

For the Cascade mode, the following generic reference to nodes in a cascaded chain is used.

north-node: For a particular O-RU in the chain of cascaded O-RUs, the north-node is the adjacent O-RU which is located closer to the O-DU. In the case of the O-RU which is adjacent to the O-DU, the north-node is the O-DU.

south-node: For the O-DU or a particular O-RU in the chain of cascaded O-RUs, the south-node is the adjacent O-RU which is located farther from the O-DU.

NOTE: For the FHM mode, the nodes are directly referred to as O-DU, FHM or O-RU, and the above terminologies are not used in the present document. The terminologies: north-node and south-node are also used for FHM in the M-Plane specification to maximize commonality for M-Plane modelling.

In this clause, the following notations are used;

N: the total number of O-RUs connected to a FHM for FHM mode, or the total number of O-RUs in a cascaded chain (including the last O-RU) for Cascade mode.

M : the total number of different cells realized by the total of N O-RUs connected to a FHM for FHM mode, or the total number of different cells realized by the total N O-RUs in a cascaded chain for Cascade mode. M equals 1 for same cell scenario for FHM mode and same cell scenario for Cascade mode. $M < N$.

N_m : the total number of O-RUs realizing a particular cell m . N_m equals N (and $m=1$) for same cell scenario for FHM mode and same cell scenario for Cascade mode. M is equal to N scenario in Cascade mode can be realized with Ethernet Forwarding and Boundary Clock capability without any Shared Cell specific configurations or capability.

n : used to denote a particular O-RU within a set of O-RUs realizing a particular cell.

13.2 Copy and Combine function

13.2.1 DL Copy function

In downlink case, FHM shall retrieve eCPRI messages coming from O-DU as payload of Ethernet frames, copy them (the entire eCPRI message including eCPRI header and eCPRI payload) without any modifications as payload into Ethernet frames and send them towards the O-RUs realizing the Shared Cell. FHM determines these O-RUs from M-Plane configuration. Figure 13.2.1-1 illustrates the same cell scenario where copy is done for O-RU#1 to O-RU# N . For the Multiple cells scenario, for a particular cell m , copy will be for O-RU#1 to O-RU# N_m .

In downlink case, Cascade O-RU shall retrieve eCPRI messages coming from the north-node as payload of Ethernet frames, copy them (the entire eCPRI message including eCPRI header and eCPRI payload) without any modifications as payload into Ethernet frames and send them towards the south-node.

DL Copy function is shown in Figure 13.2.1-1.

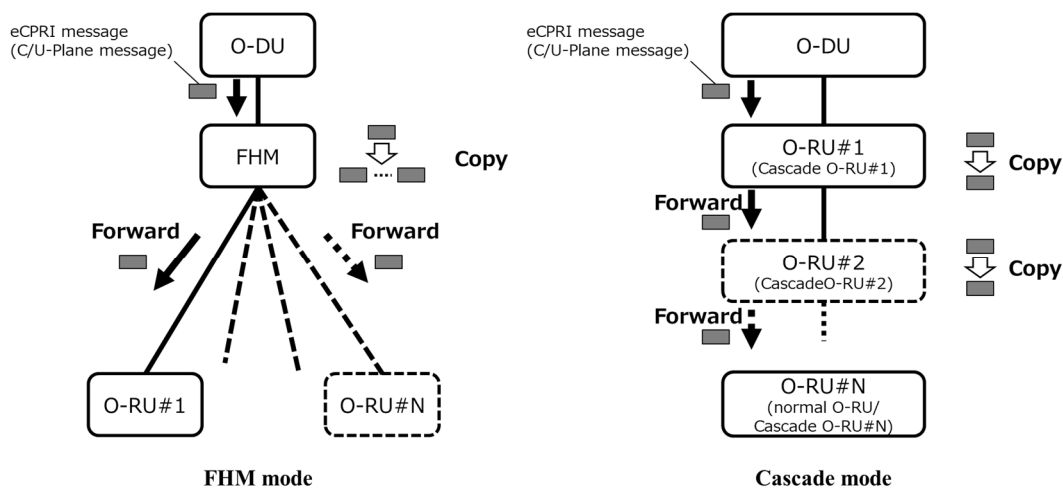


Figure 13.2.1-1: Downlink flow for Shared Cell

13.2.2 UL Combine function

FHM shall combine IQ data corresponding to the same radio resource element from the multiple eCPRI messages for UL U-Plane in the Ethernet frames transmitted from each O-RU, where the general principles are listed below;

- The FHM shall identify IQ data corresponding to the same radio resource element from the information elements in eCPRI transport header, application layer common header and application layer section field which are listed in Table 13.2.2-1. In case that transport layer fragmentation occurs, the FHM shall re-assemble the fragmented message first.

- The FHM shall retrieve compression information (if present), *iSample* and *qSample* from those corresponding eCPRI messages and calculate the combined *iSample* and *qSample* by adding *iSample* and *qSample* individually, taking compression information into account as below;

If U-Plane data compression is used then:

Combined *iSample* = *Compress (Sum (Decompress (iSample_#1), ... , Decompress (iSample_#N_m)))*,
 Combined *qSample* = *Compress (Sum (Decompress (qSample_#1), ... , Decompress (qSample_#N_m)))*,

Else

Combined *iSample* = *Sum (iSample_#1, ... , iSample_#N_m)*,
 Combined *qSample* = *Sum (qSample_#1, ... , qSample_#N_m)*,

where

iSample_#n is the *iSample* received from the O-RU#*n* and
qSample_#n is the *qSample* received from the O-RU#*n*.

- If overflow occurs from the combine operation, *iSample* and/or *qSample* should be clamped to the closest value that can be represented in the compression format used for the combined IQ data.
- The compression format for the combined IQ data shall be the same as that in the IQ data received from the O-RUs.
- In the present document, selective combining of UL signals is not supported.

After combining, the FHM shall generate eCPRI header (the *ecpriSeqid* field generated by the FHM is independent and not a copy of that received from O-RUs; other fields are just copies of those received from the O-RUs.), add combined IQ data as eCPRI message body, construct Ethernet frames carrying eCPRI messages as payload, and forward them to the O-DU.

Cascade O-RU shall combine IQ data corresponding to the same radio resource element from air and from the eCPRI messages for UL U-Plane in the Ethernet frames transmitted from the south-node, where the general principles are listed below;

- The Cascade O-RU shall identify IQ data corresponding to the same radio resource element from the information elements in eCPRI transport header, application layer common header and application layer section field which are listed in Table 13.2.2-1. In case that transport layer fragmentation occurs, the Cascade O-RU shall re-assemble the fragmented message first.
- The Cascade O-RU shall retrieve compression information (if present), *iSample* and *qSample* from those corresponding eCPRI messages and calculate the combined *iSample* and *qSample* by adding *iSample* and *qSample* individually, taking compression information into account as below;

If U-Plane data compression is used then:

Combined *iSample* = *Compress (Sum (Decompress (iSample_from_south-node) and (iSample_from air)))*,
 Combined *qSample* = *Compress (Sum (Decompress (qSample_from_south-node) and qSample_from air))*,

Else

Combined *iSample* = *Sum (iSample_from_south-node and iSample_from air)*,
 Combined *qSample* = *Sum (qSample_from_south-node and qSample_from air)*,

where

iSample_from_south-node is the *iSample* received from the south-node,
qSample_from_south-node is the *qSample* received from the south-node,
iSample_from_air is the *iSample* received from the air, and
qSample_from_air is the *qSample* received from the air.

- If overflow occurs from the combine operation, *iSample* and/or *qSample* should be clamped to the closest value that can be represented in the compression format used for the combined IQ data.
- The compression format for the combined IQ data shall be the same as that in the IQ data received from the south-node.
- In the present document, selective combining of UL signals is not supported.

After combining, the Cascade O-RU shall generate eCPRI header (the `ecpriSeqId` field generated by the Cascade O-RU is independent and not a copy of that received from the south-node; other fields are just copies of those received from the south-node.), add combined IQ data as eCPRI message body, construct Ethernet frames carrying eCPRI messages as payload, and forward them to the north-node.

Figure 13.2.2-1 illustrates the same cell scenario where combine operation is done for O-RU#1 to O-RU#N. For the Multiple cells scenario, for a particular cell m , combine operation is done for O-RU#1 to O-RU#N _{m} .

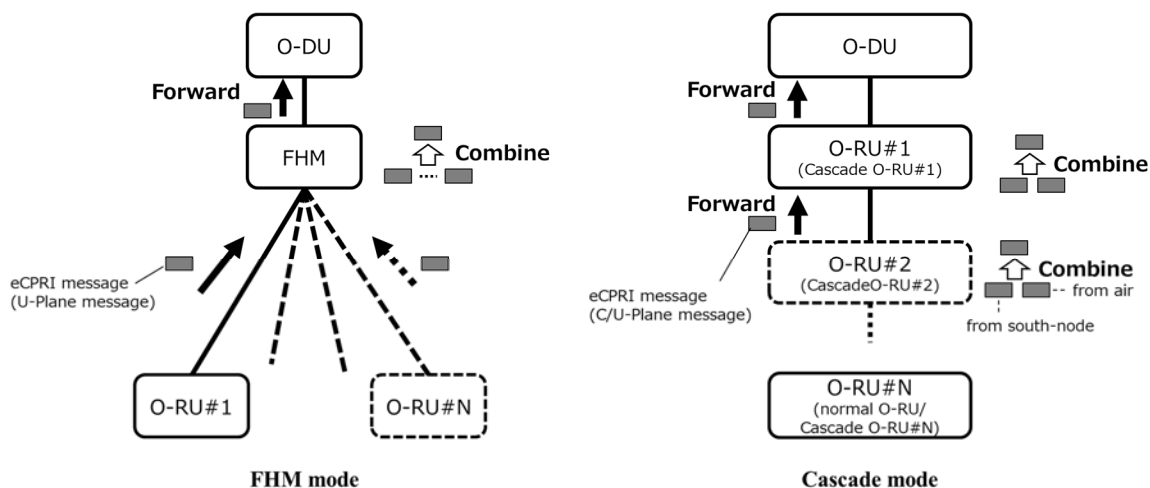


Figure 13.2.2-1: Uplink flow for Shared cell

Table 13.2.2-1: information elements to be checked when FHM/Cascade O-RU combines UL U-Plane messages

Field	Information element
eCPRI transport header	<code>ecpriPcid</code>
Application layer common header	<code>dataDirection</code> , <code>frameId</code> , <code>subframeId</code> , <code>slotId</code> , <code>symbolId</code>
Application layer section field	<code>rb</code> , <code>symInc</code> , <code>startPrbu</code> , <code>numPrbu</code>

13.3 Selective transmission and reception using beamId

13.3.1 Introduction

In order to avoid unnecessary interference and noise enhancement, selective transmission and reception function can be useful and it can be realized with minimum implementation impact on both O-DU/O-RU by using `beamId`.

For the selective transmission and reception function using `beamId`:

- Predefined-beam beamforming is assumed.
- Applicability for Cascade mode (cascaded O-RUs or cascaded FHMs) is out of scope in the present document.

The concept is to utilize the multiple O-RUs involved in a Shared Cell collectively as a single beamforming O-RU, and to realize selective transmission and reception using beamforming operations.

- In normal beamforming operation, a unique SSB and/or CSI-RS are transmitted from each beam so that the O-DU scheduler can understand which beam should be used based on L1-RSRP reports from the UE; where each is further mapped to a unique `beamId` that the O-DU signals on the fronthaul.

- For Shared Cell with selective transmission and reception using beamId:
 - In the simplest form, all O-RUs involved can be non-beamforming O-RUs, and each of them can be regarded as separate beam of one collective beamforming O-RU. In this case, a unique SSB and/or CSI-RS are transmitted from each non-beamforming O-RU; where each shall be further mapped to a unique beamId that the O-DU signals on the fronthaul. Then the FHM, based on the beamId signalled from the O-DU, shall route the messages to the appropriate non-beamforming O-RU. At this time, the FHM also shall translate the beamId so that the O-RU can transmit/receive using the appropriate beam. This assumes each O-RU handles unique user data meaning the FHM is not executing any copy/combining operation.
 - Other forms can be considered. For example:
 - All O-RUs involved can still be non-beamforming O-RUs, but multiple O-RUs illuminating overlapping/neighbouring areas can be made to transmit a common SSB/CSI-RS (and common user data) and be regarded as a common beam from the O-DU perspective. The FHM, based on the beamId signalled from the O-DU which corresponds to the common beam, shall route the messages to the appropriate set of multiple O-RUs and shall execute a copy/combining operation for this data.
 - Multiple beamforming O-RUs can collectively be considered as a single beamforming O-RU. Each beam of the individual O-RUs may be regarded as a separate beam, or multiple beams illuminating overlapping/neighbouring areas can be made to transmit a common SSB/CSI-RS (and common user data) and be regarded as a common beam from the O-DU perspective. The FHM, based on the beamId signalled from the O-DU, shall route the messages to the appropriate beamforming O-RU and shall execute a copy/combining operation for this data, and also translate the beamId value so that the O-RU can transmit/receive using appropriate beams.

NOTE: For the O-RU beams illuminating overlapping/neighbouring areas that are made to transmit a common SSB/CSI-RS and are regarded as a common beam from the O-DU perspective, coherency between the beams is not required.

From here on, the beamId used by the O-DU and signalled from the O-DU to FHM is referred to as the global beamId; and the beamId signalled from the FHM to O-RU and used by the O-RU is referred to as the local beamId.

Related M-Plane aspects are described below:

- If O-DU receives the capability from FHM via M-Plane which indicates selective transmission and reception support, O-DU can configure the FHM to use selective transmission and reception function.
- The mapping information between each global beamId, O-RU(s) and their local beamId is configured to the FHM during the M-Plane start-up procedure.
- The inter-beam relationship information such as coarse-fine-beam-relation and neighbour-beams are reported from each O-RU via M-Plane. Therefore, although O-DU considers the set of O-RUs as one beamforming O-RU, coarse-fine-beam-relation and neighbour-beams information across different O-RUs should not be available.

The details of copy combine function with selective transmission and reception function using beamId are described in clause 13.3.2 and clause 13.3.3.

13.3.2 Selective transmission function

In downlink case, FHM shall retrieve eCPRI messages coming from O-DU as payload of Ethernet frames, read C/U-Plane message-section field and separate C-Plane message section fields based on beamId and also separate U-Plane message section header fields/PRB fields based on sectionIds corresponding to each beamId. Then FHM shall create new eCPRI messages by adding common header field on each new C/U-Plane message which may includes multiple separated section fields for C-Plane and section header fields/PRB fields for U-Plane belonging to the same beam. The FHM shall send them towards appropriate O-RUs according to the mapping table between beamId value and O-RU(s). For C-Plane messages, FHM needs to overwrite beamId field according to the mapping table between global beamId and local beamId.

Figure 13.3.2-1 illustrates the same cell scenario where selective transmission is done for O-RU#1 to O-RU#N. For the Multiple cells scenario, for a particular cell *m*, selective transmission will be done for O-RU#1 to O-RU#N_{*m*}.

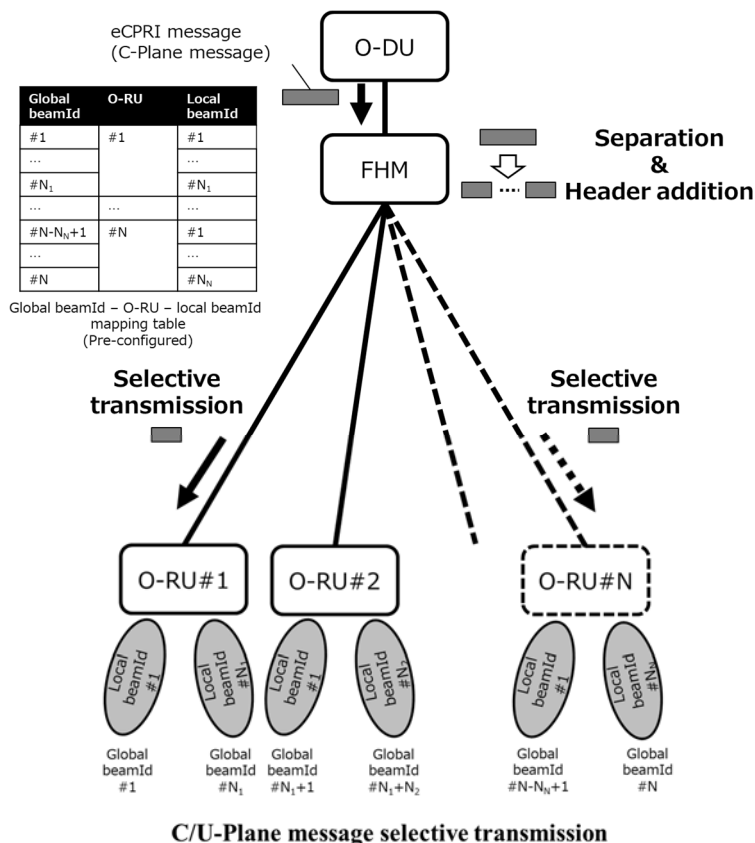


Figure 13.3.2-1: Downlink flow for selective transmission

13.3.3 Selective reception function

FHM shall send C-Plane messages as a same manner of selective transmission, i.e. the FHM shall send C-Plane messages to RUs selectively based on the beamId value in the C-Plane message received from O-DU and the mapping table between global beamId value and O-RU; and overwrite beamId value according to the mapping table between global beamId value and local beamId value for the O-RU.

Then corresponding O-RU(s) transmits UL U-Plane messages which includes UL signal received by using the beam indicated by the beamId value in the C-Plane message received from FHM.

FHM shall combine IQ data according to clause 13.2. In case UL U-Plane message received from different O-RUs do not include IQ data for the same radio resource element, combined IQ data is equivalent to IQ data received from particular O-RU.

FHM shall reconstruct eCPRI payload carrying UL U-Plane messages to be sent to the O-DU by concatenating combined IQ data and corresponding section header field with one UL U-Plane message common header.

After eCPRI payload reconstruction, the FHM shall generate eCPRI header (the ecpriSeqId field generated by the FHM is independent and not a copy of that received from O-RUs; other fields are just copies of those received from the O-RUs.), construct Ethernet frames carrying eCPRI messages as payload, and forward them to the O-DU.

Figure 13.3.3-1 illustrates the same cell scenario where selective reception is done for O-RU#1 to O-RU#N. For the Multiple cells scenario, for a particular cell, selective reception will be done for O-RU#1 to O-RU#N_{*m*}.

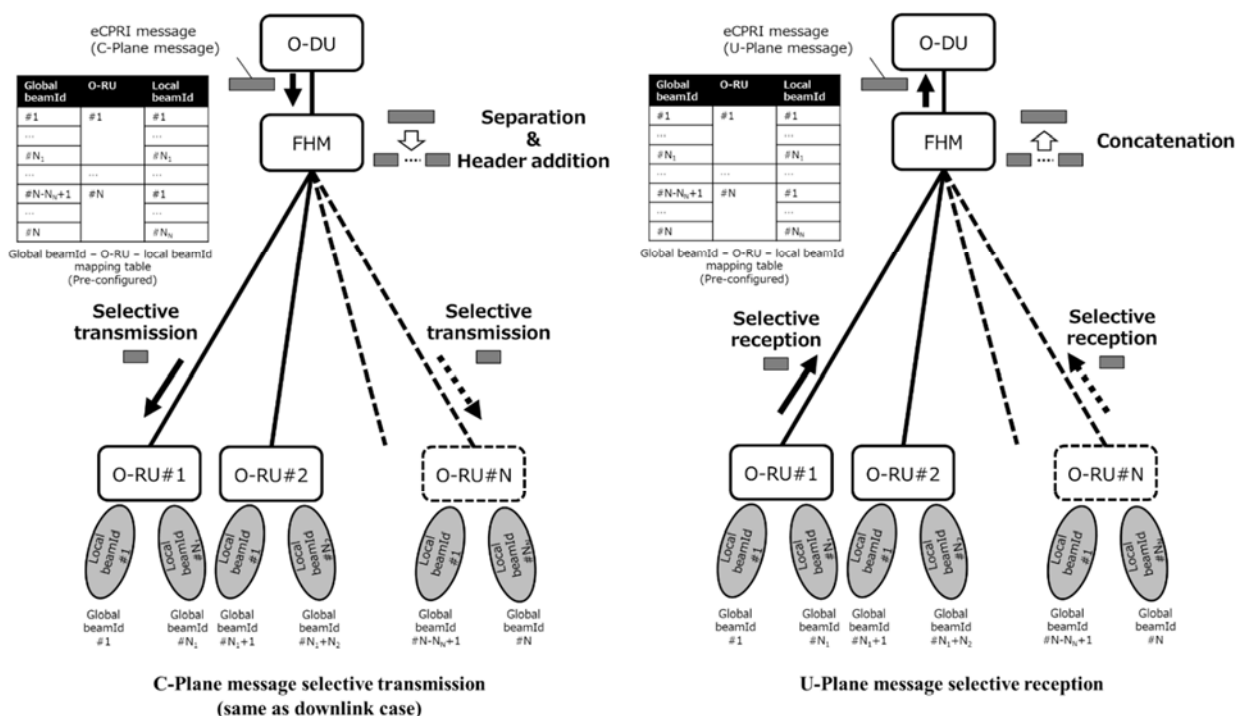


Figure 13.3.3-1: Uplink flow for selective reception

Beam mapping parameters for selective transmission and reception using beamId are summarized in Table 13.3.3-1.

Table 13.3.3-1: Beam mapping parameters for selective transmission and reception using beamId

Parameter	Description	Requirement
global beamId	This parameter indicates the beam ID to be applied to the U-Plane data on O-DU as same as existing beamId. One global beamId can be mapped to one or multiple local beamId. The mapping information between each global beamId, O-RU(s), and local beamId is configured to the FHM during M-Plane start-up procedure.	Since FHM and O-RUs are treated as a beamforming O-RU, global-beam-id = 0 shall not be used. Value range is same as that of existing beamId explained in 7.5.3.9..
local beamId	This parameter indicates the beam ID to be applied to the U-Plane data on O-RU as same as existing beamId. In one O-RU, different beams can not map to one global beamId. Local beamId shall be unique within O-RU.	

13.4 Multiple Cells Scenario in Cascade mode

13.4.1 Introduction

The details of multiple cells scenario in Cascade mode using Ethernet Transparent Bridging is described below.

13.4.2 Multiple Cell scenario in downlink

Multiple Cell scenario in downlink is shown in Figure 13.4.2-1.

O-RU shall retrieve eCPRI messages as payload of Ethernet frames, and read MAC address in Ethernet header (see clause 5.1.1). If O-RU receives packet from the north-node with 'destination MAC address' same as its own address, then O-RU shall perform DL copy function and send them towards the south-node as defined in clause 13.2 for Cascade mode. Otherwise, the O-RU shall forward whole of the Ethernet frame to the target MAC address without any

modification (Ethernet Transparent Bridging). Delay management for Shared Cell as defined in clause 13.5 shall be applied for the cascade chains serving a single cell.

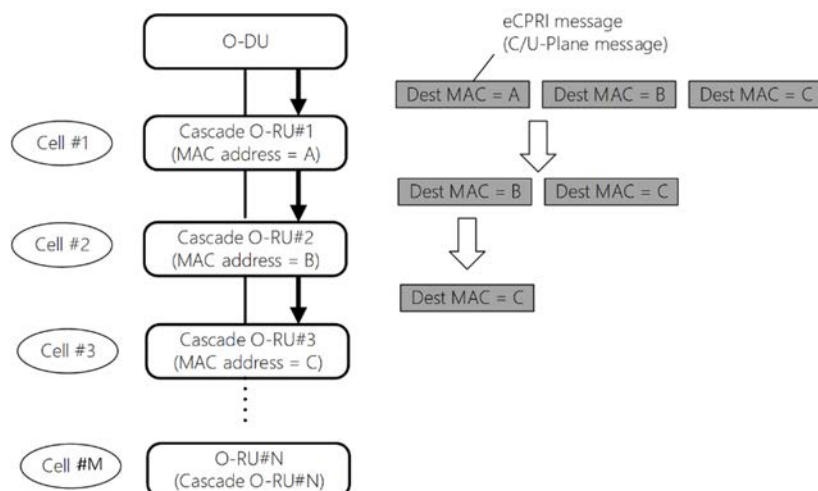


Figure 13.4.2-1: Downlink flow for Multiple Cell scenario

13.4.3 Multiple Cell scenario in uplink

Multiple Cell scenario in uplink is shown in Figure 13.4.3-1.

O-RU shall send C/U-Plane messages in the same manner as in downlink operation, i.e., if O-RU receives UL packets with "destination MAC address" different as its own address, O-RU shall forward the UL packets to the destination MAC address, otherwise O-RU shall perform UL Combine function for Cascade O-RUs and forward to its north-node as specified in clause 13.2.

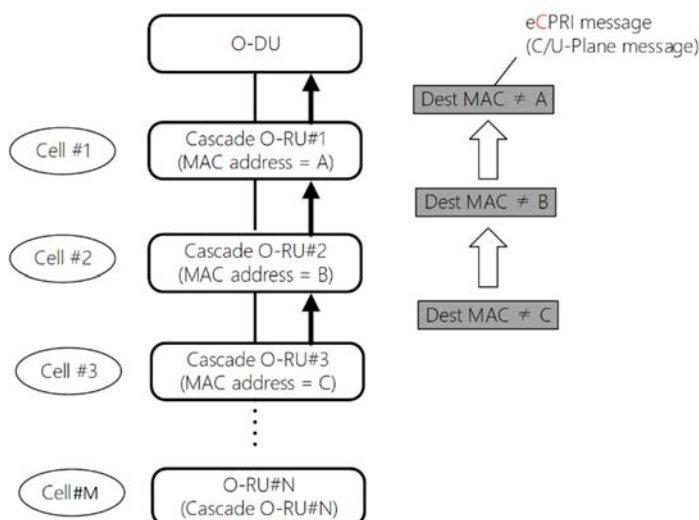


Figure 13.4.3-1: Uplink flow for Multiple Cell scenario

An O-RU shall indicate whether it can support Multiple Cell scenario in Cascade mode as a capability. The cascaded O-RU supporting this feature can support Shared Cell at the same time. As shown in the Figure 13.4.3-2 and Figure 13.4.3-3, this means that it can support mixed configuration of different cells and Shared Cells by combining the Ethernet forwarding and Shared Cell operation. However, a cascaded O-RU can be configured only for one Shared Cell for a given frequency. That is, a cascaded O-RU cannot simultaneously support cell #0 and cell #1 composed of different O-RU groups in the cascade chain if both cell #0 and cell #1 operate in same frequency. Cell #0 and cell #1 shall follow the delay management as defined in clause 13.5.

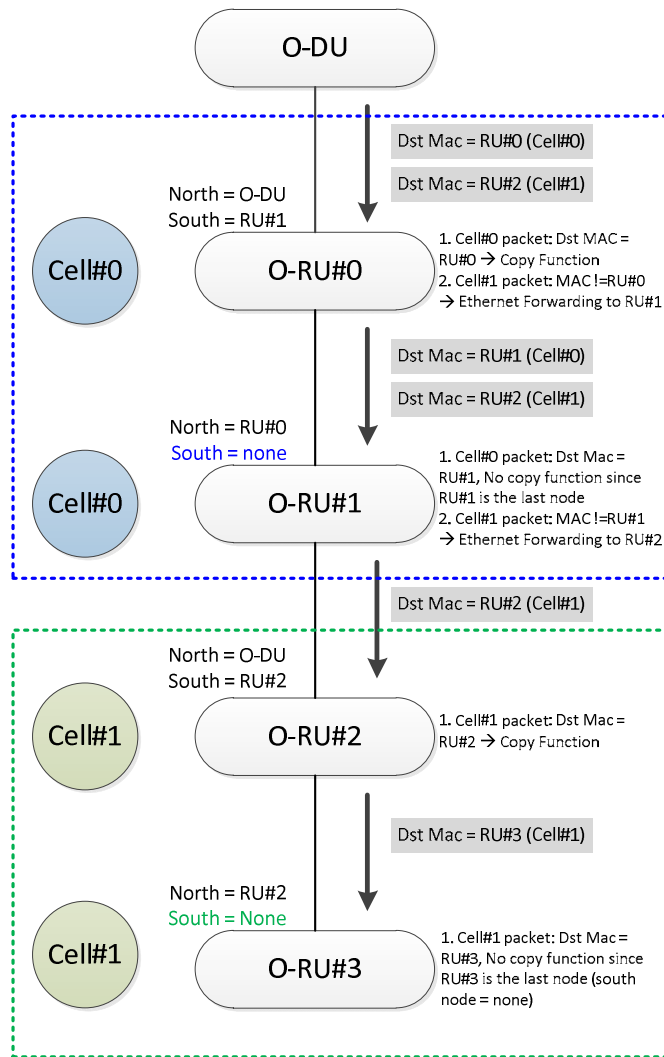


Figure 13.4.3-2: Example#1 Serving different cells and Shared Cell in a cascade O-RU chain

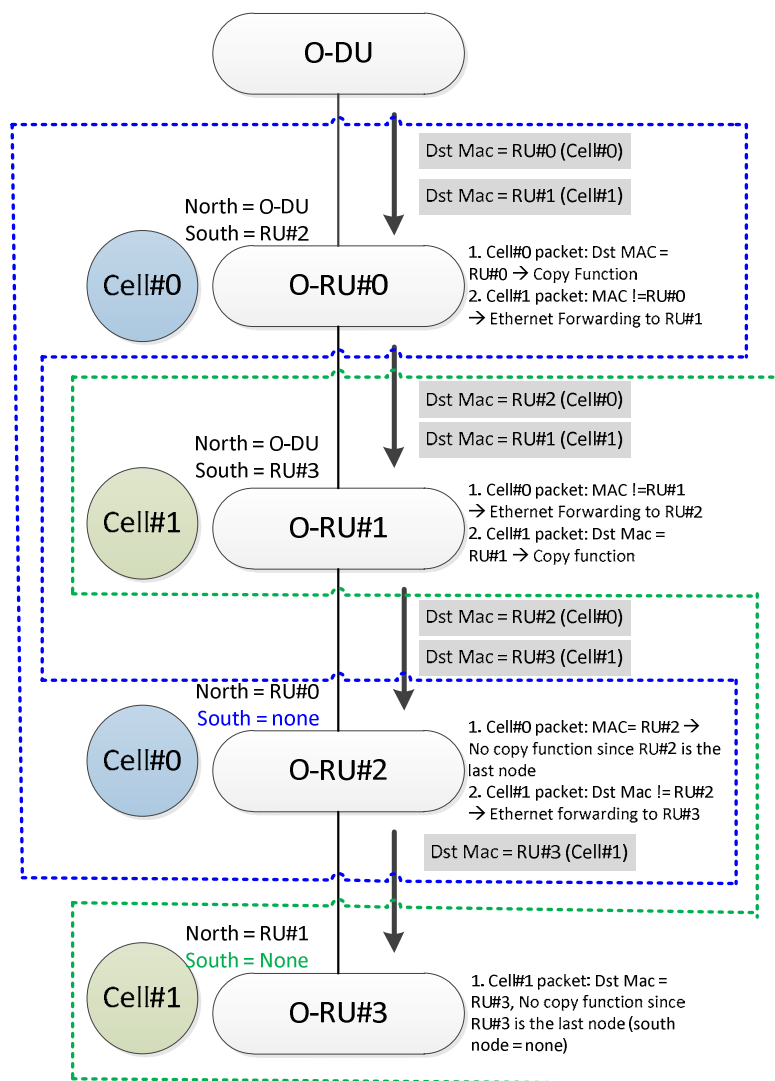


Figure 13.4.3-3: Example#2 of Serving different cells and Shared cell in a cascade O-RU chain

13.5 Delay management for Shared cell

13.5.1 Introduction

Total fronthaul distance between O-DU and O-RU shall be reduced compared to the case when there are no FHM/Cascade O-RU. This is required in order to keep the same total fronthaul delay between O-DU and O-RU even in the presence of processing time at FHM/Cascade O-RU, and to ensure UL messages arrive at O-DU within O-DU reception window as well as DL messages arrive at O-RU within O-RU reception window. Further the transmission/reception timing of FHM/Cascade O-RU needs to be clarified. So, additional requirements than those in clause 4.3 apply.

13.5.2 DL delay management for Shared cell

DL delay model parameters for FHM and Cascade mode are shown in Figure 13.5.2-1 and Figure 13.5.2-2.

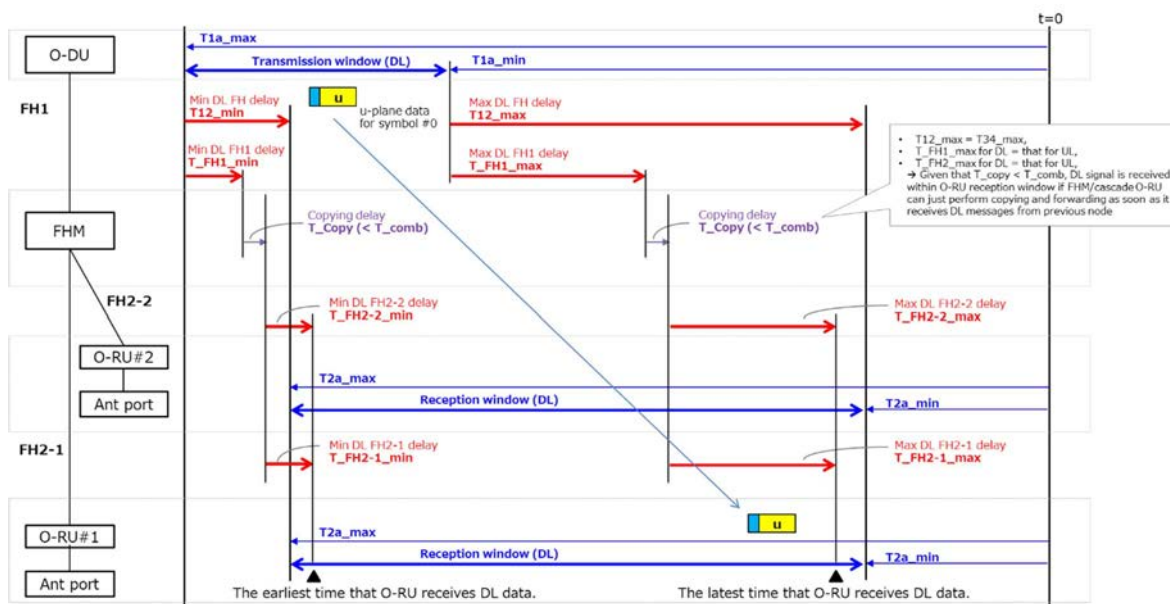


Figure 13.5.2-1: Delay model parameters for FHM mode (2 O-RUs case, i.e. $N_m = 2$, is illustrated as an example)

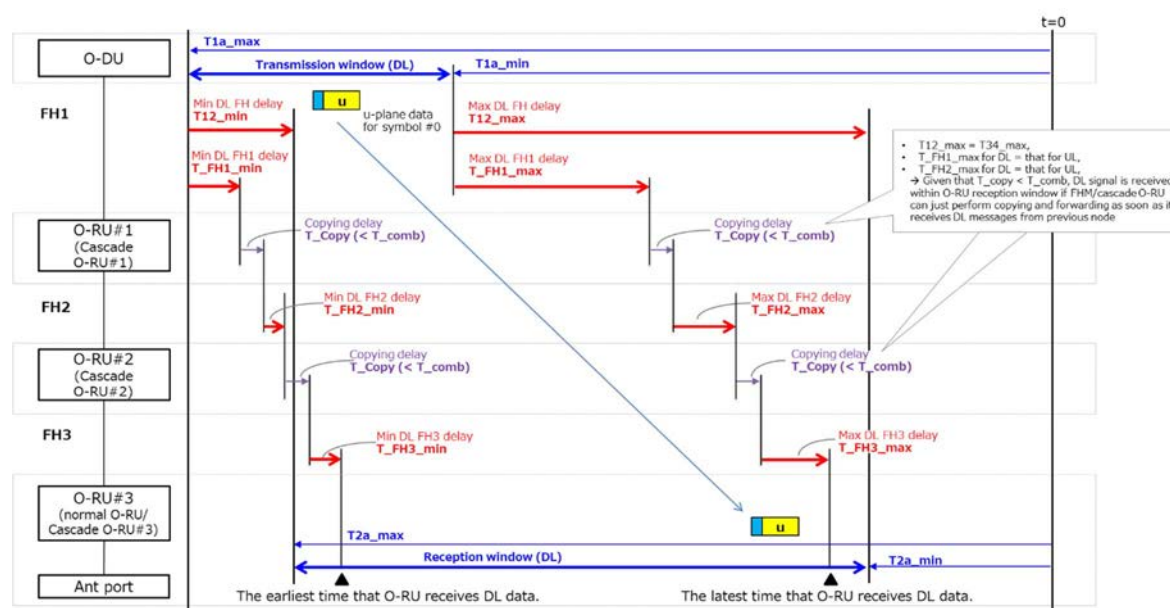


Figure 13.5.2-2: Delay model parameters for Cascade mode (3 O-RUs case, i.e. $N=3$, is shown as an example)

It is assumed that additional delay due to combining in UL case is larger than copying in DL case. Therefore, as long as fronthaul and FHM/Cascade O-RU are configured to satisfy delay management for UL case, nothing additional is needed for DL delay management with Shared cell. i.e., FHM/Cascade O-RU can just perform copying and forwarding as soon as it receives DL messages from O-DU (FHM mode) or north-node (Cascade mode), and this ensures that DL messages are received within O-RU reception window. Although the additional requirement for DL is not needed, if O-DU knows processing delay for copy which is shorter than T_{Comb} , it can be beneficial for O-DU transmission window optimization. For this reason, T_{Copy} is specified as O-RU capability.

It will need additional consideration for FDD cell (DL only band) formed of one DL carrier.

Delay parameter for Shared cell for DL is summarized in Table 13.5.2-1.

Table 13.5.2-1: Delay parameters for Shared cell for DL

Parameter	Description	Additional information
T_Copy	Corresponding to the maximum FHM or Cascade O-RU processing delay between receiving an IQ sample over the fronthaul interface from O-DU (FHM mode) or the north-node (Cascade mode), coping it and transmitting it over the fronthaul interface to O-RU (FHM mode) or the south-node (Cascade mode). In addition to the above, in FHM mode with selective transmission, T_Copy also includes reconstructing U-Plane message.	Capability parameter. FHM or Cascade O-RU reports the value by M-Plane.

13.5.3 UL delay management for Shared cell

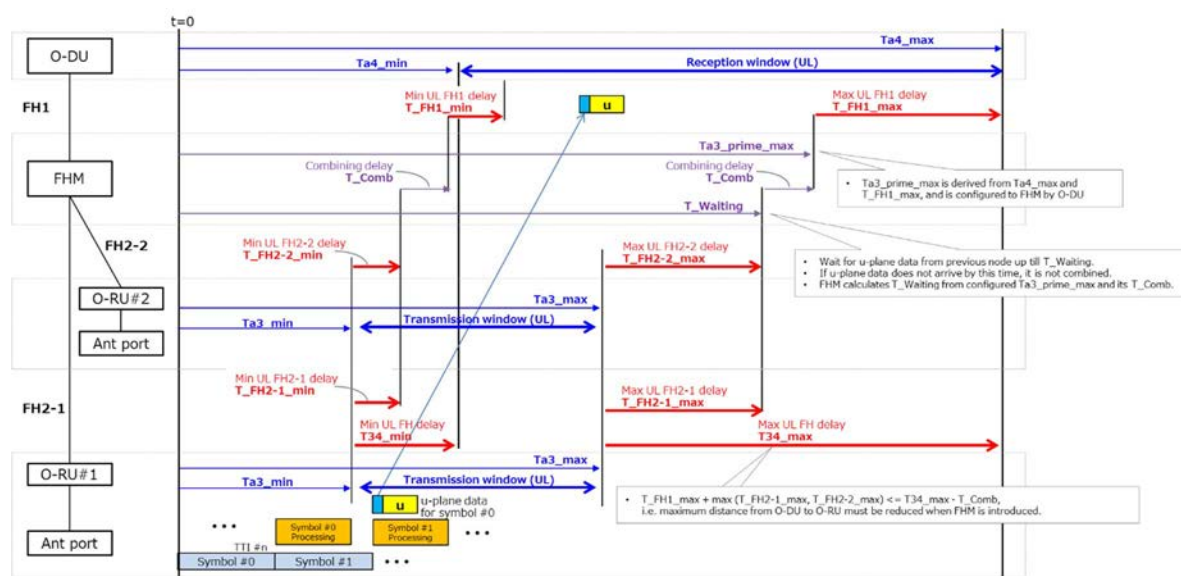


Figure 13.5.3-1: Delay model parameters for FHM mode (2 O-RUs case, i.e. $N_m = 2$, is illustrated as an example)

UL delay model parameters for FHM mode are shown in Figure 13.5.3-1. Since FHM processing delay for combining UL U-Plane messages effectively adds to the total fronthaul delay between O-DU and O-RU, total fronthaul distance between O-DU and O-RU shall be reduced compared to the case when there is no FHM in order to keep the same total fronthaul delay between O-DU and O-RU so as to ensure UL U-Plane messages arrive at O-DU within O-DU reception window. In other words, configuration of fronthaul including FHM shall meet;

$$T_{FH1_max} + \max(T_{FH2-1_max}, \dots, T_{FH2-N_m_max}) \leq T_{34_max} - T_{Comb},$$

where

T_{FH1_max} is the maximum transport delay between FHM and O-DU,

T_{FH2-n_max} is the maximum transport delay between O-RU# n and FHM ($n = 1, \dots, N_m$), and

T_{Comb} is the sum of the transmission duration over the fronthaul interface for messages within a specific time ($T_{x_Duration}$) and FHM processing delay for combining IQ samples for a message ($T_{Combine_net}$). The FHM reports maximum value as capability via M-Plane.

To ensure that UL U-Plane messages arrive at O-DU within O-DU reception window, the latest time FHM can send combined UL U-Plane messages towards O-DU shall be " $Ta4_max - T_{FH1_max}$ ". Since the value of T_{FH1_max} depends on position of FHM, FHM needs to be told about the latest time. This time is defined as $Ta3_prime_max$ and the O-DU shall configure it to FHM via M-Plane [7]. SCS associated with $eAxC_ID$ should be informed to FHM via M-Plane. Considering T_{Comb} , this means that UL U-Plane messages received by " $Ta3_prime_max - T_{Comb}$ " are subject to UL U-Plane message combining. In other words, even if FHM does not receive all UL U-Plane messages

from O-RUs by "Ta3_prime_max – T_Comb", FHM shall combine whatever received UL U-Plane messages and send it to O-DU by Ta3_prime_max. This time "Ta3_prime_max (configured) – T_Comb (FHM capability)" is defined as T_Waiting and calculated at FHM.

When IQ samples should be sent with multiple messages or jumbo frames, transmission duration for them over fronthaul interface should be considered. While T_Combine_net has the small variance upon cases, Tx_Duration usually has a large variance and more dependent upon cases. When Tx_Duration is not configured though **enhanced-t-combine-enabled** is configured via M-Plane as 'true', FHM shall calculate the amount of this parameter from the configured U-Plane parameters; herein Tx_Duration is the maximum transmission duration for messages to be sent within Ta3_prime_max over the fronthaul interface to north-node or O-DU. The time "Ta3_prime_max (configured) – T_Combine_net (FHM capability) – Tx_Duration (configured or calculated)" is defined as T_Waiting and calculated at FHM supporting ENHANCED-T-COMBINE when **enhanced-t-combine-enabled** is configured as 'true' via M-Plane.

In addition, in order to combine all UL U-Plane messages from O-RUs, T_Waiting needs to be larger than or equal to "Ta3_max + max (T_FH2-1_max, ... , T_FH2-N_m_max)". Therefore, O-DU shall configure Ta3_prime_max which meets following condition:

$$Ta3_{max} + \max (T_{FH2-1_max}, \dots , T_{FH2-N_m_max}) + T_{Comb} \leq Ta3_prime_max \leq Ta4_max - T_{FH1_max}$$

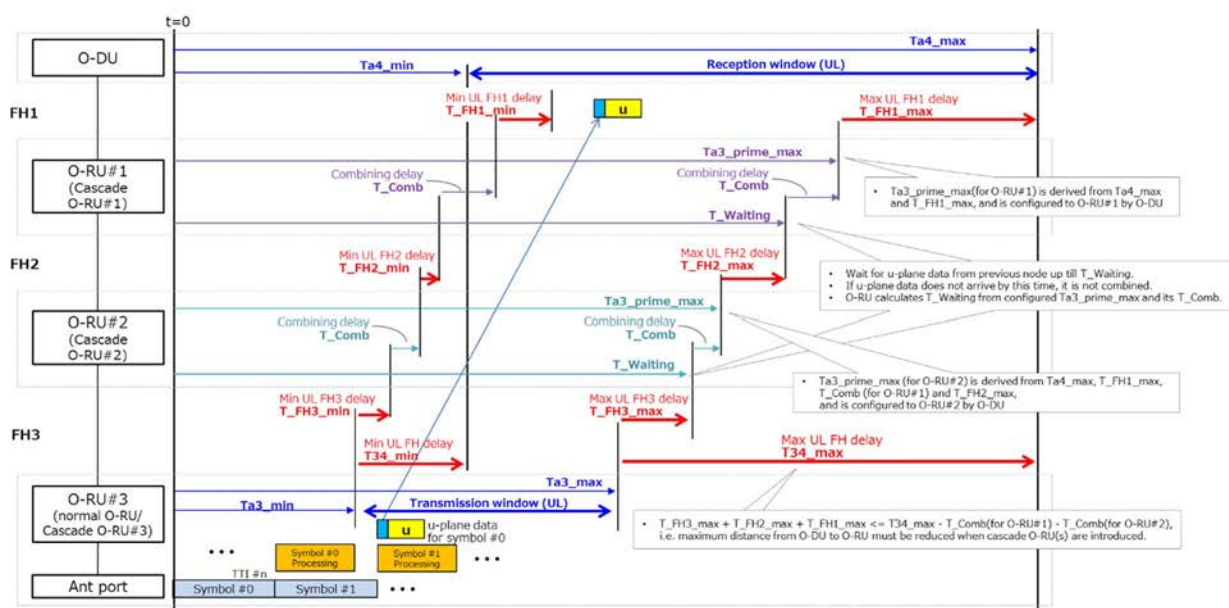


Figure 13.5.3-2: Delay model parameters for Cascade mode (3 O-RUs case, i.e. N=3, is shown as an example)

UL delay model parameters for Cascade mode are shown in Figure 13.5.3-2. In a similar manner with FHM mode, configuration of fronthaul including Cascade O-RU shall meet:

$$\sum_{i=1}^N T_{FH_i_max} \leq T_{34_max} - \sum_{i=1}^{N-1} T_{Comb} \text{ (for Cascade O-RU\#i)}$$

where:

T_FHn_max is

the maximum transport delay between O-RU#1 and O-DU (n=1),

the maximum transport delay between O-RU#n and O-RU#n-1 (n=2 ... N),

T_Comb is the sum of the transmission duration over the fronthaul interface for messages within a specific time (Tx_Duration) and Cascade O-RU processing delay for combining IQ samples for a message (T_Combine_net), The Cascade O-RU reports maximum value as capability via M-Plane.

