



ATIS-0700004.2007(S2017)

**HIGH CAPACITY – SPATIAL DIVISION MULTIPLE ACCESS (HC-SDMA)
RADIO INTERFACE STANDARD**

AMERICAN NATIONAL STANDARD FOR TELECOMMUNICATIONS



ATIS is the leading technical planning and standards development organization committed to the rapid development of global, market-driven standards for the information, entertainment and communications industry. More than 200 companies actively formulate standards in ATIS' Committees and Forums, covering issues including: IPTV, Cloud Services, Energy Efficiency, IP-Based and Wireless Technologies, Quality of Service, Billing and Operational Support, Emergency Services, Architectural Platforms and Emerging Networks. In addition, numerous Incubators, Focus and Exploratory Groups address evolving industry priorities including Smart Grid, Machine-to-Machine, Connected Vehicle, IP Downloadable Security, Policy Management and Network Optimization.

ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a member and major U.S. contributor to the International Telecommunication Union (ITU) Radio and Telecommunications' Sectors, and a member of the Inter-American Telecommunication Commission (CITEL). ATIS is accredited by the American National Standards Institute (ANSI). For more information, please visit < www.atis.org >.

AMERICAN NATIONAL STANDARD

Approval of an American National Standard requires review by ANSI that the requirements for due process, consensus, and other criteria for approval have been met by the standards developer.

Consensus is established when, in the judgment of the ANSI Board of Standards Review, substantial agreement has been reached by directly and materially affected interests. Substantial agreement means much more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered, and that a concerted effort be made towards their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards.

The American National Standards Institute does not develop standards and will in no circumstances give an interpretation of any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

CAUTION NOTICE: This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

Notice of Disclaimer & Limitation of Liability

The information provided in this document is directed solely to professionals who have the appropriate degree of experience to understand and interpret its contents in accordance with generally accepted engineering or other professional standards and applicable regulations. No recommendation as to products or vendors is made or should be implied.

NO REPRESENTATION OR WARRANTY IS MADE THAT THE INFORMATION IS TECHNICALLY ACCURATE OR SUFFICIENT OR CONFORMS TO ANY STATUTE, GOVERNMENTAL RULE 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. ATIS SHALL NOT BE LIABLE, BEYOND THE AMOUNT OF ANY SUM RECEIVED IN PAYMENT BY ATIS FOR THIS DOCUMENT, AND IN NO EVENT SHALL ATIS BE LIABLE FOR LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES. ATIS EXPRESSLY ADVISES THAT ANY AND ALL USE OF OR RELIANCE UPON THE INFORMATION PROVIDED IN THIS DOCUMENT IS AT THE RISK OF THE USER.

<p>NOTE - The user's attention is called to the possibility that compliance with this standard may require use of an invention covered by patent rights. By publication of this standard, no position is taken with respect to the validity of this claim or any patent rights in connection therewith. The patent holder has, however, filed a statement of willingness to grant license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license. Details may be obtained from the publisher.</p>
--

ATIS-0700004.2007(R2012), High Capacity – Spatial Division Multiple Access (HC-SDMA) Radio Interface Standard

Is an American National Standard developed by the **Radio Access Networks (RAN) Subcommittee** under the **ATIS Wireless Technologies and Systems Committee (WTSC)**.

Published by

**Alliance for Telecommunications Industry Solutions
1200 G Street, NW, Suite 500
Washington, DC 20005**

Copyright © 2012 by Alliance for Telecommunications Industry Solutions
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher. For information contact ATIS at 202.628.6380. ATIS is online at < <http://www.atis.org> >.

Printed in the United States of America.

American National Standard for Telecommunications

HIGH CAPACITY – SPATIAL DIVISION MULTIPLE ACCESS (HC-SDMA) RADIO INTERFACE STANDARD

Secretariat

Alliance for Telecommunications Industry Solutions

Approved March 23, 2007

American National Standards Institute, Inc.

Abstract

The HC-SDMA interface provides wide-area broadband wireless data-connectivity for fixed, portable, and mobile computing devices and appliances. The protocol is designed to be implemented with smart antenna array techniques to substantially improve the radio frequency (RF) coverage, capacity and performance for the system.

FOREWORD

The information contained in the Foreword is not part of this American National Standard (ANS) and has not been processed in accordance with ANSI’s requirements for an ANS. As such, this Foreword may contain material that has not been subjected to public review or a consensus process. In addition, it does not contain requirements necessary for conformance to the Standard

The Alliance for Telecommunication Industry Solutions (ATIS) serves the public through improved understanding between carriers, customers, and manufacturers. The Wireless Technology and Systems Committee’s (WTSC) Wireless Wideband Internet Access (WTSC-WWINA) develops, maintains, amends and enhances WTSC standards related to the radio and network aspects of WWINA systems optimized for internet data applications. WTSC develops and recommends positions on related subjects under consideration in other North American, regional and international standards bodies.

ANSI guidelines specify two categories of requirements: mandatory and recommendation. The mandatory requirements are designated by the word *shall* and recommendations by the word *should*. Where both a mandatory requirement and a recommendation are specified for the same criterion, the recommendation represents a goal currently identifiable as having distinct compatibility or performance advantages.

Suggestions for improvement of this document are welcome. They should be sent to the Alliance for Telecommunications Industry Solutions, WTSC Secretariat, 1200 G Street NW, Suite 500, Washington, DC 20005.

At the time it approved this document, WWINA, which is responsible for the development of HC-SDMA, had the following members:

- M. Younge, WTSC Chair
- D. Zelmer, WTSC Vice-Chair
- C. Underkoffler, ATIS Chief Editor
- R. Canchi, WTSC Technical Editor

Organization Represented	Name of Representative
Alcatel USA Inc.	Ken Biholar Laurent Cruchant (Alt.)
ArrayComm, Inc.	Joanne Wilson Mel Woinsky (Alt.)
AT&T	Randolph Wohlert Chuck Bailey
C.S.I. Telecommunications	Michael S. Newman Thomas G. Croda (Alt.)
Cingular Wireless LLC	Don Zelmer Marc Grant (Alt.)
Department of Defense	Chris Fitzgerald
Embarq	Carl M. Coopage John M. Heinz (Alt.)
Ericsson Incorporated	Jim Ragsdale Pierre Truong (Alt.)
FBI ESTS	Marybeth Paglino Edward Ignacio (Alt.)
Hewlett-Packard	Steve Mills
InterDigital Communications	Brian Kiernan Gil LaVean (Alt.)
Kineto Wireless	Mike Gallagher Rajeev Gupta (Alt.)
Kyocera International	Dr. R. Canchi Terry Lingren (Alt.)
LGE Mobile Research LLC	Soon Yil Kwon
Lucent Technologies	Robert B. Beeson Jane Brownley (Alt.)
Motorola Inc.	Syed Husain
National Communications System	Carol-Lyn Taylor Nicholas Andre (Alt.)
National Telecom & Info Admin – NTIA	Tim Riley

Organization Represented	Name of Representative
Navini Networks	David Maez Sai Subramanian (Alt.)
Nokia Telecommunications Inc.	Ed Ehrlich Shou Gung (Alt.)
Nortel	Amar Deol Ronald D. Ryan (Alt.)
PSEP Canada	Sim Simanis Gary Hutchinson (Alt.)
Qualcomm Incorporated	Mark Epstein Ed Tiedemann (Alt.)
Qwest	Michael Fargano
Research In Motion	Atul Asthana Adrian Buckley
Rogers Wireless Inc.	Edward O’Leary Tak Wing Wan
Siemens Communications, Inc.	Nagaraja Rao Xiaoyang Lee (Alt.)
SOMA Networks	Jonathan Jaeger
SS8 Networks Inc.	Cemal Dikman Scott Coleman (Alt.)
T-Mobile USA Inc.	Gary K. Jones Mark Younge (Alt.)
Telcordia Technologies	Jim Garrahan Cliff Halevi (Alt.)
Tellabs Operations, Inc.	Kevin Mooney William A. Walker (Alt.)
TruePosition, Inc.	Gustavo Pavan
Union Telephone Company	David Cotton, Jr.
VeriSign Inc.	Anthony M. Rutkowski
Verizon Communications	Wendy Pugh J. Martin Carroll (Alt.)

The Wireless Wideband Internet Access (WWINA) Working Group was responsible for the development of this document.

TABLE OF CONTENTS

FOREWORD	II
TABLE OF CONTENTS	III
TABLE OF FIGURES	XIII
TABLE OF TABLES	XV
1 SYSTEM OVERVIEW	1
1.1 OVERVIEW OF HC-SDMA	1
1.1.1 <i>Scope of This Document</i>	1
1.1.2 PURPOSE	1
1.2 ACRONYMS	2
1.3 CONVENTIONS	4
1.3.1 <i>Measurement Uncertainty</i>	4
1.3.2 <i>Terminology</i>	4
1.3.3 <i>Mathematical Expressions</i>	4
1.3.4 <i>Message Sequence Chart (MSC) notation</i>	4
1.3.5 <i>Reserved Fields</i>	5
1.4 APPLICATION OVERVIEW	5
1.4.1 <i>Services and Applications</i>	5
1.4.2 <i>Typical Deployment</i>	6
1.4.3 <i>User Data Transport</i>	7
1.5 PROTOCOL OVERVIEW	8
1.5.1 <i>Protocol Features</i>	8
1.6 PROTOCOL REFERENCE MODEL AND INTERFACES	8
1.6.1 <i>Layer 3 Protocol</i>	9
1.6.1.1 <i>L3 CM (Layer 3 Connection Management)</i>	9
1.6.1.2 <i>L3 RM (Layer 3 Registration Management)</i>	9
1.6.1.3 <i>L3 MMC (Layer 3 Mobility Management and Control)</i>	10
1.6.1.4 <i>L3 RRC (Layer 3 Radio Resource Control)</i>	10
1.6.2 <i>Layer 2 Protocol</i>	10
1.6.2.1 <i>L2 RLC (Layer 2 Radio Link Control)</i>	10
1.6.2.2 <i>L2 MAC (Layer 2 Medium Access Control)</i>	10
1.6.3 <i>Layer 1 Protocol</i>	10
1.6.4 <i>Layer 0 Protocol</i>	10
1.7 NORMATIVE REFERENCES	10
2 SPECTRAL LAYOUT TERMINOLOGY AND REQUIREMENTS	11
2.1 OVERVIEW	11
2.1.1 <i>Definitions for Basic HC-SDMA Networks</i>	11
2.1.1.1 <i>Network Related Definitions</i>	11
2.1.1.1.1 <i>Basic HC-SDMA Networks</i>	11
2.1.1.1.2 <i>Carrier Allocation</i>	11
2.1.1.1.3 <i>Frequency Origin and Carrier Numbering</i>	11
2.1.1.1.4 <i>Compound Networks</i>	11
2.1.1.2 <i>Base Station Related Definitions</i>	11
2.1.1.2.1 <i>Carrier Set</i>	11
2.1.1.2.2 <i>Relative Carrier Number and Base Carrier</i>	11
2.1.1.3 <i>UT Related Definitions</i>	12
2.1.1.3.1 <i>UT Addressable Space</i>	12
2.1.1.3.2 <i>UT Active Carrier Set</i>	12
2.1.2 <i>Requirements</i>	12
2.1.2.1 <i>Carrier Allocation</i>	12
2.1.2.2 <i>Frequency Origin and Preferred Frequencies</i>	12
2.1.2.3 <i>Carrier Set and Supported Frequencies</i>	12
2.1.2.4 <i>Broadcast Carriers</i>	12
2.1.2.5 <i>UT Operation in a Basic Network</i>	13
3 SLOT AND FRAME STRUCTURE	13
3.1 OVERVIEW	13