T34_max is the maximum uplink transport delay between O-DU and the last O-RU (as specified in clause 4.3).

In addition, in a similar manner with FHM mode, Ta3_prime_max configured to each Cascade O-RU shall meet following condition:

$$Ta3_{max} + \sum_{i=n+1}^N T_{FH_i_max} + \sum_{i=n}^{N-1} T_{Comb} \text{ (for Cascade O-RU\#i)} \leq Ta3_prime_max \text{ (for Cascade O-RU\#n), and}$$

$$Ta3_prime_max \text{ (for Cascade O-RU\#n)} \leq Ta4_max - \sum_{i=1}^n T_FHi_max - \begin{cases} \sum_{i=1}^{n-1} T_Comb \text{ (for Cascade O-RU\#i)}, n \geq 2 \\ 0, n = 1 \end{cases}$$

Delay parameters for Shared cell for UL are summarized in Table 13.5.3-1.

Table 13.5.3-1: Delay parameters for Shared cell for UL

Parameter	Description	Additional Information
T_Comb	Corresponding to the maximum FHM or Cascade O-RU processing delay between receiving an IQ sample over the fronthaul interface from all O-RUs (FHM mode) or the south-node (Cascade mode), combining them and transmitting it over the fronthaul interface to O-DU (FHM mode) or the north-node (Cascade mode).	Capability parameter <i>t-combine</i> . FHM or Cascade O-RU reports the value by M-Plane.
Ta3_prime_max	The latest time that FHM or Cascade O-RU is allowed to send UL U-Plane message to O-DU (FHM mode) or north-node (Cascade mode) relative to reception timing at O-RU antenna.	Configuration parameter. Value configured to FHM or Cascade O-RU by M-Plane.
T_Combine_net	The maximum FHM or Cascade O-RU processing delay between starting time to combine IQ samples from the received message(s) and starting time to transmit them by a message over the fronthaul interface to O-DU (FHM mode) or the north-node (Cascade mode).	Capability parameter. FHM or Cascade O-RU reports the value by M-Plane.
Tx_Duration	The maximum transmission duration time for messages to be sent within Ta3_prime_max over the fronthaul interface to O-DU (FHM mode) or the north-node (Cascade mode).	Configured parameter or calculated parameter. When Tx_Duration isn't configured though enhanced-t-combine-enabled is configured as 'true', FHM or Cascade O-RU calculates from configured number-of-prb and compression-method per eAxC-id, and the calculated amount doesn't exceed symbol duration without cyclic prefix designated by Ta3_prime_max.
T_Waiting	Time when FHM or Cascade O-RU shall wait UL U-Plane message sent by O-RU (FHM mode) or south-node (Cascade mode).	Calculated parameter. FHM or Cascade O-RU calculates from configured Ta3_prime_max, its T_Comb, T_Combine_net, and Tx_Duration; i.e., this parameter is not defined in M-Plane.

Regarding O-RU delay profile reporting via M-Plane:

- For FHM, only T_comb and T_Combine_net (if FHM supports it) are applicable. Other parameters (T2a and Ta3 related parameters and Tcp_adv_dl) are not applicable for FHM since it does not behave as an O-RU
- For Cascade O-RU, T_comb and T_Combine_net (if Cascade O-RU supports it) are applicable. Other parameters (T2a and Ta3 related parameters and Tcp_adv_dl) are also applicable and the values are for the case it behaves as a normal O-RU

UL message combining in shared cell is strictly processed based on its reference timing of the message. If the reference timing cannot be identified with U-Plane message only, FHM shall refer to the associated C-Plane message.

13.5.4 Section Extension 10 support for Shared Cell

One C-Plane message with Section Extension 10 conveys information for multiple eAxC_IDs in an O-RU. O-DU may use Section Extension 10 in a shared-cell use case if all the impacted O-RUs and FHMs support Section Extension 10. This relation across multiple eAxC_IDs shall be configured in all O-RUs in the shared cell by the M-Plane to support the Section Extension 10. The eAxC_ID group information in UL shall be configured in a FHM to support UL combining with the Section Extension 10 use case. The **eaxc-id-group-capability** reported by the FHM indicates the

FHM can handle U-Plane messages commanded by the C-Plane with Section Extension 10, and the FHM also reports how many rx-exac-id-groups and rx-exac-ids-per-group can be configured.

13.6 S-plane for Shared cell

13.6.1 Introduction

The same synchronization framework and requirements specified in clause 11 also apply to the FHM, O-DU and O-RU for Shared cell, where the FHM (in the FHM mode), and the cascaded O-RU(s) (in the Cascade mode), are typically regarded as Ethernet switches on the synchronization chain, meeting the requirements specified for T-BC or T-TC Ethernet clocks in Clause 11 from an S-plane point of view. Whenever an O-RU is not on the synchronization path to other O-RUs, the O-RU need not support T-BC or T-TC functions. In any case the existing limits specified in clause 11 between the S-plane input port and RF port apply to all O-RUs.

13.6.2 Node behavior at O-RU in FREERUN state

13.6.2.1 O-DU

If synchronization states on all O-RUs in a Shared cell used for operating the same cell transit to the FREERUN or HOLDOVER state, the O-DU shall stop sending C-Plane and U-Plane related data to these O-RUs. In other words, the O-DU shall continue sending unless all O-RUs in the Shared cell used for operating the same cell transit to the FREERUN or HOLDOVER state.

13.6.2.2 O-RU

If synchronization state on a cascade O-RU transits to the FREERUN state, the cascade O-RU shall autonomously stop RF transmission, switch all carriers to INACTIVE state and send a notification to the O-DU about synchronization and carriers state change. The Cascade O-RU shall enable to continue the function for copy and forward to south-node when the O-RU receives the C-Plane and U-Plane related data from the north-node and shall be enabled to continue forwarding without any combining function when the O-RU receives the U-Plane related data from the south-node.

If synchronization state on an FHM transits to the FREERUN state, the FHM shall send a notification to the O-DU about the synchronization. The FHM shall enable to continue function for copy and forward to south-nodes when the O-RU receives the C-Plane and U-Plane related data from the north-node and shall be enabled to continue the function for combine and forward to the north-node when the FHM receives the U-Plane related data from the south-nodes.

The other O-RUs in a Shared cell used for operating the same cell as cascade O-RU or FHM might be still LOCKED synchronization state. The cascade O-RU or FHM shall be able to continue copy and combine functions for the other O-RUs in chain or star topology which provides valid U-Plane and C-Plane related data from/to the O-RUs. However, it has to be considered that if one of the O-RUs in the cascaded chain or FHM in FHM mode moves to FREERUN or HOLDOVER, the clockClass values advertised by this FHM/O-RU towards south-node(s) might impact the sync plane of all the other RUs which are listening on PRTC clockClass values and might trigger state change on all O-RUs which eventually might move to FREERUN or HOLDOVER based on the sync state of that particular O-RU.

13.7 Cascade-FHM mode

13.7.1 General

Following clauses define one hybrid mode Cascade-FHM. The mode is shown in Figure 13.7.1-1.

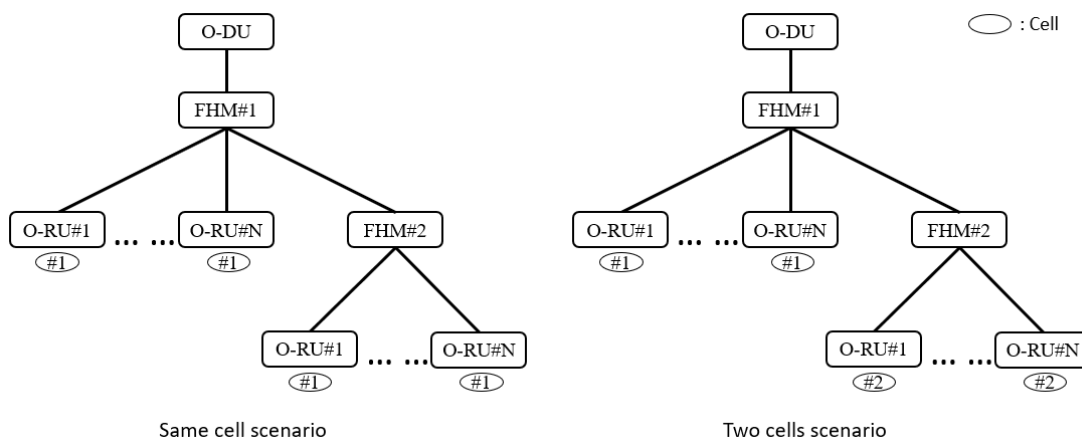


Figure 13.7.1-1: Cascade-FHM mode

Figure 13.7.1-1 shows two typical scenarios in Cascade-FHM mode: one is same cell scenario, that is O-RU(s) serving cascaded FHMs belong to same cell, and the other is Two cells scenario, that is O-RU(s) serving different FHM belong to different cell. Other different cell scenarios are also possible. The first cascade FHM nearest to O-DU is named FHM#1, the second cascade FHM is named FHM#2. For FHM#1, its north-node is O-DU and south-nodes are O-RUs and FHM#2. For FHM#2, its north-node is FHM#1 and south-nodes are O-RUs.

Cascade-FHM mode is realized by at least two chained FHMs and O-RUs connected to them. Star-like topology where many FHMs would be connected to one FHM is not in scope of the present document. In the present document maximum level of cascaded FHMs is limited to 2.

NOTE: In the present document, the reasons for limited depth of FHM chain are as follows: 1) Noise floor level may have negative impact when deeper chain is constructed. Such impact is not yet analyzed; 2) Delay management may be affected by deeper chain and also requires further analysis. 3) Cascade-FHM mode is suitable to be deployed in Indoor Distribution System, and generally the usage with two cascaded FHMs is typical and sufficient.

13.7.2 Copy and Combine function

DL Copy function and UL Combine function for Cascade-FHM mode are actually same with FHM mode with a little attention that FHM#1 regards FHM#2 as its one O-RU from aspect of Copy and Combine. Following figures show downlink and uplink flow for same cell scenario (Figure 13.7.2-1) and two cells scenario (Figure 13.7.2-2).

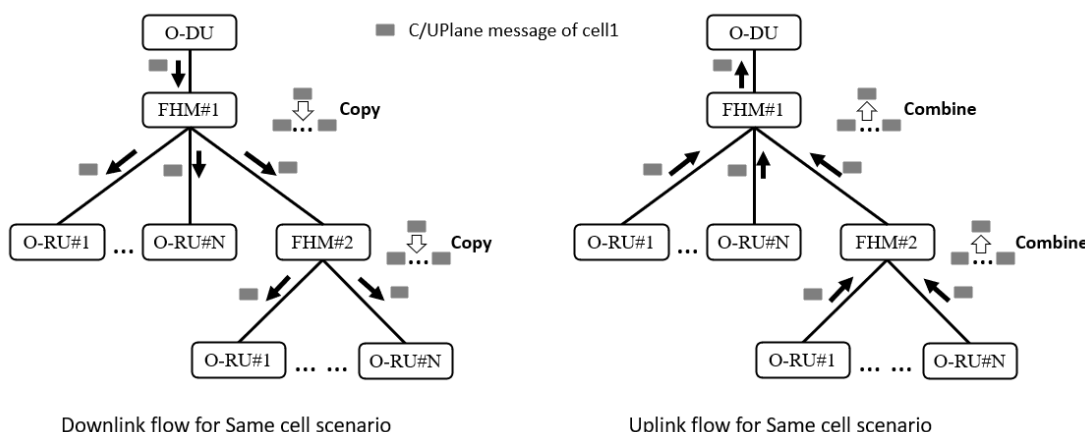


Figure 13.7.2-1: Cascade-FHM flow for same cell

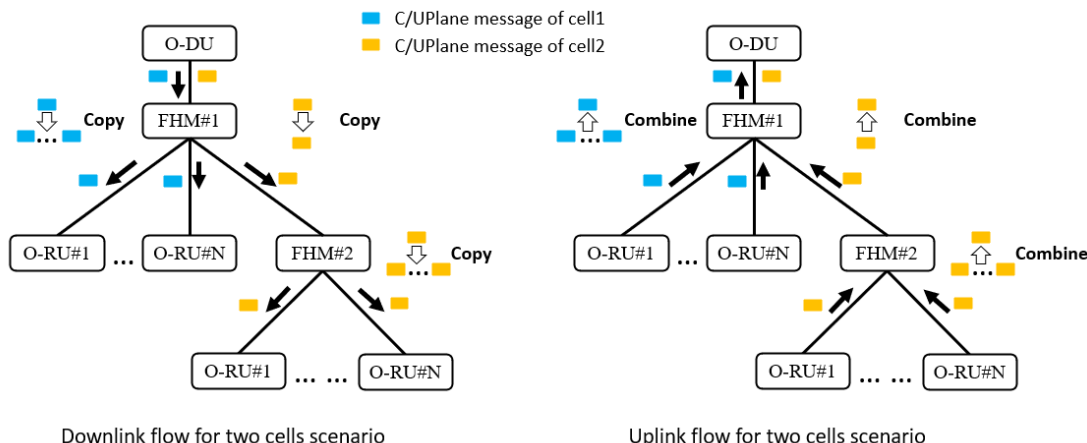


Figure 13.7.2-2: Cascade-FHM flow for two cells

13.7.3 Delay management

Delay management for Cascade-FHM mode follows the combination of FHM mode and Cascade mode while FHM#2 cannot be treated as one normal O-RU from aspect of delay management since it has own processing delay and transport delay towards its O-RUs.

Figure 13.7.3-1 shows a DL delay model for two Cascaded FHMs, and the case where FHM#1 has one O-RU and FHM#2 has two O-RUs in same cell scenario.

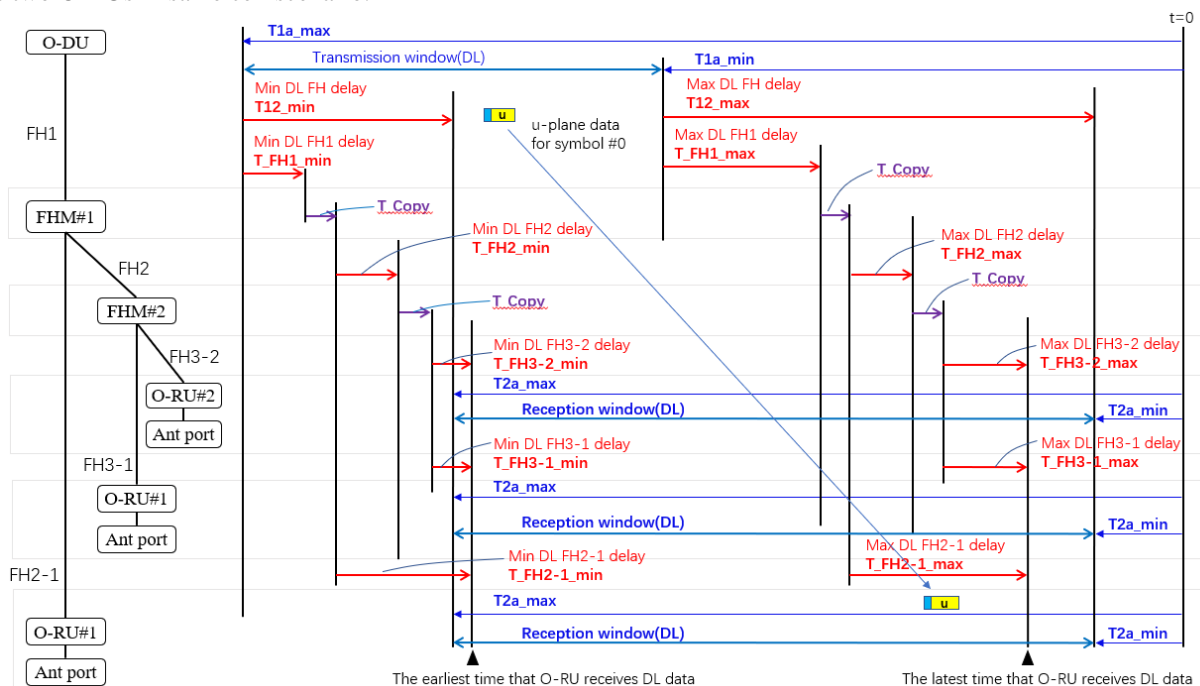


Figure 13.7.3-1: DL delay model parameters for Cascade-FHM mode

To ensure DL U-Plane messages arrive at O-RU within O-RU reception window the configuration of fronthaul including two cascaded FHMs shall satisfy the following relationships:

$$T_FH1_min + \min(\min(T_FH2_1_min, \dots, T_FH2_N_min), T_FH2_min + \min(T_FH3_1_min, \dots, T_FH3_M_min)) + T_Copy) \geq (T12_min - T_Copy)$$

where:

- T_FH1_min is the minimum transport delay between FHM#1 and O-DU,
- T_FH2_min is the minimum transport delay between FHM#2 and FHM#1,
- T_FH2-n_min is the minimum transport delay between O-RU#n and FHM#1 (n = 1, ..., N),

N is the total number of O-RUs realizing a particular cell m under FHM#1, and in this case $N=1$,
 T_{FH3-m_min} is the minimum transport delay between O-RU# m and FHM#2 ($m = 1, \dots, M$),
 M is the total number of O-RUs realizing a particular cell m under FHM#2, and in this case $M=2$,
 T_Copy is FHM processing delay for copying which depends on FHM capability and is reported via M-Plane.

$$T_FH1_max + \max(\max(T_FH2_1_max, \dots, T_FH2-N_max), T_FH2_max + \max(T_FH3_1_max, \dots, T_FH3-M_max) + T_Copy) \leq (T12_max - T_Copy)$$

where:

T_FH1_max is the maximum transport delay between FHM#1 and O-DU,
 T_FH2_max is the maximum transport delay between FHM#2 and FHM#1,
 T_FH2-n_max is the maximum transport delay between O-RU# n and FHM#1 ($n = 1, \dots, N$),
 N is the total number of O-RUs realizing a particular cell m under FHM#1, and in this case $N=1$,
 T_FH3-m_max is the maximum transport delay between O-RU# m and FHM#2 ($m = 1, \dots, M$),
 M is the total number of O-RUs realizing a particular cell m under FHM#2, and in this case $M=2$,
 T_Copy is FHM processing delay for copying which depends on FHM capability and is reported via M-Plane.

Figure 13.7.3-2 shows an UL delay model for two Cascaded FHMs, and the case where FHM#1 has one O-RU and FHM#2 has two O-RUs in same cell scenario.

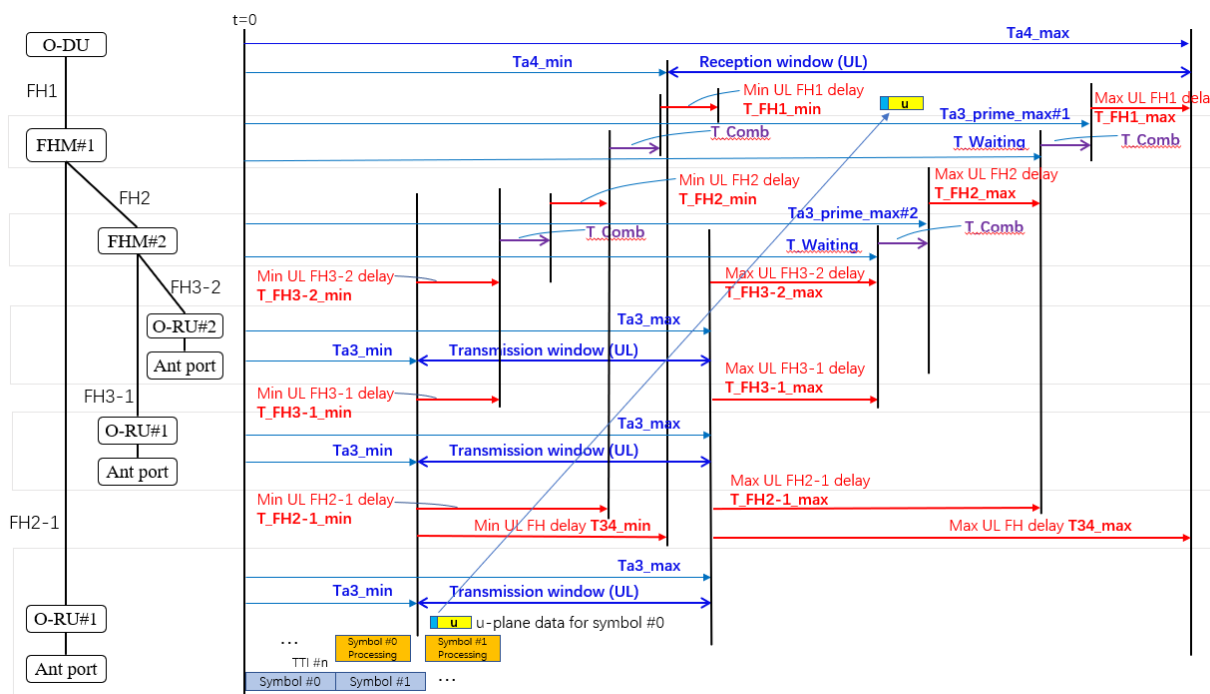


Figure 13.7.3-2: UL delay model parameters for Cascade-FHM mode

To ensure UL U-Plane messages arrive at O-DU within O-DU reception window the configuration of fronthaul including two cascaded FHMs shall satisfy the following relationships:

$$T_FH1_min + \min(\min(T_FH2_1_min, \dots, T_FH2-N_min), T_FH2_min + \min(T_FH3_1_min, \dots, T_FH3-M_min) + T_Comb) \geq (T34_min - T_Comb)$$

where:

T_FH1_min is the minimum transport delay between FHM#1 and O-DU,
 T_FH2_min is the minimum transport delay between FHM#2 and FHM#1,
 T_FH2-n_min is the minimum transport delay between O-RU# n and FHM#1 ($n = 1, \dots, N$),
 N is the total number of O-RUs realizing a particular cell m under FHM#1, and in this case $N=1$,
 T_FH3-m_min is the minimum transport delay between O-RU# m and FHM#2 ($m = 1, \dots, M$),
 M is the total number of O-RUs realizing a particular cell m under FHM#2, and in this case $M=2$,
 T_Comb is FHM processing delay for combining UL U-Plane messages which depends on FHM capability and is reported via M-Plane.

$$T_FH1_max + \max(\max(T_FH2-1_max, \dots, T_FH2-N_max), T_FH2_max + \max(T_FH3-1_max, \dots, T_FH3-M_max) + T_Comb) \leq (T34_max - T_Comb)$$

where:

- T_FH1_max is the maximum transport delay between FHM#1 and O-DU,
- T_FH2_max is the maximum transport delay between FHM#2 and FHM#1,
- T_FH2-*n*_max is the maximum transport delay between O-RU#*n* and FHM#1 (*n* = 1, ... , *N*),
- N* is the total number of O-RUs realizing a particular cell *m* under FHM#1, and in this case *N*=1,
- T_FH3-*m*_max is the maximum transport delay between O-RU#*m* and FHM#2 (*m* = 1, ... , *M*),
- M* is the total number of O-RUs realizing a particular cell *m* under FHM#2, and in this case *M*=2,
- T_Comb is FHM processing delay for combining UL U-Plane messages which depends on FHM capability and is reported via M-Plane.

In addition, in order to combine all UL U-Plane messages from O-RUs, O-DU shall configure Ta3_prime_max for two cascaded FHMs which meets following condition:

$$(Ta3_max + \max(T_FH3-1_max, \dots, T_FH3-M_max) + T_Comb) \leq Ta3_prime_max_FHM\#2,$$

$$(Ta3_max + \max(\max(T_FH2-1_max, \dots, T_FH2-N_max), T_FH2_max + \max(T_FH3-1_max, \dots, T_FH3-M_max) + T_Comb) + T_Comb) \leq Ta3_prime_max\#1 \leq (Ta4_max - T_FH1_max),$$

where:

- Ta3_prime_max_FHM#1 is Ta3_prime_max for FHM#1,
- Ta3_prime_max_FHM#2 is Ta3_prime_max for FHM#2,
- T_FH2-*n*_max is the maximum transport delay between O-RU#*n* and FHM#1 (*n* = 1, ... , *N*),
- N* is the total number of O-RUs realizing a particular cell *m* under FHM#1, and in this case *N*=1,
- T_FH3-*m*_max is the maximum transport delay between O-RU#*m* and FHM#2 (*m* = 1, ... , *M*),
- M* is the total number of O-RUs realizing a particular cell *m* under FHM#2, and in this case *M*=2.

14 Support of Shared O-RU

14.1 General

"Shared O-RU" is defined as an O-RU that supports the SHARED-ORU-MULTI-ODU and/or SHARED-ORU-MULTI-OPERATOR. For SHARED-ORU-MULTI-ODU feature, the O-RU shall be configured to operate with multiple O-DUs of same operator. For SHARED-ORU-MULTI-OPERATOR feature, the O-RU may be configured to operate with one or more O-DUs from one or more Shared Resource Operators (SRO). O-DUs of same or different operators shall connect to the Shared O-RU using existing CUS-Plane interface definitions and procedures. C-Plane and U-Plane protocol between multiple O-DUs connected to Shared O-RU is expected to be identical, however transport level parameters which may include Ethernet header parameters, IP/UDP headers shall be configured uniquely for each O-DU and Shared O-RU pair. Value of eAxC_IDs used per O-DU and Shared O-RU interfaces shall be configured by applicable M-Plane yang models.

Shared O-RU realization shall support single and/or multiple operators use cases where multiple O-DUs connected to the Shared O-RU can belong to a single operator and/or multiple operators. Shared O-RU shall support SHARED-ORU-MULTI-ODU and/or SHARED-ORU-MULTI-OPERATOR feature and based on which of these feature O-RU supports, respective restrictions for resource partitioning of the O-RU shall apply. In Figure 14.1-1, O-DU#1 and O-DU#2 may belong to a single operator or multiple operators. Policies governing partitioning of Shared O-RU resources/capabilities are out of scope of CUS and M-Plane specification.

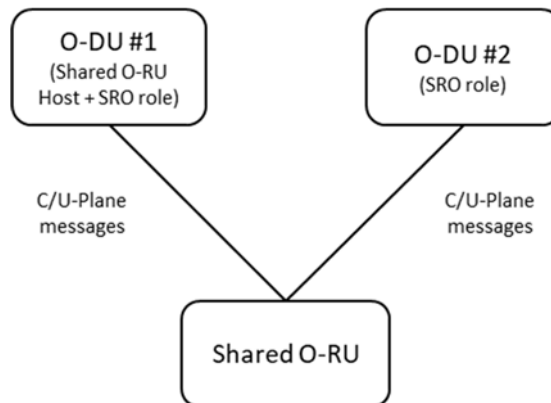


Figure 14.1-1 : Shared O-RU Concept

14.2 Delay management for Shared O-RU

14.2.1 General

The Shared O-RU Host is responsible for configuring shared O-RU aspects related to delay management. How the Shared O-RU Host determines the parameters to use in the delay management configuration of the shared O-RU is out of scope of the present document.

14.3 S-Plane for Shared O-RU

14.3.1 General

The requirements for S-Plane in the Shared O-RU application are unchanged and requirements in clause 11 shall be met. Shared O-RU Host shall ensure that the Shared O-RU has the appropriate sync source to operate within specifications. Shared O-RU Host may determine the sync topology based on the deployment scenario for Shared O-RU, e.g., for LLS-C1 and LLS-C2 topologies, Shared O-RU Host O-DU may be used as the sync source, alternatively, Shared O-RU Host can also choose multiple sync sources one from Shared O-RU Host O-DU and another sync source from a SROs O-DU (Prioritization of O-DUs as time sources can be done using BMCA in the used ITU-T profile; see clause 11.2.4.2 for applicable PTP profiles)

15 Support of Narrowband-IoT

15.1 Overview (informative)

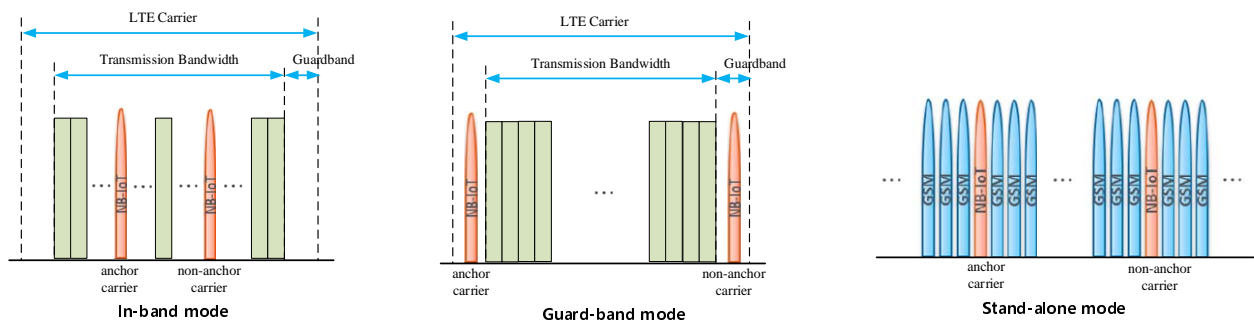
The term "NB-IoT" stands for Narrowband Internet of Things, which is a 3GPP defined low power wide area network (LPWAN) technology developed to enable extended coverage in e.g., indoor and underground scenarios and to enable support for massive number of IoT devices and services. NB-IoT is also based on Orthogonal Frequency Division Multiplexing (OFDM) like LTE and NR with similar resource grid.

An NB-IoT carrier occupies a bandwidth of 180 kHz and can be deployed as "in-band mode", "guard-band mode" and "stand-alone mode". These NB-IoT modes indicate to NB-IoT devices the channel raster location and resource element allocation in an NB-IoT subframe. For example, an NB-IoT carrier located inside an NR carrier can be operated with "guard-band mode". However, the present document uses NB-IoT carrier deployment mode as the relationship between the NB-IoT carrier and the LTE carrier or the NR carrier.

- For in-band mode, an NB-IoT carrier is located within the transmission bandwidth of an LTE/NR host carrier and utilizes one or more resource blocks of the LTE/NR host carrier.
- For guard-band mode, an NB-IoT carrier is located at the guard-band of an LTE/NR host carrier, which is inside the channel bandwidth but outside the transmission bandwidth of the host carrier.

For stand-alone mode, an NB-IoT carrier is located on channel raster (100 kHz). Support for stand-alone mode is out of scope of the present document. To support a massive number of devices in a cell and to achieve spectrum efficiency, multi-carrier operation is also supported. Two types of carriers are defined in NB-IoT, anchor and non-anchor carrier, and they are not differentiated in the fronthaul interface.

Example of the NB-IoT carrier deployment for the three modes are shown in Figure 15.1-1.



NOTE: Support for stand-alone mode is out of scope of the present document.

Figure 15.1-1: Example NB-IoT carrier deployment

Support for host carrier with non-15 kHz SCS is out of scope of the present document. Scope of the present document is limited to a common O-DU for host carriers and NB-IoT carriers, and does not cover host carriers and NB-IoT carriers on different O-DUs

15.2 NB-IoT Downlink enabling methods

15.2.1 Extended PRB Grid

The frequency-time resource for NB-IoT in-band mode is treated as one resource block of the LTE/NR host carrier, but for the guard-band mode, the NB-IoT carriers are located outside the PRBs of the host carrier. In order to have unified processing for in-band and guard-band in DL, the PRB grid of the host carrier is extended to accommodate the NB-IoT carrier in the guard-band, so that the guard-band NB-IoT carrier can also be treated as one resource block of the LTE/NR host carrier. The O-DU shall determine if the O-RU's static-low-level-tx-endpoints support the Extended PRB Grid method by checking if **max-prb-range** exceeds the number of PRBs of the host carrier. Such an endpoint may also report **supported-filter-pass-bandwidth** to inform the O-DU of the size of its transmission passband filters, and if it is reported, the O-DU shall check the NB-IoT carriers can fit within the reported **supported-filter-pass-bandwidth** and shall provide

the O-RU with occupied-bandwidth (upper and lower bound) to inform O-RU the part of the frequency spectrum that may be requested to be transmitted.

To extend the PRB grid of the host carrier, the PRB #0 starts from the lower part of the guard-band, so the PRB index of the host carrier is extended by the number of PRBs of the lower part of the guard-band, and the PRB index in the upper part of the guard-band follows the highest host carrier PRB index after extension. The PRB boundaries shall be kept same with and without PRB extension, i.e. in figure 15.2.1-1, PRB#0 before extension and PRB#2 after extension have same boundaries. The O-RU shall construct Extended PRB Grid of the host carrier by M-plane parameters **number-of-prb**, **offset-to-absolute-center** as well as **center-of-channel-bandwidth**, type in tx-array-carrier. An O-RU supporting extended PRB grid method shall support the symmetrically extended number of PRBs altered from the 3GPP definition for the LTE/NR host carrier within the boundaries of the O-RU capability reported in **filter-pass-bandwidth**. Transmission filter bandwidth also shall be widened to satisfy 3GPP TS 36.141 [40] requirement, which shall be reported by **filter-pass-bandwidth** in endpoint capability. There can be multiple NB-IoT carriers in the guard-band of the LTE/NR host carrier, and they shall be located within filter-pass-bandwidth.

An example of the extended PRB grid for a 15 MHz LTE carrier is shown in Figure 15.2.1-1.

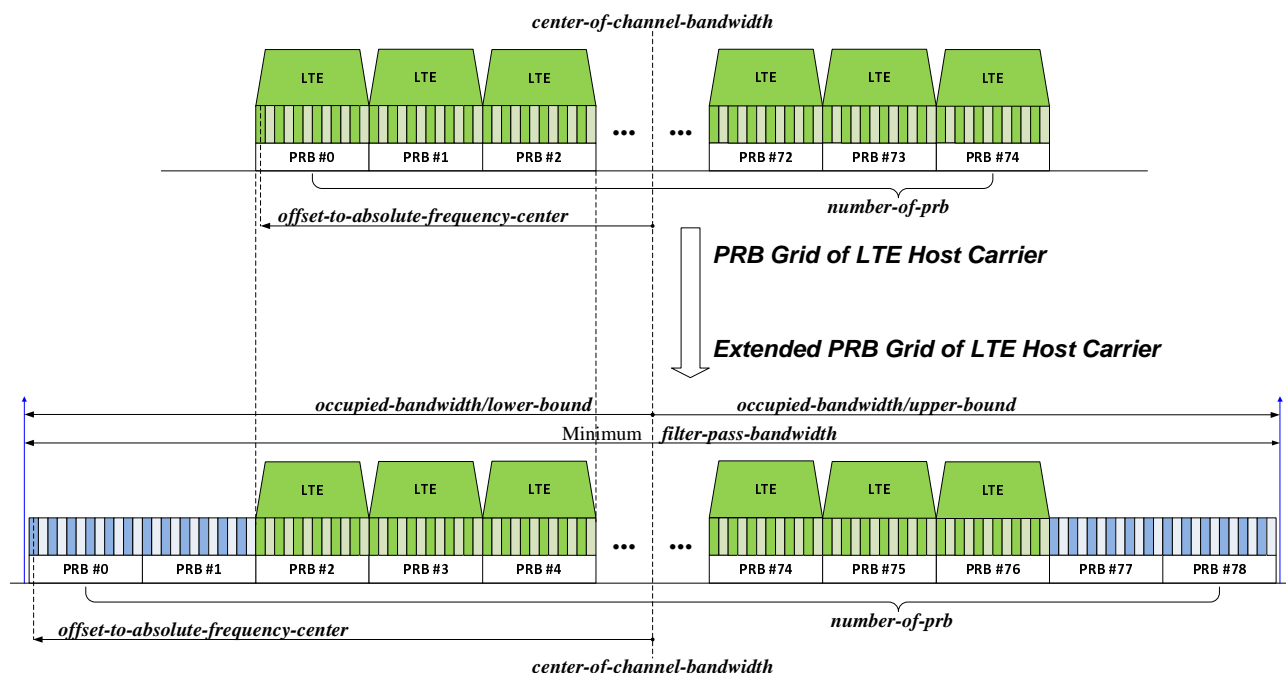


Figure 15.2.1-1: Example of extended PRB grid (DL)

The PRB grid for guard-band NB-IoT carrier need not be aligned with the extended PRB grid of the LTE/NR host carrier. Figure 15.2.1-2 shows an anchor carrier deployed in guard-band of 15MHz LTE carrier, PRB grid misalignment happens. Figure 15.2.1-3 shows an anchor carrier deployed in-band of 15MHz LTE carrier, PRB grid misalignment for non-anchor carrier does not happen in this example.




Figure 15.2.1-2: Example of PRB grid misalignment (DL)




Figure 15.2.1-3: Example of PRB grid without misalignment (DL)

15.3 NB-IoT Downlink use cases

15.3.1 Overview

Within this document, use cases of NB-IoT in-band and guard-band mode deployed with a SCS 15 kHz LTE/NR host carrier is supported. In case of in-band deployment, the DL enabling method as described in either clause 15.3.2 or clause 15.3.3 shall be applied, and in case of guard-band deployment, the DL enabling method as described in clause 15.3.3 shall be applied.

DC subcarrier of iFFT (k_0) is important for the proper signal processing between O-DU and O-RU. When either methods defined in 15.3.2 or 15.3.3 is applied, the O-RU shall follow signal processing of host carrier and O-DU shall perform phase compensation for frequency offset.

15.3.2 In-band mode with 15 kHz subcarrier spacing of host carrier

Transparent operation using Section Type 1 in a host carrier is selected as a baseline method for DL in-band mode. With this method, the NB-IoT carrier shares the same **tx-array-carrier** and endpoint as the LTE/NR host carrier, and the C-Plane and U-Plane messages for the NB-IoT carrier are treated same as the C-Plane and U-Plane messages for the LTE/NR host carrier. Phase compensation for the NB-IoT carrier shall be done in the O-DU to compensate for the phase offset caused by the frequency offset of iFFT between the NB-IoT carrier and the LTE/NR host carrier. There may be multiple NB-IoT carriers within one LTE/NR host carrier, and the maximum number of carriers is not limited by the present document.

NOTE: This approach is simple for O-RU implementation, because the existence of the NB-IoT carrier is transparent to the O-RU, and no capability report nor configuration by M-Plane is required. Also, this approach does not need additional dedicated resource for the NB-IoT carrier for signal processing such as iFFT in the O-RU.

An example for the in-band mode operation is shown in Figure 15.3.2-1.

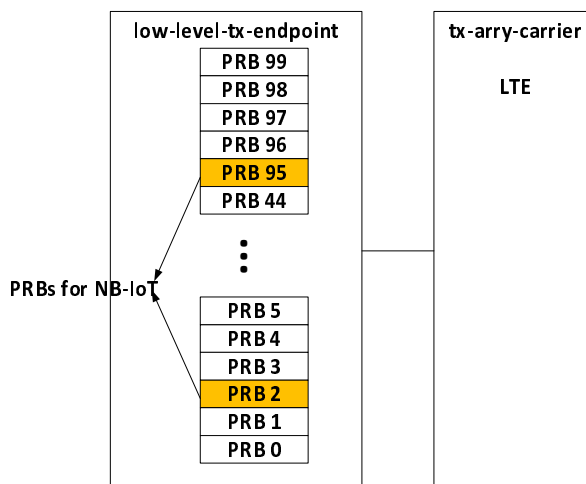


Figure 15.3.2-1: Example of in-band mode operation

15.3.3 Guard-band mode with 15 kHz subcarrier spacing of host carrier

As described in clause 15.2.1, the PRB grid for the host carrier is extended to accommodate the NB-IoT carrier in the guard-band, as a result, like in-band mode, the frequency-time resource in NB-IoT guard-band mode also can be treated as one or more resource blocks of the LTE/NR host carrier. In this use case, Section Type 1 C-Plane and the extended PRB as defined in clause 15.2.1 shall be used. For the same reason as in-band mode, phase compensation for the guard-band mode NB-IoT carrier shall be performed by the O-DU.

An O-RU supporting NB-IoT carriers in guard-band mode shall support the extended PRB grid method as specified in clause 15.2.1.

An example for the guard-band mode operation is shown in Figure 15.3.3-1.

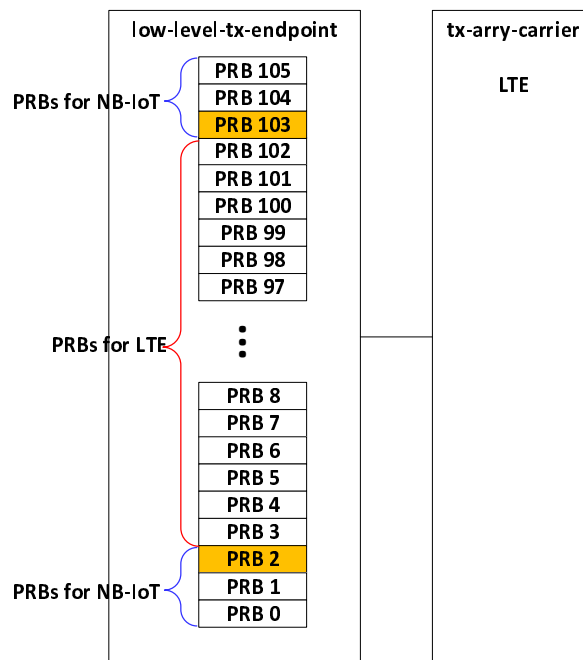


Figure 15.3.3-1: Example of configuration for guard-band mode operation

As described in clause 15.2.1, the PRB grid for guard-band NB-IoT carrier need not be aligned with the extended PRB grid of the LTE/NR host carrier. In case of PRB grid misalignment, the frequency-time resource in NB-IoT guard-band mode is treated as two PRBs instead of a single PRB. The O-DU shall populate those non-occupied REs to zero in the U-Plane message, and this operation is transparent to the O-RU. An example for this operation is shown in Figure 15.3.3-2.

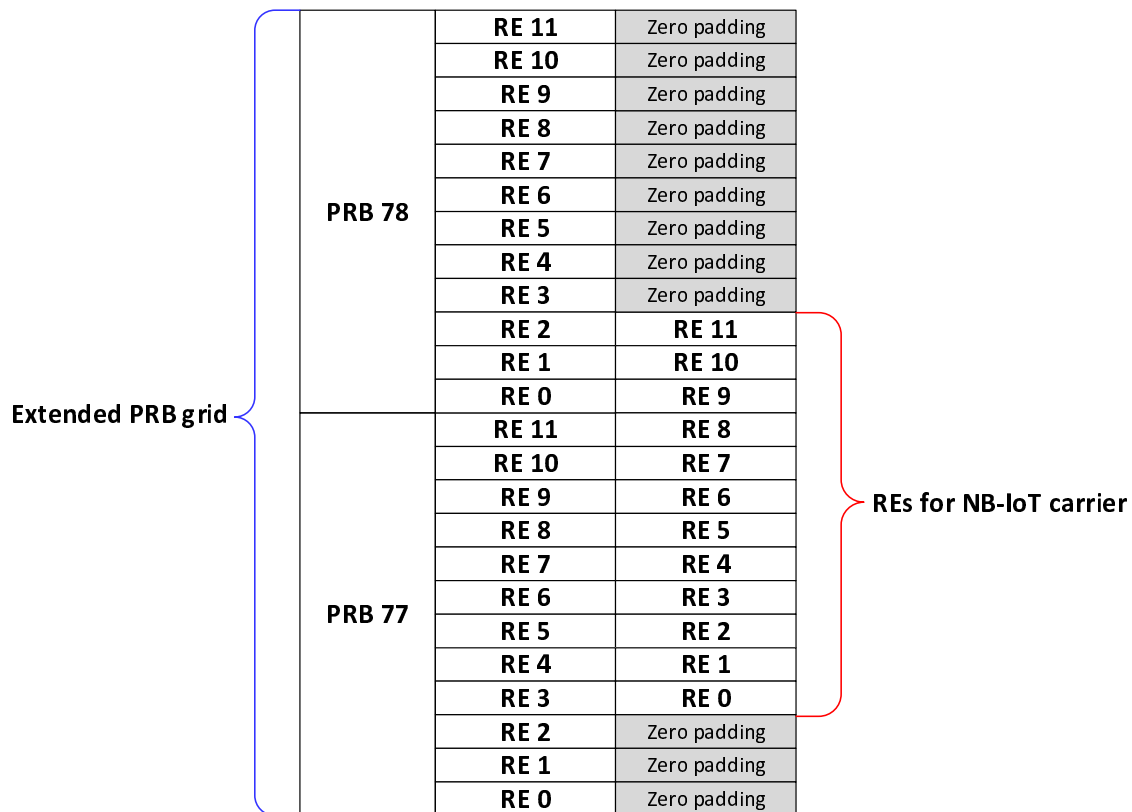


Figure 15.3.3-2: Example of 2 PRB utilization for one NB-IoT carrier in case of PRB grid misalignment

15.4 NB-IoT Uplink enabling methods

15.4.1 Single Section Single Carrier (SSSC)

Section Type 3 shall be used for the SSSC method to convey the allocation to the O-RU when the NB-IoT carriers are deployed in guard-band or in-band mode.