3.2	RF CHANNEL AND FRAME STRUCTURE.....	13
3.2.1	RF Channels.....	13
3.2.2	Frames and Slots.....	14
3.2.3	Burst Types.....	14
3.2.4	BCH/CCH Superframe.....	15
3.2.5	Base Station Synchronization and Frame Numbering.....	16
3.2.6	Base Station Color Code (BSCC).....	16
3.3	BURST FORMATS.....	16
3.3.1	Frequency Synchronization.....	16
3.3.2	Timing Synchronization.....	17
3.3.3	Broadcast Burst.....	18
3.3.4	Page Burst.....	19
3.3.5	Configuration Request Burst.....	21
3.3.6	Standard Uplink Burst (FAACH, RACH, & TCH).....	23
3.3.7	Standard Downlink Burst (CM, AA, & TCH).....	28
4	MODULATION AND CHANNEL CODING	35
4.1	OVERVIEW.....	35
4.2	STANDARD MODULATION AND CODING.....	36
4.2.1	Encryption.....	38
4.2.2	Cyclic Redundancy Check.....	38
4.2.3	Multiplexing.....	38
4.2.4	Tail Append.....	39
4.2.5	Convolutional Encoding.....	39
4.2.6	Puncturing and Repeating.....	39
4.2.7	Block Coding.....	40
4.2.7.1	Extended Hamming Code.....	40
4.2.7.2	Parity Check Code.....	41
4.2.8	Block Shaper.....	42
4.2.9	Symbol Mapping.....	43
4.2.10	Interleaving.....	44
4.2.11	Scrambling.....	45
4.2.12	$\pi/2$ Rotation and Scaling.....	47
4.3	BROADCAST CHANNEL MODULATION AND CODING.....	47
4.3.1	General Definitions.....	48
4.3.2	Timing Correction Burst.....	48
4.3.3	Broadcast Burst.....	48
4.3.4	Paging Burst.....	49
5	USER TERMINAL RADIO TRANSMISSION AND RECEPTION	51
5.1	HC-SDMA INTRODUCTION.....	51
5.1.1	Terminology.....	51
5.1.2	Measurements.....	52
5.1.3	Symbols.....	53
5.2	TIME AND FREQUENCY CHARACTERISTICS.....	54
5.2.1	RF Synchronization Characteristics.....	54
5.2.1.1	Initial Carrier Frequency Error.....	54
5.2.1.2	Carrier Frequency Error.....	55
5.2.1.3	Transmit Burst Timing Error.....	55
5.2.2	Radio Tuning Characteristics.....	55
5.2.2.1	Carrier Tuning Frequencies.....	55
5.2.2.2	Carrier Tuning Speed.....	55
5.3	TRANSMITTER CHARACTERISTICS.....	55
5.3.1	Output RF Power Characteristics.....	56
5.3.1.1	Absolute Maximum Transmit Power.....	56
5.3.1.2	Nominal Output Power.....	56
5.3.1.3	Transmit Leakage Power.....	56
5.3.2	Ramp-Up and Ramp-Down Transmit Power.....	57
5.3.3	Uplink Power Control.....	57
5.3.3.1	Open Loop Power Control Absolute Error.....	57
5.3.4	Minimum Transmit Power.....	57
5.3.4.1	Power Control Step Size.....	58
5.3.4.2	Power Control Maximum Rate.....	58
5.3.5	Transmit Modulation.....	58

5.3.5.1	Transmit Pulse Shaping Filter	58
5.3.5.2	Transmitter Modulation Classes	58
5.3.6	Modulation Accuracy	58
5.3.7	Output RF Spectrum Characteristics	59
5.3.7.1	Occupied Bandwidth	59
5.3.7.2	Adjacent Channel Power Ratio (ACPR).....	59
5.3.7.3	Out-of-Band Spurious Emissions	59
5.4	RECEIVER CHARACTERISTICS.....	60
5.4.1	Input RF Power Characteristics.....	60
5.4.1.1	Input Modulation Class	60
5.4.1.2	Reference Sensitivity Level.....	60
5.4.1.3	Maximum Non-Distorting Input Power.....	60
5.4.2	Receiver Selectivity	61
5.4.2.1	Adjacent Channel Selectivity.....	61
5.4.2.2	Out-of-Band Blocking Characteristics.....	61
5.4.3	Receiver Leakage Power.....	62
5.4.4	RF Measurement Characteristics.....	62
5.4.4.1	DSSI Estimator Accuracy.....	62
5.4.4.2	SINR Estimator Accuracy.....	62
6	BASE STATION RADIO TRANSMISSION AND RECEPTION	63
6.1	INTRODUCTION	63
6.2	CARRIER AND SPATIAL CHANNEL SUPPORT	63
6.3	TRANSMITTER CHARACTERISTICS	63
6.3.1	Output RF Power Characteristics.....	64
6.3.1.1	Rated Power Per Data Stream	64
6.3.1.2	Minimum Stream Power.....	64
6.3.1.3	Transmit Power Control Step Size	64
6.3.1.4	Downlink Power Control Maximum Rate	64
6.3.1.5	Transmit Power Control Nonlinearity.....	64
6.3.1.6	Transmit Leakage Power	64
6.3.1.7	Guard Time Transmit Power.....	65
6.3.1.8	Transmitter Power Absolute Error.....	65
6.3.1.9	Carrier Frequency Error	65
6.3.2	Transmit Modulation.....	65
6.3.2.1	Transmit Symbol Clock	65
6.3.2.2	Transmit Burst Timing Error	65
6.3.2.3	Transmit Pulse Shaping Filter	66
6.3.2.4	Phase Polarity.....	66
6.3.2.5	Ramp-Up and Ramp-Down Transmit Power.....	66
6.3.2.6	Modulation Accuracy.....	66
6.3.3	Output RF Spectrum Characteristics	66
6.3.3.1	Occupied Bandwidth.....	67
6.3.3.2	Adjacent Carrier Power.....	67
6.3.3.3	Multi-carrier Inter-modulation Products	67
6.3.3.4	Out-of-Band Spurious Emissions	67
6.4	RECEIVER CHARACTERISTICS.....	68
6.4.1	Input RF Power Characteristics.....	68
6.4.1.1	Reference Sensitivity Level.....	68
6.4.2	Adjacent Channel Selectivity.....	68
6.4.2.1	Non-Distorting Input Power.....	69
6.4.3	RF Measurement Characteristics.....	69
6.4.3.1	DSSI Estimator Accuracy.....	69
6.4.3.2	SINR Estimator Accuracy.....	69
6.5	MULTI-ANTENNA PERFORMANCE	70
6.5.1	Multi-antenna Architecture	70
6.5.2	Uplink and Downlink Multi-antenna Operation.....	71
6.5.3	Calibration Figure of Merit	72
7	L2 MAC PROTOCOL SUBLAYER SPECIFICATION	73
7.1	L2 MAC OVERVIEW	73
7.1.1	Access Overview	74
7.1.2	L2 MAC Logical Channels Overview.....	74
7.2	LAYER TO LAYER COMMUNICATION	75

7.2.1	L3 RM Communication	75
7.2.2	L3 RRC Communication.....	75
7.2.3	L3 MMC Communication.....	75
7.2.4	L2 RLC Communication.....	76
7.2.5	L1 Communication.....	76
7.3	LOGICAL CHANNELS	76
7.3.1	Conventions.....	76
7.3.1.1	Ordering of bits in message fields.....	76
7.3.1.2	Message Field to Value Conversion.....	76
7.3.2	Broadcast Channel (BCH).....	77
7.3.2.1	Basic Access Description	77
7.3.2.2	Superframe Structure	77
7.3.2.3	BCH Message Fields.....	78
7.3.3	Paging Channel (PCH).....	78
7.3.3.1	Frequency and Timeslot Hopping for Paging Messages.....	79
7.3.3.2	Base Station Page Procedures.....	79
7.3.3.3	User Terminal Page Procedures.....	81
7.3.4	Configuration Channel (CCH).....	81
7.3.4.1	Configuration Request (CR).....	82
7.3.4.2	Configuration Message (CM).....	82
7.3.5	Random Access Channel (RACH).....	84
7.3.5.1	RACH Message Definitions.....	85
7.3.5.1.1	RA Field Definitions.....	85
7.3.5.1.2	AA field definitions.....	86
7.3.5.1.2.1	Common message format and field definitions.....	86
7.3.5.1.2.2	AA subtype definitions.....	88
7.3.5.1.2.3	Subtype-dependent message formats and field definitions	89
7.3.5.1.2.3.1	Message Format for AA-cts and AA-short.....	89
7.3.5.1.2.3.2	Message Format for AA-rreg	89
7.3.5.1.2.3.3	Message Format for AA-pending	90
7.3.5.1.2.3.4	Message Format for AA-reject	90
7.3.5.1.2.3.5	Message Format for AA-invalid-id.....	90
7.3.5.1.2.3.6	Message Format for other subtypes	91
7.3.5.1.2.3.7	Field definitions for subtype-specific fields.....	91
7.3.5.1.2.4	RA to AA Mapping	93
7.3.5.2	RACH Policy and Procedures.....	93
7.3.5.2.1	Undirected RACH Retry Policy.....	94
7.3.5.2.2	Directed RACH retry policy.....	97
7.3.6	Fast Associated Control Channel (FACCH).....	98
7.3.6.1	FACCH on CM, RA, and AA Bursts	98
7.3.6.2	FACCH on TCH Bursts.....	98
7.3.6.2.1	Modulation Class	99
7.3.6.2.2	Additional Available Power	99
7.3.7	Traffic Channel (TCH).....	100
7.3.7.1	Stream Start.....	100
7.3.7.2	RMU Mux Function.....	102
7.3.7.2.1	RMU Field Definition.....	102
7.3.7.2.2	Header Field Insertion.....	102
7.3.7.2.3	UM Message Insertion	104
7.3.7.3	RMU Demux Function	106
7.3.7.4	Stream Closure.....	106
7.3.7.4.1	Stream Release Sequence.....	106
7.3.7.4.2	Graceful Stream Closure.....	107
7.3.7.4.3	Abrupt stream Closure	108
7.3.7.4.4	Stream closure parameters.....	110
7.4	COMMUNICATION ELEMENTS	110
7.5	PAGE HOPPING SEQUENCES.....	112
7.5.1	Properties.....	112
7.5.2	Page Hop Function	112
	Appendix 7A. Setting bsLoad on BCH.....	113
8	L2 RLC PROTOCOL SPECIFICATION.....	114
8.1	GENERAL.....	114
8.1.1	Definitions.....	114

8.1.1.1	Sequence Number Operands	114
8.1.2	Sublayer Overview	114
8.1.3	L2 RLC Model.....	115
8.1.3.1	AM Transmit Buffer.....	116
8.1.3.2	RMU Mux.....	116
8.1.3.3	Transmit ARQ.....	116
8.1.3.4	RMU Demux.....	117
8.1.3.5	Receive ARQ.....	117
8.1.3.6	AM Receive Buffer/Reassembly.....	117
8.2	PEER-TO-PEER COMMUNICATION ELEMENTS.....	117
8.2.1	RMU Fields and Formats.....	117
8.2.1.1	AM RMU.....	117
8.2.1.1.1	RMU Header.....	117
8.2.1.1.2	RMU Payload.....	117
8.2.2	State Variables.....	118
8.2.2.1	Transmit State Variables.....	118
8.2.2.2	Receive State Variables.....	120
8.2.2.3	Timers.....	121
8.3	ERROR HANDLING	121
8.4	LAYER-TO-LAYER COMMUNICATION	122
8.4.1	L3 RM	122
8.4.2	L2 MAC.....	122
8.5	L2 RLC PROCEDURES AND PROCESS STATES	122
8.5.1	Transmit Procedure	122
8.5.1.1	Transmit State Execution	123
8.5.1.2	mapTx Update Procedure	124
8.5.2	Receive Procedure	124
8.5.2.1	Receive Task Execution.....	124
8.5.2.2	mapTx Retransmit Update.....	126
8.5.2.3	Reset Procedure.....	126
8.5.3	Shutdown Procedure.....	126
8.5.3.1	Forced Shutdown.....	126
8.5.3.2	Graceful Shutdown.....	126
9	L3 PROTOCOL SPECIFICATION	127
9.1	OVERVIEW.....	127
9.1.1	Introduction.....	127
9.1.2	General Protocol Concepts.....	127
9.1.3	Logical Channels and Typical Ladder Diagram.....	128
9.2	LAYER 3 COMPONENTS	129
9.2.1	Connection Management.....	130
9.2.1.1	L3 CM Interface with L4.....	131
9.2.1.2	L3 CM Interface with L3 RM	131
9.2.1.3	L3 CM State Machines.....	131
9.2.1.3.1	UT L3 CM State Machine.....	131
9.2.1.3.1.1	UT L3 CM Initial State.....	132
9.2.1.3.1.2	UT L3 CM Deferred State	133
9.2.1.3.1.3	UT L3 CM Connecting State	133
9.2.1.3.1.4	UT L3 CM Connected State.....	133
9.2.1.3.1.5	UT L3 CM Closing State	134
9.2.1.3.2	BS L3 CM State Machine	134
9.2.1.3.2.1	BS L3 CM Initial State	135
9.2.1.3.2.2	BS L3 CM Unconnected State	135
9.2.1.3.2.3	BS L3 CM Connecting State.....	136
9.2.1.3.2.4	BS L3 CM Connected State	136
9.2.1.3.2.5	BS L3 CM Closing State.....	136
9.2.1.3.2.6	BS L3 CM Closed State	137
9.2.1.4	Multiple Connections per UT.....	137
9.2.1.4.1	RID versus PID.....	137
9.2.1.4.2	in-band paging	137
9.2.1.4.3	in-band queue depth reporting.....	138
9.2.1.4.4	ping renews all selected registrations	138
9.2.1.4.5	TCH renews all selected registrations.....	138
9.2.1.4.6	Network Access UT certificate contains session reference.....	138