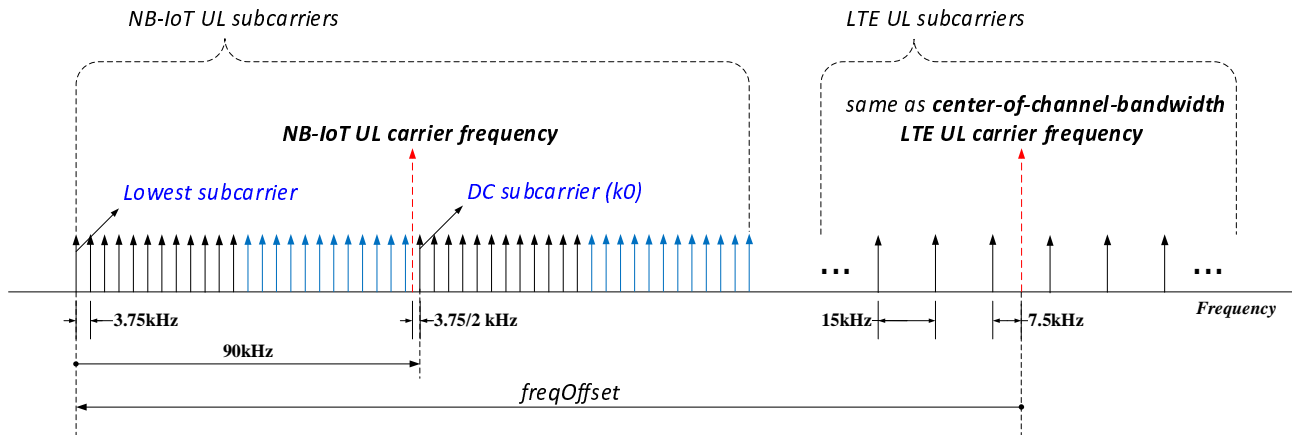
The O-DU shall determine if the O-RU's static-low-level-rx-endpoint supports SSSC method by checking if the endpoint has **center-from-freqoffset-supported** = TRUE and a config within a set has **center-from-freqoffset** = TRUE. The O-DU shall determine the max number of supported time-overlapped sections of a SSSC capable endpoint by checking **max-overlapping-instances**. To activate SSSC, the O-DU shall configure the low-level-rx-endpoints with **center-from-freqoffset-enabled** = TRUE. Such an endpoint may also report **supported-filter-pass-bandwidth** to inform O-DU the size of its reception passband filters, and if it is reported, the O-DU shall configure **occupied-bandwidth** (upper and lower bound) to inform O-RU the part of the frequency spectrum that may be requested to be received.

When using the SSSC method, the O-DU shall use a single data section for an NB IoT carrier with startPrbc limited to value 0 and numPrbc limited to values of 1, 4 or 12 for 15 kHz, 3.75 kHz and 1.25 kHz SCS, respectively. The O-DU shall indicate the DC subcarrier location of FFT (k0) of the requested NB-IoT carrier with a single section whose freqOffset field is set to the center of the lowest subcarrier of the NB-IoT carrier. The O-RU shall infer the k0 location from the formula in clause 15.4.2 with startPrbc limited to value 0 and numPrbc limited to values of 1, 4 or 12 for 15 kHz, 3.75 kHz and 1.25 kHz SCS, respectively. This will result in the term 'startPrbc × 12 × SCS + numPrbc × 6 × SCS' being equal to 90 kHz. This mode of NB-IoT section description is referred to as single-section-single-carrier (SSSC) operation. When an endpoint has **center-from-freqoffset-enabled** configured as TRUE, the same RE or PRB can be addressed multiple times in C-plane messages or in U-plane messages as an exception. When using the SSSC method, the O-DU shall use a single data section to describe all the populated PRBs in a NB-IoT carrier for any unique timeOffset and SCS applicable to the NB-IoT carrier, and freqOffset shall be set to the lowest subcarrier in the NB-IoT carrier.

EXAMPLE: While using SCS 3.75 kHz, if the O-DU require NPUSCH PRBs 1, 2, and 3 for one timeOffset value then, freqOffset would be set to RE#0 of PRB#0, startPrbc = 0, numPrbc = 4, and the all PRBs would be sent by the O-RU and the O-DU would discard the PRB#0. Another data section might require only PRB#0 at a different timeOffset (so in a

different C-Plane message) e.g. NPRACH, and would include $\text{startPrbc} = 0$, $\text{numPrbc} = 4$, and all four PRBs would be sent by the O-RU, and the O-DU would keep only PRB#0.

NOTE: The dedicated resource utilization using Section Type 3 C-Plane message is intended to align DC subcarrier (k_0) of FFT with the NB-IoT carrier and NB-IoT Carrier frequency for down-conversion, so that the O-DU is not required to perform phase compensation as is required when using the receive method in clause 15.4.3. An example of the DC frequency indication of NB-IoT UL from LTE center frequency is illustrated in Figure 15.4.1-1.



NOTE: The subcarrier arrows represent center of subcarrier frequency.

Figure 15.4.1-1: Example of DC subcarrier frequency indication with NB-IoT UL and LTE

15.4.2 Single Section Multiple Carriers (SSMC)

An O-RU supporting SSMC operation may handle the signal processing for multiple contiguous UL carriers having same SCS simultaneously when a single section in Section Type 3 that covers those carriers is received from the O-DU. This mode of NB-IoT section description is referred to as single-section-multiple-carrier (SSMC) operation. In this method, there may be a phase offset in the IQ data due to the mismatch of DC subcarrier in FFT and carrier frequency conversion, so the O-DU shall compensate the phase offset.

The O-DU shall determine if O-RU's static-low-level-rx-endpoints supports SSMC method by checking if the endpoint has **center-from-freqoffset-supported** = TRUE, and a config within a set has **center-from-freqoffset** = TRUE if the parameter is present and **max-prb-range** is larger than the number of PRBs in one NB-IoT carrier with the corresponding SCS. The O-DU shall determine the max number of supported time-overlapped sections of a SSMC capable endpoint by checking **max-overlapping-instances**.

To activate SSMC, the O-DU shall configure the low-level-rx-endpoints with **center-from-freqoffset-enabled** = TRUE. Such an endpoint may also report **supported-filter-pass-bandwidth** to inform O-DU the size of its reception passband filters, and if it is reported, the O-DU shall provide the O-RU with **occupied-bandwidth** (upper and lower bound) to inform O-RU the part of the frequency spectrum that may be requested to be received. The O-RU shall infer the DC subcarrier location of FFT from the below formula from the SSMC C-Plane.

$$\text{DC subcarrier } (k_0) \text{ frequency} = \text{center-of-channel-bandwidth} + \text{freqOffset} \times 0.5 \times \text{SCS} + \text{startPrbc} \times 12 \times \text{SCS} + \text{numPrbc} \times 6 \times \text{SCS}$$

where **center-of-channel-bandwidth** is the frequency configured via M-plane in the corresponding **rx-array-carrier**, freqOffset , startPrbc and numPrbc are provided with the Section Type 3 C-Plane section and SCS is the subcarrier spacing provided in frameStructure (see clause 7.5.2.13). The sum of startPrbc and numPrbc shall not exceed **max-prb-range** in the intended config. The O-DU shall use numPrbc values greater than 1, 4 or 12 for 15 kHz, 3.75 kHz and 1.25 kHz SCS respectively in SSMC method.

When SSMC method is used, every section description shall be counted as SSMC instance which is limited by **max-overlapping-instances**. It is recommended to combine as many NB-IoT carriers as permitted by **max-prb-range** within one section description.

15.4.3 Extended PRB Grid

With this method a PRB grid of a host carrier is extended to accommodate guard-band NB-IoT carriers.

The O-DU shall determine if O-RU's static-low-level-rx-endpoints supports the Extended PRB Grid method by checking if **max-prb-range** exceeds the number of PRBs of the host carrier. Such an endpoint may also report **supported-filter-pass-bandwidth** to inform O-DU of the size of its reception passband filters. If O-RU reports **supported-filter-pass-bandwidth**, then the O-DU shall provide O-RU with **occupied-bandwidth** (upper and lower bound) to inform O-RU the part of the frequency spectrum that may be requested to be received.

An example of the extended PRB grid for a 15MHz LTE carrier is shown in Figure 15.4.3-1.

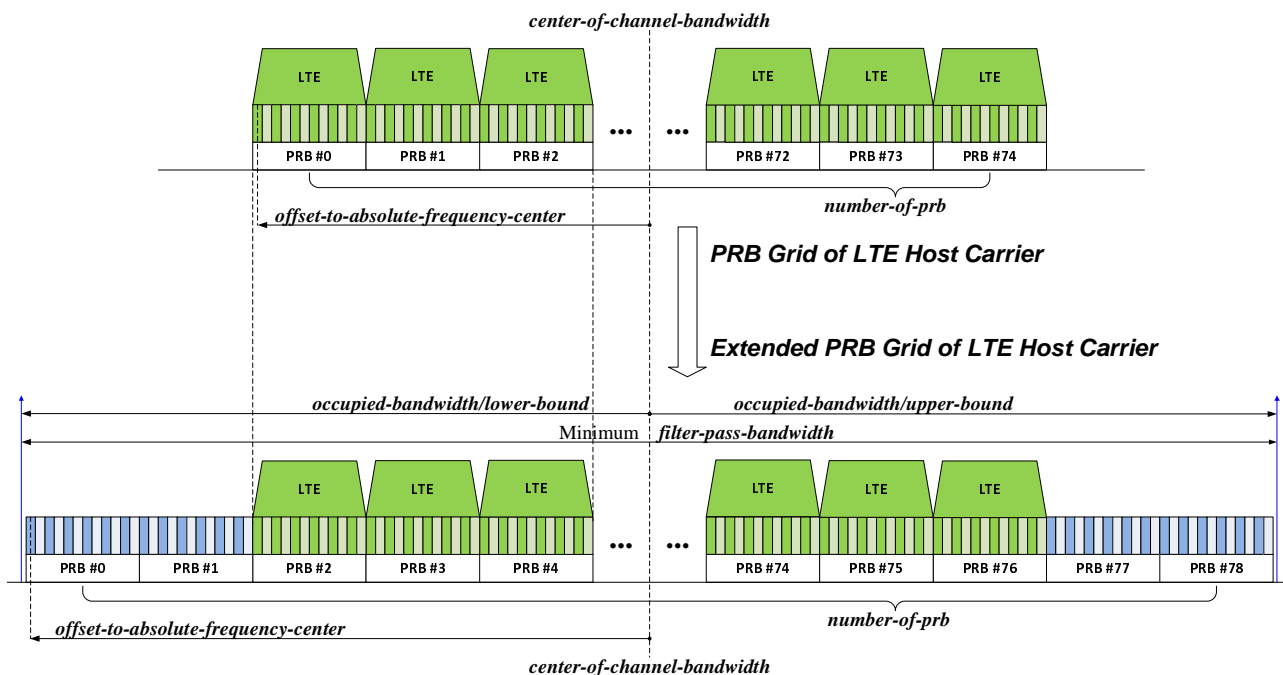


Figure 15.4.3-1: Example of extended PRB grid (UL)

The additional PRBs may be used to receive NB-IoT guard-band carriers. In this case the O-RU shall operate as if the PRB grid of the host carrier was larger; the O-DU may request reception of additional PRBs (as permitted by the extended grid). In addition, the O-DU shall handle any phase shift compensation related to misalignment of NB-IoT carrier transmitter and receiver frequencies. The PRB grid for guard-band NB-IoT carrier may be aligned with the PRB grid of the host carrier. In this case the O-DU extracts NB-IoT PRBs directly from U-plane messages. Figure 15.4.3-2 shows an example with NB-IoT PRB grid aligned to the PRB grid of the LTE host carrier.




Figure 15.4.3-2: Example of aligned PRB grid (UL)

The PRB grid for guard-band NB-IoT carrier may not be aligned with the extended PRB grid of the host carrier. In this case the O-RU operation is not changed, but the O-DU shall extract REs from adjacent PRBs in U-plane messages sent by the O-RU to collect NB-IoT carrier's PRB. An example of PRB grid misalignment is shown in Figure 15.4.3-3.




Figure 15.4.3-3: Example of not aligned PRB grid (UL)

15.5 NB-IoT Uplink use cases

15.5.1 Overview

An NB-IoT capable O-RU shall support UL NB-IoT carriers with the SCS 3.75 kHz and SCS 15 kHz for NPUSCH and the SCS 3.75 kHz or SCS 1.25 kHz for NPRACH. The following clauses describe use cases of NB-IoT carrier reception. The clause 15.5.2 describes a general solution for in-band and guard-band mode NB-IoT carrier reception for NPUSCH with SCS 15 kHz and 3.75 kHz and NPRACH with SCS 3.75 kHz and 1.25 kHz. Clause 15.5.3 describes special case of in-band and guard-band mode NB-IoT carrier reception for NPUSCH with SCS 15 kHz intended to reduce O-RU resource usage. Clause 15.5.4 describes specialised solution for NPRACH reception without C-Plane.

15.5.2 General case for in-band and guard-band mode

The reception of NB-IoT carriers in in-band and guard-band mode with the SCS 3.75 kHz and/or 15 kHz for NPUSCH and the SCS 3.75 and/or 1.25 kHz for NPRACH using the Single Section Single Carrier method (SSSC) is described in clause 15.4.1. The reception of multiple NB-IoT carriers with SCS 3.75 kHz as well as 15 kHz can be handled using individual freqOffset in the Section Type 3 as an example in Figure 15.5.2-1.

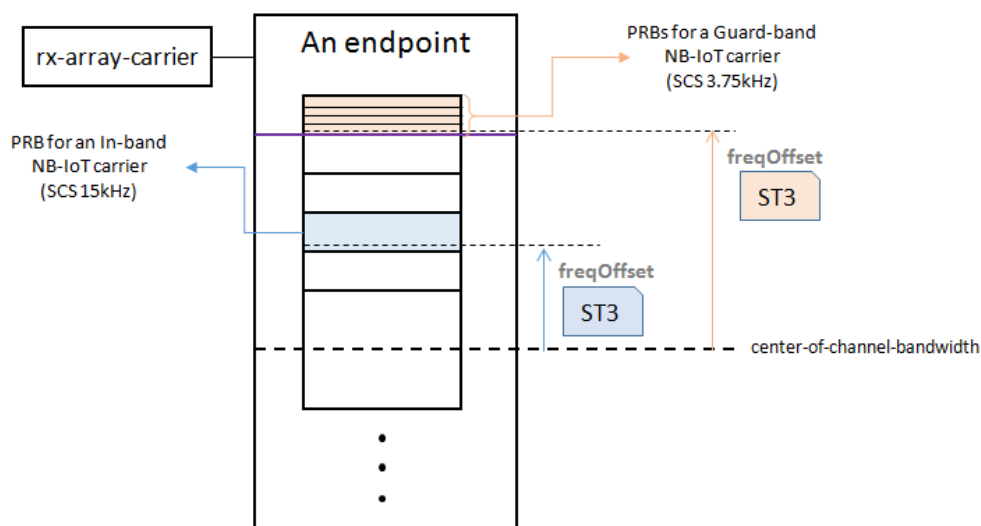


Fig. 15.5.2-1. Example for SSSC method to indicate SCSs 3.75 kHz and 15 kHz of NPUSCH using ST 3

The reception of the above-mentioned NB-IoT carriers may also be handled by the Single Section Multiple Carrier (SSMC) method described in clause 15.4.2, if both the O-DU and the O-RU support the SSMC method. The reception of multiple contiguous NB-IoT carriers with equivalent SCS may be handled with a single freqOffset in a Section Type 3 message as illustrated in example Fig. 15.5.2-2.

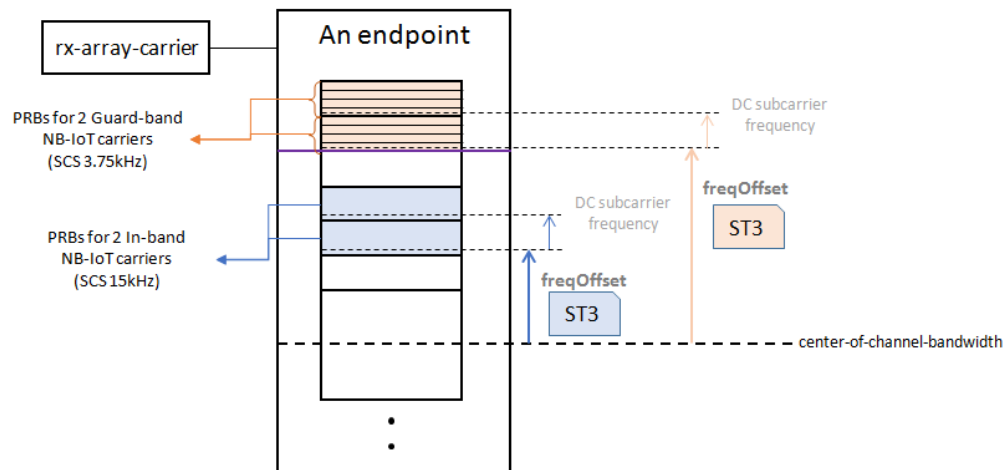


Fig. 15.5.2-2. Example for SSMC method to indicate SCSs 3.75kHz and 15kHz of NPUSCH using ST 3

Both SSSC and SSMC methods may require allocation of one or more additional RX endpoints for NB-IoT reception depending on the O-DU and O-RU capabilities.

In order to guarantee that NPUSCH and NPRACH can be scheduled without a limitation, at least 2 channels with symbols overlapping in time should be supported by an O-RU that supports NB-IoT.

15.5.3 Special case of in-band mode and guard-band mode NPUSCH with extended PRB Grid

The reception of NB-IoT NPUSCH in in-band mode with the SCS 15 kHz may be handled transparently for O-RU. In this case the O-DU shall configure the low-level-rx-endpoints via M-plane with appropriate **offset-to-absolute-frequency-center** and SCS, and send Section Type 1 C-plane message to receive a set of PRBs. The endpoint may be an endpoint used for the host carrier reception or an additional endpoint if the configuration of the endpoints used for the host carrier is not matching the configuration required for the NB-IoT carrier reception. The O-RU shall handle the reception as per the configuration provided by the O-DU. The O-DU shall handle any phase shift compensation related to misalignment of NB-IoT carrier transmitter and receiver frequencies.

Alternatively, reception of NB-IoT NPUSCH in in-band mode can be handled with the extended grid method as below especially if both in-band and guard-band NB-IoT carriers are present.

The reception of NB-IoT NPUSCH in guard-band mode with SCS 15 kHz may be handled with the extended PRB grid method described in clause 15.4.3 (handling of in-band mode is also permitted with this method). In this case, the O-DU shall configure the O-RU with **occupied-bandwidth** (the upper and the lower bound) and shall configure the O-RU's endpoint via M-plane with appropriate **offset-to-absolute-frequency-center** and SCS and send ST 1 C-plane message to receive a set of PRBs. The endpoint may be an endpoint used for the host carrier reception or an additional endpoint if the configuration of the endpoints used for the host carrier is not matching the configuration required for the NB-IoT carrier reception. The O-RU shall handle the reception as per the configuration provided by the O-DU. The O-DU handles extraction of NB-IoT PRBs from the extended PRB grid (see clause 15.4.3) and any phase shift compensation related to misalignment of NB-IoT carrier transmitter and receiver frequencies.

15.5.4 NPRACH without C-plane

Optionally, the reception of NPRACH may be handled without C-plane (see clause 7.9.9). For a subset of NPRACH configurations NPRACH reception without C-plane cannot be configured unambiguously without O-RU and O-DU supporting EXTENDED-PRACH-CONFIGURATION feature. This feature allows O-DU to provide O-RU with *snr-offset* that disambiguates NPRACH pattern start.

15.6 subframeId, slotId and symbolId numbering for NB-IoT

15.6.1 Overview

In the UL, a NB-IoT carrier may use different SCS than the 15 kHz host LTE or NR carrier driven by the physical channels NPUSCH and NPRACH as described in clause 15.6.2 and 15.6.3.

15.6.2 NPUSCH timing

Symbol numbering in a slot of NPUSCH with SCS 3.75 kHz follows the frame structure defined in 3GPP TS 36.211 [42] clause 10.1.2.3, which shall be from zero to six from an even subframe boundary with a duration of 2 ms. When the SCS value provided by "frameStructure" in the C-Plane messages is less than 15 kHz, startSymbolId shall be based on the frame structure of SCS 15 kHz, specified in clause 7.2.2. The even subframeId is applied in line with clause 7.5.2.5 for C-Plane and U-Plane messages for SCS 3.75 kHz NPUSCH. When Section Type 3 is used for NPUSCH data symbol with SCS 3.75 kHz, symbolId in the U-Plane message for the same data symbol will be based on the startSymbolId in the C-Plane message. Minimum scheduling unit for NPUSCH consists of multiple UL slots. A UL slot with NPUSCH SCS 3.75 kHz shall be commanded by C-Plane message with startSymbolId = "0" and numSymbol = "7" when Section Type 3 is used. An example for these timing parameters is shown in Fig.15.6.2-1.

NOTE: Section Type 1 frame structure for SCS 3.75 kHz is not defined in the present specification.

frameId	0											1																	
subframeId	0			2			4			6			8			0													
slotId	0						0						0						r										
symbolId	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0

Figure 15.6.2-1: Timing parameters for SCS 3.75 kHz NPUSCH

However, data symbol numbering in a slot of NPUSCH with SCS 15 kHz or downlink follows the frame structure of NR SCS 15 kHz. An example is shown in Fig.15.6.2-2.

frameId	0											1																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																													
subframeId	0	1	2	3	4	5	6	7	8	9	0																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																														
slotId	0	0	0	0	0	0	0	0	0	0	0																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																														
symbolId	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143	144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159	160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175	176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191	192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207	208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223	224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239	240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255	256	257	258	259	260	261	262	263	264	265	266	267	268	269	270	271	272	273	274	275	276	277	278	279	280	281	282	283	284	285	286	287	288	289	290	291	292	293	294	295	296	297	298	299	300	301	302	303	304	305	306	307	308	309	310	311	312	313	314	315	316	317	318	319	320	321	322	323	324	325	326	327	328	329	330	331	332	333	334	335	336	337	338	339	340	341	342	343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360	361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378	379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396	397	398	399	400	401	402	403	404	405	406	407	408	409	410	411	412	413	414	415	416	417	418	419	420	421	422	423	424	425	426	427	428	429	430	431	432	433	434	435	436	437	438	439	440	441	442	443	444	445	446	447	448	449	450	451	452	453	454	455	456	457	458	459	460	461	462	463	464	465	466	467	468	469	470	471	472	473	474	475	476	477	478	479	480	481	482	483	484	485	486	487	488	489	490	491	492	493	494	495	496	497	498	499	500	501	502	503	504	505	506	507	508	509	510	511	512	513	514	515	516	517	518	519	520	521	522	523	524	525	526	527	528	529	530	531	532	533	534	535	536	537	538	539	540	541	542	543	544	545	546	547	548	549	550	551	552	553	554	555	556	557	558	559	560	561	562	563	564	565	566	567	568	569	570	571	572	573	574	575	576	577	578	579	580	581	582	583	584	585	586	587	588	589	590	591	592	593	594	595	596	597	598	599	600	601	602	603	604	605	606	607	608	609	610	611	612	613	614	615	616	617	618	619	620	621	622	623	624	625	626	627	628	629	630	631	632	633	634	635	636	637	638	639	640	641	642	643	644	645	646	647	648	649	650	651	652	653	654	655	656	657	658	659	660	661	662	663	664	665	666	667	668	669	670	671	672	673	674	675	676	677	678	679	680	681	682	683	684	685	686	687	688	689	690	691	692	693	694	695	696	697	698	699	700	701	702	703	704	705	706	707	708	709	710	711	712	713	714	715	716	717	718	719	720	721	722	723	724	725	726	727	728	729	730	731	732	733	734	735	736	737	738	739	740	741	742	743	744	745	746	747	748	749	750	751	752	753	754	755	756	757	758	759	760	761	762	763	764	765	766	767	768	769	770	771	772	773	774	775	776	777	778	779	780	781	782	783	784	785	786	787	788	789	790	791	792	793	794	795	796	797	798	799	800	801	802	803	804	805	806	807	808	809	810	811	812	813	814	815	816	817	818	819	820	821	822	823	824	825	826	827	828	829	830	831	832	833	834	835	836	837	838	839	840	841	842	843	844	845	846	847	848	849	850	851	852	853	854	855	856	857	858	859	860	861	862	863	864	865	866	867	868	869	870	871	872	873	874	875	876	877	878	879	880	881	882	883	884	885	886	887	888	889	890	891	892	893	894	895	896	897	898	899	900	901	902	903	904	905	906	907	908	909	910	911	912	913	914	915	916	917	918	919	920	921	922	923	924	925	926	927	928	929	930	931	932	933	934	935	936	937	938	939	940	941	942	943	944	945	946	947	948	949	950	951	952	953	954	955	956	957	958	959	960	961	962	963	964	965	966	967	968	969	970	971	972	973	974	975	976	977	978	979	980	981	982	983	984	985	986	987	988	989	990	991	992	993	994	995	996	997	998	999	1000

Figure 15.6.2-2: Timing parameters for SCS 15 kHz NPUSCH

15.6.3 NPRACH timing

NPRACH preamble consists of multiple symbol groups. Each symbol group has a single cyclic prefix, one or more PRACH Symbols and is consecutive in time domain. When Section Type 3 is used for NPRACH then an entire symbol group shall be commanded by the single section description with cpLength = 0. SCS 3.75 kHz or 1.25 kHz is used for NPRACH, so that the subsequent symbolId in U-Plane messages following startSymbolId and numSymbol can be increased to value more than 13. This is in line with symbolId value definition in clause 8.3.3.7. An example for timing parameters in C-Plane and U-Plane messages which are associated with NPRACH preamble format 0 in FDD frame structure is shown in Figure 15.6.3-1.

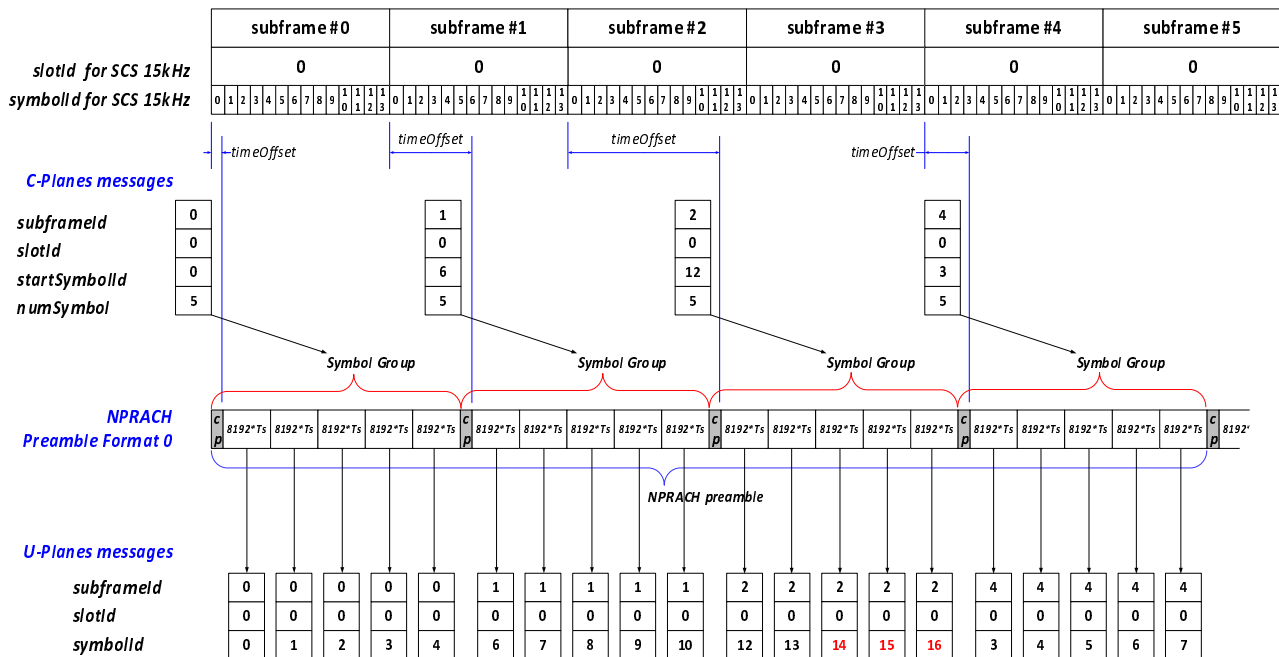


Figure 15.6.3-1: Example of Timing parameters in NPRACH U-Plane messages

Annex A (normative): Compression methods

A.1 Block floating point compression

A.1.1 Overview

The compressed data representation shall be as follows. For each PRB, In-phase (I) and Quadrature (Q) samples shall be converted to floating point format. The samples shall be represented as signed mantissa and a shared exponent (see Figure A.1.1-1). The compression procedure receives 12 subcarriers with 24 uncompressed I and Q samples. The I and Q samples shall be subsequently compressed to a signed, fixed bitwidth integer mantissa and 4-bit unsigned integer exponent. Further, the exponent shall be included for each compression block to be sent per PRB (see Figure A.1.1-2).

NOTE: Mantissa bitwidths are specified in the compHdr field of the relevant U-Plane or C-Plane message (range 1-16 bitwidth). O-RU support for particular mantissa bitwidths is to be defined on individual basis, see Table 10.2-1.




Figure A.1.1-1: Block Floating Point Compression data representation




Figure A.1.1-2: Block Floating Point Compression process overview

A.1.2 Block floating point compression algorithm

The following pseudo code depicts an example implementation of the compression algorithm.

Inputs:

- fPRB – Original physical resource block (PRB), 12 complex resource elements with "native" word length of the implementation e.g., 24 bits (UL should use as accurate value as possible from FFT & beamforming)
- iqWidth – Word length after compression (includes sign bit)

Outputs:

- cPRB – Compressed PRB, 12 complex resource elements with word length iqWidth
- exponent – Common exponent for compressed PRB

```

// Find max and min
maxV = max(Re(fPRB), Im(fPRB)), minV = min(Re(fPRB), Im(fPRB))
// Determine max absolute value
maxValue = max(maxV, |minV|-1) /* msb of negative value can be one higher */
// Calculate exponent
rawExp = [floor(log2(maxValue) + 1)] /* msb of maxValue */
// Calculate shift value and limit to positive
exponent = max(rawExp - iqWidth + 1, 0)
// Determine right shift value
scaler = 2-exponent
For iRe = 1:length(fPRB)
    // Scale and round, mult. may be bit-shift, Quantize may be or-round
    Re(cPRB(iRE)) = Quantize (scaler × Re(fPRB(iRE)))
    Im(cPRB(iRE)) = Quantize (scaler × Im(fPRB(iRE)))
End

```

A.1.3 Block floating point decompression algorithm

The following pseudo code depicts an example implementation of the decompression algorithm.

Inputs:

- cPRB – Compressed PRB, 12 complex resource elements with word length WL
- exponent – Common exponent for compressed PRB

Outputs:

- fPRB – Decompressed physical resource block (PRB), 12 complex resource elements with "native" word length for further processing. For example, 24 bits or 32 bits

```

//Determine scaler
scaler = 2exponent
For iRe = 1:length(cPRB)
    //Scale, this may be replaced with a bit-shift operation
    Re(fPRB(iRE)) = scaler × Re(cPRB(iRE))
    Im(fPRB(iRE)) = scaler × Im(cPRB(iRE))
End

```

A.2 Block scaling compression

A.2.1 Overview

A block scaling algorithm is specified, which is similar in concept to the block floating point representation, except that instead of data being represented by mantissa values and exponent shared within the block, data is instead represented by post-scaled values and a multiplicative scale value shared within the block. The data block size for block scaling compression shall be a single PRB, same as for the block floating point representation. Figure A.2.1-1 shows the algorithm in principle (assuming an 8-bit scaler value).

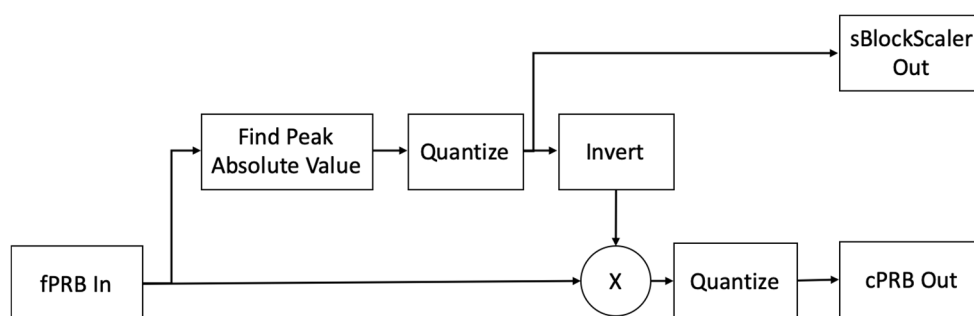


Figure A.2.1-1: Block scaling compression process diagram

A.2.2 Block scaling compression algorithm

The following pseudo code depicts an example implementation of the compression algorithm.

Inputs:

- fPRB – Original physical resource block (PRB), 12 complex resource elements with "native" word length of the implementation e.g., 16-bits I + 16-bits Q is assumed in this definition.
- iqWidth – Word length after compression (includes sign bit)

Outputs:

- cPRB – Compressed PRB, 12 complex resource elements with word length iqWidth
- sblockScaler – Common scaler for compressed PRB

```
// Find max and min of resource element real and imaginary parts
maxV = max(max(real(fPRB)), max(imag(fPRB)));
minV = min(min(real(fPRB)), min(imag(fPRB)));

// Determine maximum absolute value, which will be 15-bits
maxValue = max(maxV, abs(minV)-1); // Negative values offset by 1 to fit 15-bits

// Map sBlockScaler to 8-bits
sblockScaler = ceil(maxValue/2^8); // Q1.7, 0->128 represents 0->1.000

// Calculate inverse of sBlockScaler - this can be implemented as a look-up table
if(sblockScaler==0)
    sBlockScaler = 1; // Trap divide by zero
end
inverseBlockScaler = 2^7/sblockScaler; // Scaled 128->1

// Max output value, used for saturating compressed samples to iqWidth
qs = 2^(iqWidth-1);

// Scale each RE. Scaling assumes 16-bit uncompressed data width.
for iRe =1:length(fPRB)
    cRe_re = round(inverseBlockScaler * real(fPRB(iRe)) / 2^(16-iqWidth)); // Scale real part
    cRe_re = min(max(cRe_re, -qs), qs-1); // Saturate
    cRe_im = round(inverseBlockScaler * imag(fPRB(iRe)) / 2^(16-iqWidth)); // Scale imaginary part
    cRe_im = min(max(cRe_im, -qs), qs-1); // Saturate
    cPRB(iRe) = complex(cRe_re, cRe_im); // Complex output value
end
```

A.2.3 Block scaling decompression algorithm

The following pseudo-code depicts an example implementation of the block scaling decompression algorithm.

Inputs:

- cPRB – Compressed PRB, 12 complex resource elements with word length iqWidth
- blockScaler – Common scaler for compressed PRB

Outputs:

- fPRB – Decompressed physical resource block (PRB), 12 complex resource elements with "native" word length for further processing. For example, 16-bits I + 16-bits Q is assumed in this definition.

```
// Re-scale each resource element back to 16-bit uncompressed width
for iRe = 1:length(cPRB)
    fRe_re = round(sblockScaler * real(cPRB(iRe)) / 2^(iqWidth-9));
    fRe_im = round(sblockScaler * imag(cPRB(iRe)) / 2^(iqWidth-9));
    fPRB(iRe) = complex(fRe_re, fRe_im); // Complex output value
end
```

A.3 μ -Law compression

A.3.1 μ -Law compression algorithm

Inputs:

- `prbI` & `prbQ` – Original physical resource block (PRB), 12 complex resource elements with a word length of 16-bits I and 16-bits Q. The input bit width is fixed to 16-bits.
- `compBitWidth` – the length of I bits and the length of Q bits after compression over the entire PRB.

Outputs:

- `compI` & `compQ` – compressed PRB, 12 complex resource elements with word length `compBitWidth`, including sign, exponent and mantissa.
- `compShift` – the shift applied to the entire PRB.

The μ -law compression is a "block" compression and the block size shall be one PRB (same as for block floating point compression and block scaling compression). The O-RAN μ -law compression method combines a bit shift operation (for dynamic range) with a nonlinear piece wise approximation of μ -law compression. Parameters values shall be $\mu = 8$, 1 bit sign, and 2-bit mantissa. These values have been chosen for implementation efficiency.

Computations of the of μ -law compression algorithm are shown in Figure A.3.1-1.

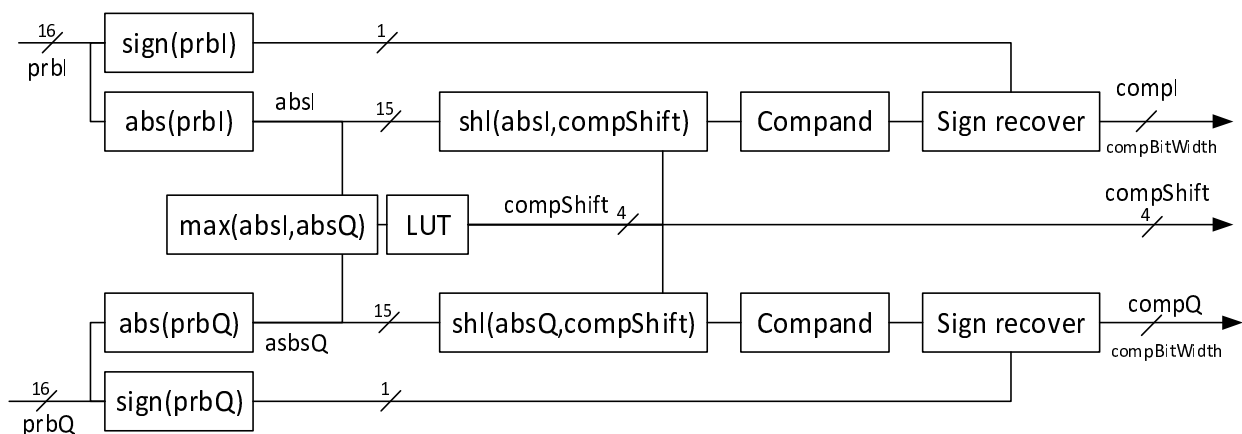


Figure A.3.1-1: μ -Law compression algorithm

The following pseudo-code depicts an example implementation of the μ -law compression algorithm.

```
// extract the sign bit and absolute value for the PRB
signI=sign(prbI), signQ=sign(prbQ), absI=abs(prbI), absQ=abs(prbQ)

// Find the maximum in absI and absQ
maxVal=max(absI, absQ)

// Determine compShift, the shift to be applied to the entire PRB
if maxVal>=2^14 then compShift=0
if maxVal<2^14 then compShift=1
if maxVal<2^13 then compShift=2
if maxVal<2^12 then compShift=3
if maxVal<2^11 then compShift=4
if maxVal<2^10 then compShift=5
if maxVal<2^9 then compShift=6
if maxVal<2^8 then compShift=7

// Apply round and shift left (shl - make greater)
absI = shl(absI, compShift)
absQ = shl(absQ, compShift)

// compand each sample, absBitWidth=15
if absI(i) > (2^absBitWidth-1) then absI(i) = (2^absBitWidth-1) // saturate
if absI(i) <= 2^(absBitWidth-2) then
  compI(i) = absI(i)/2^(absBitWidth- compBitWidth )
elseif absI(i) <= 2^(absBitWidth-1) then
  compI(i) = absI(i)/2^(absBitWidth- compBitWidth +1) + 2^(compBitWidth-3)
else
  compI(i) = absI(i)/2^(absBitWidth- compBitWidth +2) + 2^(compBitWidth-2)
end
if absQ(i) > (2^absBitWidth-1) then absQ (i) = (2^absBitWidth-1)
if absQ(i) <= 2^(absBitWidth-2) then
```

```

    compQ(i) = absQ(i)/2^(absBitWidth- compBitWidth )
  elseif absQ(i) <= 2^(absBitWidth-1) then
    compQ(i) = absQ(i)/2^(absBitWidth- compBitWidth +1) + 2^(compBitWidth-3)
  else
    compQ(i) = absQ(i)/2^(absBitWidth- compBitWidth +2) + 2^(compBitWidth-2)
  end
  // re-apply sign
  compI = round(signI .* compI);
  compQ = round(signQ .* compQ);

```

A.3.2 μ -Law decompression algorithm

Inputs:

- compI & compQ – compressed physical resource block (PRB), 12 complex resource elements with a word length of compBitWidth bits I and compBitWidth bits Q.
- compShift – the shift applied to the entire PRB.

Outputs:

- decompI & decompQ – the decompressed PRB, 12 complex resource elements with word length fixed to 16-bits I and 16-bits Q, including sign, exponent and mantissa.

The following pseudo code provides an example implementation of μ -law decompression algorithm.

```

// extract the sign bit and absolute value for the PRB
signI=sign(compI), signQ=sign(compQ), absI=abs(compI), absQ=abs(compQ)

// decompand each sample, absBitWidth=15
if absI(i) > (2^(compBitWidth -1)-1) then absI(i) = (2^(compBitWidth -1)-1)
if absI(i) <= 2^(compBitWidth -2) then
  decompI(i) = absI(i)*2^(absBitWidth- compBitWidth)
elseif absI(i) <= (2^(compBitWidth -2) + 2^(compBitWidth -3)) then
  decompI(i) = absI(i)*2^(absBitWidth- compBitWidth +1) - 2^13
else
  decompI(i) = absI(i)*2^(absBitWidth- compBitWidth +2) - 2^15
end

if absQ(i) > (2^(compBitWidth -1)-1) then
  absQ(i) = (2^(compBitWidth -1)-1)
if absQ(i) <= 2^(compBitWidth -2) then
  decompQ(i) = absQ(i) * 2^(absBitWidth- compBitWidth)
elseif absQ(i) <= (2^(compBitWidth -2) + 2^(compBitWidth -3)) then
  decompQ(i) = absQ(i) * 2^(absBitWidth- compBitWidth +1) - 2^13
else
  decompQ(i) = absQ(i) * 2^(absBitWidth- compBitWidth +2) - 2^15
end;

// Apply sign and shift
decompI = signI .* decompI
decompQ = signQ .* decompQ
decompI = decompI/2^compShift
decompQ = decompQ/2^compShift

```

A.3.3 μ -law udCompParam and IQ data format

PRB fields (see clause 8.3.3) are populated as follows:

- udCompParam (8 bits)
 - compBitWidth, 4 bits, (MSB)
 - compShift, 4-bits, (LSB)
- IQ samples, total bits = 12x 2x compBitWidth
 - 1st sample I, compBitWidth-bits

- 1st sample Q, compBitWidth-bits
- 2nd sample I, compBitWidth-bits
- 2nd sample Q, compBitWidth-bits
- ...
- 12th sample I, compBitWidth-bits
- 12th sample Q, compBitWidth-bits

A.4 Beamspace compression type I and type II

A.4.1 Overview

Two types of beamspace compression methods, i.e., beamspace compression type I and beamspace compression type II are specified. The beamspace compression methods are specific to beamforming weights and are not suitable for user or control IQ data. Hence these compression methods are only used as part of the `bfwCompMeth` in the C-Plane. Table A.4.1-1 lists the basic characteristics of the two methods.

Table A.4.1-1: Beamspace compression type I and II

	Transformation	IQ value compression
Beamspace compression type I	Only use 1D-DFT-based transformation	Block scaling
Beamspace compression type II	Linear antenna array uses 1D-DFT-based transformation and rectangular antenna array uses 2D-DFT-based transformation, according to the antenna array configuration	Block floating point

The pseudo codes for beamspace compression type I and beamspace compression type II are provided in clauses A.4.2 and A.4.3, respectively, as example implementations.

A.4.2 Beamspace compression type I

A.4.2.1 Overview

In beamspace compression type I, only 1D-DFT-based transformation is used for the transformation between antenna-space and beamspace, while block scaling compression is used to represent the IQ values of beamforming weights. In the following, the pseudo codes of the compression and decompression algorithms for beamspace compression type I are provided in following subclauses.

A.4.2.2 Beamspace compression algorithm

The following pseudo code depicts an example implementation of the beamspace compression type I algorithm.

Inputs:

- `fBV` – Original beamforming vector of K complex elements. K is the number of digital antenna ports supported by the O-RU (see clause 12.5.3 and Annex J) and is communicated to the DU during startup by the OAM subsystem. Each element is a complex number with a native bitwidth, e.g., 16-bit I, and 16-bit Q. Element `fBV(n)` of this vector corresponds to beamforming vector element $k = n-1$ (see clause 12.5.4).
- `iqWidth` – Word length of each I and Q value after compression (includes sign bit)

Outputs:

- `cBV` – Compressed beamforming coefficients

- `blockScaler` – Common scaler for compressed beamforming coefficients
- `activeBeamspaceCoefficientMask` – Vector of bits `activeBeamspaceCoefficientMask(n)` for $1 \leq n \leq K$. Bit value `activeBeamspaceCoefficientMask(n)=1` indicates presence of beamspace coefficient associated with the beamforming vector element k (see clause 12.5.4) in the compressed beamforming vector. The bit `activeBeamspaceCoefficientMask(1)` is conveyed in the most significant bit of the first octet of the field `bfwCompParam`.