9.2.1.4.7	<i>concerted handover of all connections</i>	138
9.2.1.4.8	<i>connection-based QoS</i>	138
9.2.2	<i>Registration Management</i>	139
9.2.2.1	<i>L3 RM Interface with L3 CM</i>	139
9.2.2.2	<i>L3 RM Interface with L2 RLC</i>	139
9.2.2.3	<i>L3 RM Interface to L2 MAC</i>	140
9.2.2.4	<i>L3 RM Interface to L3 MMC</i>	140
9.2.2.5	<i>L3 RM State Machines</i>	141
9.2.2.5.1	<i>UT L3 RM State Machine</i>	141
9.2.2.5.1.1	<i>UT L3 RM Initial State</i>	141
9.2.2.5.1.2	<i>UT L3 RM Discovering State</i>	141
9.2.2.5.1.3	<i>UT L3 RM Negotiating State</i>	142
9.2.2.5.1.4	<i>UT L3 RM Registered State</i>	142
9.2.2.5.1.5	<i>UT L3 RM Closing State</i>	143
9.2.2.5.1.6	<i>UT L3 RM Closed State</i>	143
9.2.2.5.2	<i>BS L3 RM State Machine</i>	143
9.2.2.5.2.1	<i>BS L3 RM Initial State</i>	144
9.2.2.5.2.2	<i>BS L3 RM Discovering State</i>	144
9.2.2.5.2.3	<i>BS L3 RM Negotiating State</i>	144
9.2.2.5.2.4	<i>BS L3 RM Registered State</i>	145
9.2.2.5.2.5	<i>BS L3 RM Closing State</i>	146
9.2.2.5.2.6	<i>BS L3 RM Closed State</i>	146
9.2.2.5.3	<i>L3 RM Stream Control Summary</i>	146
9.2.3	<i>Radio Resource Control</i>	147
9.2.3.1	<i>L3 RRC Interfaces</i>	147
9.2.3.2	<i>UT Class and Capabilities</i>	147
9.2.3.3	<i>Modulation Class Selection</i>	147
9.2.3.3.1	<i>Initial Modulation Class Selection</i>	147
9.2.3.3.2	<i>Modulation Class Adaptation</i>	148
9.2.3.4	<i>Power Control</i>	148
9.2.3.4.1	<i>Open-loop Power Control</i>	149
9.2.3.4.1.1	<i>CR Burst</i>	149
9.2.3.4.1.2	<i>CM Burst</i>	149
9.2.3.4.1.3	<i>RA Burst</i>	149
9.2.3.4.1.4	<i>AA Burst</i>	149
9.2.3.4.1.5	<i>Initial Uplink TCH Burst Power</i>	150
9.2.3.4.1.6	<i>Initial Downlink TCH Burst Power</i>	150
9.2.3.4.2	<i>Closed-loop Power Control</i>	150
9.2.3.4.2.1	<i>Uplink Power Control</i>	150
9.2.3.4.2.2	<i>Downlink Power Control</i>	153
9.2.4	<i>Mobility Management and Control</i>	155
9.2.4.1	<i>L3 MMC Interfaces</i>	156
9.2.4.2	<i>L3 MMC Management</i>	156
9.2.4.2.1	<i>L3 MMC State Machine</i>	156
9.2.4.2.2	<i>BS Selection</i>	158
9.2.4.2.3	<i>BS Monitoring</i>	160
9.2.4.3	<i>L3 MMC Control</i>	161
9.2.4.3.1	<i>Downlink Timing Measurements</i>	161
9.2.4.3.2	<i>Uplink Timing Measurements</i>	162
9.3	<i>PROTOCOL EVENTS</i>	162
9.3.1	<i>Introduction</i>	162
9.3.2	<i>BCH Acquisition and Selection</i>	162
9.3.3	<i>Session Start</i>	164
9.3.3.1	<i>Proactive Registration</i>	164
9.3.3.2	<i>Reactive Registration</i>	164
9.3.3.3	<i>Registration Progression</i>	164
9.3.4	<i>Data Delivery</i>	168
9.3.4.1	<i>BS Side</i>	168
9.3.4.1.1	<i>BS-initiated Extended Stream Start</i>	168
9.3.4.1.2	<i>Paging States and Transitions</i>	169
9.3.4.1.3	<i>Paging Selection</i>	170
9.3.4.1.4	<i>Stream Aggregation (Downlink Backlog)</i>	170
9.3.4.1.5	<i>In-band Page</i>	172
9.3.4.2	<i>UT Side</i>	173

9.3.4.2.1	UT-initiated Extended Stream Start	174
9.3.4.2.2	Stream Retry Strategy.....	175
9.3.4.2.3	RA Latency Messages.....	177
9.3.4.2.4	Stream Aggregation (Uplink Backlog).....	178
9.3.4.2.5	Resource Map Message.....	179
9.3.4.2.6	Intra-UT Stream Conflicts.....	180
9.3.4.3	L3 initiated Stream Closure	181
9.3.4.3.1	Graceful Stream Closure.....	181
9.3.4.3.2	Abrupt Stream Closure.....	182
9.3.4.3.3	Stream Duration Limit	183
9.3.5	Registration Maintenance	183
9.3.5.1	Registration Renewal (Ping).....	184
9.3.5.2	Ping Promotion to TCH.....	184
9.3.5.3	Deregistration	185
9.3.5.3.1	Bilateral Deregistration.....	186
9.3.5.3.2	Unilateral Deregistration.....	186
9.3.5.4	RID Reuse Detection.....	188
9.3.6	Handover.....	190
9.3.6.1	Proactive Registration.....	190
9.3.6.2	Reactive Registration	192
9.3.6.3	Registration Teardown During Handover.....	194
9.3.7	Session Stop	195
9.3.8	Session Abort	197
9.4	DATA FLOW.....	199
9.4.1	Introduction.....	199
9.4.2	Outgoing Data Flow	200
9.4.3	Incoming Data Flow	201
9.4.4	Control vs. Data.....	203
9.4.5	Packet Delimiting and Reconstruction.....	203
9.4.6	Packet Data Conversion.....	203
9.4.6.1	Packet Checksumming.....	204
9.4.6.2	Packet Fragmentation and Recombination	204
9.4.6.3	Packet Sequencing and Reordering.....	205
9.4.7	Packet Discarding.....	205
9.4.8	Packet Routing	205
9.4.9	Data Recovery from L2 RLC.....	206
9.4.10	Stream Aggregation	206
9.5	MESSAGE DEFINITIONS	207
9.5.1	Acknowledged Mode (AM) Messages.....	207
9.5.1.1	Message Structure.....	208
9.5.1.1.1	Message Top-level Structure	208
9.5.1.1.2	Message TLVG Structure	208
9.5.1.1.3	Message Field Structure.....	209
9.5.1.2	Message Formatting.....	209
9.5.1.3	Individual Messages.....	211
9.5.1.3.1	BS Params Message.....	212
9.5.1.3.2	UT Capability Message.....	214
9.5.1.3.3	UT Params Message	216
9.5.1.3.4	Reg Params Message	222
9.5.1.3.5	Net Sess Start Message.....	228
9.5.1.3.6	Net Sess Stop Message	229
9.5.1.3.7	In Band Page Message.....	229
9.5.1.3.8	Rsrc Map Message.....	230
9.5.1.3.9	RA Latency Request Message.....	233
9.5.1.3.10	RA Latency Report Message.....	233
9.5.1.3.11	Dereg Message.....	235
9.5.1.3.12	QoS Message.....	235
9.5.1.3.13	AFN Verify Message.....	236
9.5.1.3.14	Co Reg Queue Report Message.....	237
9.5.1.3.15	RACH Retry Params Message.....	238
9.5.2	Unacknowledged Mode (UM) Messages.....	239
9.5.2.1	queue.rpt.....	240
9.5.2.2	aggregate.cmd.....	240
9.5.2.3	starve.cmd.....	240