In the following example implementation, if the absolute value of a beamspace coefficient is less than 'threshold', it is deemed inactive, i.e., the `activeBeamspaceCoefficientMask` shall be '0' at that coefficient index and this index shall not be sent across the fronthaul. The decompression algorithm shall assume a value of 0 for that coefficient. The value of threshold may be chosen by the implementer. Other methods to determine active or inactive beamspace coefficients may also be used.

```
// Generate DFT basis matrix
for k = 1 to K
  for l = 1 to K
    W(k,l) = 1/sqrt(K)*exp(i*2*pi*(k-1)*(l-1)/(K)) // W is a KxK complex matrix
  end for
end for
// Transform into beamspace
cBV = W*fBV // multiplication of a KxK complex matrix with a Kx1 complex vector yields another
//complex vector.

/* The algorithm is initialized to assume that all Beamspace Coefficients are transmitted across the
fronthaul link. */
for k = 1 to K
  activeBeamspaceCoefficientMask(k) = 1
end for

/* At this stage some of the beamspace coefficients may be removed from the vector of coefficients
to transmit across the fronthaul. */

t = 0
for k = 1 to K
  if abs(cBV(k)) < threshold
    activeIndex(k) = 0
    cBV(k) = null // remove the element from the vector
  else
    activeIndex(k) = 1
    t = t + 1
  end if
end for

T = t

// Calculate scaler
maxValue = max(abs(Re(cBV)), abs(Im(cBV)))

blockScaler = maxValue /* scaler can be chosen to be larger than maxValue. */
For iRe = 1 to T
  //Scale and round.
  //Quantize may be truncate or round
  //Inverse may be implemented via look up table or other methods
  Re(cBV(iRe)) = Quantize (Inverse(blockScaler) × Re(cBV(iRe)))
  Im(cBV(iRe)) = Quantize (Inverse(blockScaler) × Im(cBV(iRe)))
End
```

A.4.2.3 Beamspace decompression algorithm

The following pseudo-code depicts an example implementation of the block scaling decompression algorithm of beamspace compression type I.

Inputs:

- `cBV` – Compressed beamforming coefficients
- `blockScaler` – Common scaler for compressed beamforming coefficients
- `activeBeamspaceCoefficientMask` – see definition in clause A.4.2.2

Outputs:

- fBV – Decompressed beamforming vector of K (see K in clause 12.5.3) complex elements. Element fBV(n) of this vector corresponds to beamforming vector element $k = n-1$ (see clause 12.5.4)

```

m = 0
for k = 1 to K
  if activeBeamspaceCoefficientMask (k) = 1
    //Scale
    fBSC(k) = blockScaler × cBV(m)
    m = m + 1
  else
    fBSC(k) = 0
  end if
end for

// Generate DFT basis matrix
for k = 1 to K
  for l = 1 to K
    W (k,l) = 1/sqrt(K)*exp(-i*2*pi*(k-1)*(l-1)/(K)) // W is a K×K complex matrix
  end for
end for

fBV = W * fBSC

```

A.4.3 Beamspace compression type II

A.4.3.1 Overview

Beamspace compression type II uses 1D-DFT or 2D-DFT based transformation configuration, depending on array configuration. In this case, different beamspace transformation configurations are used for different antenna array configurations, i.e., the number of polarization P , the number of rows M , and the number of columns N . The total number of antenna ports $K = P \times M \times N$. Table A.4.3.1-1 shows how transformation is configured in O-RU and O-DU based on the antenna array configuration. Further, the block floating point compression shall be used to represent the IQ values of beamforming weights. More details are provided in the pseudo codes of the compression and decompression algorithms for beamspace compression type II in following subclauses.

Table A.4.3.1-1: type II beamspace transformation configuration

Antenna array	Transformation configuration
Single-polarized linear array ($P=1, M>1, N=1$ or $P=1, M=1, N>1$)	Transformation based on size K 1D-DFT
Dual-polarized linear array ($P=2, M>1, N=1$ or $P=2, M=1, N>1$)	Transformation based on size $K/2$ 1D-DFT per polarization
Single-polarized rectangular planar array ($P=1, M>1, N>1$)	Transformation based on size $M \times N$ 2D-DFT
Dual-polarized rectangular planar array ($P=2, M>1, N>1$)	Transformation based on size $M \times N$ 2D-DFT per polarization

A.4.3.2 Beamspace compression algorithm

The following pseudo code depicts an example reference implementation of the compression algorithm of beamspace compression type II.

Inputs:

- fBV – Original beamforming vector of K complex elements. $K = P \times M \times N$ is the number of digital antenna ports comprising P polarizations, M rows of antenna ports and N columns of antenna ports in the antenna array supported by the O-RU (see clause 12.5.3 and Annex J) and is communicated to the DU during startup by the OAM subsystem. Each element is a complex number with a native bitwidth e.g., 16-bit I, and 16-bit Q. Element fBV(n) of this vector corresponds to beamforming vector element $k = n-1$ (see clause 12.5.4).
- bwIqWidth – Word length of each I and Q value after compression (includes sign bit)

Outputs:

- cBV – Compressed beamforming coefficients
- exponent- Exponent for compressed beamforming coefficients

- **activeBeamspaceCoefficientMask** – Vector of bits **activeBeamspaceCoefficientMask(n)** for $1 \leq n \leq K$. Bit value **activeBeamspaceCoefficientMask(n)=1** indicates presence of beamspace coefficient associated with the beamforming vector element $k = n-1$ (see clause 12.5.4) in the compressed beamforming vector. The bit **activeBeamspaceCoefficientMask(1)** is conveyed in the most significant bit of the first octet of the field **bfwCompParam**.

In the following example implementation, if the absolute value of a beamspace coefficient is less than 'threshold', it is deemed inactive, i.e., the **activeBeamspaceCoefficientMask** shall be '0' at that coefficient index and this index shall not be sent across the fronthaul. The decompression algorithm shall assume a value of 0 for that coefficient. The value of threshold may be chosen by the implementer. Other methods to determine active or inactive beamspace coefficients may also be used.

For details of Kronecker product of matrices see clause A.4.3.4.

```
// Generate two DFT basis matrices and one identity matrix according to antenna array parameters of
//P, M and N.
for k = 1 to M
  for l = 1 to M
    Dm(k,l) = 1/sqrt(M)*exp(i*2*pi*(k-1)*(l-1)/M) // Dm is an MxM complex matrix
  end for
end for

for k = 1 to N
  for l = 1 to N
    Dn(k,l) = 1/sqrt(N)*exp(i*2*pi*(k-1)*(l-1)/N) // Dn is an NxN complex matrix
  end for
end for

for k = 1 to P
  for l = 1 to P
    if k==l
      Ip(k,l) = 1 // Ip is a PxP identity matrix
    else
      Ip(k,l) = 0
    end if
  end for
end for

// W is a KxK complex matrix, generated by Kronecker product of matrices Ip, Dn and Dm
W = kron(kron(Ip, Dn), Dm)

// Transform from antenna-space to beamspace
cBV = W*fBV // multiplication of a KxK complex matrix with a Kx1 complex vector yields another
//complex vector.

/* The algorithm is initialized to assume that all Beamspace Coefficients are transmitted across the
fronthaul link. */
for k = 1 to K
  activeBeamspaceCoefficientMask(k) = 1
end for

/* At this stage some of the beamspace coefficients may be removed from the vector of coefficients
to transmit across the fronthaul. */

t = 0
for k = 1 to K
  if abs(cBV(k)) < threshold
    activeIndex(k) = 0
    cBV(k) = null // remove the element from the vector
  else
    activeIndex(k) = 1
    t = t + 1
  end if
end for
T = t
// Block floating point compression of beamspace coefficients
// Calculate scaler. To simplify expressions, it is assumed that abs(cBV) < 1 (equality to be
//handled by some exception)
maxValue = max(abs(Re(cBV)), abs(Im(cBV)))

// Calculate exponent and limit
raw_exp = 1 + floor(log2(maxValue)) // less or equal to zero since maxValue < 1
MAXEXPBFW = 15 // The largest exponent, 24-1. This is a constant.
minExpBfw = 0 /* Limit dynamic range to correspond to (MAXEXPBFW - minExpBfw) + bfwIqWidth bits of
fixed-point. This could be configurable, e.g. to match 12-bit or 16-bit fixed point*/
```



```

exponent = max(rawExp + MAXEXPBFW, minExpBfw) /* Exponent can be chosen to be larger than this value
(but not larger than maxExp), e.g. to avoid mantissa overflow after rounding */

// Calculate shift value and inverse scaler
shiftValue = exponent - MAXEXPBFW - bfwIqWidth + 1
invScaler = 2-shiftValue

For iRe = 1 to T
  // Scale and round: Quantize could be truncate or round, also ensuring that mantissa values are
  //within bounds, multiplication by invScaler could be bit-shift with the shiftValue.
  Re(cBV(iRE)) = Quantize (invScaler × Re(cBV(iRE)))
  Im(cBV(iRE)) = Quantize (invScaler × Im(cBV(iRE)))
End

```

A.4.3.3 Beamspace decompression algorithm

The following pseudo code depicts an example implementation of the decompression algorithm of beamspace compression type II.

Inputs:

- cBV – Compressed beamforming coefficients
- exponent – exponent for compressed beamforming coefficients
- activeBeamspaceCoefficientMask – see definition in clause A.4.3.2
- bfwIqWidth – Word length of each I and Q value after compression (includes sign bit)

Outputs:

- fBV – Decompressed beamforming vector of K (see K in clause 12.5.3) complex elements. $K = P \times M \times N$ is the number of digital antenna ports comprising P polarizations, M rows of antenna ports and N columns of antenna ports. Element fBV(n) of this vector corresponds to beamforming vector element $k = n-1$ (see clause 12.5.4)

```

MAXEXPBFW = 15 // The largest exponent, 24-1. This is a constant.
shiftValue = exponent - MAXEXP - bfwIqWidth + 1
scaler = 2shiftValue // scale back compressed values to the proper beamforming weight range

m = 0
for k = 1 to K
  if activeBeamspaceCoefficientMask (k) = 1
    // Block floating point decompression of beamforming coefficients
    fBSC(k) = scaler × cBV(m)
    m = m + 1
  else
    fBSC(k) = 0
  end if
end for

// Generate two DFT basis matrices and one identity matrix according to antenna array parameters of
// P, M and N.
for k = 1 to M
  for l = 1 to M
    Dm (k,l) = 1/sqrt(M)*exp(-i*2*pi*(k-1)*(l-1)/M) // Dm is an M×M complex matrix
  end for
end for
for k = 1 to N
  for l = 1 to N
    Dn (k,l) = 1/sqrt(N)*exp(-i*2*pi*(k-1)*(l-1)/N) // Dn is an N×N complex matrix
  end for
end for
for k = 1 to P
  for l = 1 to P
    if k==l
      Ip(k,l) = 1 // Ip is a P×P identity matrix
    else
      Ip(k,l) = 0
    end if
  end for
end for

// W is a K×K complex matrix, generated by Kronecker product of matrices Ip, Dn and Dm.
W = kron(kron(Ip, Dn), Dm)

```

```
// Transform from beamspace back to antenna-space
fBV = W * fBSC
```

A.4.3.4 Kronecker product algorithm

The Kronecker product of an $L \times M$ matrix A and an $N \times P$ matrix B is denoted as $C = A \otimes B$, where C is an $(L \times N) \times (M \times P)$ matrix:

$$C = A \otimes B = \begin{bmatrix} a_{1,1}B & \cdots & a_{1,M}B \\ \vdots & \ddots & \vdots \\ a_{L,1}B & \cdots & a_{L,M}B \end{bmatrix}$$

The following pseudo code gives an example of the Kronecker product function $C = \text{kron}(A, B)$:

Inputs:

- A – an $L \times M$ matrix
- B – an $N \times P$ matrix

Output:

- C – Kronecker product of matrices A and B

```
function C = kron(A, B)
for l = 1 to L
  for m = 1 to M
    for n = 1 to N
      for p = 1 to P
        C((l-1)*N+n, (m-1)*P+p) = A(l, m)*B(n, p)
      end for
    end for
  end for
end for
```

A.5 Modulation compression

Modulation compression is an IQ data compression method that may be applied to DL data only and depends on the observation that modulated data symbols are represented by a very limited number of I and Q bits. For example, a QPSK modulated symbol has only two potential states of I and two potential states of Q, so such a symbol is representable with no loss of information with a single bit of I and a single bit of Q. Likewise, a 64QAM constellation point (8x8 constellation) is representable by at most 3 bits of I and 3 bits of Q. This allows for a reduction in DL throughput. See the two figures below for a description of this concept (Figure A.5-1 and Figure A.5-2 assume a single modulation type, MCS, is used for the data section).

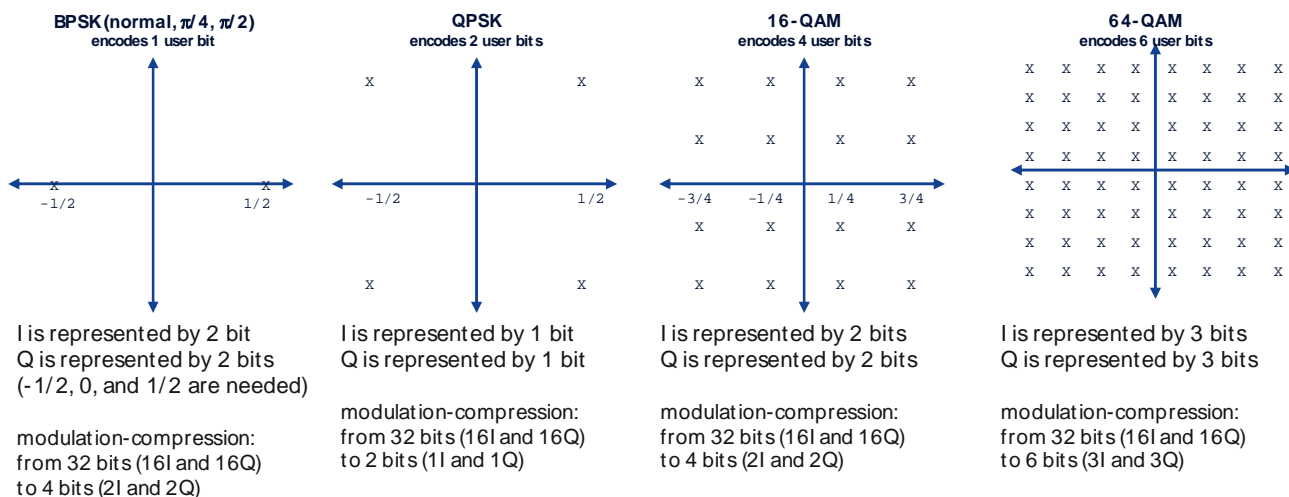


Figure A.5-1: Example constellation types

To represent the constellation points as I and Q values that also overlap allowing multiple constellation sizes to be represented by a single word-width, the constellations shall be "shifted" to allow a two's-complement I and Q value to represent any constellation point. Figure A.5-2 shows the same constellations after shifting.




Figure A.5-2: Shifted constellation points

Once the constellations are shifted, the I and Q values shall be encoded in a limited number of bits, being the larger number needed to represent the largest constellation possible in the compression block (the data section). This means that if some data in the section use 64QAM and others use QPSK (e.g., a reference RE) all REs would use the largest needed representation which in this example is 3 bits for I and 3 bits for Q (for 64QAM). This spoils the compression efficiency a little bit but because reference REs are a small fraction of the total number of REs, the efficiency degradation is small (there is further clarification below regarding mixing MCS in a data section). In general, every user has its own data section (and own beamforming index), so users with high-order modulation need and use more bits of I and Q while users with lower-order modulation need and use fewer bits of I and Q.

Some constellations need not be shifted. For example, BPSK needs I and Q data to take the values -1/2, zero, and +1/2 (different varieties of BPSK rotate these as with $\pi/4$ BPSK and $\pi/2$ BPSK). For this reason, BPSK would use two bits for I and 2 bits for Q; while this seems counterintuitive (BPSK using more bits than QPSK) this is a small penalty given the rarity of BPSK as a modulation type. Here, BPSK shall not be shifted. Likewise, PHICH constellations encode 3 states for each I and Q: -1, zero and 1. For this constellation the representation shall be -1/2, 0 and 1/2 with no constellation shift needed. However, constellations of all QAM modulations shall be shifted (except in mixed-MCS cases). Applied constellation shift shall be indicated by the "csf" field, where for every bit set (=1) in the reMask "csf" indicates whether to shift (csf=1) or not (csf=0) the associated RE.

Mixed-MCS cases represent another example when constellations shall not be shifted. This occurs when user data REs (at high MCS) and signaling data REs (at low MCS) are in the same PRB – hence in the same data section. The reMask discriminates the REs at high MCS from the REs at the lower MCS (and provides different modCompScaler or

mcScaleOffset values for the different-MCS data), but all the REs in the PRB shall use the same number of I and Q bit-widths. In this case only the high-MCS constellation shall be shifted, the lower-MCS constellation shall not be shifted, because its data points already overlap with the shifted high-MCS data points. Figure A.5-3 shows an example of 16QAM data overlain with 64QAM data.




Figure A.5-3: Multiple-MCS example for modulation compression

In the Figure A.5-3 example, for overlain 16QAM (green) and 64QAM (red), the high-MCS (red) points are shifted by $\frac{1}{2}$ the high-MCS resolution (here, $-\frac{1}{8}$) to allow all points to share the same "grid", as shown in the middle subfigure, wherein the red and green points overlay. All I and Q values in the given example are represented by 3 bits each on the fronthaul interface.

The O-DU shall use the constellation shift flag (csf) to tell the O-RU which data (red points) to "unshift" by adding $\frac{1}{8}$ to them, thereby restoring the original constellation values. After that, modCompScaler (or mcScaleOffset) shall be applied to set the data to the correct power levels (separate modCompScaler or mcScaleOffset values may be used for the differing MCS data).

When decompressing, the O-RU shall "unshift" the constellation depending on "csf" value and apply a scale factor for the constellation types represented in the section. There shall be either one or two modulation types in one section. The modulation type shall be inferred from the reMask bits, where each "one" bit indicates the shift command ("csf") and scale factor ("modCompScaler" when using Section Extension 4, and "mcScaleOffset" when using Section Extension 5) for the REs in the subject PRB. The scale factor allows not only for correcting for different constellation scaling (e.g., for multiplexed channel data in a PRB including QPSK and 16QAM, QPSK involves a $2/\sqrt{2}$ factor while 16QAM involves a $4/\sqrt{10}$ factor), but also allows for different channel power scaling which is permitted as a 3GPP option.

NOTE: Modulation compression method is essentially lossless, except that the scale factors, being 15 bits, impose a limit on the accuracy of representation. 15 bits is considered sufficient for all LTE and NR data representations.

When compressing, constellation points shall be shifted by the shift value defined in Table 7.7.4.2-1, and the I and Q values shall be represented as signed two's complement fractional notation and included in the U-Plane message as udIqWidth bit vectors, where udIqWidth is dependent on the modulation constellation type and is defined in Table 7.7.4.2-1.

The following pseudo code depicts an example implementation of the modulation decompression algorithm:

1. Read *iqSample* as an udIqWidth bit vector in the U-Plane message [this is all the IQ data in the data section]
2. Map *iqSample* $[0, 2^{\text{udIqWidth}} - 1]$ to *iqSampleFx* $[-1, 1)$ assuming that the udIqWidth bits are represented as Q1.(udIqWidth-1) [this is the normal two's-complement representation of the I and Q samples represented in fractional notation].
- 3X. For each RE in the PRB (using Section Extension 4):
 - 3Xa: fetch the "csf" and "modCompScaler" values for which this RE has a "1" in the reMask
 - 3Xb. If "csf" == 1 then $iqSampleFx = iqSampleFx + 2^{-\text{udIqWidth}}$ [this is "unshifting" the constellation point].
 - 3Xc. $iqSampleScaled = \text{modCompScaler} \times iqSampleFx \times \sqrt{2}$ [this scales the constellation point]
- 3Y. For each RE in the PRB (using Section Extension 5):
 - 3Ya: fetch the "csf" and "mcScaleOffset" values for which this RE has a "1" in the relevant mcScaleReMask

3Yb. If "csf" == 1 then $iqSampleFx = iqSampleFx + 2^{-udIqWidth}$ [this is "unshifting" the constellation point].

3Yc. $iqSampleScaled = mcScaleOffset \times iqSampleFx \times \sqrt{2}$ [this scales the constellation point]

After decompression, $|iqSampleScaled|$ shall be ≤ 1 and a value of $|iqSampleScaled| = 1.0$ matches 0 dBFS.

A.6 Selective RE sending

Selective RE sending may be used together with either block floating point compression (see clause A.1) or modulation compression (see clause A.5). The compression algorithm shall be performed on all 12 IQ-samples in the PRB, including any possible zero valued IQ-samples.

For each uncompressed IQ-sample where both the I- and the Q-value are 0, corresponding bit shall be set to 0b in the Selective RE sending bitmask (sReSMask) and the corresponding compressed IQ-sample shall not be transmitted over the interface in the U-Plane message. Remaining bits in sReSMask shall be set to 1b, see clause 8.3.3.18.

If the total number of bits for the transmitted IQ-samples is not a multiple of 8 (1 byte), then bit-padding shall be performed per PRB after the last part of the last IQ-samples of the corresponding PRB so that each set of U-Plane PRB data or new U-Plane data section shall be byte aligned.

Figure A.6-1 illustrates an example where 9 out of the 12 IQ-samples in a PRB are zero and thus not sent in the corresponding U-Plane message.

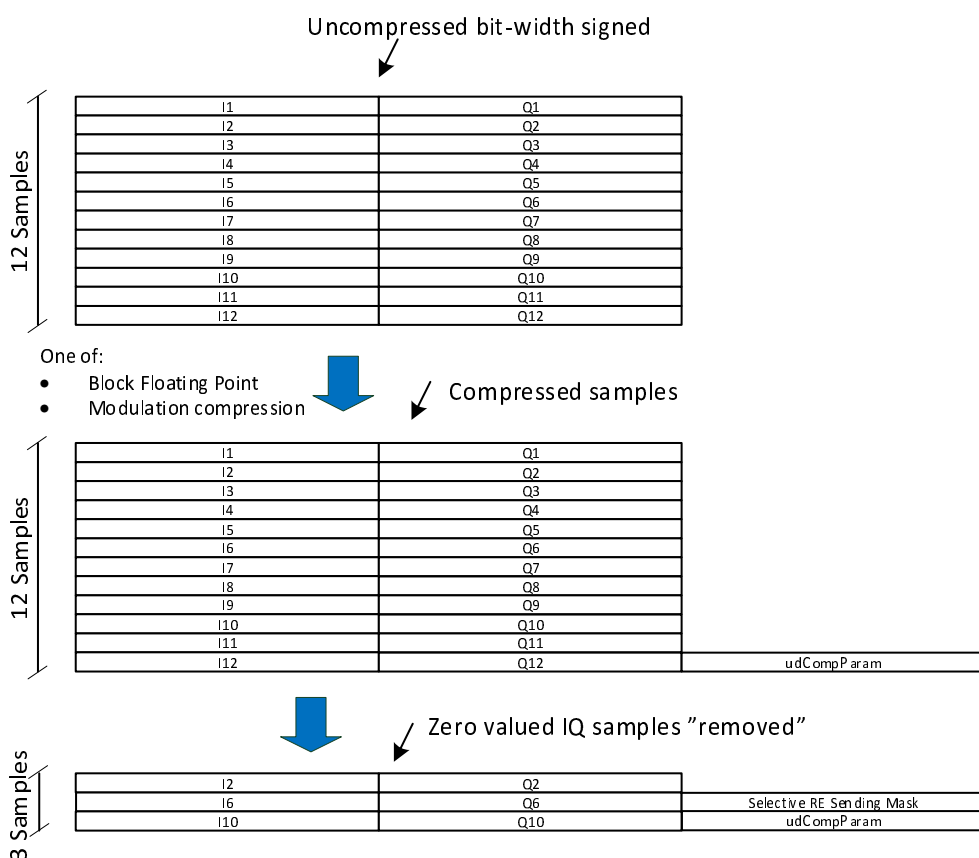


Figure A.6-1: Example of selective RE sending

Annex B (informative): Delay management use cases

B.1 Introduction

Intra-PHY lower layer fronthaul split has characteristic of a stringent bandwidth and tight latency requirement. This implies use of a special "Fronthaul Service Profile" to be supported by the transport network, and which may differ depending on the operating environment, topology and target use cases. The general concept and latency boundaries is based on eCPRI reference points for delay management definitions (See Figure B.1-1). However, the present document additionally differentiates between DL (See Figure B.1-2) and UL (See Figure B.1-3) latency boundaries. The parameters and how these are determined for a pre-defined latency configuration are explained below (an actual example of the parameters for a specific use case are presented in Annex B.2). Pre-defined latency is necessary when actual latency measurements are not provided; both the use of pre-defined latency value and use of a method for measuring actual network latency in the DL and UL are supported in the present document.

The following assumptions are considered for the delay boundaries definitions:

- $T_{cp_adv_dl}$: Smallest time advance to receive Downlink Data C-Plane message before the first IQ data can be processed
- The fronthaul transmission delay behavior for C-Plane for DL data, C-Plane for UL data, and U-Plane for DL data is equal. Thus, there is common usage of T_{12_min} and T_{12_max} parameters.
- The transmission window ($T_{1a_max} - T_{1a_min}$) for C-Plane for DL data, C-Plane for UL data, and U-Plane for DL data all have the same length.
- The reception window ($T_{2a_max} - T_{2a_min}$) for C-Plane for DL data, C-Plane for UL data, and U-Plane for DL data all have the same length.
- $T_{2a_min_cp_ul}$: Latest availability at O-RU of C-Plane for UL data message before reception of the first IQ data sample of the respective user's U-Plane UL data packet is received over the air interface.




Figure B.1-1: Definition of reference points for delay management (adapted from [2])

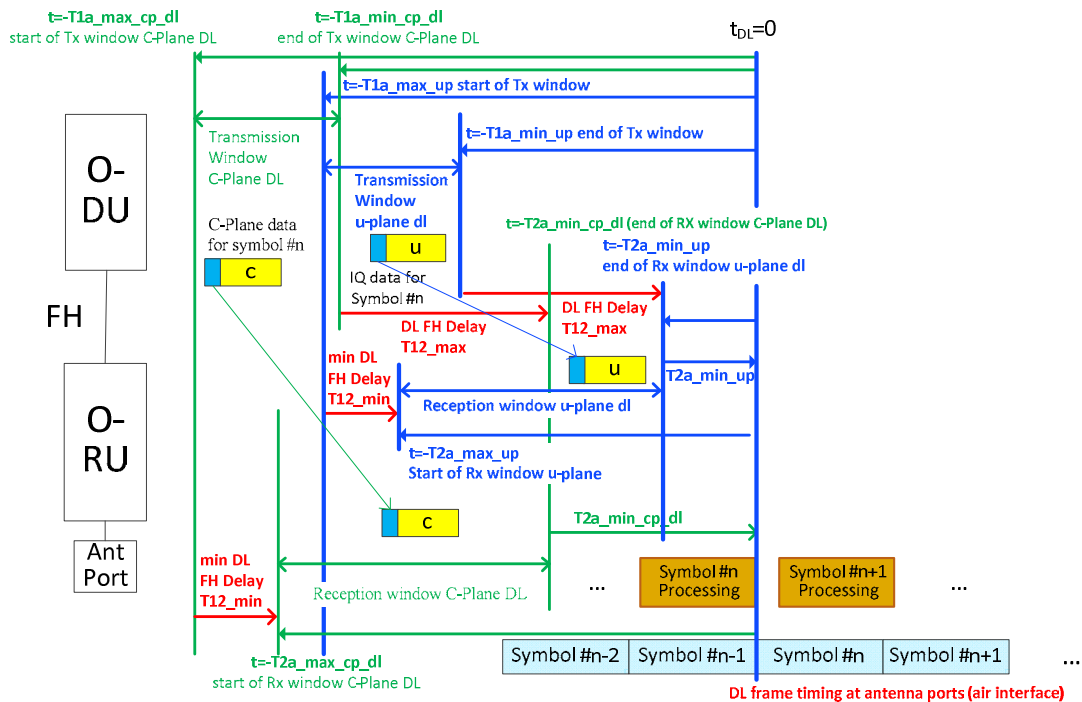


Figure B.1-2: Timing relations per symbol IQ in DL direction (U-Plane and C-Plane)

Figure B.1-2 is based on the eCPRI delay measurement model on timing relations in DL direction. More detail is added to illustrate the following data transfer timing relations:

- U-Plane DL data: Blue path
- C-Plane for DL data: Green path

To understand this timing diagram, it is easier to work backward in timeline.

For U-Plane DL data flow (use symbol #0 transmission as an example):

- $t = 0$: time of transmission (at air interface) of the first sample for symbol #0 (see $t_{DL} = 0$)
- $t = -T2a_{min_up}$: O-RU has a fixed data processing delay ($T2a_{min_up}$). In order to meet air interface time for symbol#0 transmission at $t = 0$, symbol#0 data is presented on time for the processing unit (yellow block).
 - For symbol#0, start of processing time $t = (0 - T2a_{min_up}) = -T2a_{min_up}$
- End of reception window: The latest time that O-RU can accept U-Plane DL data for a specific symbol prior to start of data processing. If U-Plane DL data arrives earlier than this time and is within reception window range, DL data may wait inside reception window buffer until the start of processing time, hence - the end of reception window is the same time point as the start of processing time as described immediately above.
 - For symbol#0, end of reception window time $t = (0 - T2a_{min_up}) = -T2a_{min_up}$.
- Start of reception window: This earliest time that O-RU can accept U-Plane DL data for a specific symbol prior to start of data processing. If U-Plane DL data arrives later than this time and is within reception window range, DL data will wait inside reception window buffer.
 - For symbol#0, start of reception window time $t = (0 - T2a_{max_up}) = -T2a_{max_up}$
- O-RU reception window range = $T2a_{max_up} - T2a_{min_up}$
- End of transmission window: The latest time that O-DU can send U-Plane DL data for a specific symbol out to transport interface.
 - For symbol#0, end of transmission window time $t = (0 - T1a_{min_up}) = -T1a_{min_up}$.

- Start of transmission window: The earliest time that O-DU can send U-Plane DL for a specific symbol out to transport interface.
 - For symbol#0, start of transmission window time $t = (0 - T1a_max_up) = - T1a_max_up$.
- O-DU transmission window range = $(T1a_max_up - T1a_min_up)$
- DL FH transport delay: $T12_min$ and $T12_max$ is the min and max transport delay. When FH transport is behaving properly, it should guarantee each packet transport (carrying C-Plane or U-Plane DL packet) delay is within this min and max boundary.
- O-DU transmission window, FH transport delay and O-RU reception window satisfy the following criteria for proper behavior:
 - Start of reception window can accept a packet sent at start of transmission window AND experienced min FH transport delay
 - $T2a_max_up < (-T1a_max_up + T12_min)$
 - End of reception window can accept a packet sent at end of transmission window AND experienced max FH transport delay
 - $T2a_min_up > (-T1a_min_up + T12_max)$
 - Combining both criteria, the logical conclusion is as follows:
 - **O-RU reception window range > (O-DU transmission window + FH DL transport max-min)**
 - **$(T2a_max_up - T2a_min_up) > ((T1a_max_up - T1a_min_up) + (T12_max - T12_min))$**

In other words, the reception window at O-RU is greater than or equal to the total of the O-DU transmission variability and the Transport time variability.

For C-Plane to support DL data flow (use symbol #n = 0 transmission as an example):

- Focus on the green path, same principle is applied to relate reception window, transmission window and FH transport delay.
- the logical conclusion is as follows:
 - **Reception window range > (Transmission window + FH DL transport max-min)**
 - **$(T2a_max_cp_dl - T2a_min_cp_dl) > ((T1a_max_cp_dl - T1a_min_cp_dl) + (T12_max - T12_min))$**

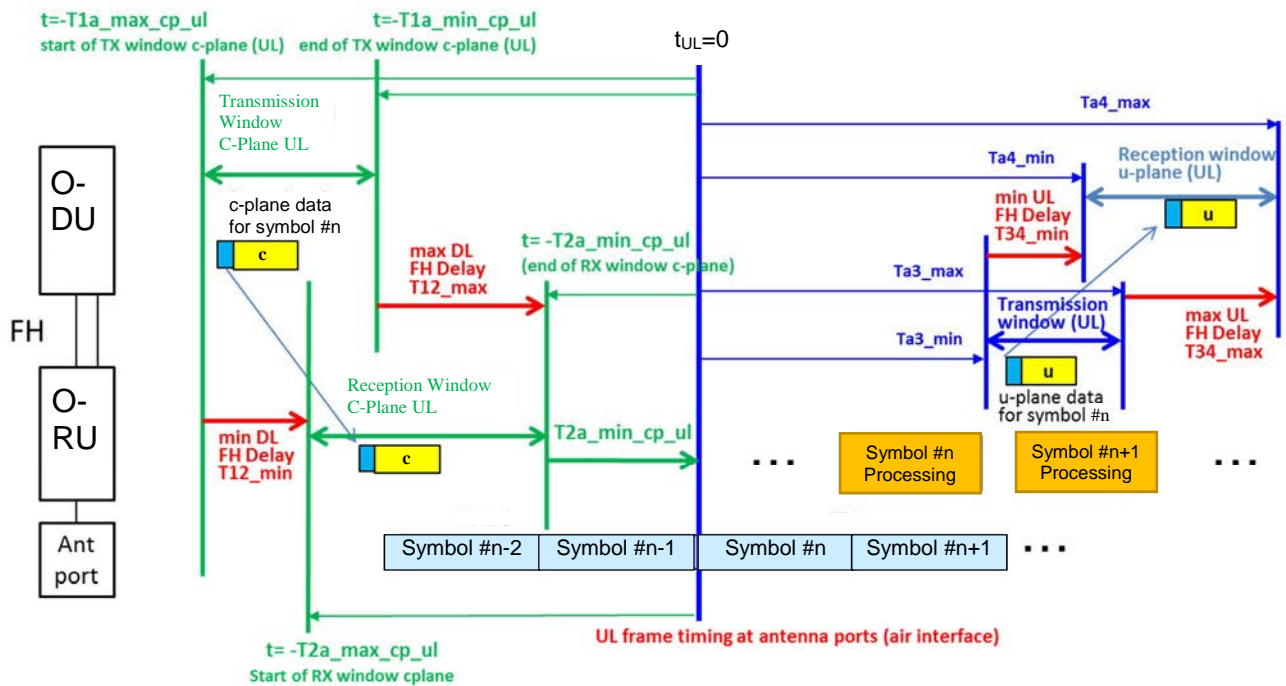


Figure B.1-3: Timing relations per symbol IQ in UL direction (U-Plane and C-Plane)

Figure B.1-3 is based on the eCPRI delay measurement model on timing relations in UL direction. More detail is added to illustrate the following data transfer timing relations:

- U-Plane UL data: Blue path.
- C-Plane for UL data: Green path.

To understand this timing diagram, it is easier to work forward in timeline for the U-Plane UL data and work backward in timeline for the C-Plane to support UL data flow. This example is applicable for non-PRACH channel.

For U-Plane UL data flow (use symbol #n = 0 reception as an example):

- $t = 0$: time of reception (at air interface) of the first sample for symbol #0 (see $t_{UL} = 0$).
- $t = (0 + Ta3_{min_up})$: O-RU has a fixed data processing delay ($Ta3_{min_up}$). Air interface data is immediately presented to data processing unit (yellow block)..
 - For symbol#0, end of processing = $0 + Ta3_{min_up} = Ta3_{min_up}$
- Start of transmission window: The earliest time that O-RU can send U-Plane UL for a specific symbol out to transport interface. The earliest time is immediately after data processing.
 - For symbol#0, start of transmission window = $(0 + Ta3_{min_up} = Ta3_{min_up})$.
- End of transmission window: The latest time that O-RU can send U-Plane UL data for a specific symbol out to transport interface.
 - For symbol#0, end of transmission window = $(0 + Ta3_{max_up} = Ta3_{max_up})$.
- O-RU transmission window range = $(Ta3_{max_up} - Ta3_{min_up})$
- Start of Reception window: This earliest time that O-DU can accept U-Plane UL data for a specific symbol prior to start of O-DU data processing. If U-Plane DL data arrives later than this time and is within reception window range, UL data will wait inside reception window buffer.
 - For symbol#0, start of reception window = $(0 + Ta4_{min}) = Ta4_{min}$
- End of Reception window: The latest time that O-DU can accept U-Plane UL data for a specific symbol prior to start of O-DU data processing.

- For symbol#0, end of reception window = $(0 + Ta4_max) = Ta4_max$.
- O-DU reception window range = $(Ta4_max - Ta4_min)$
- UL FH Transport delay: $T34_min$ and $T34_max$ is the min and max transport delay. When FH transport is behaving properly, it should guarantee each packet transport (carrying U-Plane UL packet) delay is within this min and max boundary.
- Transmission window, FH transport delay and reception window satisfies the following criteria for proper behavior:
 - Start of reception window can accept a packet sent at start of transmission window AND experienced min FH transport delay

$$Ta4_min < (Ta3_min_up + T34_min)$$
 - End of reception window can accept a packet sent at end of transmission window AND experienced max FH transport delay

$$Ta4_max > (Ta3_max_up + T34_max)$$
 - Combining both criteria, the logical conclusion is as follows:
 - **O-DU reception window range > (O-RU transmission window + FH UL transport max-min)**
 - **$(Ta4_max - Ta4_min) > ((Ta3_max_up - Ta3_min_up) + (T34_max - T34_min))$**

In other words, the reception window at O-DU is greater than or equal to the total of the O-RU transmission variability and the Transport time variability.

For C-Plane to support UL data flow (use symbol #0 transmission as an example):

- Focus on the green path, same principle to C-Plane to support DL data is applied to relate reception window, transmission window and FH transport delay.
- the logical conclusion is as follows:
 - **Reception window range > (Transmission window + FH transport min-max difference)**
 - **$(T2a_max_cp_ul - T2a_min_cp_ul) > ((T1a_max_cp_ul - T1a_min_cp_ul) + (T12_max - T12_min))$**

It is useful to consider several different delay management use cases because delay management, in particular the need to buffer large amounts of data within O-RUs, can have a significant cost and complexity impact on O-RU design. More specifically, accommodating long network delays (allowing long fiber lengths and/or many switch hops) can impose a significant buffering requirement on the O-RU while use cases involving short network delays e.g. an in-building application may allow much smaller in-O-RU buffering thereby allowing a lower-cost and lower-power design. Additionally, the SCS may affect the transmission windows described above, and thus use cases may be SCS specific as well. Here it is assumed the DL buffering will be done in the O-RU, otherwise the O-DU would have to always implement extensive buffering to accommodate any O-RU use case which would increase overall costs.

In the case of using pre-defined network latency values per use case, the following is pre-defined (for measured-network-latency cases these same values are determined via the measurement process):

- $T12_min$, $T12_max$, $T34_min$ and $T34_max$ are determined based on the desired network configuration. Determination of $T12_min$ and $T34_min$ values includes transmission paths, both fiber and minimum switching delays. In addition, longest fibers, switching delays, and PDV introduced due to the variable delay nature of ethernet is accounted for in $T12_max$ and $T34_max$.
- $T2a_min_up$, $T2a_min_cp_ul$, and $Ta3_min$ are determined across all O-RU equipment to be used in the system. The determined values are greater than or equal to the largest of these values across all supported equipment. Equipment with lower values need to have additional delay added to align with the determined values.

Similarly, the maximum transmission windows is determined across all O-DU and O-RU equipment. The transmission windows for DL UP, DL CP and UL CP are all assumed to be the same. However, this is the maximum transmission window. Equipment may use less time for transmission of any symbol.

Additionally, the advance between DL UP/ CP are determined. ($T_{cp_adv_dl}$)

Once the above parameters are determined, the remaining values are calculated as shown in the following Table B.1-1 through B.1-4:

Downlink Data Direction

Table B.1-1: U-Plane DL delay boundaries

Downlink	Method
$T1a_max_up$	$\leq (T12_min + T2a_max_up)$
$T1a_min_up$	$\geq (T12_max + T2a_min_up)$
$T2a_max_up$	$\geq (T2a_min_up + (T12_max - T12_min) + \text{O-DU Transmission Window})$
$T2a_min_up$	Specified per Use Case
$T12_max$	Specified per Use Case
$T12_min$	Specified per Use Case
O-DU Transmission Window	Specified per Use Case

Table B.1-2: C-Plane DL delay boundaries

Downlink	Method
$T_{cp_adv_dl}$	Specified per Use Case
$T1a_max_cp_dl$	$\geq (T1a_max_up + T_{cp_adv_dl})$
$T1a_min_cp_dl$	$(T1a_min_up + T_{cp_adv_dl})$
$T2a_max_cp_dl$	$\geq (T2a_max_up + T_{cp_adv_dl})$
$T2a_min_cp_dl$	$(T2a_min_up + T_{cp_adv_dl})$
$T12_max$	Same as U-Plane DL
$T12_min$	Same as U-Plane DL

Uplink Data Direction (need not be the same as the Downlink values)

Table B.1-3: U-Plane Uplink delay boundaries

Uplink	Method
$Ta3_max$	$\leq (Ta3_min + \text{O-RU Transmission Window})$
$Ta3_min$	Specified per Use Case
$Ta4_max$	$\geq (Ta3_max + T34_max)$
$Ta4_min$	$\leq (Ta3_min + T34_min)$
$T34_max$	Specified per Use Case
$T34_min$	Specified per Use Case
O-RU Transmission Window	Specified per Use Case

Table B.1-4: C-Plane Uplink delay boundaries

Downlink	Method
$T1a_max_cp_ul$	$\leq (T12_min + T2a_max_cp_ul)$
$T1a_min_cp_ul$	$\geq (T12_max + T2a_min_cp_ul)$
$T2a_max_cp_ul$	$\geq (T2a_min_cp_ul + (T12_max - T12_min) + \text{O-DU uplink C-Plane Transmission Window duration})$
$T2a_min_cp_ul$	Specified per Use Case
$T12_max$	Specified per Use Case
$T12_min$	Specified per Use Case
O-DU Transmission Window	Specified per Use Case

B.2 Latency categories and sub-categories

The O-DU and O-RU are defined to be in "delay categories" and "delay sub-categories" for the purposes of allowing a matching of O-DU and O-RU units that will operate properly together from the point of view of accommodating a

customer's network delay. Network delay comprises the "time-of-flight" of signals through (typically) a fiber-optic cable (so can be known from the fiber length) added to the signal traversal latency through any switches in the network.

The delay category and delay sub-category values depend in part on the processing latency within the O-DU and O-RU. It may be expected especially for an O-RU that the processing latency may depend on the specific frequency band and sub-carrier spacing that is used. Further, a multi-band radio may experience different processing latencies for its different bands. Therefore, it may be expected that an O-RU (and perhaps more rarely an O-DU) will have different delay category and delay sub-category ratings for different bands.

O-DU categories are defined as A.##[###] (per direction, i.e. O-DU UL and O-DU DL are different categories)

- Prefix(blank/f): The prefix distinguishes the type of O-DU with dynamic/fixed timing advance/retard
- A: Category from [AAAA-N] per table below
- .##[###]: sub-category from [.00 - .1000] per tables below

O-RU categories are defined as X.##[###] (per direction, i.e. O-RU UL and O-RU DL are different categories)

- X: Category from [O-ZZ] per table below
- .##[###]: sub-category from [.00-.1000] per tables below

Table B.2-1 is used by equipment vendors to assign categories to their equipment based on design characteristics.

Table B.2-1: O-DU and O-RU Delay Categories

O-DU Category		O-RU Category	
Category	<ul style="list-style-type: none"> • T1a_max_up_{O-DU} - Txmax_{O-DU} • Ta4_max_up_{O-DU} [μsec]	Category	<ul style="list-style-type: none"> • T2a_min_up OR • Ta3_max [μsec]
AAAA	≥ 30000	O	0 to 50
AAA	10000 to 29999	P	51 to 70
AA	3000 to 9999	Q	71 to 90
A	400 to 2999	R	91 to 110
B	380 to 399	S	111 to 130
C	360 to 379	T	131 to 150
D	340 to 359	U	151 to 170
E	320 to 339	V	171 to 190
F	300 to 319	W	191 to 210
G	280 to 299	X	211 to 230
H	260 to 279	Y	231 to 250
I	240 to 259	Z	251 to 270
J	220 to 239	ZZ	≥ 271
K	200 to 219		
L	180 to 199		
M	160 to 179		
N	0 to 159		

NOTE: Categories are defined to group endpoints with similar delay characteristics for easy evaluation relative to use cases. Specifically, categories AA-AAAA and ZZ are defined for non-ideal fronthaul. The calculated value for O-DU or O-RU falling anywhere within the range for the category indicates that the endpoint is classified as that category. It is NOT required that the endpoint be able to support the full range of the category.

Table B.2-2 and Table B.2-3 are intended for use by network providers to determine the best and worst case T12_max/ T34_max values that can be supported by a given equipment combination (excluding non-ideal fronthaul). Alternatively, network providers may locate the desired T12_max/ T34_max and select from the equipment

combinations meeting that criteria. Common criteria are identified by different colors on diagonals through the Table B.2-2 and Table B.2-3.

Table B.2-2: Latency_min (Minimum supported T12_max/ T34_max in μ s)

RU	O-DU													
	A	B	C	D	E	F	G	H	I	J	K	L	M	N
O	350	330	310	290	270	250	230	210	190	170	150	130	110	0
P	330	310	290	270	250	230	210	190	170	150	130	110	90	0
Q	310	290	270	250	230	210	190	170	150	130	110	90	70	0
R	290	270	250	230	210	190	170	150	130	110	90	70	50	0
S	270	250	230	210	190	170	150	130	110	90	70	50	30	0
T	250	230	210	190	170	150	130	110	90	70	50	30	10	0
U	230	210	190	170	150	130	110	90	70	50	30	10	0	0
V	210	190	170	150	130	110	90	70	50	30	10	0	0	0
W	190	170	150	130	110	90	70	50	30	10	0	0	0	0
X	170	150	130	110	90	70	50	30	10	0	0	0	0	0
Y	150	130	110	90	70	50	30	10	0	0	0	0	0	0
Z	130	110	90	70	50	30	10	0	0	0	0	0	0	0

Table B.2-3: Latency_max (Maximum supported T12_max/ T34_max in μ s)

RU	O-DU													
	A	B	C	D	E	F	G	H	I	J	K	L	M	N
O	3000	399	379	359	339	319	299	279	259	239	219	199	179	159
P	2949	348	328	308	288	268	248	228	208	188	168	148	128	108
Q	2929	328	308	288	268	248	228	208	188	168	148	128	108	88
R	2909	308	288	268	248	228	208	188	168	148	128	108	88	68
S	2889	288	268	248	228	208	188	168	148	128	108	88	68	48
T	2869	268	248	228	208	188	168	148	128	108	88	68	48	28
U	2849	248	228	208	188	168	148	128	108	88	68	48	28	8
V	2829	228	208	188	168	148	128	108	88	68	48	28	8	0
W	2809	208	188	168	148	128	108	88	68	48	28	8	0	0
X	2789	188	168	148	128	108	88	68	48	28	8	0	0	0
Y	2769	168	148	128	108	88	68	48	28	8	0	0	0	0
Z	2749	148	128	108	88	68	48	28	8	0	0	0	0	0

Table B.2-4: Use Case Mapping (Based on Latency_min)

Range (KM)	T12max T34max (μ sec)	Minimum T12Max Guaranteed for Category Combinations (Combinations may support larger T12max)													
		AO - AT	BO - BS	CO - CR	DO - DQ	EO - EP	FO								
50	250	AO - AT	BO - BS	CO - CR	DO - DQ	EO - EP	FO								
45	225	AU	BT	CS	DR	EQ	FP	GO							
40	200	AV	BU	CT	DS	ER	FQ	GP	HO						
35	175	AW	BV	CU	DT	ES	FR	GQ	HP	IO					
30	150	AX - AY	BW - BX	CV - CW	DU - DV	ET - EU	FS - FT	GR - GS	HQ - HR	IP - IQ	JO - JP	KO			
25	125	AZ	BY	CX	DW	EV	FU	GT	HS	IR	JQ	KP	LO		
20	100		BZ	CY	DX	EW	FV	GU	HT	IS	JR	KQ	LP	MO	
15	75			CZ	DY	EX	FW	GV	HU	IT	JS	KR	LQ	MP	
10	50				DZ	EY - EZ	FX - FY	GW - GX	HV - HW	IU - IV	JT - JU	KS - KT	LR - LS	MQ - MR	
5	25						FZ	GY	HX	IW	JV	KU	LT	MS	
0	0							GZ	HY	IX	JW	KV	LU	MT	

NOTE: The Use Case Mapping Table B.2-4 can be used to determine the O-DU/ O-RU delay category combinations which are guaranteed to support a given latency use case. This is defined based on Latency_min for the category. It may be possible for equipment from a lower category to support a higher use case, but cannot be guaranteed. To determine the exact maximum latency for an O-DU/ O-RU pair, it is necessary to calculate based on the delay parameter values for the paired equipment.

Table B.2-5: O-DU and O-RU Tx/Rx Window Sub-Categories

Receive Window Sub-Category		Transmit Window Sub-Category	
Sub-Category	<ul style="list-style-type: none"> • T2a_max_up - T2a_min_up • RXmax_{O-DU} [in usec] 	Sub-Category	<ul style="list-style-type: none"> • TXmax_{O-DU} • Ta3_max - Ta3_min [in usec]
.1000	≥ 10000	.20	≥ 200
.300	3000 to 9999	.19	190 to 199
.100	1000 to 2999	.18	180 to 189
.40	400 to 999	.17	170 to 179
.38	380 to 399	.16	160 to 169
.36	360 to 379	.15	150 to 159
.34	340 to 359	.14	140 to 149
.32	320 to 339	.13	130 to 139
.30	300 to 319	.12	120 to 129
.28	280 to 299	.11	110 to 119
.26	260 to 279	.10	100 to 109
.24	240 to 259	.09	90 to 99
.22	220 to 239	.08	80 to 89
.20	200 to 219	.07	70 to 79
.18	180 to 199	.06	60 to 69
.16	160 to 179	.05	50 to 59
.14	140 to 159	.04	40 to 49
.12	120 to 139	.03	30 to 39
.10	100 to 119	.02	20 to 29
.08	80 to 99	.01	10 to 19
.06	60 to 79	.00	0 to 9
.04	40 to 59		
.02	20 to 39		
.01	10 to 19		
.00	0 to 9		

NOTE: Sub-categories (see Table B.2-5) are defined to group endpoints with similar delay characteristics for easy evaluation relative to use cases. Specifically, sub-categories .100 to .1000 are defined for non-ideal fronthaul. The calculated value for O-DU or O-RU falling anywhere within the range for the category indicates that the endpoint is classified as that sub-category. It is NOT required that the endpoint be able to support the full range of the sub-category.

Table B.2-6 is used by service providers to identify equipment sub-category combinations which meet the desired network variability (excluding non-ideal fronthaul). Variability is shown in km in Table B.2-6. This range is based on 5 μs per km.

Table B.2-6: Dynamic Range (in km) for sub-category pairs

TxMax	Receive Window																				
	.40	.38	.36	.34	.32	.30	.28	.26	.24	.22	.20	.18	.16	.14	.12	.10	.08	.06	.04	.02	.01
.20	40	36	32	28	24	20	16	12	8	4											
.19	42	38	34	30	26	22	18	14	10	6	2										
.18	44	40	36	32	28	24	20	16	12	8	4										
.17	46	42	38	34	30	26	22	18	14	10	6	2									
.16	48	44	40	36	32	28	24	20	16	12	8	4									
.15	50	46	42	38	34	30	26	22	18	14	10	6	2								
.14	52	48	44	40	36	32	28	24	20	16	12	8	4								
.13	54	50	46	42	38	34	30	26	22	18	14	10	6	2							
.12	56	52	48	44	40	36	32	28	24	20	16	12	8	4							
.11	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2						
.10	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4						
.09	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2					
.08	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4					
.07	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2				
.06	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4				
.05	70	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2			
.04	72	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4			
.03	74	70	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2		
.02	76	72	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4		
.01	78	74	70	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2	
.00	80	76	72	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4	2

B.3 Example case: Evaluating O-DU / O-RU combinations

This example illustrates an O-RU and an O-DU, each supporting multiple SCS. It illustrates how the actual O-DU/ O-RU delay profiles map to categories, and then how the corresponding categories are used to evaluate the resulting transport network limitations. The values represented are exemplary only. The example is provided to illustrate:

- Interoperability across multiple SCS
 - O-RU and O-DU will have multiple delay profiles
 - Delay profiles are different for uplink and downlink
- applicability service provider use case

Table B.3-1 and Table B.3-2, respectively, list an example O-RU and O-DU delay profile. This results in 6 different category combinations as described in Table B.3-3.

Table B.3-1: Delay profiles for the example O-RU

	O-RU Parameters	30 kHz	120 kHz	15 kHz
Downlink	T2a_min_up,	50	50	205
	T2a_max_up,	250	180	435
	T2a_min_cp_dl,	175	175	330
	T2a_max_cp_dl,	375	205	460
	Tcp_adv_dl	125	125	125
	Category	O.20 (250-50)=200	O.12	W.22
Uplink	Ta3_min,	50	50	70
	Ta3_max	100	70	235
	T2a_min_cp_ul,	125	125	125
	T2a_max_cp_ul	325	255	360
	Category	R.05	P.02	Y.16

Table B.3-2: Delay profiles for the example O-DU.

O-DU Parameters		30 kHz	120 kHz	15 kHz
Downlink	T1a_max_up _{O-DU} ,	250	180	435
	TXmax _{O-DU}	40	20	70
	Category	K.04	M.02	C.07
Uplink	Ta4_max _{O-DU} ,	250	180	395
	RXmax _{O-DU}	200	130	325
	Category	I.20	L.12	B.32

Table B.3-3: Resulting 6 different category combinations

	30 kHz	120 kHz	15 kHz
Downlink	KO	MO	CW
Uplink	IR	LP	BY

The respective minimum and maximum T12max values are highlighted in the Table B.3-4 and Table B.3-5. Using the 30 kHz as an example, it can be seen that the range on the transport is limited by the uplink (IR) to between 130 μs and 168 μs. This means that this combination can be guaranteed to be able to support at least 130 μs of delay, and may be capable of supporting up to 168 μs of delay. Using the actual delay values for the combination at 30 kHz, the uplink is limited to (Ta4_max – Ta3_max) = (250 – 100) = 150 μs. The downlink value is higher, so uplink becomes the limiting factor for this combination.

Table B.3-4: Delay Category O-DU and O-RU with highlighted valid options for this example, minimum T12max

RU	lls-DU													
	A	B	C	D	E	F	G	H	I	J	K	L	M	N
O	350	330	310	290	270	250	230	210	190	170	150	130	110	0
P	330	310	290	270	250	230	210	190	170	150	130	110	90	0
Q	310	290	270	250	230	210	190	170	150	130	110	90	70	0
R	290	270	250	230	210	190	170	150	130	110	90	70	50	0
S	270	250	230	210	190	170	150	130	110	90	70	50	30	0
T	250	230	210	190	170	150	130	110	90	70	50	30	10	0
U	230	210	190	170	150	130	110	90	70	50	30	10	0	0
V	210	190	170	150	130	110	90	70	50	30	10	0	0	0
W	190	170	150	130	110	90	70	50	30	10	0	0	0	0
X	170	150	130	110	90	70	50	30	10	0	0	0	0	0
Y	150	130	110	90	70	50	30	10	0	0	0	0	0	0
Z	130	110	90	70	50	30	10	0	0	0	0	0	0	0

Table B.3-5: Delay Category O-DU and O-RU with highlighted valid options for this example, maximum T12max

RU	IIs-DU													
	A	B	C	D	E	F	G	H	I	J	K	L	M	N
O	3000	399	379	359	339	319	299	279	259	239	219	199	179	159
P	2949	348	328	308	288	268	248	228	208	188	168	148	128	108
Q	2929	328	308	288	268	248	228	208	188	168	148	128	108	88
R	2909	308	288	268	248	228	208	188	168	148	128	108	88	68
S	2889	288	268	248	228	208	188	168	148	128	108	88	68	48
T	2869	268	248	228	208	188	168	148	128	108	88	68	48	28
U	2849	248	228	208	188	168	148	128	108	88	68	48	28	8
V	2829	228	208	188	168	148	128	108	88	68	48	28	8	0
W	2809	208	188	168	148	128	108	88	68	48	28	8	0	0
X	2789	188	168	148	128	108	88	68	48	28	8	0	0	0
Y	2769	168	148	128	108	88	68	48	28	8	0	0	0	0
Z	2749	148	128	108	88	68	48	28	8	0	0	0	0	0

Evaluating the overall combination across all SCS, it can be seen that the maximum possible range for this combination if all 3 SCS are to be used is 148 μs (limited by the 120 kHz uplink) and the minimum guaranteed range is 110 μs. (In actuality, the limit is 110 μs for this combination.)

Similarly, there are 6 combinations of sub-categories to evaluate as shown in Table B.3-6. In this case, the sub-category combinations happen to be the same for 120 kHz uplink and downlink. In this case the different between T12max and T12min is limited by the 120 kHz delay profile, with a maximum range of ~20 km (~100 μs).

Table B.3-6: Delay Sub-Category O-DU and O-RU with highlighted (in red) valid options for this example

TxMax	Receive Window																				
	.40	.38	.36	.34	.32	.30	.28	.26	.24	.22	.20	.18	.16	.14	.12	.10	.08	.06	.04	.02	.01
.20	40	36	32	28	24	20	16	12	8	4											
.19	42	38	34	30	26	22	18	14	10	6	2										
.18	44	40	36	32	28	24	20	16	12	8	4										
.17	46	42	38	34	30	26	22	18	14	10	6	2									
.16	48	44	40	36	32	28	24	20	16	12	8	4									
.15	50	46	42	38	34	30	26	22	18	14	10	6	2								
.14	52	48	44	40	36	32	28	24	20	16	12	8	4								
.13	54	50	46	42	38	34	30	26	22	18	14	10	6	2							
.12	56	52	48	44	40	36	32	28	24	20	16	12	8	4							
.11	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2						
.10	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4						
.09	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2					
.08	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4					
.07	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2				
.06	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4				
.05	70	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2			
.04	72	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4			
.03	74	70	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2		
.02	76	72	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4		
.01	78	74	70	66	62	58	54	50	46	42	38	34	30	26	22	18	14	10	6	2	
.00	80	76	72	68	64	60	56	52	48	44	40	36	32	28	24	20	16	12	8	4	2

Using only the O-DU/ O-RU category/ sub-category combinations, if the corresponding O-DU/ O-RU is to be used across all 3 SCS, the resulting delay constraints limit the implementation to:

- $110 \mu s \leq T12max \leq 148 \mu s$
- $T12min = (T12max - 100) \mu s$

The service provider can then use this analysis to determine if the resulting combination suits their target use case.

B.4 Example case: Non-ideal transport O-DU/O-RU categories

This example illustrates the usage of the non-ideal transport delay categories and sub-categories included in Table B.2-1 and Table B.2-5 respectively. These categories and sub-categories were created to allow for the use of transport networks which have delay introduced at the MAC layer that is larger than the propagation delay e.g. DOCSIS networks, E/GPON networks, microwave links, or G.Fast networks. In such systems, delay is decoupled from the transport network length and therefore better specified in microseconds or milliseconds.

Non-ideal transport is defined by 3GPP in document TR 36.932 [i.2] per Table B.4-1.

Table B.4-1: Non-Ideal transport categorization from 3GPP TR 36.932 [i.2] Table 6.1-1

Backhaul Technology	Latency (One way)	Throughput	Priority (1 is the highest)
Fiber Access 1	10-30 ms	10M-10 Gbps	1
Fiber Access 2	5-10 ms	100-1000 Mbps	2
Fiber Access 3	2-5 ms	50M-10 Gbps	1
DSL Access	15-60 ms	10-100 Mbps	1
Cable	25-35 ms	10-100 Mbps	2
Wireless Backhaul	5-35 ms	10 Mbps – 100 Mbps typical, maybe up to Gbps range	1

In Table B.4-2, example deployment use cases are used to illustrate O-DU/O-RU delay profile mapping. The values, while meant to be representative, are exemplary only.

Table B.4-2: Non-Ideal fronthaul O-DU/O-RU Delay Profiles by link type

		DOCSIS		Microwave		PON		G.Fast	
		T12 (DL)	T34 (UL)	T12 (DL)	T34 (UL)	T12 (DL)	T34 (UL)	T12 (DL)	T34 (UL)
	Latency (RTT)	< 1 ms	< 29 ms	< 10 ms	<10 ms	< 500 μ s	< 2.5 ms	< 2 ms	< 2 ms
	Jitter/PDV	< 320 μ s	< 9 ms	< 2.5 ms	<2.5 ms	< 1 ms	<1 ms	< 1 ms	< 1 ms
O-DU	T1a_max_up_O-DU	1,500 μ s	-	12,500 μ s	-	750 μ s	-	2,250 μ s	-
	TXmax_O-DU	50 μ s	-	100 μ s	-	50 μ s	-	50 μ s	-
	Ta4_max_O-DU	-	29,500 μ s	-	12,500 μ s	-	2,750 μ s	-	2,250 μ s
	Rx_max_O-DU	-	9,500 μ s	-	3,250 μ s	-	1,250 μ s	-	1,250 μ s
	Category	A.05	AAA.300	AAA.10	AAA.300	A.05	A.100	A.05	A.100
O-RU	T2a_min_up	100 μ s	-	50 μ s	-	50 μ s	-	70 μ s	-
	T2a_max_up	500 μ s	-	2,650 μ s	-	1,100 μ s	-	1,200 μ s	-
	Ta3_min	-	50 μ s	-	50 μ s	-	100 μ s	-	70 μ s
	Ta3_max	-	150 μ s	-	200 μ s	-	200 μ s	-	200 μ s
	Category	R.40	T.10	O.100	W.15	O.40	W.10	P.100	W.13

The use of non-ideal transport links for fronthaul traffic comes with system performance tradeoffs. See Appendix L for further discussion of these considerations.

Annex C (informative):
Void

Annex D (informative): IQ sample and exponent packetization for different bitwidths

D.1 IQ sample ordering for big-endian byte order

As specified in clause 8.3.2, bit-ordering and packetization for I and Q samples and compression parameters follow patterns that repeat after every transmission of 12 resource elements. Tables D.1-1 through D.1-11 illustrate this bit-ordering and repetition for big-endian byte order. The cells in the tables indicate the bit ordering for the following IEs:

- udCompParam (assumed to be one byte in the tables)
- I samples denoted by $I_{\text{bitwidth}-1} \dots I_0$
- Q samples denoted by $Q_{\text{bitwidth}-1} \dots Q_0$

Table D.1-1: IQ data samples bit-ordering (6-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I_5	I_4	I_3	I_2	I_1	I_0	Q_5	Q_4	1	N+1
Q_3	Q_2	Q_1	Q_0	...				1	N+2
...									
			...	I_5	I_4	I_3	I_2	1	N+17
I_1	I_0	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	N+18
udCompParam								1	Octet M
I_5	I_4	I_3	I_2	I_1	I_0	Q_5	Q_4	1	M+1
Q_3	Q_2	Q_1	Q_0	...				1	M+2
...									
			...	I_5	I_4	I_3	I_2	1	M+17
I_1	I_0	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	M+18
...									

Table D.1-2: IQ data samples bit-ordering (7-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I_6	I_5	I_4	I_3	I_2	I_1	I_0	Q_6	1	N+1
Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	...		1	N+2
...									
	...	I_6	I_5	I_4	I_3	I_2	I_1	1	N+20
I_0	Q_6	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	N+21
udCompParam								1	Octet M
I_6	I_5	I_4	I_3	I_2	I_1	I_0	Q_6	1	M+1
Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	...		1	M+2
...									
	...	I_6	I_5	I_4	I_3	I_2	I_1	1	M+20
I_0	Q_6	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	M+21
...									

Table D.1-3: IQ data samples bit-ordering (8-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+1
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+2
...									
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+23
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+24
udCompParam								1	Octet M
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+1
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+2
...									
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+23
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+24
...									

Table D.1-4: IQ data samples bit-ordering (9-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	1	N+1
I ₀	Q ₈	Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	N+2
Q ₁	Q ₀	...						1	N+3
...									
					...	I ₈	I ₇	1	N+25
I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	Q ₈	1	N+26
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+27
udCompParam								1	Octet M
I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	1	M+1
I ₀	Q ₈	Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	M+2
Q ₁	Q ₀	...						1	
...									
					...	I ₈	I ₇	1	M+25
I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	Q ₈	1	M+26
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+27
...									

Table D.1-5: IQ data samples bit-ordering (10-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	1	N+1
I ₁	I ₀	Q ₉	Q ₈	Q ₇	Q ₆	Q ₅	Q ₄	1	N+2
Q ₃	Q ₂	Q ₁	Q ₀	...				1	N+3
...									
			...	I ₉	I ₈	I ₇	I ₆	1	N+28
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	Q ₉	Q ₈	1	N+29
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+30
udCompParam								1	Octet M
I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	1	M+1
I ₁	I ₀	Q ₉	Q ₈	Q ₇	Q ₆	Q ₅	Q ₄	1	M+2
Q ₃	Q ₂	Q ₁	Q ₀	...				1	M+3
...									
			...	I ₉	I ₈	I ₇	I ₆	1	M+28
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	Q ₉	Q ₈	1	M+29
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+30
...									

Table D.1-6: IQ data samples bit-ordering (11-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	N+1
I ₂	I ₁	I ₀	Q ₁₀	Q ₉	Q ₈	Q ₇	Q ₆	1	N+2
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	...		1	N+3
...									
	...	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	1	N+31
I ₄	I ₃	I ₂	I ₁	I ₀	Q ₁₀	Q ₉	Q ₈	1	N+32
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+33
udCompParam								1	Octet M
I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	M+1
I ₂	I ₁	I ₀	Q ₁₀	Q ₉	Q ₈	Q ₇	Q ₆	1	M+2
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	...		1	M+3
...									
	...	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	1	M+31
I ₄	I ₃	I ₂	I ₁	I ₀	Q ₁₀	Q ₉	Q ₈	1	M+32
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+33
...									

Table D.1-7: IQ data samples bit-ordering (12-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₁₁	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	1	N+1
I ₃	I ₂	I ₁	I ₀	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+2
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+3
...									
I ₁₁	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	1	N+34
I ₃	I ₂	I ₁	I ₀	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+35
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+36
udCompParam								1	Octet M
I ₁₁	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	1	M+1
I ₃	I ₂	I ₁	I ₀	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+2
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+3
...									
I ₁₁	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	1	M+34
I ₃	I ₂	I ₁	I ₀	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+35
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+36
...									

Table D.1-8: IQ data samples bit-ordering (13-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	1	N+1
I ₄	I ₃	I ₂	I ₁	I ₀	Q ₁₂	Q ₁₁	Q ₁₀	1	N+2
Q ₉	Q ₈	Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	N+3
Q ₁	Q ₀	...						1	N+4
...									
					...	I ₁₂	I ₁₁	1	N+36
I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	N+37
I ₂	I ₁	I ₀	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+38
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+39
udCompParam								1	Octet M
I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	1	M+1
I ₄	I ₃	I ₂	I ₁	I ₀	Q ₁₂	Q ₁₁	Q ₁₀	1	M+2
Q ₉	Q ₈	Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	M+3
Q ₁	Q ₀	...						1	M+4
...									
					...	I ₁₂	I ₁₁	1	M+36
I ₁₀	I ₉	I ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	M+37
I ₂	I ₁	I ₀	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+38
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+39
...									

Table D.1-9: IQ data samples bit-ordering (14-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
l_{13}	l_{12}	l_{11}	l_{10}	l_9	l_8	l_7	l_6	1	N+1
l_5	l_4	l_3	l_2	l_1	l_0	Q_{13}	Q_{12}	1	N+2
Q_{11}	Q_{10}	Q_9	Q_8	Q_7	Q_6	Q_5	Q_4	1	N+3
Q_3	Q_2	Q_1	Q_0	...				1	N+4
...									
...				l_{13}	l_{12}	l_{11}	l_{10}	1	N+39
l_9	l_8	l_7	l_6	l_5	l_4	l_3	l_2	1	N+40
l_1	l_0	Q_{13}	Q_{12}	Q_{11}	Q_{10}	Q_9	Q_8	1	N+41
Q_7	Q_6	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	N+42
udCompParam								1	Octet M
l_{13}	l_{12}	l_{11}	l_{10}	l_9	l_8	l_7	l_6	1	M+1
l_5	l_4	l_3	l_2	l_1	l_0	Q_{13}	Q_{12}	1	M+2
Q_{11}	Q_{10}	Q_9	Q_8	Q_7	Q_6	Q_5	Q_4	1	M+3
Q_3	Q_2	Q_1	Q_0	...				1	M+4
...									
...				l_{13}	l_{12}	l_{11}	l_{10}	1	M+39
l_9	l_8	l_7	l_6	l_5	l_4	l_3	l_2	1	M+40
l_1	l_0	Q_{13}	Q_{12}	Q_{11}	Q_{10}	Q_9	Q_8	1	M+41
Q_7	Q_6	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	M+42
...									

Table D.1-10: IQ data samples bit-ordering (15-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
l_{14}	l_{13}	l_{12}	l_{11}	l_{10}	l_9	l_8	l_7	1	N+1
l_6	l_5	l_4	l_3	l_2	l_1	l_0	Q_{14}	1	N+2
Q_{13}	Q_{12}	Q_{11}	Q_{10}	Q_9	Q_8	Q_7	Q_6	1	N+3
Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	...		1	N+4
...									
...		l_{14}	l_{13}	l_{12}	l_{11}	l_{10}	l_9	1	N+42
l_8	l_7	l_6	l_5	l_4	l_3	l_2	l_1	1	N+43
l_0	Q_{14}	Q_{13}	Q_{12}	Q_{11}	Q_{10}	Q_9	Q_8	1	N+44
Q_7	Q_6	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	N+45
udCompParam								1	Octet M
l_{14}	l_{13}	l_{12}	l_{11}	l_{10}	l_9	l_8	l_7	1	M+1
l_6	l_5	l_4	l_3	l_2	l_1	l_0	Q_{14}	1	M+2
Q_{13}	Q_{12}	Q_{11}	Q_{10}	Q_9	Q_8	Q_7	Q_6	1	M+3
Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	...		1	M+4
...									
...		l_{14}	l_{13}	l_{12}	l_{11}	l_{10}	l_9	1	M+42
l_8	l_7	l_6	l_5	l_4	l_3	l_2	l_1	1	M+43
l_0	Q_{14}	Q_{13}	Q_{12}	Q_{11}	Q_{10}	Q_9	Q_8	1	M+44
Q_7	Q_6	Q_5	Q_4	Q_3	Q_2	Q_1	Q_0	1	M+45
...									

Table D.1-11: IQ data samples bit-ordering (16-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam								1	Octet N
I ₁₅	I ₁₄	I ₁₃	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	N+1
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+2
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+3
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+4
...									
I ₁₅	I ₁₄	I ₁₃	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	N+45
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+46
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+47
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+48
udCompParam								1	Octet M
I ₁₅	I ₁₄	I ₁₃	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	M+1
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+2
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+3
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+4
...									
I ₁₅	I ₁₄	I ₁₃	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	M+45
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+46
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+47
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+48
...									

D.2 IQ sample ordering for little-endian byte order

When the optional "little endian byte order" format is used, the complex IQ data order is changed relative to that shown in Annex D.1. The Q data will be sent before I data, while within IQ data, the low byte is sent before high byte. All other conventions are observed same as in (Annex D.1) big-endian format. The usage of "little endian byte ordering" can be negotiated between the O-DU and O-RU via the M-Plane. Tables D.2-1 through **D.2-11** illustrate different IQ bitwidths when data is sent in little-endian byte order. The cells in the tables indicate the bit ordering for the following IEs:

- udCompParam (assumed to be one byte)
- I samples denoted by $I_{\text{bitwidth}-1} \dots I_0$
- Q samples denoted by $Q_{\text{bitwidth}-1} \dots Q_0$

Table D.2-1 through Table D.2-11 also apply to different IQ bitwidths for beamforming weights (Section Type 1 or 3, with Section Extension 1, 11 or 19) and channel information value (Section Type 6) when the optional "little endian byte order" is chosen via M-Plane. For beamforming weights, the udCompParam in the tables should be treated as bfwCompParam. $I_{\text{bitwidth}-1} \dots I_0$ denotes bfw $I_{\text{bitwidth}-1} \dots \text{bfw}I_0$, and $Q_{\text{bitwidth}-1} \dots Q_0$ denotes bfw $Q_{\text{bitwidth}-1} \dots \text{bfw}Q_0$. For channel information value, the udCompParam in the tables should be treated as ciCompParam. $I_{\text{bitwidth}-1} \dots I_0$ denotes ci $I_{\text{samplebitwidth}-1} \dots \text{ci}I_{\text{sample}0}$, and $Q_{\text{bitwidth}-1} \dots Q_0$ denotes ci $Q_{\text{samplebitwidth}-1} \dots \text{ci}Q_{\text{sample}0}$.

Table D.2-1: Bit order of IQ data samples in little endian mode (6-bit IQ bit width example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	I ₅	I ₄	1	N+1
I ₃	I ₂	I ₁	I ₀	...				1	N+2
...									
			...	Q ₅	Q ₄	Q ₃	Q ₂	1	N+17
Q ₁	Q ₀	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+18
udCompParam									Octet M
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	I ₅	I ₄	1	M+1
I ₃	I ₂	I ₁	I ₀	...				1	M+2
...									
			...	Q ₅	Q ₄	Q ₃	Q ₂	1	M+17
Q ₁	Q ₀	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+18
...									

Table D.2-2: Bit order of IQ data samples in little endian mode (7-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	I ₆	1	N+1
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	...		1	N+2
...									
	...	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	1	N+20
Q ₀	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+21
udCompParam									Octet M
Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	I ₆	1	M+1
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	...		1	M+2
...									
	...	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	1	M+20
Q ₀	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+21
...									

Table D.2-3: Bit order of IQ data samples in little endian mode (8-bit IQ bit width example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+2
...									
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+23
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	N+24
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+2
...									
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+23
I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	1	M+24
...									

Table D.2-4: Bit order of IQ data samples in little endian mode (9-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	1	N+2
I ₀	I ₈	...						1	N+3
...					...	Q ₇	Q ₆	1	N+25
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	Q ₈	I ₇	1	N+26
I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	I ₈	1	N+27
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	1	M+2
I ₀	I ₈	...						1	
...					...	Q ₇	Q ₆	1	M+25
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	Q ₈	I ₇	1	M+26
I ₆	I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	I ₈	1	M+27
...									

Table D.2-5: Bit order of IQ data samples in little endian mode (10-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	1	N+2
I ₁	I ₀	I ₉	I ₈	...				1	N+3
...				...	Q ₇	Q ₆	Q ₅	1	N+28
Q ₃	Q ₂	Q ₁	Q ₀	Q ₉	Q ₈	I ₇	I ₆	1	N+29
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	I ₉	I ₈	1	N+30
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	1	M+2
I ₁	I ₀	I ₉	I ₈	...				1	M+3
...				...	Q ₇	Q ₆	Q ₅	1	M+28
Q ₃	Q ₂	Q ₁	Q ₀	Q ₉	Q ₈	I ₇	I ₆	1	M+29
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	I ₉	I ₈	1	M+30
...									

Table D.2-6: Bit order of IQ data samples in little endian mode (11-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	N+2
I ₂	I ₁	I ₀	I ₁₀	I ₉	I ₈	...		1	N+3
...									
	...	Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	N+31
Q ₁	Q ₀	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	1	N+32
I ₄	I ₃	I ₂	I ₁	I ₀	I ₁₀	I ₉	I ₈	1	N+33
udCompParam								1	Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	M+2
I ₂	I ₁	I ₀	I ₁₀	I ₉	I ₈	...		1	M+3
...									
	...	Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	M+31
Q ₁	Q ₀	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	1	M+32
I ₄	I ₃	I ₂	I ₁	I ₀	I ₁₀	I ₉	I ₈	1	M+33
...									

Table D.2-7: Bit order of IQ data samples in little endian mode (12-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	1	N+2
I ₃	I ₂	I ₁	I ₀	I ₁₁	I ₁₀	I ₉	I ₈	1	N+3
...									
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+34
Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	1	N+35
I ₃	I ₂	I ₁	I ₀	I ₁₁	I ₁₀	I ₉	I ₈	1	N+36
udCompParam								1	Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	1	M+2
I ₃	I ₂	I ₁	I ₀	I ₁₁	I ₁₀	I ₉	I ₈	1	M+3
...									
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+34
Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	1	M+35
I ₃	I ₂	I ₁	I ₀	I ₁₁	I ₁₀	I ₉	I ₈	1	M+36
...									

Table D.2-8: Bit order of IQ data samples in little endian mode (13-bit bitwidth mantissa example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	1	N+2
I ₄	I ₃	I ₂	I ₁	I ₀	I ₁₂	I ₁₁	I ₁₀	1	N+3
I ₉	I ₈	...						1	N+4
...					...	Q ₇	Q ₆	1	N+36
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	Q ₁₂	Q ₁₁	1	N+37
Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	N+38
I ₂	I ₁	I ₀	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	N+39
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	1	M+2
I ₄	I ₃	I ₂	I ₁	I ₀	I ₁₂	I ₁₁	I ₁₀	1	M+3
I ₉	I ₈	...						1	M+4
...					...	Q ₇	Q ₆	1	M+36
Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	Q ₁₂	Q ₁₁	1	M+37
Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	1	M+38
I ₂	I ₁	I ₀	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	M+39
...									

Table D.2-9: Bit order of IQ data samples in little endian mode (14-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	1	N+2
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	I ₁₃	I ₁₂	1	N+3
I ₁₁	I ₁₀	I ₉	I ₈	...				1	N+4
...				Q ₇	Q ₆	Q ₅	Q ₄	1	N+39
Q ₃	Q ₂	Q ₁	Q ₀	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	1	N+40
Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	1	N+41
I ₁	I ₀	I ₁₃	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	N+42
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	I ₇	I ₆	1	M+2
I ₅	I ₄	I ₃	I ₂	I ₁	I ₀	I ₁₃	I ₁₂	1	M+3
I ₁₁	I ₁₀	I ₉	I ₈	...				1	M+4
...				Q ₇	Q ₆	Q ₅	Q ₄	1	M+39
Q ₃	Q ₂	Q ₁	Q ₀	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	1	M+40
Q ₉	Q ₈	I ₇	I ₆	I ₅	I ₄	I ₃	I ₂	1	M+41
I ₁	I ₀	I ₁₃	I ₁₂	I ₁₁	I ₁₀	I ₉	I ₈	1	M+42
...									

Table D.2-10: Bit order of IQ data samples in little endian mode (15-bit IQ bitwidth example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	l ₇	1	N+2
l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	l ₀	l ₁₄	1	N+3
l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	...		1	N+4
...									
...		Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	N+42
Q ₁	Q ₀	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	1	N+43
Q ₈	l ₇	l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	1	N+44
l ₀	l ₁₄	l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	1	N+45
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	l ₇	1	M+2
l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	l ₀	l ₁₄	1	M+3
l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	...		1	M+4
...									
...		Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	1	M+42
Q ₁	Q ₀	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	1	M+43
Q ₈	l ₇	l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	1	M+44
l ₀	l ₁₄	l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	1	M+45
...									

Table D.2-11: Bit order of IQ data samples in little endian mode (16-bit IQ bit width example)

0 (msb)	1	2	3	4	5	6	7 (lsb)	Number of Octets	
...									
udCompParam									Octet N
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+1
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+2
l ₇	l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	l ₀	1	N+3
l ₁₅	l ₁₄	l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	1	N+4
...									
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	N+45
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	N+46
l ₇	l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	l ₀	1	N+47
l ₁₅	l ₁₄	l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	1	N+48
udCompParam									Octet M
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+1
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+2
l ₇	l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	l ₀	1	M+3
l ₁₅	l ₁₄	l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	1	M+4
...									
Q ₇	Q ₆	Q ₅	Q ₄	Q ₃	Q ₂	Q ₁	Q ₀	1	M+45
Q ₁₅	Q ₁₄	Q ₁₃	Q ₁₂	Q ₁₁	Q ₁₀	Q ₉	Q ₈	1	M+46
l ₇	l ₆	l ₅	l ₄	l ₃	l ₂	l ₁	l ₀	1	M+47
l ₁₅	l ₁₄	l ₁₃	l ₁₂	l ₁₁	l ₁₀	l ₉	l ₈	1	M+48
...									

Annex E (informative): OFDM phase compensation

Consider the time interval $t_{start,l}^\mu \leq t < t_{start,l}^\mu + (N_u^\mu + N_{CP,l}^\mu)T_c$ as defined in 3GPP TS 38.211 [4], clause 5.3.1, which corresponds to the transmission of the l th OFDM symbol with baseband waveform $s_l^{(p,\mu)}(t)$, and subcarrier spacing configuration μ .

Suppose that the transmitter performs up-conversion of the signal $s_l^{(p,\mu)}(t)$ to RF centre frequency f_0 , and the receiver performs down-conversion from RF centre frequency f_1 . Unlike LTE, in NR it is possible that the transmitter and receiver have different centre frequencies, $f_0 \neq f_1$, as for example can happen when a bandwidth part is configured, which motivates the introduction of OFDM phase compensation in NR.

Ignoring channel distortions and noise for simplicity, the baseband signal at the receiver in the same time interval can be written

$$s_l^{(p,\mu)}(t) \cdot e^{j2\pi(f_0-f_1)t} = s_l^{(p,\mu)}(t) \cdot e^{j\theta_l} \cdot e^{j2\pi(f_0-f_1)(t-t_{CP,l}^\mu)}$$

where $\theta_l = 2\pi(f_0 - f_1)t_{CP,l}^\mu$ and $t_{CP,l}^\mu = t_{start,l}^\mu + N_{CP,l}^\mu T_c$ is the time at the end of the cyclic prefix of OFDM symbol l . It can be seen that each OFDM symbol is subjected to a phase shift θ_l , which can result in large phase changes between successive OFDM symbols. This may cause problems for receiver processes such as channel estimation and frequency error tracking, which expect the channel to vary smoothly between symbols.

The solution adopted by 3GPP is to apply a phase pre-compensation term $e^{-j2\pi f_0 t_{CP,l}^\mu}$ at the transmitter and a phase post-compensation term $e^{j2\pi f_1 t_{CP,l}^\mu}$ at the receiver. These two terms together provide the required correction $e^{-j\theta_l}$. The phase pre-compensation requirement for the transmitter is specified in 3GPP TS 38.211 [4], clause 5.4. The phase compensation depends only on the starting time of each OFDM symbol, and is common for all subcarriers in that symbol transmitted using a given numerology μ , regardless of which NR physical channels they belong to, except for PRACH.

Annex F (normative): Beamforming attributes frame of reference

The beamforming attributes involving pointing angle (bfAzPt and bfZePt) specified in clause 7.7.2 are defined following 3GPP TR 38.901 [43], clause 7.1. Two coordinate systems are defined, the Global Coordinate System (GCS) and Local Coordinate System (LCS). The GCS applies across multiple BS and UT locations, while the LCS applies to a single array antenna. The GCS is defined as shown below in Figure F-1. In the GCS, ϕ refers to the azimuth angle and θ refers to the zenith angle.

The LCS is defined by a 3-parameter rotation of the GCS. The rotation with parameters α, β, γ is shown in Figure F-2 (left). The rotation parameters are defined as follows:

- α is defined as the bearing angle (sector pointing angle)
- β is defined as the downtilt angle
- γ is defined as the slant angle

A two-dimensional array antenna is defined such that the x' -axis is broadside to the array antenna. If an antenna architecture such as in 3GPP TR 38.901 [43] clause 7.3 is used, the horizontal direction is defined as y' and the vertical direction is defined as z' .

All parameters which relate to "peak" refer to the principal beam pointing in a particular direction. For a given beam configuration, the peak direction is the angle corresponding to maximum gain. Thus, bfAzPt specifies the pointing angle of maximum gain.

The pointing parameter bfAzPt is defined as ϕ' . The pointing parameter bfZePt is defined as θ' . The angles are shown in Figure F-2 (right).

The beamwidth parameters bfAz3dd and bfZe3dd are defined as the angular widths at which the beam falls to 3 dB below the peak beam gain on both sides of the peak direction, in azimuth and zenith, respectively.

The sidelobe parameters bfAzSl and bfZeSl are defined as the suppression level of the highest sidelobe relative to the peak gain in the azimuth and zenith principal planes, respectively.




Figure F-1: Global Coordinate System definition




Figure F-2: Orienting the LCS (blue) with respect to the GCS (gray) by a sequence of 3 rotations (left); Definition of spherical coordinates and unit vectors in both the GCS and LCS (right)

Annex G (informative): LAA Algorithms and Examples

An example PDSCH transmission algorithm showing message flow at the O-DU following clause 7.2.5.4.1 is shown in Figure G-1. The symbols used in Figure G-1 are defined in clause 7.2.5.3. The notation $\lceil a \rceil$ is the ceiling of number a to the nearest x , where x is 1 ms for normal SF and 0.5 ms for partially filled SF.

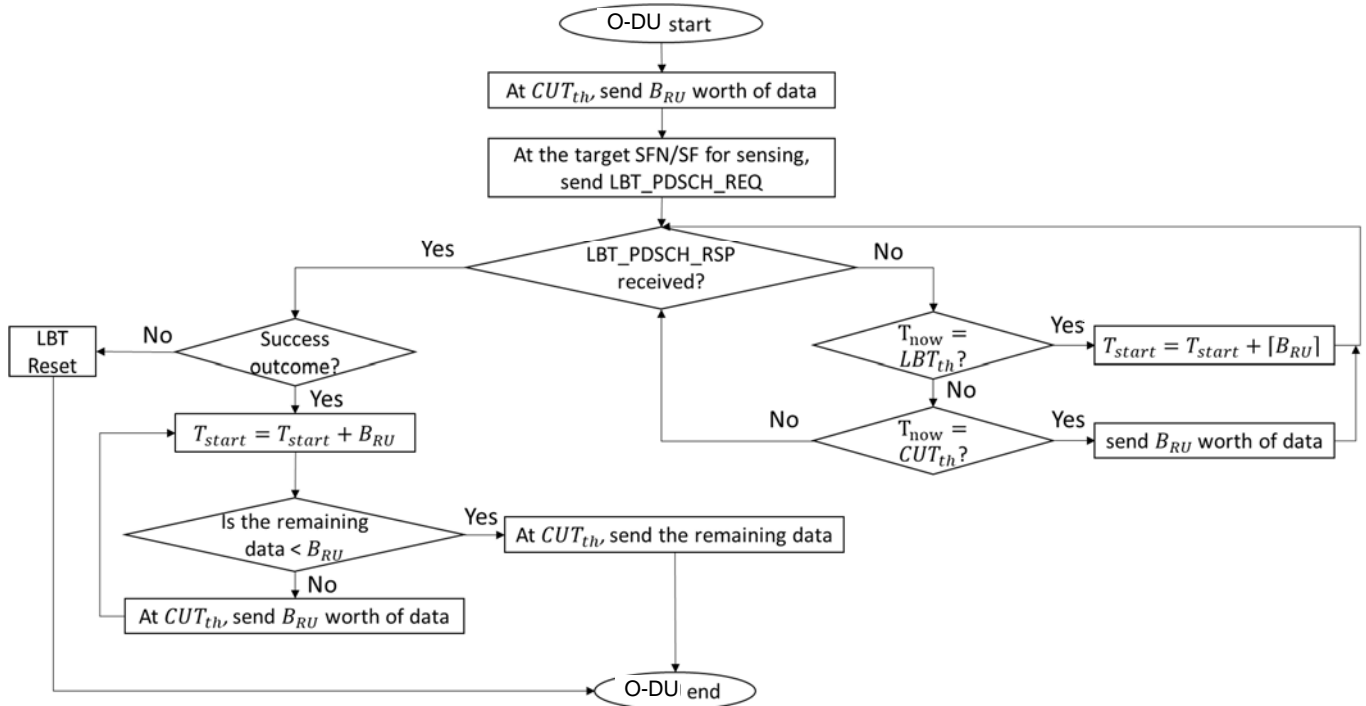


Figure G-1: PDSCH transmission algorithm O-DU flowchart

Similarly, an example PDSCH transmission algorithm showing message flow at the O-RU is shown in Figure G-2.

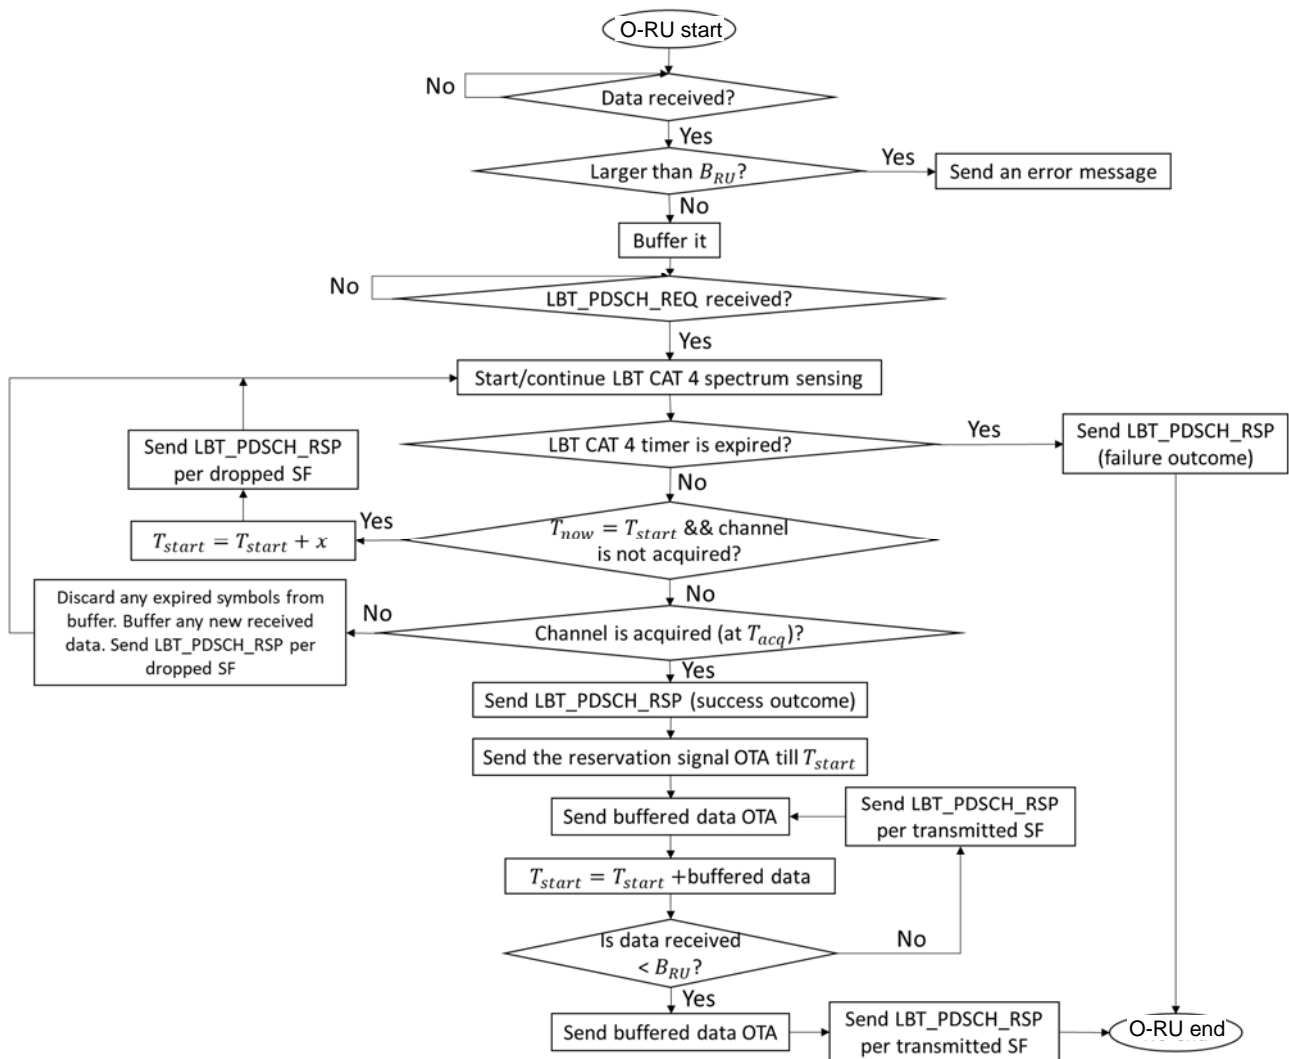


Figure G-2: PDSCH transmission algorithm O-RU flowchart

For DRS Transmission Procedure, Figure G-3 shows an example O-DU flowchart. The corresponding O-RU example flow chart when using $25 \mu\text{s}$ sensing is depicted in Figure G-4. An example flow chart for DRS transmission procedure at O-RU for LBT CAT 4 is shown in Figure G-5.