9.5.2.4	<i>hold.cmd</i>	241
9.5.2.5	<i>sinrErr.rpt</i>	241
9.5.2.6	<i>pwrCtrl.cmd</i>	241
10	PROTOCOL LAYER PRIMITIVES (INFORMATIVE)	241
10.1	OVERVIEW	241
10.2	DATA TYPES	241
10.3	PRIMITIVE LISTS AND DESCRIPTIONS	248
10.3.1	<i>Interface List</i>	248
10.3.2	<i>Individual Interfaces</i>	249
10.3.2.1	<i>L1 to L2 MAC Interface Primitives</i>	249
10.3.2.2	<i>L2 MAC to L1 Interface Primitives</i>	253
10.3.2.3	<i>L3 MMC to L1 Interface Primitives</i>	254
10.3.2.4	<i>L2 MAC to L2 RLC Interface Primitives</i>	255
10.3.2.5	<i>L2 RLC to L2 MAC Interface Primitives</i>	256
10.3.2.6	<i>L2 MAC to L3 RRC Interface Primitives</i>	256
10.3.2.7	<i>L3 RRC to L2 MAC Interface Primitives</i>	265
10.3.2.8	<i>L2 MAC to L3 MMC Interface Primitives</i>	268
10.3.2.9	<i>L3 MMC to L2 MAC Interface Primitives</i>	271
10.3.2.10	<i>L2 MAC to L3 RM Interface Primitives</i>	273
10.3.2.11	<i>L3 RM to L2 MAC Interface Primitives</i>	278
10.3.2.12	<i>L2 RLC to L3 RM Interface Primitives</i>	287
10.3.2.13	<i>L3 RM to L2 RLC Interface Primitives</i>	288
10.3.2.14	<i>L3 RRC to L3 MMC Interface Primitives</i>	290
10.3.2.15	<i>L3 MMC to L3 RRC Interface Primitives</i>	290
10.3.2.16	<i>L3 MMC to L3 RM Interface Primitives</i>	291
10.3.2.17	<i>L3 RM to L3 MMC Interface Primitives</i>	291
10.3.2.18	<i>L3 RM to L3 CM Interface Primitives</i>	292
10.3.2.19	<i>L3 CM to L3 RM Interface Primitives</i>	294
10.3.2.20	<i>L3 CM to L4 Interface Primitives</i>	296
10.3.2.21	<i>L4 to L3 CM Interface Primitives</i>	298
11	PRIVACY AND AUTHENTICATION	299
11.1	OVERVIEW	300
11.2	HC-SDMA HANDSHAKE AND AUTHENTICATION PROTOCOL, I-HAP	302
11.2.1	<i>User Terminal Requests for Registration</i>	303
11.2.2	<i>Base Station Transmits Its Digital Certificate</i>	304
11.2.3	<i>User Terminal Analyzes Base Station Certificate Message</i>	307
11.2.4	<i>User Terminal Encrypted Shared Secret Transmission</i>	308
11.2.5	<i>Base Station Decryption of Encrypted Shared Secret</i>	309
11.2.6	<i>User Terminal Verifies BS Certificate Signature</i>	309
11.3	HC-SDMA TERMINAL AUTHENTICATION PROTOCOL, I-TAP	309
11.3.1	<i>UT Certificates</i>	312
11.3.1.1	<i>Certificate Usage Rules</i>	312
11.3.1.2	<i>UT Certificate Scheme</i>	313
11.3.1.2.1	<i>UT Identity Certificate</i>	314
11.3.1.2.2	<i>UT Service Authorization Certificate</i>	315
11.3.1.2.3	<i>UT Network Access Certificate</i>	315
11.3.1.3	<i>Operational Recommendations</i>	316
11.3.2	<i>UT Authenticator Message and String Generation</i>	317
11.3.3	<i>AFN Verification with a Centralized Time Server</i>	318
11.3.4	<i>UT Parameter Scrambling</i>	321
11.3.5	<i>UT Params Message Transmission</i>	321
11.3.6	<i>BS Analysis of UT Certificate and Authenticator String</i>	322
11.3.7	<i>Failed UT Authentication</i>	323
11.4	HC-SDMA SECURE COMMUNICATIONS PROTOCOL, I-SEC	323
11.4.1	<i>TCH Streams</i>	323
11.4.2	<i>Symmetric Key Stream Cipher Algorithm</i>	323
11.4.2.1	<i>Initialization Vector Selection</i>	324
11.4.2.2	<i>Stream Counter Test</i>	324
11.4.2.3	<i>Determination of Encryption Key from Shared Secret</i>	325
11.4.2.4	<i>Stream Cipher State Initialization</i>	326
11.4.2.5	<i>Keystream Generation</i>	327
11.4.2.6	<i>Burst Counter Test</i>	327