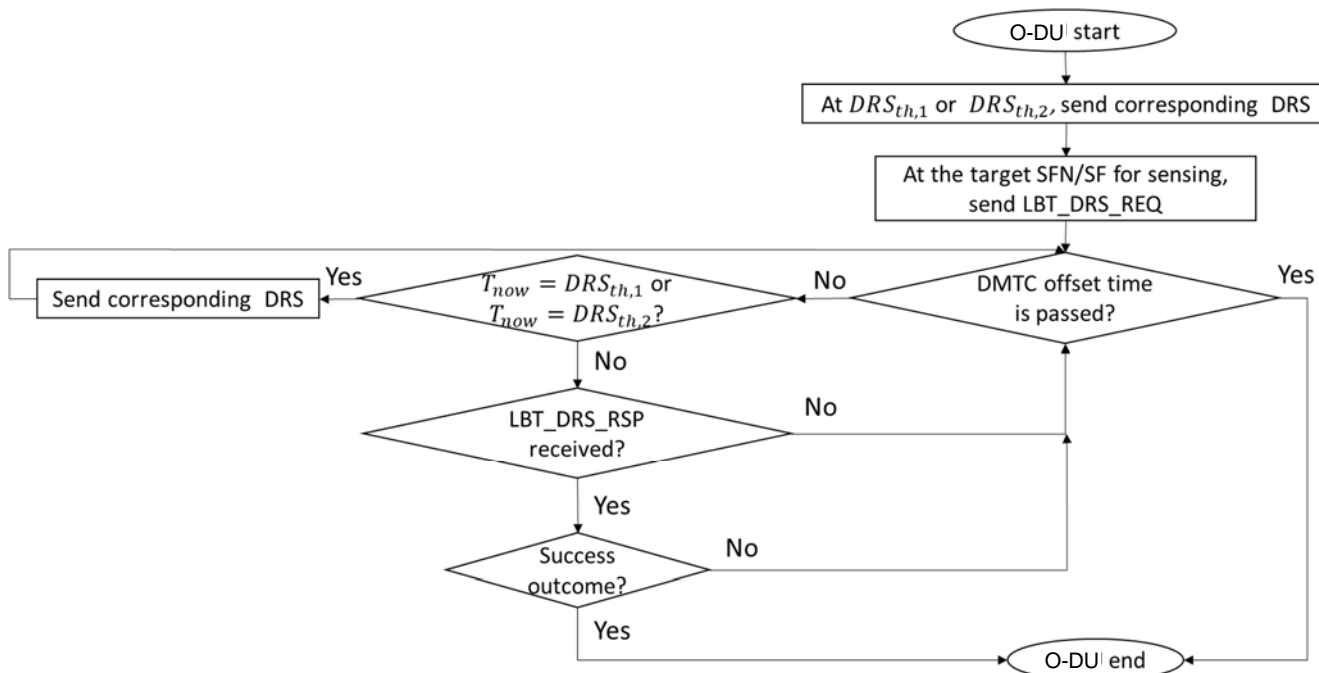


Figure G-3: DRS transmission procedure O-DU flowchart

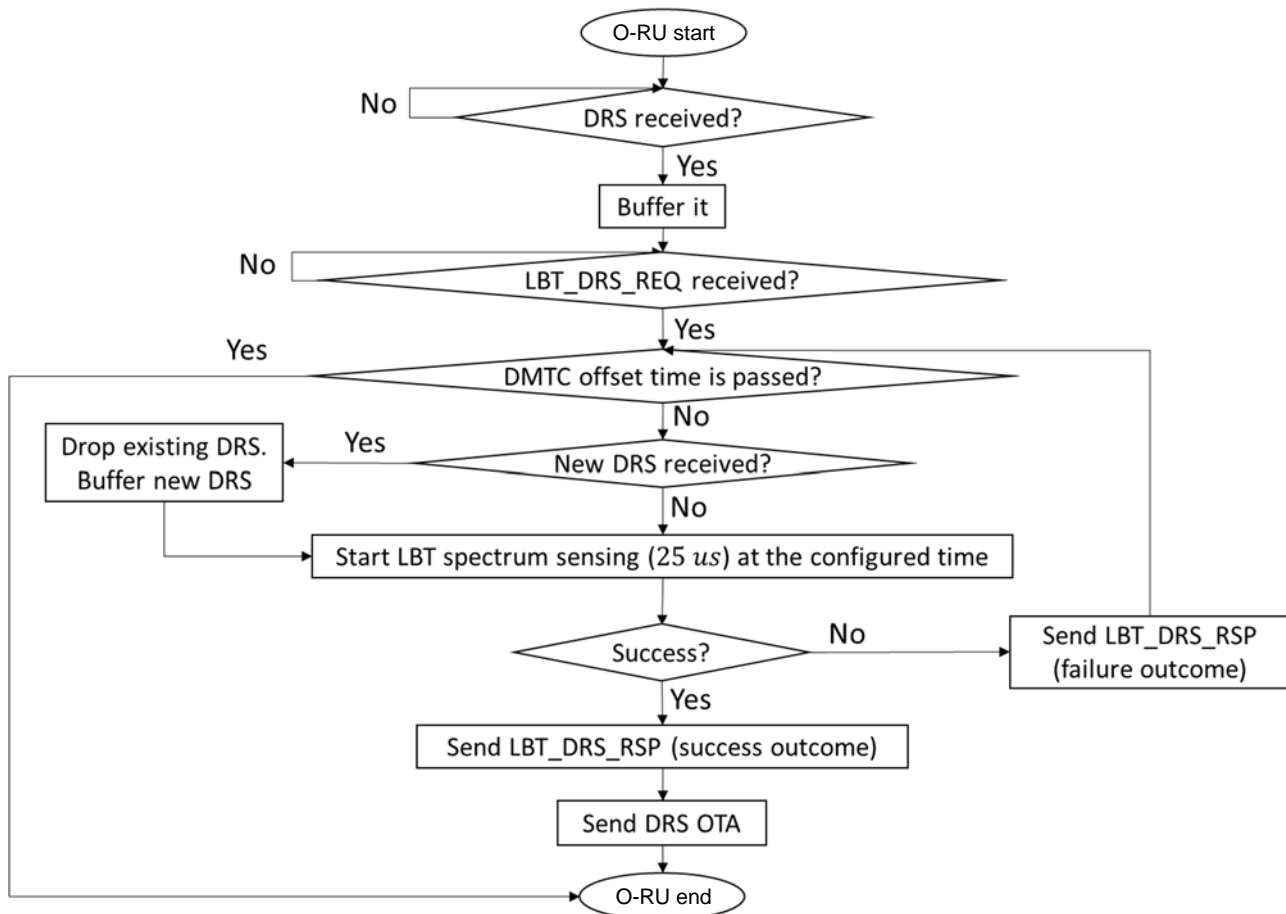


Figure G-4: DRS transmission procedure O-RU flowchart – 25 μs sensing

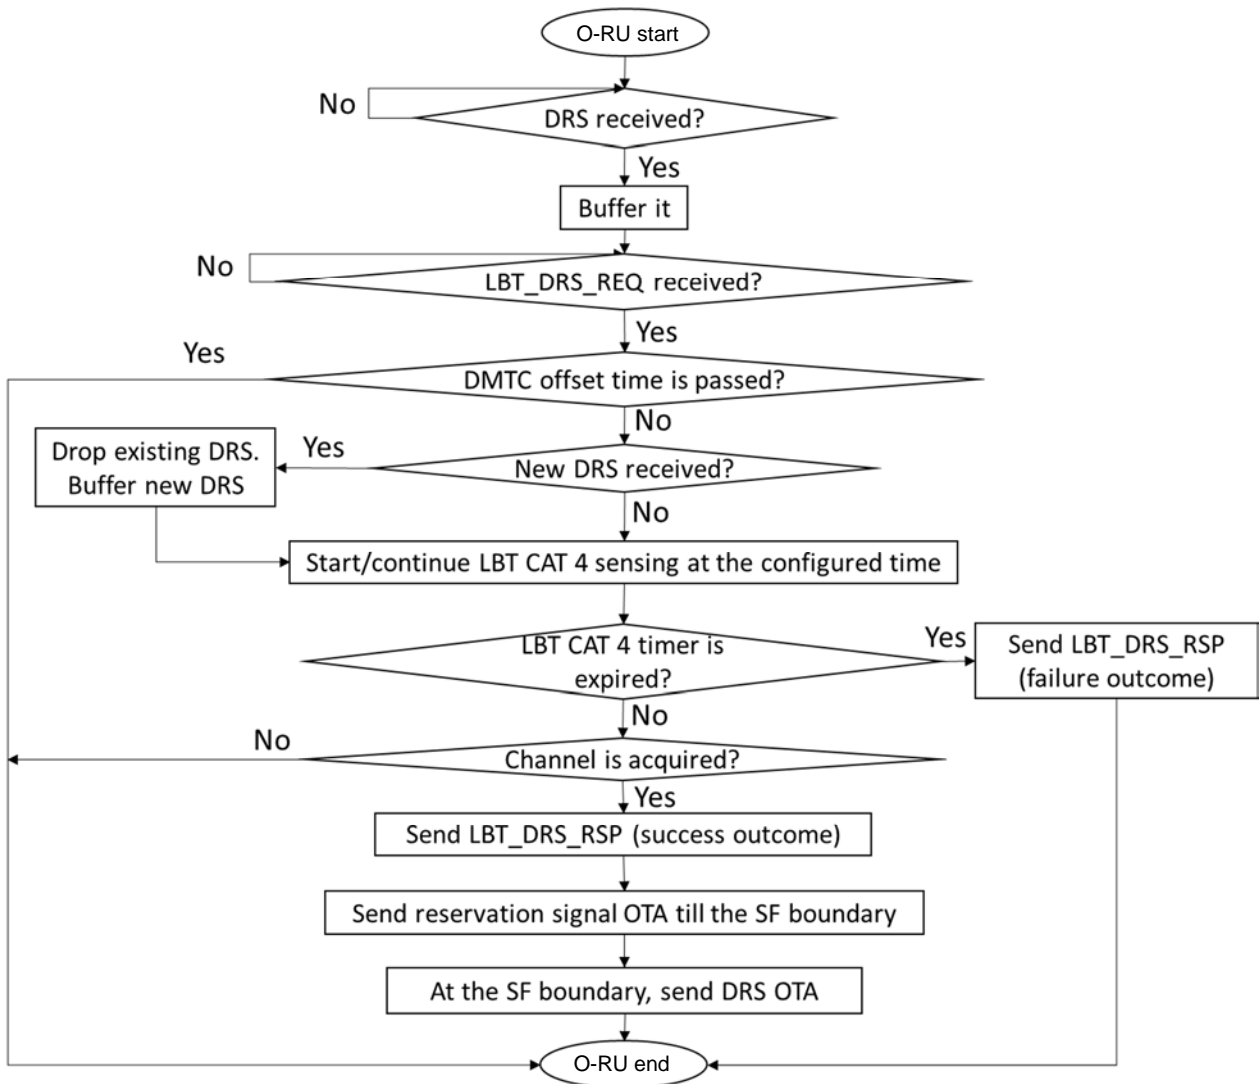


Figure G-5: DRS transmission procedure O-RU flowchart – LBT Cat 4

Annex H (informative): S-Plane detailed frequency and phase error budget analysis, and future ITU-T clock types and classes reference

H.1 Void

H.2 Frequency and time error budget analysis

H.2.1 Overview

This Annex provides the informative analysis to support budget allocation in clause 11.3.2 for a Full Timing Support network (as per Recommendation ITU-T G.8271.1 [22] for the limits, Recommendation ITU-T G.8273.2 [28] for the clock definition). The analysis serves 2 purposes:

Considering T-BC Class B and C switches (as per Recommendation ITU-T G.8273.2 [28]) in a Recommendation ITU-T G.8271.1 [22] compliant network, the number of allowed switches to satisfy the allocated network limit is computed in detail as an example.

NOTE: the following configurations are outside the scope of this annex:

- Recommendation ITU-T G.8271.1 [22] compliant networks using class D T-BCs.
- Recommendation ITU-T G.8271.1 [22] compliant networks using T-TCs instead of T-BCs.
- Non-ITU-T G.8271.1 [22]-compliant networks, such as Recommendation ITU-T G.8271.2 [24] ones.
- O-RU using SyncE/eSyncE. In the example below, only PTP is considered as synchronization source for O-RU.

Each network element in the fronthaul clock chain generates time error (including constant cTE and dynamic dTE_H , dTE_L), which accumulates through the entire clock chain and is present at the O-RU UNI, as described in Recommendation ITU-T G.8271.1 [22] Appendix IV. This Annex consider the accumulation of centered, symmetrical noise. In particular, accumulated dynamic time error causes O-RU synchronization receiver clock FFO (fractional frequency error) after clock recovery and filtering. Given that O-RU meets the 3GPP air interface frequency accuracy target (± 50 ppb), O-RU filtering is needed to filter the accumulated dynamic time error and reduce the frequency error down to an acceptable level. The allowed network limit (i.e. dynamic time error), reasonable O-RU filter bandwidth and acceptable frequency error after filtering are the result of a compromise exercise as shown in the following analysis.

The value of the O-RU filtering bandwidth is a key compromise, combined with the local oscillator thermal sensitivity under variable temperature profiles (such as the ones specified in Recommendation ITU-T G.8273 [27]):

- The higher filtering bandwidth, the faster frequency correction of the local oscillator thermal sensitivity and therefore the lower temporary accumulated time error under thermal variations, but the poorer efficiency in low pass filtering the dynamic noise seen on the UNI.
- The lower filtering bandwidth, the better efficiency in low pass filtering the dynamic noise seen on the UNI, but the poorer frequency correction of the local oscillator thermal sensitivity and therefore the higher temporary accumulated time error under thermal variations.

H.2.2 Network limit guidelines for frequency error budget (LLS-C1 and LLS-C2):

Based on the above compromise explanation, a practical expectation of O-RU filtering max BW is set to 75mHz to start the analysis as shown in Table H.2.2-1.

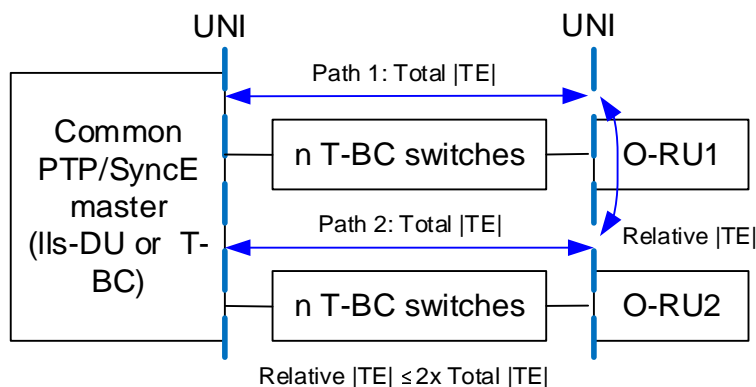
Table H.2.2-1: Network (LLS-C1/C2) Frequency Error Budget

Frequency error budget analysis considering O-DU class:	A	B
Consider O-DU PTP/SyncE master max frequency error budget = (refer to NOTE 1 in Table 11.3.2.1-1)	15 ppb	5 ppb
Consider Fronthaul and O-RU total max frequency error budget based on O-DU frequency error budget taken away from the 3GPP air interface (± 50 ppb) budget =	35 ppb	45 ppb
Further split the Fronthaul and O-RU total frequency error budget as follows as an example of O-RU design: <ul style="list-style-type: none"> max FFO_FH (see NOTE 1) = max FFO_RU (O-RU internal additive frequency noise generation) = 	21 ppb 14 ppb	27 ppb 18 ppb
The relationship between max FFO_FH and the acceptable dynamic time error at the input of the O-RU is the following (NOTE 2): <ul style="list-style-type: none"> max FFO_FH (in ppb) = $\pi \cdot \max(\text{dTE})$ (in ns) \cdot filter BW (in Hz) max(dTE) = max FFO_FH / ($\pi \cdot$ filter BW) = <p>max(dTE) is therefore the maximum allowed network Fronthaul network noise generation limit (between O-DU UNI and O-RU UNI) guaranteeing the 3GPP FFO limit at the air interface with such a 75mHz BW O-RU. It is up to O-RU vendor implementation to select filter BW (not necessarily 75mHz) to trade off the internal budget split between max FFO_FH and max FFO_RU as long as the frequency error budget (35 ppb or 45 ppb) is still met.</p>	90 ns	114 ns
Based on Recommendation ITU-T G.8271.1 [22] Appendix IV guidance and working assumptions to calculate accumulated error: Total dynamic noise = RSS (dTE) <ul style="list-style-type: none"> max(dTE) = RSS (max(dTE_L)) of all nodes excluding O-RU's T-TSC, and max(dTE_H) of the last clock of the chain) = 	90 ns	114 ns
Considering the model of clock chain of n T-BC clocks (between O-DU UNI to RU UNI) Recommendation ITU-T G.8273.2 [28] (class B) switch: max(dTE _L) = 40 ns, max(dTE _H) = 70 ns <ul style="list-style-type: none"> max(dTE) = sqrt [n*40² + 70²] = n = (max(dTE)² - 70²)/40², the maximum number of class B T-BCs in each chain (excluding O-DU) <p>Considering the model of clock chain of n T-BC clocks (between O-DU UNI to RU UNI) Recommendation ITU-T G.8273.2 [28] (class C) switch: max(dTE_L) = 10 ns, max(dTE_H) = 20 ns (NOTE 3) <ul style="list-style-type: none"> max(dTE) = sqrt [n*10² + 20²] = n = (max(dTE)² - 20²)/10², the maximum number of class C T-BCs in each chain (excluding O-DU) </p>	90 ns 2	114 ns 5
NOTE 1: This FFO_FH contributor is generated by the fronthaul network elements under ideal O-DU output and filtered by ideal O-RU PTP clock. Such measurement output signal (eg. 1PPS or PTP port) at the O-RU PTP clock output is optional and may therefore not exist on all O-RUs. NOTE 2: The noise is generally assumed to be mainly centered. The impact of rare highly asymmetric events on the FFO is outside the scope of this edition of the document. As per clause 11.2.5.1 max(dTE) is peak-to-peak value. This limit mainly applies to the higher frequency range; for the lower frequency range, which depends on O-RU vendor-specific low pass filtering characteristics, a higher limit may be allowed. NOTE 3: This max(dTE _H) limit is not yet specified by Recommendation ITU-T G.8273.2 [28] and is therefore an estimation.		

H.2.3 Network limit guidelines for time error budget (LLS-C1 and LLS-C2):

Using existing class B T-BCs, and considering no time error contribution by the fiber asymmetry nor from two master ports of the same T-BC, then:

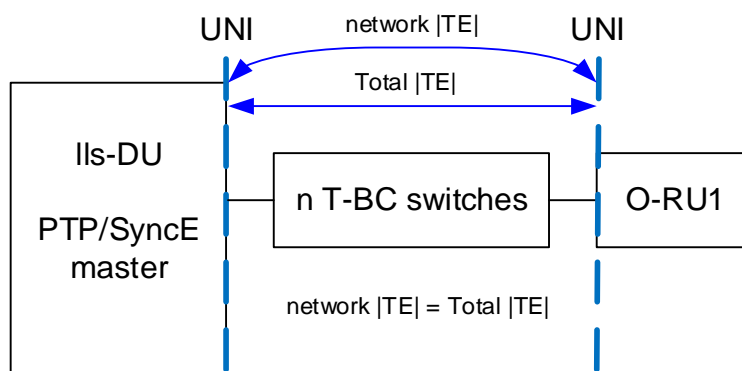
Following Recommendation ITU-T G.8271.1 [23], Appendix IV guidance to calculate accumulated error with the following clock chain models as shown in Figure H.2.3-1, and Figure H.2.3-2:



NOTE 1: This picture does not show the port-to-port relative TE between two Master ports of the common master. Ils-DU is considered as having no such port-to-port relative time error. However, T-BCs as specified by Recommendation ITU-T G.8273.2 does not specify it, and it may be more prudent to consider the common T-BC inside the « n » chain.

NOTE 2: This picture does not show the relative TE caused by asymmetry on the links, therefore some margin should be considered in the TE budget to take them into account

Figure H.2.3-1: Clock chain model for analysis for two paths



NOTE: This picture does not show the relative TE caused by asymmetry on the links, therefore some margin should be considered in the TE budget to take it into account

Figure H.2.3-2: Clock chain model for analysis

As per Recommendation ITU-T G.8271.1 [23] Appendix IV formula IV-13 "case 1" with centered noise:

Total $\max|\text{TE}| = \text{sum}(\max|\text{cTE}| \text{ of } n \text{ nodes}) + \frac{1}{2} \cdot \text{RSS} \text{ sum}(\max(\text{dTE}_L) \text{ of } n \text{ nodes and } \max(\text{dTE}_H) \text{ of last node})$

$$= n \cdot \max|\text{cTE}| + \frac{1}{2} \cdot \text{sqrt}(n \cdot \max(\text{dTE}_L)^2 + \max(\text{dTE}_H)^2)$$

However, the O-RU's time error budget ($\max|\text{cTE}|, \max(\text{dTE}_L)^2 + \max(\text{dTE}_H)^2$) has already been considered in the budget (see tables below), so only the T-BCs in the network are included in "n" and their contribution is limited to $\max|\text{cTE}|$ and $\max(\text{dTE}_L)$. As a result, the above general formula can be further simplified for the chain of n T-BCs excluding the O-RU's T-TSC:

$$\text{Total } \max|\text{TE}_L| = n \cdot \max|\text{cTE}| + \frac{1}{2} \cdot \text{sqrt}(n \cdot \max(\text{dTE}_L)^2),$$

Where a node is based on T-BC Class B switch with the following noise generation specification:

Constant time error = $\max|\text{cTE}| = 20 \text{ ns}$ for class B, 10 ns for class C

Low-band dynamic error = $\max(\text{dTE}_L) = 40 \text{ ns}$ for class B, 10 ns for class C (considering centered noise)

Using the above results in an O-RU TE budget as shown in Table H.2.3-1.

Table H.2.3-1: O-RU Time Error Budget

For O-RU type= and overall relative $ TE_L $ limit = to meet timing category (NOTES 1, 2)	Enhanced 60 ns Cat A	Regular 100 ns Cat B	Enhanced 190 ns Cat B	Regular 95 ns Cat C	Enhanced 140 ns Cat C
$\max TE_L $ per branch (NOTE 2) $n \cdot \max cTE + \frac{1}{2} \cdot \sqrt{(n \cdot \max(dTE_L))^2} < \text{limit per branch}$	30 ns	50 ns	95 ns	95 ns	140 ns
Maximum number of class B T-BCs (n) on each branch from common network element (either O-DU or T-BC) to O-RU: Noise accumulation in each branch, calculated using $n \cdot \max cTE + \sqrt{(n \cdot \max(dTE_L))^2}$: Remaining relative TE margin overall (NOTE 3):	0 0 ns 60 ns	1 40 ns 20 ns	2 68 ns 54 ns	2 68 ns 27 ns	4 120 ns 20 ns
Maximum number of class C T-BCs (n) on each branch from common network element (either O-DU or T-BC) to O-RU: Noise accumulation in each branch, calculated using $n \cdot \max cTE + \sqrt{(n \cdot \max(dTE_L))^2}$: Remaining relative TE margin overall (NOTE 3):	1 15 ns 30 ns	3 39 ns 22 ns	7 83 ns 24 ns	7 83 ns 12 ns	11 127 ns 13 ns
<p>NOTE 1: Proposed maximum n values are conservative and leave at least 10 ns margin for fiber asymmetry or relative $\max TE_L$ between two ports of the branching clock. It is therefore recommended to limit the number of fronthaul clocks to this value, although an additional one may still allow meeting the expected limits.</p> <p>NOTE 2: For IEEE802.1CM [11,12] clause 6 timing categories A and B, the overall limit corresponds to the maximum relative TE_L between two O-RU input ports (the end of two branches). The table calculates the noise accumulation in each branch separately, with the limit per branch equal to half of the overall limit. The two branches are then added linearly to calculate the remainder. A less conservative assumption, such as treating both branches as a single chain, could result in a larger margin for the remaining $\max TE_L$. For IEEE802.1CM [11, 12] clause 6 timing category C, the overall limit corresponds to $\max TE_{RL}$ between the O-DU output and any O-RU input ports, and therefore the limit per branch is the same as the overall limit.</p> <p>NOTE 3: For IEEE802.1CM [11, 12] clause 6 timing categories A and B, the remaining relative TE should be assigned to fiber asymmetry in each of the branches from the common network element to the O-RU input ports, plus the relative TE between ports of the common network element For category C, the remaining relative TE should be assigned to fiber asymmetry in the chain between the O-DU output port and the O-RU input ports</p>					

H.2.4 Network limit guidelines for LLS-C1 and LLS-C2 based on the budget

The above budget can be translated into network limits at various points in the network. This is useful for measurement purposes. For the LLS-C1 and LLS-C2 configurations, a simplified version of the synchronization network can be viewed as shown in Figure H.2.4-1 below. The numbers included are all quoted from Table 11.3.2.1-1, which is the normative source for LLS-C1 and LLS-C2.

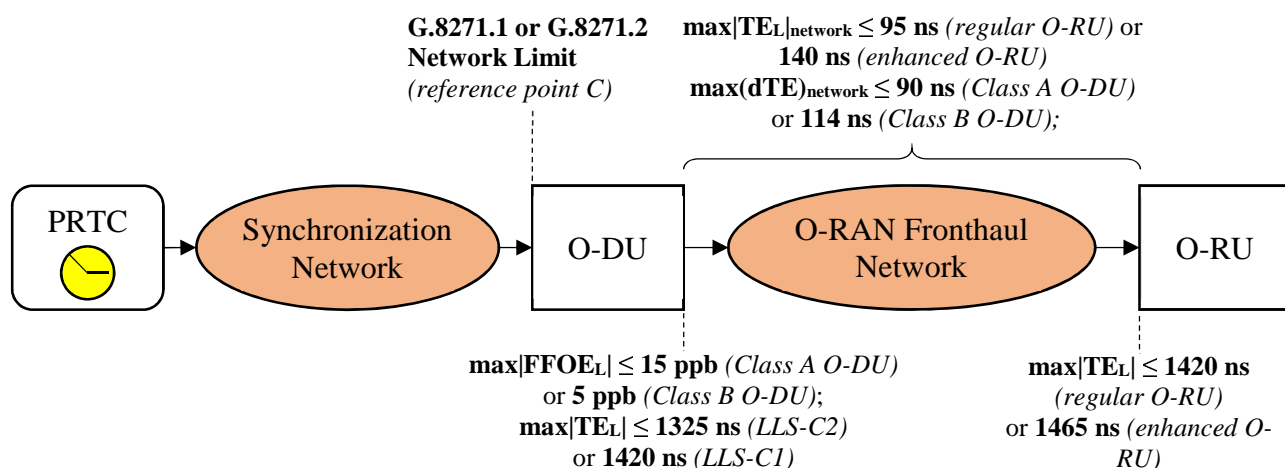


Figure H.2.4-1: Simplified view of LLS-C1 and LLS-C2 configurations, with network limits

The limit at the output of the O-DU can be described in terms of MTIE (Maximum Time Interval Error). This is given in Table H.2.4-1 and Figure H.2.4-2 (for a Class A O-DU) and Table H.2.4-2 and Figure H.2.4-3 (for a Class B O-DU). In both cases, these are measured after filtering by a first-order low-pass filter of bandwidth 0.1 Hz (see Table 11.3.2.1-1, NOTE 1).

The maximum MTIE of 1300 ns at long observation intervals is derived from the network limit in Recommendation ITU-T G.8271.2 [24] clause 7 of 1100 ns (for an APTS network at the input to the O-DU), plus 200 ns for the noise generation of the O-DU (based on the noise generation of a T-TSC-A in Recommendation ITU-T G.8273.4 [35] clause 7). This is the maximum MTIE value that should occur if the O-DU is synchronized using APTS. If the synchronization network at the input to the O-DU is FTS, while the slope of the MTIE graph at lower observation intervals will be the same, the maximum output should be much lower than the limit.

An unfiltered limit is also shown in each case, calculated by adding a constant value to the filtered limit representing the amount of noise that is expected to be removed by the measurement filter. The value to be added is calculated by using the equation quoted in Table H.2.2-1-1 above:

$$\max(\text{dTE}) = \frac{\max|\text{FFO}|}{\pi f_m}$$

where $\max|\text{FFO}|$ is the frequency accuracy budget of the O-DU, and f_m is the bandwidth of the measurement filter. For an O-DU class A with a 15 ppb frequency accuracy budget, and a measurement filter of bandwidth 0.1 Hz, the $\max(\text{dTE})$ is 48 ns, resulting in an addition of 48 ns to the filtered mask. For an O-DU class B, the addition to the mask is 16 ns.

Table H.2.4-1: MTIE Limit at the output of an O-DU Class A

Observation interval (s)	MTIE limit (0.1 Hz filtered) (ns)	MTIE limit (unfiltered) ns
$0.1 < \tau \leq 86.67$	15τ	$48 + 15 \tau$
$\tau > 86.67$	1300	1348




Figure H.2.4-2: MTIE Limit at the output of an O-DU Class A

Table H.2.4-2: MTIE Limit at the output of an O-DU Class B

Observation interval (s)	MTIE limit (0.1 Hz filtered) (ns)	MTIE limit (unfiltered) ns
$0.1 < \tau \leq 260$	5τ	$16 + 5 \tau$
$\tau > 260$	1300	1316

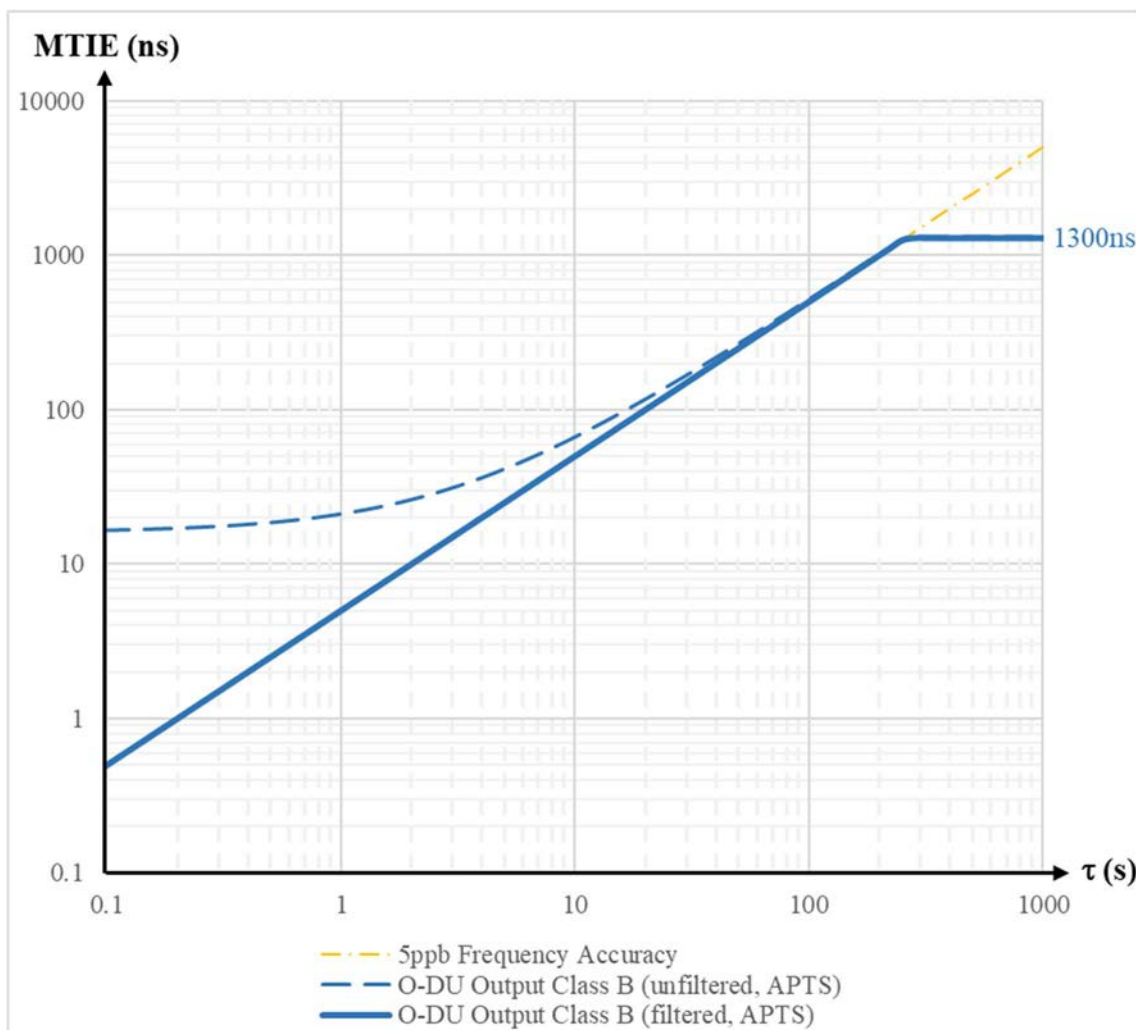


Figure H.2.4-3: MTIE Limit at the output of an O-DU Class B

Table H.2.2-1 makes assumptions about the bandwidth of the O-RU and the budget allocation for the FFO; however, designers are free to change these assumptions provided the overall performance requirements at the RF output are met. The only figures for the amount of noise that are quoted in the normative section (Table 11.3.2.1-1) are the two max(dTE) figures of 90 ns for an O-DU Class A, and 114 ns for an O-DU Class B.

The maximum limit at the input to an O-RU can then be calculated by adding these figures to the unfiltered limit at the output of the O-DU. For a worst-case input, linear addition is assumed, adding 90 ns for an O-DU Class A, and 114 ns for an O-DU Class B. Tables H.2.4-3, and H.2.4-4, show an unfiltered measurement:

Table H.2.4-3: MTIE Limit at the input to an O-RU (connected to an O-DU Class A)

Observation interval (s)	MTIE limit (unfiltered) ns
$0.1 < \tau \leq 86.67$	$138 + 15 \tau$
$\tau > 86.67$	1438

Table H.2.4-4: MTIE Limit at the input to an O-RU (connected to an O-DU Class B)

Observation interval (s)	MTIE limit (unfiltered) ns
$0.1 < \tau \leq 260$	$130 + 5 \tau$
$\tau > 260$	1430

These unfiltered limits are shown in Figure H.2.4-4 below in the dotted green and blue lines. The 15 ppb and 5 ppb lines are shown for information only. It is not possible to calculate a filtered version of these limits without assuming something about the frequency accuracy of the various noise contributions (e.g., the budget figures used in Table H.2.2-1), therefore the unfiltered limits are the primary limits.

The limit for a Class A O-DU is in all cases slightly higher than that for a Class B O-DU. Since when an O-RU is designed it is not known what class of O-DU it will be deployed with, the maximum noise that an O-RU tolerates at its input is that for a Class A O-DU (the dashed green curve in Figure H.2.4-4 below):




Figure H.2.4-4: MTIE limit at the input to an O-RU

The MTIE limits shown in Figure H.2.4-2, Figure H.2.4-3 and Figure H.2.4-4 are not new requirements. They are derived from the existing requirements documented in Table 11.3.2.1-1, and hence are a way of expressing those requirements using MTIE. These limits should not create compatibility issues for existing equipment tested to the parameters in Table 11.3.2.1-1.

H.2.5 Network limit guidelines for frequency error budget (LLS-C3):

- Based on the above compromise explanation, a practical expectation of O-RU filtering max BW is set to 75 MHz to start the analysis
- Based on Recommendation ITU-T G.8272, PRTC/T-GM MTIE (during lock) specification can be used to describe PRTC/T-GM dynamic noise generation, see Tables H.2.5-1 and H.2.5-2:

Table H.2.5-1: Wander generation (MTIE) for PRTC-A

MTIE limit (μs)	Observation interval (s)
$0.275 \times 10^{-3} \tau + 0.025$	$0.1 < \tau \leq 273$
0.10	$\tau > 273$

Table H.2.5-2: Wander generation (MTIE) for PRTC-B

MTIE limit (μs)	Observation interval (s)
$0.275 \times 10^{-3} \tau + 0.025$	$0.1 < \tau \leq 54.5$
0.04	$\tau > 54.5$

- Given O-RU filtering max BW = 75 mHz, it corresponds to observation interval $\tau = 1/(\pi \cdot 75 \text{ mHz}) = 4.2 \text{ s}$. From the above table, MTIE limit (with $\tau = 4.2 \text{ s}$) = 26.2 ns pk-pk for both PRTC-A and PRTC-B.
- Besides MTIE, which can be treated as dynamic noise during lock condition, there is an additional consideration of PRTC/T-GM during HOLDOVER condition. Potential semi-static frequency drift could happen during HOLDOVER, $\pm 2 \text{ ppb}$ is reserved based on Recommendation ITU-T G.8271.1 Appendix V PRTC failure scenario (b) which permits 400 ns HOLDOVER limit for short period of 5 minutes.

Based on the above, the error budget is shown below in Table H.2.5-3.

Table H.2.5-3: Network (LLS-C3) frequency error budget

Frequency error budget analysis for PRTC class:	A
Consider PRTC PTP/SyncE master max frequency error budget (refer to NOTE 1 in Table 11.3.2.2-1)	2 ppb
Consider Fronthaul and O-RU total frequency error budget based on PRTC frequency error budget taken away from the 3GPP air interface ($\pm 50 \text{ ppb}$) budget =	48 ppb
Further split the Fronthaul and O-RU total frequency error budget as follows as an example of O-RU design: <ul style="list-style-type: none"> max FFO_FH (see NOTE 1) = max FFO_RU (O-RU internal additive frequency noise generation) = 	30 ppb 18 ppb
The relationship between max FFO_FH and the acceptable dynamic time error at the input of the O-RU is the following (NOTE 2): <ul style="list-style-type: none"> max FFO_FH (in ppb) = $\pi \cdot \text{max}(dTE)$ (in ns) \cdot filter BW (in Hz) max(dTE) = max FFO_FH / ($\pi \cdot$ filter BW) = <p>max(dTE) is therefore the maximum allowed Fronthaul network noise generation limit (between PRTC UNI and O-RU UNI) guaranteeing the 3GPP max frequency error limit at the air interface with such a 75mHz BW O-RU.</p> <p>It is up to O-RU vendor implementation to select filter BW (not necessarily 75mHz) to trade off the internal budget split between max FFO_FH and max FFO_RU as long as the O-RU total max frequency error budget (48 ppb) is still met.</p>	126 ns
Based on Recommendation ITU-T G.8271.1 [22] Appendix IV guidance and working assumptions to calculate accumulated error: <ul style="list-style-type: none"> Total dynamic noise = RSS (dTE) max(dTE) = RSS (max(dTE_L) of all nodes including PRTC/T-GM but not O-RU's T-TSC, and max(dTE_H) of the last clock in the chain) <p>Considering T-BC class B with : max(dTE_L) = 40 ns, max(dTE_H) = 70ns</p> <ul style="list-style-type: none"> max(dTE) = $\text{sqrt}[26^2 + n \cdot 40^2 + 70^2]$ = n = $(\text{max}(dTE)^2 - 70^2 - 26^2) / 40^2$ = the maximum number of class B T-BCs in each chain (after PRTC) <p>Considering T-BC class C with: max(dTE_L) = 10 ns, max(dTE_H) = 20 ns (NOTE 3)</p> <ul style="list-style-type: none"> max(dTE) = $\text{sqrt}(26^2 + n \cdot 10^2 + 20^2)$ = n = $(\text{max}(dTE)^2 - 20^2 - 26^2) / 10^2$ = the maximum number of class C T-BCs in each chain (after PRTC) 	126 ns 6 126 ns >>10

NOTE 1: This FFO_FH contributor is generated by the fronthaul network elements under ideal PRTC/T-GM output and filtered by ideal O-RU PTP clock. Such measurement output signal (eg. 1PPS or PTP port) at the O-RU PTP clock output is optional and may therefore not exist on all O-RUs
NOTE 2: The noise is generally assumed to be mainly centered. The impact of rare highly asymmetric events on the FFO is outside the scope of the present document. As per clause 11.2.5.1, max(dTE) is peak-to-peak value. This limit mainly applies to the higher frequency range; for the lower frequency range, which depends on O-RU vendor-specific low pass filtering characteristics, a higher limit may be allowed (still within the limits specified by Recommendation ITU-T G.8271.1 [23] clause 7.3 as worst case).
NOTE 3: This max(dTE _H) limit is not yet specified by Recommendation ITU-T G.8273.2 [28] and is therefore an estimation.

H2.6 Time error budget for network limit (LLS-C3)

Recommendation ITU-T G.8271.1 [22] Appendix V (Example of design options) and Appendix XII (Examples of design options for fronthaul and clusters of base stations) provide guidelines on the number of switches that can be deployed in case of LLS-C3 for the different target requirements.

Appendix V is focusing on the absolute Time Error Requirement (Category C), while Appendix XII addresses also relative time error requirements applicable in fronthaul (Category A and B).

H.3 Guidelines for the allowed number of switches:

The recommended guidance for maximum number of switches is shown in Table H.3-1 and H.3-2 and is determined based on the smallest allowed number constrained by

- Frequency error budget.
- Operator-chosen most constraint time error budget category.
- The class of network elements (The O-RU classes are examples proposed by IEEE802.1 CM [11, 12]).

Table H.3-1: Network frequency error budget

Frequency Error Network limit	LLS-C1 and LLS-C2, class A O-DU	LLS-C1 and LLS-C2, class B O-DU	LLS-C3	Comment
Absolute Frequency error budget between time source and O-RU	2 (class B T-BC) >>10 (class C T-BC)	5 (class B T-BC) >>10 (class C T-BC)	NOTE 3	No branch exceeds this number of T-BCs from O-DU or PRTC/T-GM to meet 50 ppb frequency accuracy at the air interface.

Table H.3-2: Network time error budget

Time Error Network limit	LLS-C1 and LLS-C2, class A O-DU	LLS-C1 and LLS-C2, class B O-DU	LLS-C3	Comment
Cat A Relative Time error budget (with enhanced O-RUs) between O-RUs	0 (class B T-BC) 1 (class C T-BC)	0 (class B T-BC) 1 (class C T-BC)	NOTE 3 NOTE 4	No branch exceeds this number of T-BCs from common T-BC to meet target relative time error limit at the air interface. Value 0 means that only IIs-C1 is supported
Cat B Relative Time error budget (with regular O-RUs) between O-RUs	1 (class B T-BC) 3 (class C T-BC)	1 (class B T-BC) 3 (class C O-RU)	NOTE 3	
Cat B Relative Time error budget (with enhanced O-RUs) between O-RUs	2 (class B T-BC) 7 (class C T-BC)	2 (class B T-BC) 7 (class C T-BC)	NOTE 3	
Cat C Absolute Time error budget (with regular O-RUs) between time source and O-RU	2 (class B T-BC) 7 (class C T-BC)	2 (class B T-BC) 11 (class C T-BC)	NOTE 3	No branch exceeds this number of T-BCs from O-DU or PRTC/T-GM to meet 1500 ns absolute time error limit at the air interface. NOTE 1.
Cat C Absolute Time error budget (with enhanced O-RUs) between time source and O-RU	4 (class B T-BC) (further limit to 2 due to freq. limit) 7 (class C T-BC)	4 (class B T-BC) 11 (class C T-BC)	NOTE 3	
<p>NOTE 1: Only applicable to IIs-C1 and IIs-C2: As indicated in Table 11.3.2.1-1, the $\max TE_L _{O-DU}$ at the output of the O-DU is 1420 ns for LLS-C1 and 1325 ns for LLS-C2. This limit considers that the input of the O-DU stays within the limits at Reference point C defined by Recommendation ITU-T G.8271.1 [22] or Recommendation ITU-T G.8271.2 [24].</p> <p>NOTE 2: The analysis on the number of switches (for time error budget) is meant to be a rough estimate and excludes both fiber asymmetry factor and relative TE_R between two ports of the common network element. It is operator's responsibility to control these two parameters. The analysis for each time error budget (specific category) has some left-over margin that could be used to cover them. If the left-over margin is not enough, the alternative is to reduce the allowed number of switches.</p> <p>NOTE 3: Network design guidelines for configuration LLS-C3 are provided in Recommendation ITU-T G.8271.1 [22] Appendix V (addressing IEEE802.1CM [11, 12] synchronization Category C) and Appendix XII (addressing IEEE802.1CM [11, 12] synchronization Category A and B) of Recommendation ITU-T G.8271.1 [22]. The guidelines in Appendix V includes an indication on allowed number of switches between the PRTC/T-GM and the O-RU. The guidelines in Recommendation ITU-T G.8271.1 [22] Appendix XII also includes an indication on allowed number of switches between a clock that is common for the cooperating O-RUs, and these O-RUs.</p> <p>As an example:</p> <ul style="list-style-type: none"> - In order to meet the Cat B requirements with a regular O-RU, the number of T-BC class C switches, after the common switch (itself a class C T-BC), should be 3 or less, while there is no room for additional switches with T-BC class B clocks after the common switch (itself a class B T-BC). - In order to meet the Cat B requirements with an enhanced O-RU, the number of T-BC class C switches should be 6 or less after the common switch (itself a class C T-BC), and the number of T-BC class B switches after the common switch (itself a class B T-BC), should be 2 or less. <p>The minor differences between the maximum number of switches supported in LLS-C2 and LLS-C3 are due to different characteristics of the common clock (as indicated in Figure H.2.3-1, in LLS-C2, the O-DU has no relative time error between ports specified, while in LLS-C3, the Recommendation ITU-T G.8273.2 [28] T-BC takes this into account). Recommendation ITU-T G.8271.1 [22] also presents the case of an alternative deployment with a short clock chain that has a maximum of 4 Class C T-BC, or 1 Class B T-BC between the PRTC/T-GM and the O-RU (see reference network model in Figure II.6 of Recommendation ITU-T G.8271.1 [22] with a PRTC-B/T-GM directly connected to the common T-BC). For this case the regular O-RU was considered as it represents the worst-case scenario. This deployment, in addition to meeting IEEE802.1CM [11,12] synchronization Category C, is also suitable to support IEEE802.1CM [11, 12] synchronization Category B.</p> <p>NOTE 4: Cat A requirements concerns co-located O-RUs. It is assumed that the cooperating O-RUs are connected to the same switch (therefore there is no switch after the common T-BC).</p>				

Annex I (informative): Precoding and examples

I.1 Case 1: Tx Diversity 1-CRS Port Ant0, 1 PRB:

Case 1 is depicted in Figure I.1-1.