11.4.2.7	Encryption and Decryption Using i-SEC Keystream	328
APPENDIX 11A	REFERENCES	328
APPENDIX 11B	ELLIPTIC CURVE PARAMETERS AND ARITHMETIC.....	329
11B.1	ELLIPTIC CURVE DOMAIN PARAMETERS (ECDP)	329
11B.2	ELLIPTIC CURVE ARITHMETIC	330
APPENDIX 11C. I-HAP ALGORITHMS		330
11C.1	I-HAP CERTIFICATE SIGNATURES.....	330
11C.1.1	<i>i-HAP Certificate Authority (CA) Key Pair Generation</i>	<i>331</i>
11C.1.2	<i>i-HAP Certificate Signature Generation</i>	<i>332</i>
11C.1.3	<i>i-HAP Certificate Signature Verification.....</i>	<i>332</i>
11C.2	I-HAP BASE STATION PUBLIC KEY ENCRYPTION ALGORITHM.....	333
11C.2.1	<i>i-HAP BS Key Pair Generation</i>	<i>333</i>
11C.2.2	<i>i-HAP UT Encryption.....</i>	<i>333</i>
11C.2.3	<i>i-HAP BS Decryption.....</i>	<i>334</i>
11C.3	CYCLIC REDUNDANCY CODE FOR SHARED SECRET EXCHANGE.....	334
APPENDIX 11D. I-TAP ALGORITHMS.....		335
11D.1	I-TAP UT CERTIFICATE SIGNATURES.....	335
11D.1.1	<i>i-TAP UT Certificate Authority (CA) Key Pair Generation.....</i>	<i>335</i>
11D.1.2	<i>i-TAP UT Certificate Signature Generation.....</i>	<i>336</i>
11D.1.3	<i>i-TAP UT Certificate Signature Verification.....</i>	<i>336</i>
11D.2	I-TAP UT PUBLIC KEY ENCRYPTION ALGORITHM.....	336
11D.2.1	<i>i-TAP UT Key Pair Generation.....</i>	<i>337</i>
11D.2.2	<i>i-TAP UT Private Key Encryption.....</i>	<i>337</i>
11D.2.3	<i>i-TAP BS Verification of UT Signature</i>	<i>338</i>
11D.3	I-TAP SCRAMBLING ALGORITHM	338
11D.4	I-TAP AUTOMATIC CERTIFICATE RETRIEVAL	339
12 PERFORMANCE REQUIREMENTS AND TEST PROCEDURES.....		341
12.1	INTRODUCTION	341
12.1.1	<i>Definitions.....</i>	<i>342</i>
12.2	FREQUENCY (F) SYNCHRONIZATION CHANNEL RECEIVER REQUIREMENTS	342
12.2.1	<i>Performance in Additive White Gaussian Noise</i>	<i>342</i>
12.2.1.1	<i>Definition</i>	<i>342</i>
12.2.1.2	<i>Performance Specifications.....</i>	<i>343</i>
12.2.1.3	<i>Measurement Procedure</i>	<i>343</i>
12.2.2	<i>PERFORMANCE IN FADING.....</i>	<i>344</i>
12.2.2.1	<i>Definition</i>	<i>344</i>
12.2.2.2	<i>Performance Specifications.....</i>	<i>344</i>
12.2.2.2.1	<i>Input Specifications.....</i>	<i>344</i>
12.2.2.2.2	<i>Output Specifications</i>	<i>344</i>
12.2.2.3	<i>Measurement Procedure</i>	<i>345</i>
12.2.3	<i>PERFORMANCE IN NOISE-ONLY CONDITION</i>	<i>345</i>
12.2.3.1	<i>Definition</i>	<i>345</i>
12.2.3.2	<i>Performance Specifications.....</i>	<i>345</i>
12.2.3.2.1	<i>Input Specifications.....</i>	<i>345</i>
12.2.3.2.2	<i>Output Specifications</i>	<i>346</i>
12.2.3.3	<i>Measurement Procedure</i>	<i>346</i>
12.3	TIMING (T) CHANNEL RECEIVER REQUIREMENTS	346
12.3.1	<i>Performance in Additive White Gaussian Noise</i>	<i>346</i>
12.3.1.1	<i>Definition</i>	<i>346</i>
12.3.1.2	<i>Performance Specifications.....</i>	<i>346</i>
12.3.1.2.1	<i>Input Specifications.....</i>	<i>346</i>
12.3.1.2.2	<i>Output Specifications</i>	<i>346</i>
12.3.1.3	<i>Measurement Procedure</i>	<i>347</i>
12.3.2	<i>Performance in Fading</i>	<i>347</i>
12.3.2.1	<i>Definition</i>	<i>347</i>
12.3.2.2	<i>Performance Specifications.....</i>	<i>347</i>
12.3.2.2.1	<i>Input Specifications.....</i>	<i>347</i>
12.3.2.2.2	<i>Output Specifications</i>	<i>347</i>
12.3.2.3	<i>Measurement Procedure</i>	<i>348</i>

12.4 BROADCAST (B) CHANNEL RECEIVER REQUIREMENTS	348
12.4.1 Performance in Additive White Gaussian Noise	349
12.4.1.1 Definition	349
12.4.1.2 Performance Specifications.....	349
12.4.1.2.1 Input Specifications.....	349
12.4.1.2.2 Output Specifications	349
12.4.1.3 Measurement Procedure	349
12.4.2 Performance in Fading	350
12.4.2.1 Definition	350
12.4.2.2 Performance Specifications.....	350
12.4.2.2.1 Input Specifications.....	350
12.4.2.2.2 Output Specifications	350
12.4.2.3 Measurement Procedure	351
12.4.3 Performance in Noise-only Condition.....	351
12.4.3.1 Definition	351
12.4.3.2 Performance Specifications.....	351
12.4.3.2.1 Input Specifications.....	351
12.4.3.2.2 Output Specifications	351
12.4.3.3 Measurement Procedure	351
12.4.4 Performance in a Multi-basestation environment	352
12.4.4.1 Definition	352
12.4.4.2 Performance Specifications.....	352
12.4.4.2.1 Input Specifications.....	352
12.4.4.2.2 Output Specifications	352
12.4.4.3 Measurement Procedure	352
12.5 BCH ACQUISITION REQUIREMENTS	352
12.5.1 F-Burst Acquisition	352
12.5.1.1 Definition	352
12.5.1.2 Specifications	352
12.5.2 T-BURST ACQUISITION	353
12.5.2.1 Definition	353
12.5.2.2 Specifications	353
12.5.3 B-BURST ACQUISITION	353
12.5.3.1 Definition	353
12.5.3.2 Specifications	353
12.5.4 Measurement Procedure	353
12.5.5 Overall BCH Acquisition	353
12.5.5.1 Definition	353
12.5.5.2 Specifications for AWGN Channel.....	354
12.5.5.3 Measurement Procedure	354
12.5.5.4 Specifications for Fading Channel.....	354
12.5.5.5 Measurement Procedure	354
12.5.6 Overall False Alarm Rate Specification.....	355
12.5.6.1 Definition	355
12.5.6.2 Specification.....	355
12.5.6.3 Measurement Procedure	355
12.6 PAGING CHANNEL RECEIVER REQUIREMENTS	355
12.6.1 Performance in Additive White Gaussian Noise	355
12.6.1.1 Definition	355
12.6.1.2 Performance Specifications.....	355
12.6.1.2.1 Input Specifications.....	355
12.6.1.2.2 Output Specifications	355
12.6.1.3 Measurement Procedure	356
12.6.2 Performance in Fading	356
12.6.2.1 Definition	356
12.6.2.2 Performance Specifications.....	356
12.6.2.2.1 Input Specifications.....	356
12.6.2.2.2 Output Specifications	356
12.6.2.3 Measurement Procedure	356
12.6.3 Performance in Noise-only Condition.....	357
12.6.3.1 Definition	357
12.6.3.2 Performance Specifications.....	357
12.6.3.2.1 Input Specifications.....	357
12.6.3.2.2 Output Specifications	357