At the O-DU

- For single Tx case, $d^{(q)}(0), \dots, d^{(q)}(M_{\text{sym}}^{(q)} - 1)$ $M_{\text{q_sym}}$ modulation bits belonging to PDSCH ANT0 are packed into a PRB. All CRS REs for ANT0 are packed into a PRB for transmission and are unpacked at the O-RU.

At the O-RU

- At the O-RU, for single antenna port single Tx a single layer is used and mapping is defined as $x^{(0)}(i) = d^{(0)}(i)$
- O-RU needs to map CRS REs into antenna ports using `crsReMask`, `crsSymbolNumber` and `crsShift`. Clause I.5 describes CRS RE mapping into antenna ports.

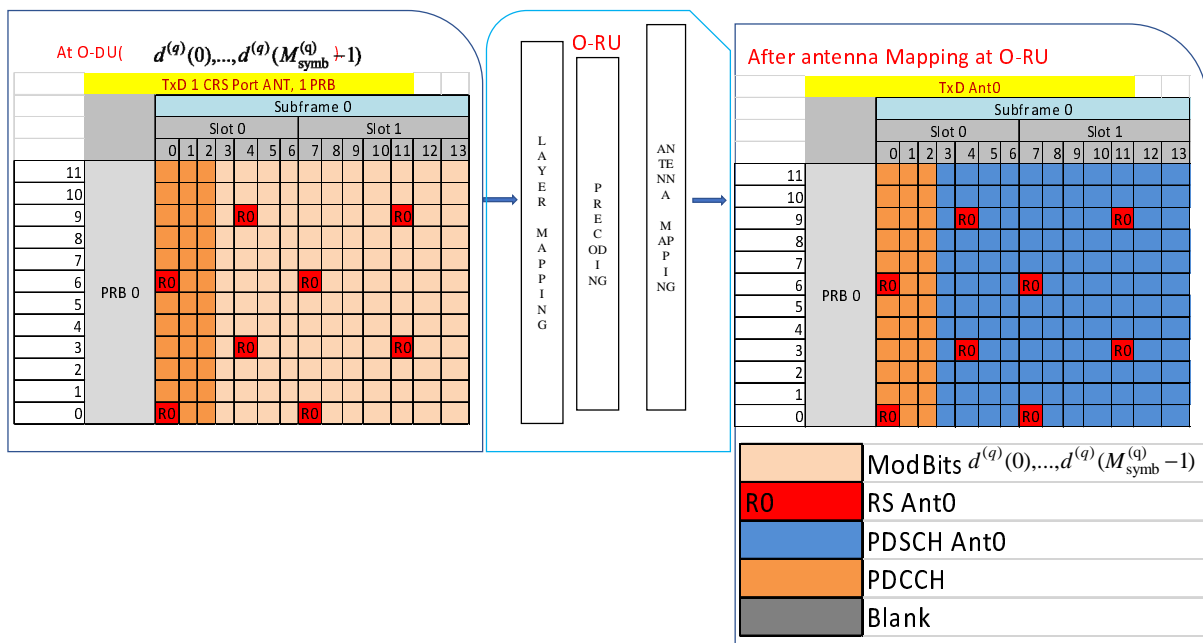


Figure I.1-1: Single Tx 1-CRS Port Ant0, 1 PRB

I.2 Case 2: Tx Diversity 2-CRS Port Ant0, Ant1, and 1 PRB

Case 2 is depicted in Figure I.2-1.

At the O-DU

- For TxD case, $d^{(q)}(0), \dots, d^{(q)}(M_{\text{sym}}^{(q)} - 1)$ $M_{\text{q_sym}}$ modulation bits belonging to PDSCH are packed into a PRB. All CRS REs for ANT0 and ANT1 are packed into the same PRB for transmission and are unpacked at the O-RU.

At the O-RU

- At the O-RU, for two antenna port TxD 2 layers are used and mapping is defined as $x^{(0)}(i) = d^{(0)}(2i)$
 $x^{(1)}(i) = d^{(0)}(2i+1)$
- For TxD, information for 2 layers are packed into a PRB for transmission and are unpacked at the O-RU. At the O-RU, after layer mapping and, precoding, CRS REs for 2 antenna ports are mapped to the appropriate RE positions and rest are left blank (gray-shaded REs) as illustrated below. Clause I.5 describes CRS RE mapping into antenna ports.

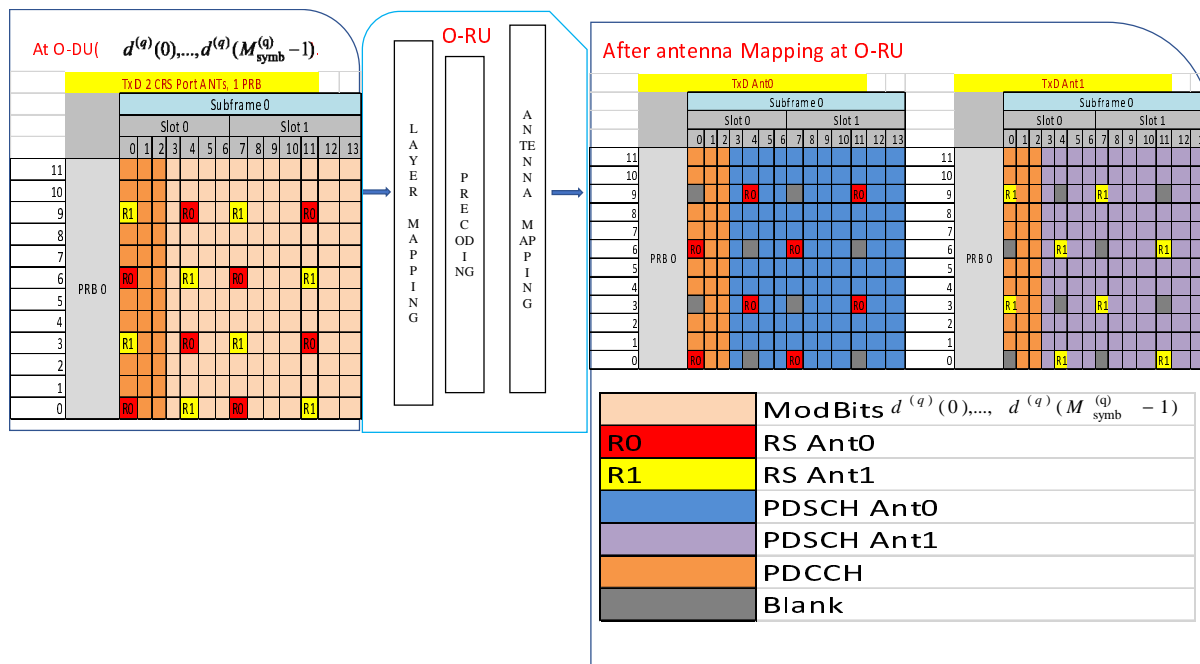


Figure I.2-1: Tx Diversity 2-CRS Port Ant0, Ant1, and 1 PRB

I.3 Case 3: Tx Diversity 4-CRS Port Ant0,1,2,3 and 1 PRB

Case 3 layer mapping is depicted in Figure I.3-1.

Table 6.3.3.3-1: Codeword-to-layer mapping for transmit diversity

Number of layers	Number of codewords	Codeword-to-layer mapping $i = 0, 1, \dots, M_{\text{sym}}^{\text{layer}} - 1$
2	1	$x^{(0)}(i) = d^{(0)}(2i)$ $x^{(1)}(i) = d^{(0)}(2i+1)$ $M_{\text{sym}}^{\text{layer}} = M_{\text{sym}}^{\text{codeword}} / 2$
4	1	$x^{(0)}(i) = d^{(0)}(4i)$ $x^{(1)}(i) = d^{(0)}(4i+1)$ $x^{(2)}(i) = d^{(0)}(4i+2)$ $x^{(3)}(i) = d^{(0)}(4i+3)$ $M_{\text{sym}}^{\text{layer}} = \begin{cases} M_{\text{sym}}^{\text{codeword}}/4 & \text{if } M_{\text{sym}}^{\text{codeword}} \bmod 4 = 0 \\ (M_{\text{sym}}^{\text{codeword}}+2)/4 & \text{if } M_{\text{sym}}^{\text{codeword}} \bmod 4 = 2 \end{cases}$ If $M_{\text{sym}}^{\text{layer}} \bmod 4 \neq 0$ two null symbols shall be appended to $d^{(0)}(M_{\text{sym}}^{\text{codeword}} - 1)$

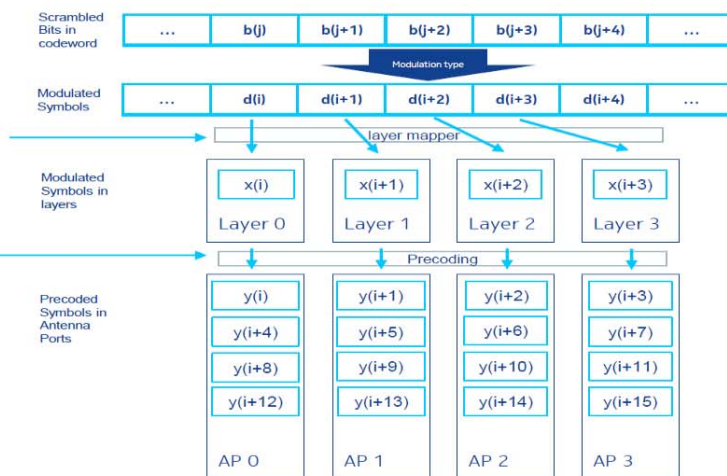
$$\begin{bmatrix}
 y^{(0)}(4i) \\
 y^{(1)}(4i) \\
 y^{(2)}(4i) \\
 y^{(3)}(4i) \\
 y^{(0)}(4i+1) \\
 y^{(1)}(4i+1) \\
 y^{(2)}(4i+1) \\
 y^{(3)}(4i+1) \\
 y^{(0)}(4i+2) \\
 y^{(1)}(4i+2) \\
 y^{(2)}(4i+2) \\
 y^{(3)}(4i+2) \\
 y^{(0)}(4i+3) \\
 y^{(1)}(4i+3) \\
 y^{(2)}(4i+3) \\
 y^{(3)}(4i+3)
 \end{bmatrix}
 = \frac{1}{\sqrt{2}}
 \begin{bmatrix}
 1 & 0 & 0 & 0 & j & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & -1 & 0 & 0 & 0 & j & & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 & 0 & j & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 1 & 0 & 0 & 0 & -j & & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 0 & 0 & 0 & j & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & -1 & 0 & 0 & 0 & j \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 1 & 0 & 0 & 0 & j \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 0 & 0 & 0 & -j & 0
 \end{bmatrix}
 \begin{bmatrix}
 \text{Re}\{x^{(0)}(i)\} \\
 \text{Re}\{x^{(1)}(i)\} \\
 \text{Re}\{x^{(2)}(i)\} \\
 \text{Re}\{x^{(3)}(i)\} \\
 \text{Im}\{x^{(0)}(i)\} \\
 \text{Im}\{x^{(1)}(i)\} \\
 \text{Im}\{x^{(2)}(i)\} \\
 \text{Im}\{x^{(3)}(i)\}
 \end{bmatrix}$$


Figure I.3-1: Case 3 Layer Mapping

At the O-DU

- For TxD case, $d^{(q)}(0), \dots, d^{(q)}(M_{\text{symb}}^{(q)} - 1)$ $M_{\text{q_sym}}$ modulation bits belongs to PDSCH are packed to into a PRB. All CRS RE's for ANT0, ANT1, ANT2 and ANT3 are packed into that same PRB for transmission and are unpacked at the O-RU.

At the O-RU

- At the O-RU, for two antenna port TxD 4 layers are used and mapping is defined as

$$x^{(0)}(i) = d^{(0)}(4i)$$

$$x^{(1)}(i) = d^{(0)}(4i + 1)$$

$$x^{(2)}(i) = d^{(0)}(4i + 2)$$

$$x^{(3)}(i) = d^{(0)}(4i + 3)$$
- For TxD, user data RE's for 4 layers are packed into a PRB for transmission and are unpacked at the O-RU. At the O-RU, after layer mapping and precoding, CRS REs for 4 antenna ports are mapped to the appropriate RE position and rest are left blank as illustrated in Figure I.3-2.

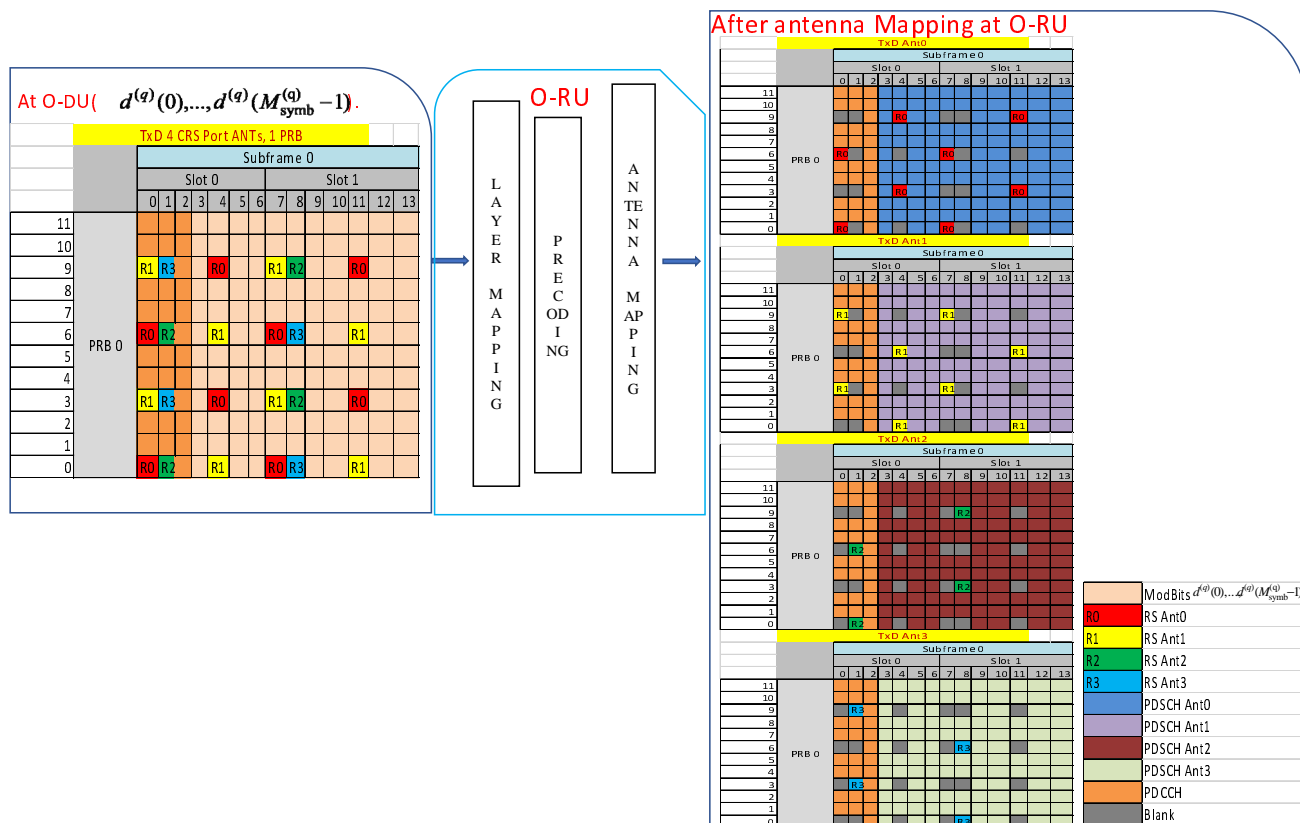


Figure I.3-2: Tx Diversity 4-CRS Port Ant0,1,2, 3, and 1 PRB

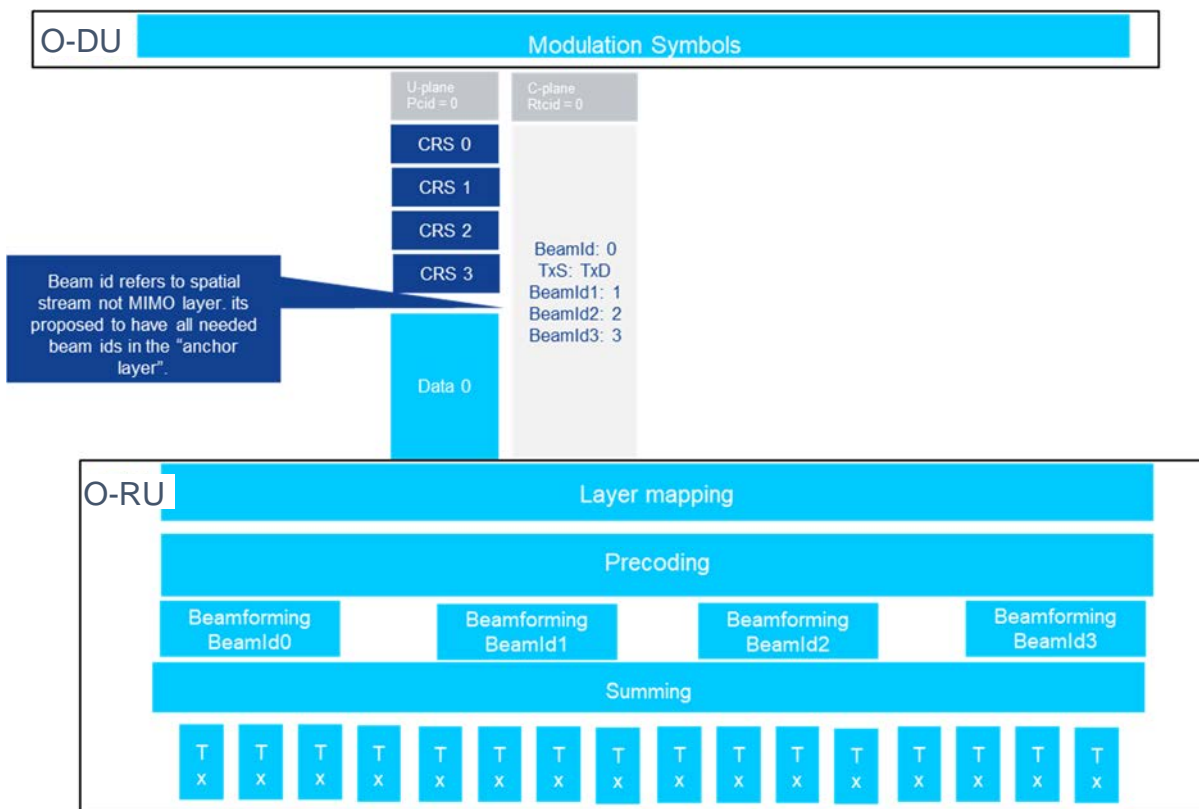


Figure I.3-3: Tx Diversity – 4 CRS Ports

I.4 Case 4: TM3/TM4 3 layers, 4 antenna ports

One symbol from each of layers is linearly mapped to each antenna port as shown in Figure I.4-1.




Figure I.4-1: Case 4 Layer Mapping

For TM3/TM4, REs belonging to all the antenna ports are mapped to the zeroth layer at the O-DU. The O-RU should consider layer zero's CRS REs for CRS mapping for all the antenna ports using reMask bit field; CRS RE positions in the non-zero layers can be ignored as illustrated below (gray-shared REs in the O-DU represent CRS RE positions that are not populated by the O-DU) in Figure I.4-2.

$$\begin{aligned}
 x^{(0)}(i) &= d^{(0)}(3i) \\
 x^{(1)}(i) &= d^{(0)}(3i+1) \quad M_{\text{symp}}^{\text{layer}} = M_{\text{symp}}^{(0)} / 3 \\
 x^{(2)}(i) &= d^{(0)}(3i+2)
 \end{aligned}$$

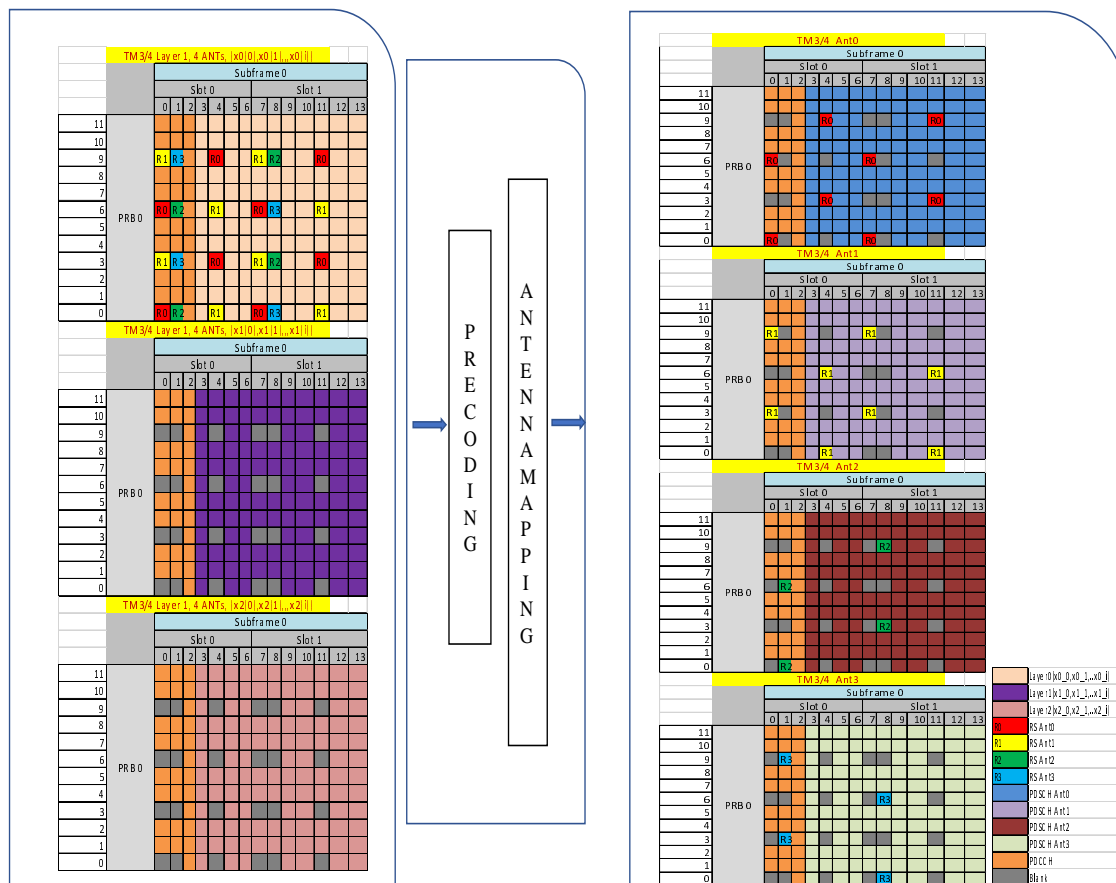


Figure I.4-2: TM3/TM4, 3 Layers 4 Antenna Ports

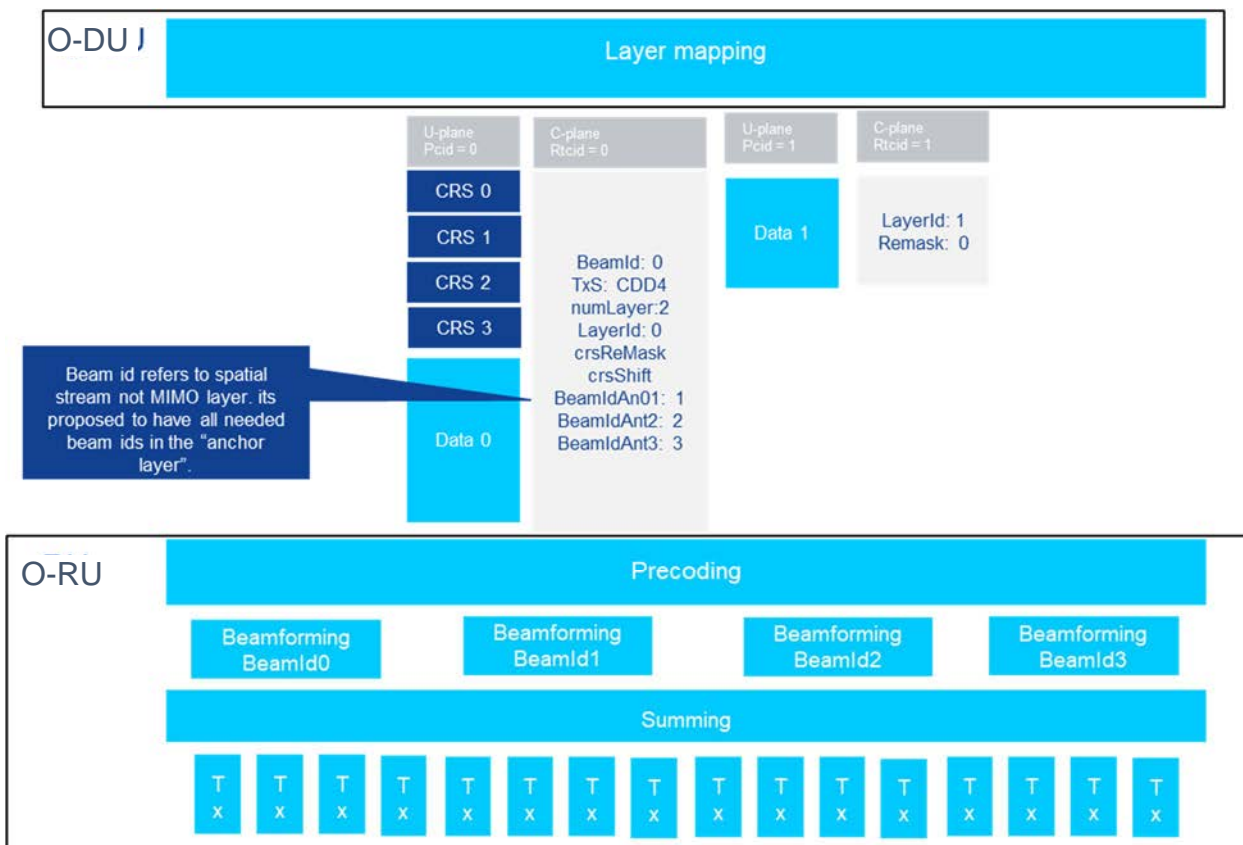


Figure I.4-3: TM3 rank 2/4 CRS ports

I.5 CRS location assignment

CRS shift pattern for antenna port 0 when $3 \leq vshift \leq 5$ is same as $0 \leq vshift \leq 3$ for antenna port 1 resulting in only three possible vshift for a 2-antenna port case and this holds good for 4-antenna port case as well. Hence the crsReMask is differentiated with the crsShift field. The CRS location for the different layers are described in Table I.5-1, Table I.5-2 and Table I.5-3.

Table I.5-1: CRS location assignment for layer1

Layer1					
		crsSymNum 0 and 7		crsSymNum 4 and 11	
vshift	crsShift	crsReMask	PosInd in Bit Ant 0	crsReMask	PosInd in Bit Ant 0
0	0	0000 0100 0001	0,6	0010 0000 1000	3,9
1	0	0000 1000 0010	1,7	0100 0001 0000	4,10
2	0	0001 0000 0100	2,8	0100 0010 0000	5,11
3	0	0010 0000 1000	3,9	0000 0100 0001	6,0
4	0	0100 0001 0000	4,10	0000 1000 0010	7,1
5	0	1000 0010 0000	5,11	0001 0000 0100	8,2

crsReMask is not repeated for all for vShift combinations, hence crsShift is always indicated as 0.

Table I.5-2: CRS Location Assignment For layer2

Layer2							
		crsSymNum 0 and 7			crsSymNum 4 and 11		
vshift	crsShift	crsReMask	PosInd in Bit Ant 0	PosInd in Bit Ant 1	crsReMask	PosInd in Bit Ant 0	PosInd in Bit Ant 1
0	0	0010 0100 1001	0,6	3,9	0010 0100 1001	3,9	0,6
1	0	0100 1001 0010	1,7	4,10	0100 1001 0010	4,10	1,7
2	0	1001 0010 0100	2,8	5,11	1001 0010 0100	5,11	2,8
3	1	0010 0100 1001	3,9	6,0	0010 0100 1001	6,0	3,9
4	1	0100 1001 0010	4,10	7,1	0100 1001 0010	7,1	4,10
5	1	1001 0010 0100	5,11	8,2	1001 0010 0100	8,2	5,11

crsReMask for vShift 0 is same as for vShift 3 case but RE locations are different for Ant 0 and 1. Hence indicating 1 for 3, 4, 5 vShift cases, thereby O-RU picks the positions as interpreted here.

Table I.5-3: CRS Location Assignment For layer4

Layer4							
		crsSymNum 0 and 7			crsSymNum 4 and 11		
vshift	crsShift	crsReMask	PosInd in Bit Ant 0	PosInd in Bit Ant 1	crsReMask	PosInd in Bit Ant 0	PosInd in Bit Ant 1
0	0	0010 0100 1001	0,6	3,9	0010 0100 1001	3,9	0,6
1	0	0100 1001 0010	1,7	4,10	0100 1001 0010	4,10	1,7
2	0	1001 0010 0100	2,8	5,11	1001 0010 0100	5,11	2,8
3	1	0010 0100 1001	3,9	6,0	0010 0100 1001	6,0	3,9
4	1	0100 1001 0010	4,10	7,1	0100 1001 0010	7,1	4,10
5	1	1001 0010 0100	5,11	8,2	1001 0010 0100	8,2	5,11
Layer4							
		crsSymNum 1			crsSymNum 8		
vshift	crsShift	crsReMask	PosInd in Bit Ant 2	PosInd in Bit Ant 3	crsReMask	PosInd in Bit Ant 3	PosInd in Bit Ant2
0	0	0010 0100 1001	0,6	3,9	0010 0100 1001	0,6	3,9
1	0	0100 1001 0010	1,7	4,10	0100 1001 0010	1,7	4,10
2	0	1001 0010 0100	2,8	5,11	1001 0010 0100	2,8	5,11
3	1	0010 0100 1001	3,9	6,0	0010 0100 1001	3,9	6,0
4	1	0100 1001 0010	4,10	7,1	0100 1001 0010	4,10	7,1
5	1	1001 0010 0100	5,11	8,2	1001 0010 0100	5,11	8,2

crsReMask for vShift 0 is same as for vShift 3 case but RE locations are different for Ant 0 and 1. Hence indicating 1 for 3, 4, 5 vShift cases, thereby O-RU picks the positions as interpreted here.

Pseudo code to determine the CRS belongs to ANT port-N

```
// Determine vShift = N_CELL_ID % 6
// Below pseudo code for Num Layer = 4
Switch(crsSymNum)
{
Case 0 || case 7:
    if( vShift == 0 )
    {
        if (bitSet_0 && bitSet_6)
        {
            mapIQ of 0 and 6 to Ant0 //rest blank
        }
        if (bitSet_3 && bitSet_9)
        {
            mapIQ of 0 and 6 to Ant1 //rest blank
        }
    }
    Break;
Case 1:
    if( vShift == 0 )
    {
        if (bitSet_0 && bitSet_6)
        {
            mapIQ of 0 and 6 to Ant2
        }
        if (bitSet_3 && bitSet_9)
        {
            mapIQ of 3 and 9 to Ant3
        }
    }
    Break;
Case 4 || case 11:
    if( vShift == 0 )
    {
        if (bitSet_0 && bitSet_6)
        {
            mapIQ of 0 and 6 to Ant1
        }
        if (bitSet_3 && bitSet_9)
        {
```



```

        mapIQ of 0 and 6 to Ant0
    }
}
Break;
Case 8:
if( vShift == 0 )
{
    if (bitSet_0 && bitSet_6)
    {
        mapIQ of 0 and 6 to Ant3
    }
    if (bitSet_3 && bitSet_9)
    {
        mapIQ of 3 and 9 to Ant2
    }
}
Break;
}
}

```

I.6 First data layer and non-first data layer association using static eAxC_ID and layerID mapping

This clause shows an example with 3 UEs using LTE TM4, each scheduled with different number of layers.. The example shows a specific order of C-Plane messages, where UE with maximum number of scheduled layers is sent first, in practical cases this order can be implementation specific and subject to reordering. The example is described in more details in Table I.6-1.

Table I.6-1: C-Plane messages order

UE1 scheduled with numLayers=4, TM4		UE2 scheduled with numLayers=1, TM4	UE3 scheduled with numLayers=2, TM4
C-Plane Section + long-format SE 3 <ul style="list-style-type: none"> layerID = 0 → eAxC_ID: 10 C-Plane Section + short-format SE 3 <ul style="list-style-type: none"> layerID = 1 → eAxC_ID: 20 C-Plane Section + short-format SE 3 <ul style="list-style-type: none"> layerID = 2 → eAxC_ID: 30 C-Plane Section + short-format SE 3 <ul style="list-style-type: none"> layerID = 3 → eAxC_ID: 40 		C-Plane Section + long-format SE 3 <ul style="list-style-type: none"> layerID=0 → eAxC_ID: 10 	C-Plane Section+long-format-SE 3 <ul style="list-style-type: none"> layerID =0→ eAxC_ID: 10 C-Plane Section+short-format SE 3 <ul style="list-style-type: none"> layerID =1→ eAxC_ID: 20
eAxC_ID: 10	layerID 0	Re-use the mapping created for UE1	Re-use the mapping created for UE1
eAxC_ID: 20	layerID_1		
eAxC_ID: 30	layerID_2		
eAxC_ID: 40	layerID_3		

Annex J (informative): Beamforming methods description

J.1 Introduction

Beams are RF energy directed in specific angular directions in space. Beamforming can generate energy lobes (maxima, or "peaks") and/or energy nulls (minima, or "valleys") in the spatial dimension. They can be formed using various methods in the analog domain, the digital domain, or a combination thereof. Beams and beamforming can further be used to re-utilize temporo-spectral (Time-Frequency) resources to achieve Spatial Multiplexing.

The present document has four distinct methods supported at the O-RU for beamforming.

J.2 Predefined-beam beamforming

In this method, beam indices ("beamId" values) are conveyed from the O-DU to the O-RU to indicate which beam to apply to the DL or UL data. The beams are expected to be pre-defined within the O-RU and the method of beamforming is not specified. However, the O-DU still needs to know whether the beamforming is frequency-domain, time-domain or "hybrid" so the beamId values can be correctly applied e.g., to not apply different time-domain beamIds to subcarriers in the same OFDM symbol.

J.3 Weight-based dynamic beamforming

In this method, beamforming weights are transmitted across the interface using C-Plane messages.

For frequency-domain beamforming the operation may be considered as follows:

$X = [x_0 \ x_1 \ \dots \ x_{P-1}]$ represents the frequency domain IQ data in one data section, where P is the number of REs in the data section and x_i is an RE, $x_i \in \mathbb{C}^{L \times 1}$, L is the number of streams.

$W = [w_0 \ w_1 \ \dots \ w_{P-1}]$ represents the beamforming weights associated with the tx-array, where w_i is a beamforming weight of RE x_i , $w_i \in \mathbb{C}^{K \times L}$, K is the number of array elements.

NOTE: O-DU may send W per PRB or PRB bundle, and O-RU may apply W per PRB or PRB bundle.

The output after beamforming $Y = [y_0 \ y_1 \ \dots \ y_{P-1}]$ is given by

$$y_i = w_i x_i$$

where y_i is a beamforming output for the RE x_i , $y_i \in \mathbb{C}^{K \times 1}$

The equation above refers to DL beamforming, but the same principle is applied for UL as well.

Time-domain beamforming operation is similar to frequency-domain beamforming, however, instead of being applied per data section the beamforming is applied for the entire OFDM symbol:

$X = [x_0 \ x_1 \ \dots \ x_{P-1}]$ represents the frequency domain IQ data in one OFDM symbol, where P is the number of REs in the OFDM symbol and x_i is an RE, $x_i \in \mathbb{C}^{L \times 1}$, L is the number of streams.

$W = [w_0 \ w_1 \ \dots \ w_{P-1}]$ represents the beamforming weights associated with one tx-array, where w_i is a beamforming weight of RE x_i , $w_i \in \mathbb{C}^{K \times L}$, K is the number of array elements.

The output after beamforming $Y = [y_0 \ y_1 \ \dots \ y_{P-1}]$ is given by

$$y_i = w_i x_i$$

where y_i is a beamforming output for the RE x_i , $y_i \in \mathbb{C}^{K \times 1}$

The equation above refers to DL beamforming, but the same principle is applied for UL as well.

Hybrid beamforming operation involves the multiplication by frequency-domain weights and time-domain weights and may be considered as follows:

There are K number of time-domain weights (that is the same as the number of array elements) and K' number of frequency-domain weights ($K' \leq K$, see clause 12.5.3). The K' frequency-domain weights are applied for the REs in a data section, and the K time-domain weights are applied for all the REs in the entire OFDM symbol.

$X = [x_0 \ x_1 \ \dots \ x_{P-1}]$ represents the frequency domain IQ data in one OFDM symbol, where P is the number of REs in the OFDM symbol and x_i is an RE, $x_i \in \mathbb{C}^{L \times 1}$, L is the number of streams.

$W = [w_0 \ w_1 \ \dots \ w_{P-1}]$ represents the beamforming weights in one OFDM symbol, where w_i is a time-domain beamforming weight of RE x_i , $w_i \in \mathbb{C}^{K \times K'}$, K is the total number of array elements.

$W' = [w'_0 \ w'_1 \ \dots \ w'_{P-1}]$ represents the beamforming weights in one OFDM symbol, where w'_i is a frequency-domain beamforming weight of RE x_i , $w'_i \in \mathbb{C}^{K' \times L}$.

The output after beamforming $Y = [y_0 \ y_1 \ \dots \ y_{P-1}]$ is given by

$$y_i = (w_i w'_i) x_i$$

where y is a beamforming output for the RE x_i , $y_i \in \mathbb{C}^{K \times 1}$

The equation above refers to DL beamforming, but the same principle is applied for UL as well.

J.4 Attribute-based dynamic beamforming

In this method, beamforming attributes and/or their indices (if already known to the O-RU) are transmitted across the interface.

Whereas a beam index provides a pointer to a beamforming vector already known to the O-RU, and beamforming weights specify an important method for how to form the beam, beamforming attributes specify the what, an inherently compact characterization of the desired beam pattern itself, to be formed directionally in space.

These beamforming attributes include:

- bfAzPt: the azimuth beamforming pointing angle in degrees
- bfZePt: the zenith beamforming pointing angle in degrees
- bfAz3dd: the azimuth beamforming 3dB down beam width in degrees
- bfZe3dd: the zenith beamforming 3dB down beam width in degrees
- bfAzSl: the azimuth beamforming sidelobe suppression value in dB
- bfZeSl: the zenith beamforming sidelobe suppression value in dB

Multiple methods of forming the beam per the beam attributes are possible and are left as O-RU implementation choices. Some potential schemes are: beamforming phased array weights (analog or digital), holographic beamforming, Butler matrices, lenses, and other known and emerging techniques. These schemes may also be hybridized with "Channel-information-based beamforming" method.

J.5 Channel-information-based beamforming

In this method, beamforming weights are calculated at the O-RU based on the channel estimates that are transmitted across the interface.

For UL:

Assuming L streams which are jointly scheduled in UL, let $X = [x_0 \ x_1 \ \dots \ x_{P-1}]$ represent the frequency domain IQ data in one data section, where P is the number of REs in the data section and x_i is an RE, $x_i \in \mathbb{C}^{K \times 1}$, a beamforming matrix $W = [w_1, \dots, w_K] \in \mathbb{C}^{L \times K}$ is applied to the frequency domain IQ data of RE x_i for L streams, K is the number of array elements, and L is the number of streams.

$H \triangleq [h_1, \dots, h_L] \in \mathbb{C}^{K \times L}$ is defined as the UL channel estimates of RE x for the L streams, and h_i is the channel matrix conveyed via the Section Type 6 message and is identified via ueId.

When the C-Plane Section Extension = 16 is used, the UL channel for each stream is fetched according to the 'ueId' and the 'antMask' fields.

Multiple methods to calculate the beamforming weights for UL are possible and are left as the O-RU implementation choices. Some potential schemes are:

- MRC: $W = H^H / \|H\|$, $\|H\|$ is the norm of H .
- Zero-forcing: $W = (H^H H)^{-1} H^H$
- Regularized zero-forcing/ MMSE: $W = (H^H H + \xi I_L)^{-1} H^H$,

ξ is the regularization parameter. C-Plane Section Extension 8 allows configuration of regularization factor per ueId in UL.

The output after beamforming $Y = [y_0 y_1 \dots y_{P-1}]$ is given by

$$y_i = Wx_i$$

where y is a beamforming output for the RE x_i , $y_i \in \mathbb{C}^{L \times 1}$.

For DL:

Assuming L layers which are jointly scheduled in DL, let $X = [x_0 x_1 \dots x_{P-1}]$ represent the frequency domain IQ data in one data section, where P is the number of REs in the data section and x_i is an RE, $x_i \in \mathbb{C}^{L \times 1}$, L is the number of layers, a beamforming matrix $W = [w_1, \dots, w_L] \in \mathbb{C}^{K \times L}$ is applied to the frequency domain IQ data of RE x for L layers, K is the number of array elements.

$H_{\text{eff}} \triangleq [h_{\text{eff},1}^T, \dots, h_{\text{eff},L}^T]^T \in \mathbb{C}^{L \times K}$ is defined as the DL channel estimates of RE x_i for the L selected layers where $h_{\text{eff},i} \in \mathbb{C}^{1 \times K}$ is the channel matrix conveyed via the Section Type 6 message and is identified via ueId.

Multiple methods to calculate the beamforming weights for DL are possible and are left as the O-RU implementation choices. Some potential schemes are:

- Zero-forcing: $W = \frac{1}{\sqrt{\Psi}} H_{\text{eff}}^H (H_{\text{eff}} H_{\text{eff}}^H)^{-1}$
- Regularized zero-forcing/ MMSE: $W = \frac{1}{\sqrt{\Psi}} H_{\text{eff}}^H (H_{\text{eff}} H_{\text{eff}}^H + \xi I_L)^{-1}$,

ξ is the regularization parameter, and the normalization parameter Ψ can be chosen to satisfy the total power constraint $\{WW^H\} \leq M$.

The output after beamforming $Y = [y_0 y_1 \dots y_{P-1}]$ is given by

$$y_i = Wx_i$$

where y is a beamforming output for the RE x_i , $y_i \in \mathbb{C}^{K \times 1}$.

C-Plane Section Extension 8 allows configuration of regularization factor per ueId in DL.

Assuming Y users' channel matrices each transmitting multiple layers' data considering transmit antenna switching (TAS), so H_{eff} can be selected by using a proper sub-spacing function $f(\cdot)$ among Y user's channel matrices H_i where $i = 1, \dots, Y$, and $f(\cdot) \in \mathbb{C}^{N_i \times K}$ where N_i is the number of antennas of i -th user. Define S_i as the number of selected layers for i -th user, and then $S = \sum_{i=1}^Y S_i$ and $H_{\text{eff}} = [H_{\text{eff},S_1}^T, \dots, H_{\text{eff},S_Y}^T]^T$ where $H_{\text{eff},S_i} = f(H_i) \in \mathbb{C}^{S_i \times K}$ and $S_i \leq N_i$.

Annex K (informative): Layers of array elements

K.1 General

Figure K.1-1 shows an example of a rectangular array with 12 rows and 4 columns of array elements. There are four data converters (not shown in Figure K.1-1). Each data converter connects to all the 48 array elements (also known as the full-connection model in 3GPP). The array element contains 4 gain and phase control elements each connecting to one of the data converters. The gain and phase control element is used to apply time domain (TD) beamforming weights. This type of connection creates four overlapping arrays ("array-layers") by reusing the same array elements.




Figure K.1-1: An example array

On the transmit side each DAC can correspond to one layer, which is then connected to one of the overlapping Tx arrays. Hence the structure of the tx-array so constructed can support 4 layers. The 48 elements corresponding to layer $q = 0 \dots 3$ receive the same input from DAC q and this input can be beamformed with a 48 element TD beamforming vector $(\theta_{q,0}, \theta_{q,47})$ and a Frequency Domain (FD) beamforming weight φ_q corresponding to layer q . Each array element hence can receive 4 equivalent beamforming weights $\varphi_0 \theta_{0,0}$ to $\varphi_3 \theta_{3,0}$ corresponding to 4 layers.

K.2 Use Case A

Use case A is to send an eAxC flow via 48 elements. In this case, the eAxC flow is beamformed in frequency domain (FD) by φ_0 and 48 complex weights corresponding to time domain (TD) using $(\theta_{0,0}, \theta_{0,47})$ as shown in Figure K.2-1. In this case, the O-RU has 4 simple tx-arrays with 48 elements each.




Figure K.2-1: Four tx-arrays with 48 elements each

In the above case the 48 array elements each have one power amplifier which is shared across the four tx-arrays, conveyed over the M-Plane as the array elements being "shared" according to clause 12.5.5.

K.3 Use Case B

Use case B is when one eAxC flow is sent over all the 4x48 elements. This can be accomplished by using 4 digital weights in FD and 4x48 TD weights (see Figure K.3-1). In this case, O-RU has one tx-array with 4x48 elements.

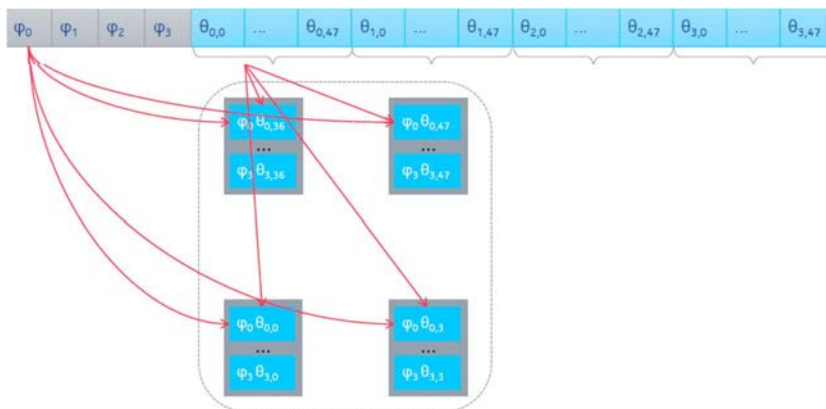


Figure K.3-1: Use case B

In the above case the 48 array elements each have one power amplifier which is shared across the four array-layers $q = 0, \dots, 3$, conveyed over the M-Plane as independent-power-budget being false according to clause 12.5.3.

Annex L (informative): Considerations when operating in non-ideal transport environments

L.1 General

The Intra-PHY lower layer fronthaul split has been optimized for operation using transport systems that are able to support tight latency and bandwidth requirements. However, this does not preclude operation of the lower layer split in environments where the transport systems are not able to meet tight latency and bandwidth requirements. This appendix is informative and describes those additional capabilities that may be included in an O-DU and/or O-RU to assist when operating in such circumstances.

Many of the following discussion points and examples assume an FDD cell configuration. FDD examples are used here as the timing aspects are easier to conceptualize. The same principles and approaches apply to TDD cells.

L.2 Deployment scenarios and performance considerations

L.2.1 Low impact

There is a class of low mobility deployment scenarios characterized by slow fading, which provide a relatively stable radio channel at small time scales. When combined with low UE density, these scenarios see low (or no) performance impact from non-ideal fronthaul vRAN deployments.

Example deployment use cases include indoor femto cells, fixed wireless access, and outdoor pico cells in non-dense urban settings.

In such scenarios the benefit from HARQ is negligible, thus the performance impact of disabling HARQ retransmissions or using predictive HARQ is negligible. Further, the impact of a diminished UE attach ramp rate is negligible as well because UE ramp rate is not a performance design attribute for such deployments.

L.2.2 Medium impact

There is a second class of low mobility deployment scenarios where again, radio channels are predominately slow fading, but where UE density may be higher. In these scenarios, the UE attach ramp rate may be a performance attribute of higher importance and thus a larger impact from non-ideal fronthaul may be expected.

Example deployment use cases include venue deployments (e.g., stadiums or arenas), shopping malls, or airport terminals.

In such scenarios inter-cell interference may become significant and the benefit from HARQ may be greater and thus a noticeable degradation of system throughput would be observed when using non-ideal fronthaul. In addition, delays in NAS attach may become noticeable to users depending on the user behavior given the increased UE density.

L.2.3 High impact

There also exists a class of deployments which are characterized by high mobility and/or fast fading. In these scenarios the performance impacts of using non-ideal fronthaul may be significant.

Example deployment use cases include freeway coverage macro cells, dense urban macro or small cells, or rural large ISD macros.

In these scenarios, HARQ often provides a benefit and thus prediction algorithms will likely result in a noticeable performance impact. In addition, the impact of delayed CSI feedback will further reduce system efficiency. Similarly, depending on the UE behavior (i.e., traffic movement) UE attach ramp rate may be a metric of higher importance.

L.3 HARQ

L.3.1 Introduction

Long latency fronthaul links may break traditional Hybrid ARQ timelines. The following clauses discuss the timelines in detail and cover techniques which can be used to minimize the impact of the timeline being broken.

L.3.2 Synchronous HARQ

Since FDD LTE offers 8 HARQ processes for each UE, one UE can be scheduled for PDSCH traffic in up to 8 different subframes simply by associating a different HARQ process to each subframe. Hence, the most demanding scenario in terms of processing timeline is given by a "full-buffer" UE continuously scheduled in consecutive subframes, which results in 8 subframes (i.e., 8 ms) of timeline budget for each HARQ process. Since the standard mandates that each ACK/NACK HARQ response is sent in uplink 4 subframes after the relevant downlink transmission, effectively only 4 ms of timeline budget are left for eNB-side operations.

The eNB can be expected to complete all tasks within the 4-ms timeline budget, with many eNB implementations taking more than 3 ms in the worst load conditions. It is this argument that has generated the common perception according to which all sources of "extra" latency cannot exceed 0.5-1.0 ms, see Figure L.3.2-1.

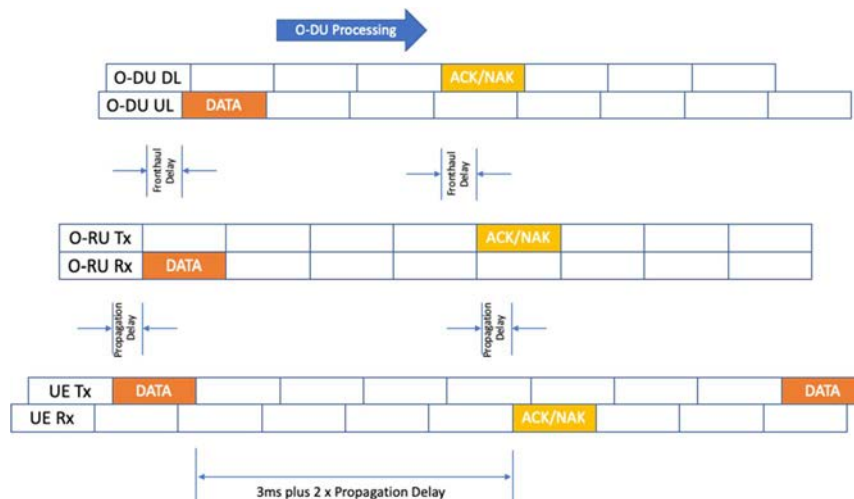


Figure L.3.2-1: Synchronous Up-link HARQ Processing Budget

In the strictest interpretation of this argument, the sub-ms upper bound includes occasional, short-term spikes, which effectively rules out virtual implementations of LTE L1 on commodity operative systems (OSs) and forces adoption of real-time OSs. Even in the broader interpretation, this argument rules out disruptive deployments such as split-7 architectures described in 3GPP TR 38.801 [13], clause 11.1.2.7, over high-latency fronthaul.

The following analysis challenges these often-cited assumptions, showing that, while they make sense for many legacy deployments, different scenarios exist in which forcing tight latency constraints is unnecessary and may be counterproductive to the exploration of innovative low-cost deployment scenarios.

L.3.3 HARQ interleaving

One option often mentioned in the context of latency-resistant LTE implementations is "HARQ (process) interleaving", which in essence requires waiting for the ACK/NACK response to be available before reusing a certain HARQ process, see Figure L.3.3-1.

Using this approach, the tolerated latency for O-DU processing time plus round trip fronthaul delay can be increased to $3 + n8$ milli-seconds, where n is number of autonomous HARQ responses sent with the New Data Indication (NDI) not toggled, ensuring that the UE re-transmits the data until it receives ACK or NACK from the O-DU.

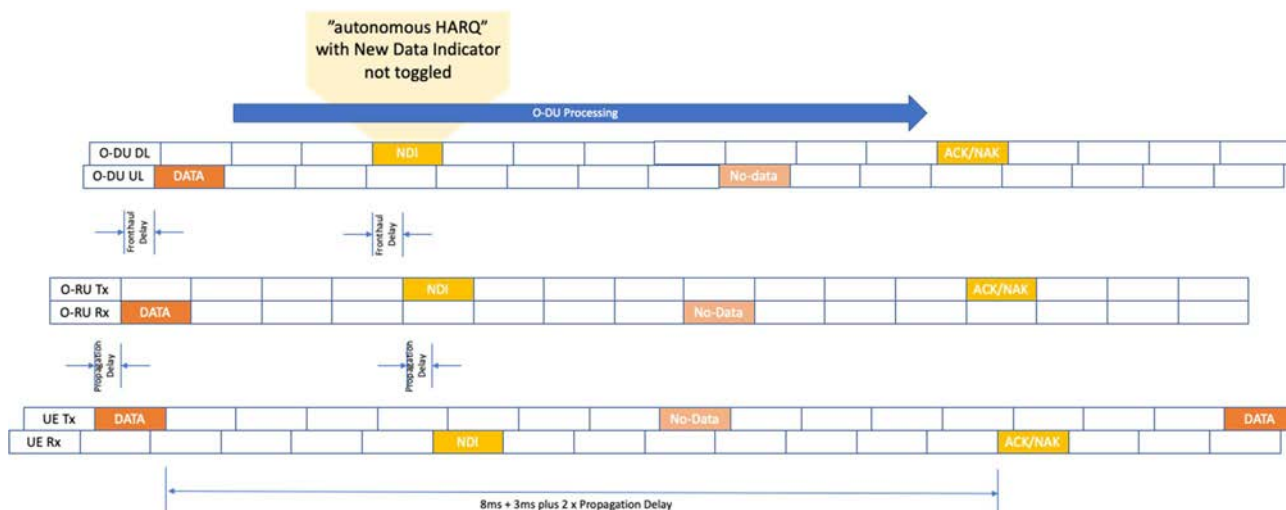


Figure L.3.3-1: Interleaved HARQ operation

In the worst case in terms of timeline budget -- that of a single, full-buffer UE -- this option is very inefficient, resulting in only 8 every (8+n) subframes being used when the overall latency exceeds the 8-ms budget by n ms. For example, 8 ms of extra latency would lead to about 50% of the subframes being unused. Hence, while possibly appealing in deployments with latency typically below the desired limit except for a few occasional spikes, this option is certainly not viable for deployments in which high latency is the norm, such as split-7 architectures over high-latency fronthaul.

L.3.4 Predictive HARQ

Another option consists of, effectively, turning off HARQ, relying on upper layers for retransmissions (e.g., RLC-AM, or transport/application layer for RLC-TM and RLC-UM). In this technique, the eNB schedules each user as it would if it had received a timely ACK response for each PDSCH transmission, which results in no subframe left unnecessarily unused. see Figure L.3.4-1.

When the late HARQ response is actually received, two cases are possible. In one case, when the response is ACK, nothing else needs to be done since "prediction of ACK" was correct -- the eNB did everything correctly, without wasting time waiting for the actual response.

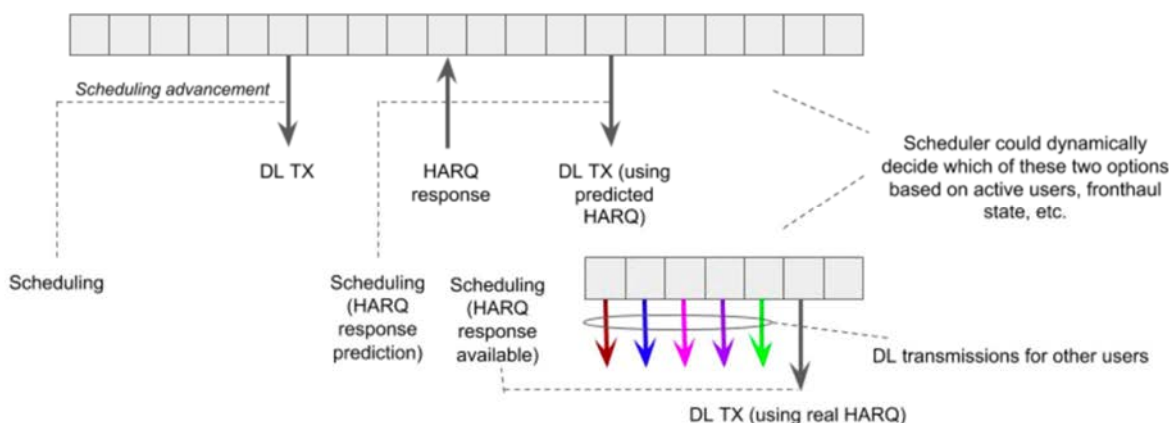


Figure L.3.4-1: DL Predictive HARQ operation

In the other case, when the response is NACK and the prediction was wrong, nothing else needs to be done, since this event is equivalent to an ACK-for-NACK decoding error, hence equivalent to an event that is typically caused by radio problems (e.g., fading, interference) against which LTE already has built-in robustness.

This simple technique performs well when NACK events are rare, most notably in peak-rate conditions, when NACK events are essentially absent even at maximum MCS. With proper tuning of the link adaptation parameters (e.g., lower target block error rate), it is possible to achieve full-buffer throughput performance within 10-15% of what an optimal, low-latency eNB can achieve in any given static/pedestrian radio conditions.

It is evident that, for these low-mobility scenarios, ruling out high-latency deployments is unnecessary: Many use cases exist in which a 10-15% throughput degradation does not justify more expensive infrastructures.

More advanced algorithms may be required for optimization of scenarios in which full-buffer throughput is not the key metric (e.g., RRC messages not protected by RLC-AM). For example, "prediction of NACK" may be used until an ACK response is actually received for a certain critical message.

While discussions of advanced techniques are beyond the scope of this clause -- and the scope of fronthaul protocol specifications -- it is critical to remark that innovative eNB vendors have the opportunity to unlock new use cases, using the longer latency O-DU and O-RU categories.

L.3.5 Asynchronous HARQ considerations

In the LTE down-link and in both 5G down-link and up-link, HARQ has been defined to be asynchronous. This means that a separate HARQ process identifier is included in the HARQ messages, thus avoiding the processing time limitations associated with synchronous up-link LTE operation.

L.4 RACH considerations

L.4.1 Introduction

The LTE attach procedure starts with UEs performing a preamble transmission on the random access channel to identifying itself to the network. This call flow is usually described as "message 1" through "message 4" as:

- 1) RACH preamble from UE
- 2) Random Access Response (RAR) from MAC
- 3) PUSCH message from UE containing CCCH or MAC signaling
- 4) Contention resolution Identity MAC CE for contention-based RACH procedure

Following are the timers used in RACH procedure (specified as per 3GPP TS 36.331 [i.3]) apart from the PRACH configuration.

```

ra-ResponseWindowSize      ENUMERATED {
    sf2, sf3, sf4, sf5, sf6, sf7, sf8, sf10},
mac-ContentionResolutionTimer  ENUMERATED {
    sf8, sf16, sf24, sf32, sf40, sf48, sf56, sf64}

```

PRACH configuration allows the RACH occasions for UEs to be available in every UL subframe to every 20 UL subframes.

Even if the above timers are set to maximum value, timers would pose an issue with RACH handling for long latency fronthaul transport links. In particular, the ra-ResponseWindowSize timer maximum value is 10ms (with timer starting 3 subframes after RACH transmission) from RACH preamble transmission occasion, see Figure L.4.1-1. During this time the vRAN system needs to handle the steps including

- 1) Complete L1 processing of RACH detection,
- 2) Sending decoded RACH message to MAC (including fronthaul transport latency)
- 3) MAC processing to allocate DL (msg2, RAR) and UL (msg3 grant) resources
- 4) Relaying of messages back to L1 (including fronthaul transport latency)

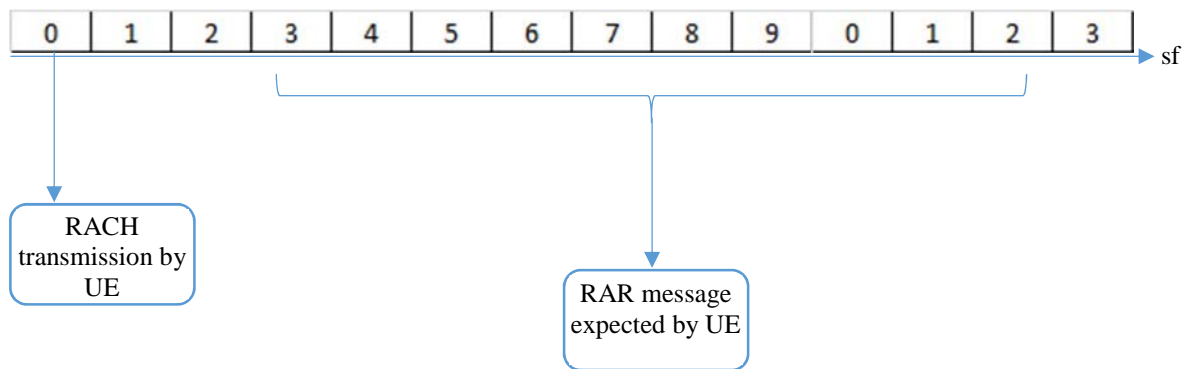


Figure L.4.1-1: Traditional RACH timing diagram

Therefore, in this process, there will be a 2-way delay (first in UL for RACH reception at MAC and second for sending RAR from MAC to L1) apart from the L1/MAC processing time. Without adjust to accommodate long latency fronthaul, this would lead to RACH procedure failure and UE will reattempt the RACH again (and would lead to same problem again) leading to UE not able to access the network.

L.4.2 Non-Ideal fronthaul RACH designs

L.4.2.1 Assumptions

When designing the RACH process for non-ideal, long latency fronthaul the following points are considered:

- 1) The method should work assuming existing 3GPP specification
- 2) There should not be any impact to UE and 3GPP compliant UE should work with the method
- 3) Currently PRACH Format-0 is considered in RACH analysis to consider FH latency requirements. Other preamble formats are not considered in this analysis.
- 4) Current clause considers the msg2 handling to meet the FH latency requirements
- 5) Contention Based RACH procedure is considered in the present document.
- 6) BI (Backoff indicator) not considered
- 7) HARQ timings for Msg3 and msg4 HARQ are considered in clause L.3.
- 8) FDD Duplexing mode is considered

With high fronthaul latency, alternative options are required to handle the RACH procedure. Two such options are discussed below.

L.4.2.2 Option 1: Semi-persistent resource allocation

Option 1 presents a semi-persistent allocation of resources that allows the O-DU to make RACH decisions without waiting for the reception of a RACH preamble. In this alternative, MAC allocates three types of resource for RAR messages:

- 1) PDCCH resources for RAR:
 - a. RA-RNTI (0 – 9) being fixed and one-to-one mapping for RACH occasions in FDD mode
 - b. Common search space resources for PDCCH carrying RAR (corresponding RA-RNTI)
 - c. Based on configured RACH configuration (RACH occasions), MAC reserves the PDCCH resources corresponding to RAR message in all required TTIs (For example, n+3rd TTI from every RACH preamble occasion).

- d. The reservation of resources will be done at MAC regardless of the RACH triggered by UE
 - e. DCI Format 1A
- 2) PDSCH resources for RAR:
 - a. Configured number of RACH to be entertained by MAC in one subframe
 - b. Based on "Number of RACH" to be processed, MAC estimates the size of RAR PDU
 - c. MAC reserves the PDSCH resources in all required TTIs similar to PDCCH
 - d. Layer-1 updates the TC-RNTI (A pool of TC-RNTI available with L1 for CBR RACH) in RAR along with RAP-ID (Preamble ID) and TA value in RAR message (decoding of MAC PDU sent by MAC)
 - e. Layer-1 maintains the TC-RNTI pool
 - 3) PUSCH resources for Msg3:
 - a. Based on "Number of RACH" to be processed in a subframe, MAC reserves the PUSCH resources
 - 4) All the above information (reserved by MAC) to L1 at cell configuration time. MAC will directly receive the PUSCH (Rx_PUSCH.Indication) with TCRNTI (and corresponding subframe number) and TA value applied in RAR
 - 5) Considering maximum value of contention resolution timer, MAC will then normally schedule msg4.
 - 6) After successful msg4 procedure, MAC will start using C-RNTI
 - 7) In case of contention failure or UE release, MAC will inform Layer-1 to free the TC-RNTI

Disadvantage: This would waste PDCCH/PDSCH/PUSCH resources if there is no RACH procedure triggered. In addition, hard coding of the timing advance is an acceptable approach for small cell deployments, but in macro cells this will not work.

L.4.2.3 Option 2: RACH retransmission estimation

The following description addresses the flow driving Msg2 scheduling in a basic configuration (more complex flows also exist).

For simplicity, the description refers to an FDD LTE deployment with RACH configured with the greatest possible values for the maximum number of preamble retransmissions (50), the preamble periodicity (20 subframes) and the RAR window size (10 subframes).

- The O-DU stack detects that a RACH preamble was transmitted at UL subframe t .
- If the scheduling in advance (which is a function of the fronthaul latency) is such that the relevant message 2 can be transmitted over the air at DL subframe $t+12$ or earlier, the "regular" flow is possible. In particular, a message 2 with content corresponding to the preamble detected in UL subframe t is transmitted by DL subframe $t+12$.
- Else, the relevant message 2 cannot be delivered within the valid RAR window, in which case the terminal that transmitted a preamble at UL subframe t is mandated to transmit another preamble at UL subframe $t+20$, for which the valid RAR window ends at DL subframe $t+32$. In this case, O-DU skips the scheduling of the message 2 corresponding to the preamble transmitted at UL subframe t (since it would not be received within the valid RAR window) and directly schedules for DL subframe $t+32$ the message 2 corresponding to the preamble transmitted at UL subframe $t+20$, even before receiving and processing the samples for said UL subframe.

The flow described above assumes that the "effective fronthaul round-trip delay" is lower than 32 ms, that is, assumes that it is possible for O-DU to process the samples relative to UL subframe t and to compute the relevant response, and for the O-RU(s) to transmit said response at DL subframe $t+32$.

A simple extension of this flow makes it viable to complete the RACH procedure in the presence of an "effective fronthaul round-trip delay" as high as 52 ms, or even 72 ms, but deployments over such high-latency fronthauls are not recommended.

Some details are omitted in the description above to maintain the flow presentation as simple as possible. In practice, the latency-resistant RACH procedure is made slightly more complex by the fact that the terminal chooses the preamble ID randomly at each retransmission, and that the content of message 2 has to include the preamble ID of the last transmission.

In the example above, at the time of scheduling message 2 for DL subframe $t+32$ (which refers to the preamble transmitted in UL subframe $t+20$), O-DU doesn't know what preamble ID to include in that message, since it has not processed yet the samples relative to said preamble transmission. The simplest method to this is to populate message 2 with the last seen preamble, see Figure L.4.2.3-1.

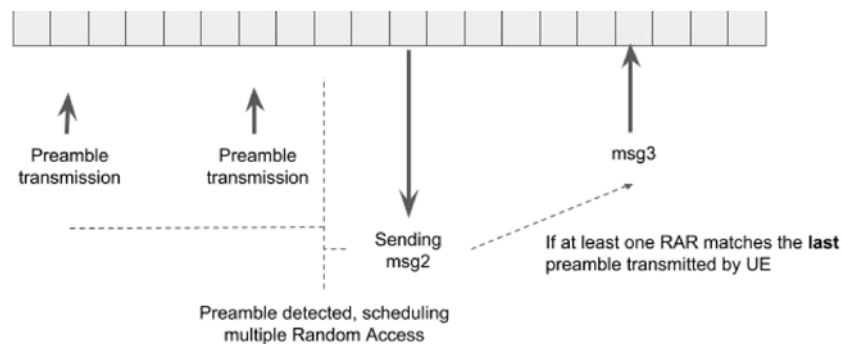


Figure L.4.2.3-1: Long latency RAR "collision" operation

By limiting the number of available preambles, the likelihood two subsequent UE RACH attempts using the same preamble can be increased sufficiently to ensure successful RACH in a timely manner.

L.5 Other latency related considerations

Supporting high latency requires careful system design to ensure seamless device connectivity while maintaining spectrum efficiency. In this clause, other system aspects that are impacted by high fronthaul latency are discussed. Some of the issues can be resolved with upper layer protocols while some may not be relevant for certain deployment use cases. Proposed solutions for these issues are for future study to see if they require any additional information to be transmitted over the fronthaul interface:

- Link adaptation and scheduling:
 - Link adaptation is implemented at layer 2, where the eNB adapts in real-time the UEs' MCS based on the CQI. Scheduling (e.g., proportional fair scheduling) depends heavily on the channel quality between the eNB and the UE. Although this depends on the channel model, it has to be done on a relatively short time-scale. Link adaptation and scheduling have to be conservative for non-ideal fronthaul with split 7 since the channel can change significantly in the time period. This may reduce the benefit of frequency selective scheduling and high mobility support since it takes longer for the scheduler to respond to the channel variations under fading conditions. However, this may not be an issue for low mobility scenarios in a small cell environment.
- UL power control:
 - Another challenging issue for non-ideal fronthaul is UL power control. This also becomes challenging under high mobility and large latencies since it takes longer time to feedback this information to the scheduler to adjust the UE transmit power. However, for use cases with limited mobility, this may not be an issue.
- Beamforming and MU-MIMO:
 - Beamforming is becoming an important use case for interference mitigation and for multi-user MIMO support for LTE and NR for macros. This requires close interaction with the scheduler in order to

suppress interference and pair users. If there is significant latency on the fronthaul, the beamforming and multi-user MIMO support becomes challenging under non-ideal conditions. This can be explored further. However, there are use cases with 2TRX/4TRX radios where beamforming support is not critical.

- COMP and advanced receiver support:
 - One of the main advantages of split 7 is centralized processing for features such as UL COMP and advanced receivers since signals from multiple base-stations can be combined and processed jointly. However, under split 7 non-ideal fronthaul, these gains are not straightforward to attain in general since the latencies on the different links can have significant variations. Hence, innovation is needed to help achieve these gains. However, there can be mechanisms to exploit these gains via other means (e.g. solutions developed for inter-site COMP which can accommodate more latency). This is also not a requirement for all use cases.
- UE attach latency:
 - The UE attach process involves 5 steps of message exchanges between the eNB and the UE. While there are messages also dependent on the core network and RRC latency, most of these messages are handled by the MAC/PHY layers and this has an impact on UE attach time, specially under large latencies. Again, this may not be relevant for low mobility use cases such as small cells.

To summarize, there are several other factors to consider for high latency support. However, such factors need not be relevant for all use cases and can be explored further in future versions.

L.6 Bandwidth limitation considerations

In addition to accommodating extra latency in the fronthaul transport link, implementations may also support variable bandwidth as another aspect of non-ideal fronthaul transport. At a high level, two bandwidth related scenarios can exist; peak rate may be limited, or bandwidth resources may be shared and thus fluctuate over time. In either case, fronthaul implementations can be designed to handle these conditions.

To support limited or variable bandwidth transport links, the O-DU will need to produce estimates of the uplink and downlink available bandwidth at periodic intervals e.g., on a per slot or per TTI interval. The O-DU can then use these bandwidth estimates to inform the scheduling decisions.

In particular a bandwidth estimation function can simply return available bandwidth values in terms of bits per second in the next time interval. Alternatively, available bandwidth could also be expressed as a function of the following parameters including but not limited to:

- Real-time IQ sample variable-bit-width
- IQ Compression
- Variable bit width per channel
- MCS / Constellation of data
- Beam count / compression
- Beam forming related signaling overhead
- C-Plane signaling overhead (reMask, symInc support, etc)
- Number of PRBs scheduled

An example API implementation may take as inputs some or all of the above parameters and return simply *numPrbu* available in the next interval.

The design goal in such implementations is to keep the cell up and active but limit the air interface user plane bandwidth to match the available fronthaul resources. Fronthaul data associated with air interface broadcast and control/reference signaling occupies a small fraction of the bandwidth needed to achieve full cell throughput. An

example implementation of a fronthaul aware scheduler could prioritize broadcast and control signaling, allocating PDSCH and PUSCH resources only after fronthaul bandwidth resources for all broadcast and control signaling has been removed from the available bandwidth estimate.

The exact fronthaul aware scheduling algorithm used is out of scope for the present document.

Annex M (informative): Use case of selective transmission and reception

M.1 General

In order to clarify how to transmit/receive signal when selective transmission and reception is used, some concrete examples of selective transmission and reception with beamId are illustrated in this clause.

- 1) Selective transmission and reception with non beamforming O-RUs and 1-to-1 mapping of global beamId and local beamId
- 2) Selective transmission and reception with beamforming O-RUs and 1-to-N mapping of global beamId and local beamId

M.2 Selective transmission and reception with non beamforming O-RUs and 1-to-1 mapping of global beamId and local beamId

In this clause, it is assumed that the total number of beams is seven. Figure M.2-1 shows comparison of C/U-Plane processing and RF signal transmission/reception procedure between normal beamforming O-RU case with seven beam and selective transmission and reception with seven non beamforming O-RUs case.

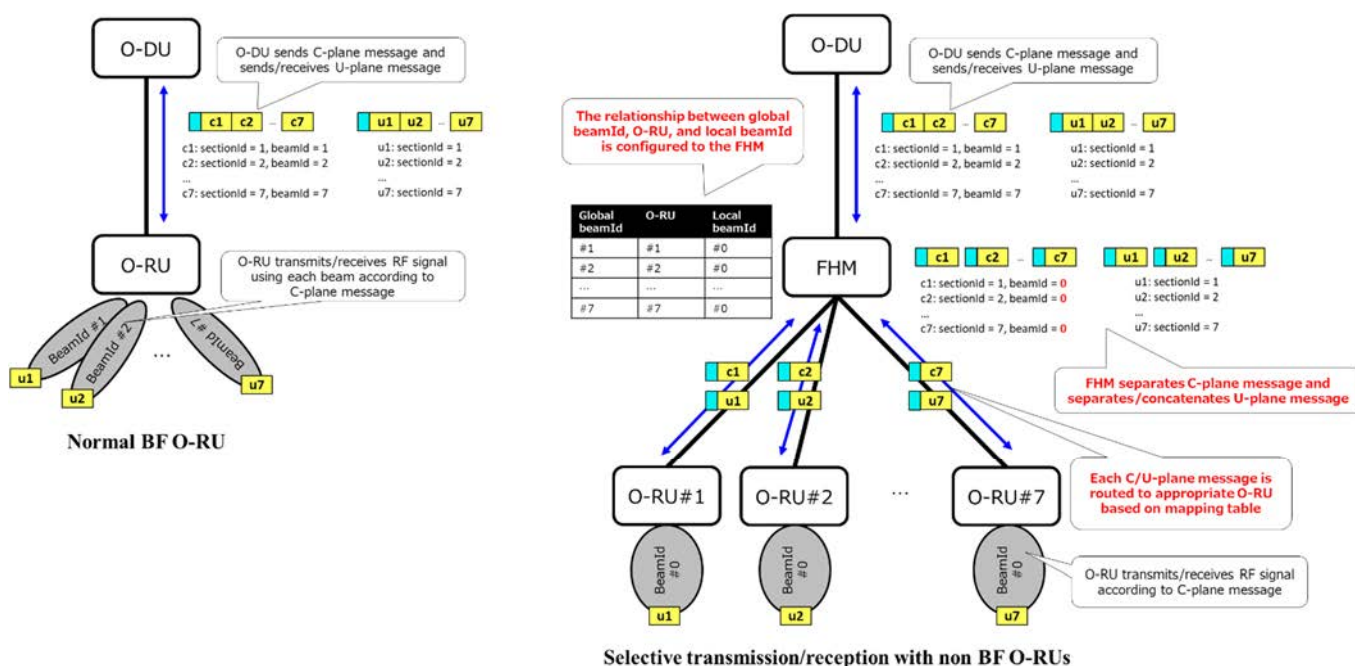


Figure M.2-1: C/U-Plane processing and RF signal transmission/reception procedure comparison (non BF O-RU case)

First of all, the mapping table between global beamId/O-RU/local beamId is configured via M-Plane during start-up procedure. Table M.2-1 shows the mapping table for this example scenario. The mapping table is assumed to be created by operator based on what an operator would like to operate.

Table M.2-1: Global beamId/O-RU/local beamId mapping table

Global beamId	O-RU	Local beamId
Beam ID #1	O-RU #1	Beam ID #0
Beam ID #2	O-RU #2	Beam ID #0
Beam ID #3	O-RU #3	Beam ID #0
Beam ID #4	O-RU #4	Beam ID #0
Beam ID #5	O-RU #5	Beam ID #0
Beam ID #6	O-RU #6	Beam ID #0
Beam ID #7	O-RU #7	Beam ID #0

When SSB is transmitted, O-DU sends C-Plane message with beamId corresponds to a transmitting SSB before U-Plane transmission. FHM reads C-Plane header and routes it to corresponding O-RU(s). After that, U-Plane packet is routed in same manner and then O-RU transmits SSB. When the UE receives SSB, UE can obtain each SSB ID by decoding PBCH included in the SSB. If UE is before attaching, UE will transmit PRACH preamble on the PRACH occasion correspondig to the best SSB ID. Since PRACH occasion is linked to SSB ID and SSB ID is linked to global beamId one by one, O-DU knows which beam is used to receive PRACH preamble on each PRACH occasion. Therefore, O-DU can send C-Plane packet with appropriate beamId to receive PRACH preamble on each PRACH occasion. The C-Plane packet is routed to O-RU(s) by FHM and then the O-RU(s) try to receive PRACH preamble and sends U-Plane packet to FHM. After UE attaching, other shared channel and reference signal are transmitted/received as same manner.

M.3 Selective transmission and reception with beamforming O-RUs and 1-to-N mapping of global beamId and local beamId

In this clause, it is assumed that the total number of beams is seven. Figure M.3-1 shows comparison of C/U-Plane processing and RF signal transmission/reception procedure between normal beamforming O-RU case with seven beam and selective transmission and reception with two beamforming O-RUs case. In this case, each beamforming O-RU for selective transmission and reception has 4 beams. Local beamId #4 of O-RU #1 and local beamId #1 of O-RU #2 are mapped to the same global beamId #4.




Figure M.3-1: C/U-Plane processing and RF signal transmission/reception procedure comparison (BF O-RU case)

First of all, the mapping table between global beamId/O-RU/local beamId is configured via M-Plane during start-up procedure. Table M.3-1 shows the mapping table for this example scenario. The mapping table is assumed to be created by operator based on what an operator would like to operate.

Table M.3-1: Global beamId/O-RU/local beamId mapping table

Global beamId	O-RU	Local beamId
Beam ID #1	O-RU #1	Beam ID #1
Beam ID #2	O-RU #1	Beam ID #2
Beam ID #3	O-RU #1	Beam ID #3
Beam ID #4	O-RU #1	Beam ID #4
	O-RU #2	Beam ID #1
Beam ID #5	O-RU #2	Beam ID #2
Beam ID #6	O-RU #2	Beam ID #3
Beam ID #7	O-RU #2	Beam ID #4

As shown in Table M.3-1, local beamId #4 of O-RU #1 and local beamId #1 of O-RU #2 are mapped to the same global beamId #4. This means O-RU #1 and O-RU #2 partially construct Shared cell area which is covered by beam ID #4 of O-RU #1 and beam ID #1 of O-RU #2. Therefore copy and combine function is applied for the signal transmit to/receive from this area. The detailed procedure of fronthaul interface and physical layer for transmitting SSB, receiving PRACH preamble, and for transmitting/receiving other shared channel are same as described in clause M.2.

Change history

Date	Version	Information about changes
11.03.2019	01.00	Final version 01.00
02.06.2019	01.00.01	Updated with NOK-0006, NOK-0007, NOK-0008, NOK-0009, NOK-0010, NOK-0011, NOK-0012, NOK-0013, ERI-0001, ERI-0002, INT-0001, SAM-0003, SAM-0004, CAB-0001, CAB-0002 and CAB-0003, which include WI non-ideal fronthaul and several mandatory / optional modifications, as well as adding Section Extension 6 and 7.
10.06.2019	01.00.02	Accommodates review comments against v01.00.01.
20.06.2019	01.00.03	Updated with changes from Tokyo F2F meeting (see comment sheet), as well as NEC-0003.
27.06.2019	01.00.04	Updated with comments #22 - #30 in the review comment sheet as well as based on discussions at the June 27 meeting. Most changes are related to mandatory/optional statements and conditions, and in 3.1.3.1.6 regarding endpoint use.
02.08.2019	02.00	Updated with CR FJT-0001 and NOK-0015 and NOK-0016 and comments #31 - #34 which are the final comments; updated version to v02.00 and created "tc" and clean versions for WG4 approval.
14.02.2020	02.00.01	Updated using the following CRs: NOK-0014, NOK-0017, NOK-0018, NOK-0019, NOK-0023, NOK-0024, INT-0002, INT-0003, INT-0004, MAV-0001, MAV-0002, MAV-0003, MAV-0004, ERI-0003, ERI-0006, ETR-0001, ETR-0002, ETR-0003 and MCS-0001, as well as made minor grammatical corrections.
21.02.2020	02.00.02	Updated using the following CRs: DCM-0002, DCM-0003, NOK-0020, NOK-0026, NOK-0029, ERI-0004, ETR-0006, SAM-0012, SAM-0013 and two comments received from Ericsson.
28.02.2020	02.00.03	Updated using the following CRs: ERI-0007, COM-0001, COM-0002 and SAM-0014 ... as well as comments from the comment review sheet.
10.03.2020	02.00.04	Updated using the following CRs: NOK-0021, MAV-0005, ERI-0009, SAM-0016, SAM-020, SAM-0021, SAM-0022 plus all comments from the comment sheet up to comment #42
26.03.2020	03.00	Update from comments #43 - #55, accept all changes and fix figure and table numbering and table of contents as needed. Document approved by WG4.
19.06.2020	03.00.01	Updated doc version and this table (comment #1), included CRs: PIC-0001, KEY-0002, ATT-0001, INT-0005, QCM-0001, NOK-0027, NOK-0028, NOK-0031.
26.06.2020	03.00.02	Updated doc version and this table (comment #1), including CRs PIC-0003, NOK-0033, NOK-0035 and ERI-0010. Also included review comments 2-5 (all that were received).
03.07.2020	03.00.03	Updated this table including CRs NOK-0025, NOK-0032, NOK-0036, MAV-0006, MAV-0007, SAM-0024, ERI-0011 & INT-0006. Also included review comments 6-55.
30.07.2020	04.00	Updated this table including review comments 56-77; accept all changes and update table of contents as needed; review is complete. Review comments 78-108, all typographical or formatting corrections, were received after the vote to approve was started and then approved for inclusion by WG4.
15.10.2020	04.00.01	Updated doc version and this table (comment #0), included CRs: NOK-0034, NOK-0038, MAV-0009, NOK-0041, NOK-0046, QCM-0002, fixed some grammatical errors.
23.10.2020	04.00.02	Updated doc version and this table (comment #0), included CRs: NOK-0042, SAM-0035, SAM-0036, SAM-0037, SAM-0039, ERI-0012, KEY-0002, CIS-0018.
29.10.2020	04.00.03	Updated doc version and this table. Added CRs: NOK-0047, ALS-0001, QCM-0003, DCM-0005, DCM-0008, CUC-0001, NEC-0005, MAV-0014.
19.11.2020	05.00	Included review comments 4-122; accept all changes and formatting as needed. Document approved by WG4.
19.02.2021	05.00.01	Updated doc version and this table (review comment #0). Included review comment 1-14; included CRs: NOK-0048, NOK-0049, NOK-0054, MAV-0015, ERI-0014, ERI-0015, QCM-0005, INT-0008, SAM-0045
26.02.2021	05.00.02	Updated doc version and this table (review comment #18). Included CRs: KEY-0012, ERI-0018, ERI-0017, QCM-0006, MAV-0016, SAM-0044, SAM-0046, DCM-0010, NOK-0053, NOK-0056.
01.03.2021	05.00.03	Updated doc version and this table (review comment #29). Included CRs: ERI-0016.
12.03.2021	06.00	Included review comments 19-62; accept all changes and formatting as needed. Document approved by WG4.
25.06.2021	06.00.01	Updated doc version and this table (review comment #0). Added CRs: ADI-0001, MAV-0018, MAV-0019, ERI-0020, QCM-0007, INT-0012, DCM-0011, SAM-0049, SAM-0050, NOK-0060, NOK-0064, NOK-0062, NOK-0071, QCM-0008, NOK-0069, INT-0010, INT-0011, CAL-001, ALS-0002, NOK-0070, QCM-0011, SAM-0053, SAM-0054, SAM-0058, ETR-0013, MAV-0023, ERI-0031.
28.06.2021	06.00.02	Updated doc version and this table (review comment #2). Included review comment 1; included CRs: NOK-0063, QCM-0012.

Date	Version	Information about changes
02.07.2021	06.00.03	Included review comments 3-146.
20.07.2021	07.00	Included review comments 147-153; accept all changes and formatting as needed. Document approved by WG4.
18.10.2021	07.00.01	Created initial version, with change-marks from previous content. Implemented CRs: NOK-0085, NOK-0086.
27.10.2021	07.00.02	Resolution of review comments: #1,3, 8-61, 63-66, 68, and 69.
28.10.2021	07. Jan	Incorporated review comments #5, 62, 67 and 71. Document approved by WG4.
29.10.2021	07.01.2001	Created initial version, with change-marks enabled from version 07.01. Implemented CRs: NOK-0055, QCM-0014, QCM-0017, NEC-0012, SAM-0051, CAL-0003, INT-0017, CAL-0002, ALS-0003, SAM-0056, MAV-0021, SAM-0057, QCM-0013, ERI-0034, NOK-0081, SOL-0005, CHM-003, ERI-0035.
03.11.2021	07.01.2002	Implemented CRs: SOL-0004, QCM-0016
19.11.2021	08.00	Incorporated review comments 1-93. Document approved by WG4.
14.03.2022	08.00.01	Created initial version, with change marks, from previous content. Implemented CRs: NOK-0091, QCM-0034, QCM-0033, NOK-0089, SAM-0079, SAM-0080, SAM-0081, NOK-0092, NOK-0097, DCM-0020, QCM-0035, ERI-0042, ERI-0043, INT-0020, NOK-0094, NOK-0095, DCM-0019, DCM-0018, NOK-0093, NOK-0104, MAV-0024, QCM-0015
18.03.2022	08.00.02	Incorporated CRs: QCM-0026, QCM-0027, ERI-0044, NOK-0103, MAV-0029, NOK-0106, QCM-0036, NOK-0108. Incorporated review comments 2-8, 10, 12, 14-16, 18-20.
31.03.2022	09.00	Incorporated review comments 11,13,17,21-472. Review comments 473-501 were received after the vote to approve was started and then approved for inclusion by WG4.
04.07.2022	09.00.01	Incorporated CRs: NOK-0096, ALS.AO-0004, SOL-0006, NOK-0088, SAM-0070, SAM-0082, SOL-0011, NOK-0111, NOK-0112
07.07.2022	09.00.02	Document adapted to new template. Incorporated CRs: QCM-0042, SAM-0071. Incorporated review comments #1-44, 46-79, 81-84, 88-92, 94-106
15.07.2022	09.00.03	Incorporated review comments #107-124, 126-146, 148-175
21.07.2022	09.00.04	Incorporated review comments #176-185, 187-194, 199, 201-217, 219-221
23.07.2022	10.00	Incorporated review comments #222-229.
27.10.2022	10.00.01	Created initial draft, with change marks from v10.00. Incorporated CRs: ADI-0004, INT.AO-0021, NEC-0019, NOK-0110, NOK-0117, QCM-0029, QCM-0039, QCM-0041, QCM-0043, QCM-0044, QCM-0053, SAM-0083, SAM-0089, SOL-0012
04.11.2022	10.00.02	Incorporated review comments #2-58, 60-64. Incorporated CR QCM-0031
12.11.2022	11.00	Incorporated review comments #59, 65-120, 121-136. Document approved by WG4
24.02.2023	11.00.01	Created initial draft from v11.00. Incorporated CRs: ADI-0006, ATT.AO-0022, ERI-0069, NOK-0125, NOK-0128, NOK-0129, NOK-0131, QCM-0047, SAM.AO-0085, SAM-0078, SAM-0090, SAM-0092
02.03.2023	11.00.02	Incorporated review comments #2-49
10.03.2023	11.00.03	Incorporated review comments #50-120
18.03.2023	12.00	Incorporated review comments #121-160. Document approved by WG4
14.03.2024	12.01	Incorporated CRs: NOK-0170, NOK-0171, NOK-0172. Document approved by WG4

History

Document history		
V7.0.2	September 2022	Publication
V12.0.1	April 2024	Publication