12.6.3.3	Measurement Procedure	357
12.6.4	Performance under Multiuser Paging.....	357
12.6.4.1	Definition	357
12.6.4.2	Performance Specifications.....	357
12.6.4.2.1	Input Specifications	357
12.6.4.2.2	Output Specifications	357
12.6.4.3	Measurement Procedure	357
12.7	CONFIGURATION CHANNEL RECEIVER REQUIREMENTS.....	358
12.7.1	Performance in Additive White Gaussian Noise	358
12.7.1.1	Definition	358
12.7.1.2	Uplink Performance Specifications (CR Burst).....	358
12.7.1.2.1	Input Specifications	358
12.7.1.2.2	Output Specifications	358
12.7.1.3	Uplink Measurement Procedure.....	359
12.7.1.4	Downlink Performance Specifications (CM Burst)	359
12.7.1.4.1	Input Specifications	359
12.7.1.4.2	Output Specifications	359
12.7.1.5	Downlink Measurement Procedure.....	360
12.8	ACCESS CHANNEL (RACH) RECEIVER REQUIREMENTS	360
12.8.1	Performance in Additive White Gaussian Noise	360
12.8.1.1	Definition	360
12.8.1.2	Uplink Performance Specifications (RA Burst).....	360
12.8.1.2.1	Input Specifications	360
12.8.1.2.2	Output Specifications	360
12.8.1.3	Uplink Measurement Procedure.....	361
12.8.1.4	Downlink Performance Specifications (AA Burst)	361
12.8.1.4.1	Input Specifications	361
12.8.1.4.2	Output Specifications	361
12.8.1.5	Downlink Measurement Procedure.....	361
12.9	TRAFFIC CHANNEL RECEIVER REQUIREMENTS	362
12.9.1	Performance in Additive White Gaussian Noise	362
12.9.1.1	Definition	362
12.9.1.2	Uplink Performance Specifications.....	362
12.9.1.2.1	Input Specifications	362
12.9.1.2.2	Output Specifications	362
12.9.1.3	Uplink Measurement Procedure.....	362
12.9.1.4	Downlink Performance Specifications	363
12.9.1.4.1	Input Specifications	363
12.9.1.4.2	Output Specifications	363
12.9.1.5	Downlink Measurement Procedure.....	363
12.10	FAST ASSOCIATED CONTROL CHANNEL RECEIVER REQUIREMENTS	364
12.10.1	Performance in Additive White Gaussian Noise.....	364
12.10.1.1	Definition	364
12.10.1.2	Uplink Performance Specifications.....	364
12.10.1.2.1	Input Specifications	364
12.10.1.2.2	Output Specifications	364
12.10.1.3	Uplink Measurement Procedure.....	364
12.10.1.4	Downlink Performance Specifications	365
12.10.1.4.1	Input Specifications	365
12.10.1.4.2	Output Specifications	365
12.10.1.5	Downlink Measurement Procedure.....	365
A	GLOSSARY	366

TABLE OF FIGURES

FIGURE 1: MSC NOTATION	5
FIGURE 2: ELEMENTS IN A TYPICAL HC-SDMA DEPLOYMENT.....	6
FIGURE 3: END-USER TRAFFIC ENCAPSULATION	7
FIGURE 4: SUBLAYERS AND INTERFACES IN THE HC-SDMA PROTOCOL REFERENCE MODEL.....	9
FIGURE 5: RF CHANNELIZATION FOR A 5 MHZ FREQUENCY ALLOCATION	14
FIGURE 6: THE HC-SDMA FRAME STRUCTURE (NOT TO SCALE)	14

FIGURE 7: HC-SDMA BCH/CCH CARRIER SUPERFRAME STRUCTURE	16
FIGURE 8: FREQUENCY SYNCHRONIZATION BURST STRUCTURE	17
FIGURE 9: TIMING SYNCHRONIZATION BURST STRUCTURE	18
FIGURE 10: BROADCAST BURST STRUCTURE	19
FIGURE 11: DIVISION OF A DOWNLINK SLOT INTO TWO PAGE BURSTS	21
FIGURE 12: CONFIGURATION REQUEST BURST STRUCTURE	22
FIGURE 13: STANDARD UPLINK BURST STRUCTURE	23
FIGURE 14: STANDARD DOWNLINK BURST STRUCTURE	29
FIGURE 15: BLOCK DIAGRAM OF ERROR CONTROL CODING SCHEME	36
FIGURE 16: CONVOLUTIONAL ENCODER	39
FIGURE 17: SCRAMBLING SEQUENCE GENERATOR	47
FIGURE 18: QUATERNARY SEQUENCE GENERATOR	50
FIGURE 19: RAMP-UP AND RAMP-DOWN TRANSMIT POWER	57
FIGURE 20: EXAMPLE OF BASE STATION MULTI-ANTENNA ARCHITECTURE	71
FIGURE 21: FOM DEFINITION	73
FIGURE 22: BCH BASIC ACCESS DIAGRAM	77
FIGURE 23: BCH SUPERFRAME	78
FIGURE 24: PAGING SUBSLOT DIVERSITY	80
FIGURE 25: RA TIME OF ARRIVAL RELATIVE TO BASE STATION FRAME TIMING	87
FIGURE 26: UT L2 MAC UNDIRECTED RACH RETRY POLICY FOR L2MacUtStart . req	96
FIGURE 27: DIRECTED RACH RETRY POLICY (FOR $X > 0$)	97
FIGURE 28: UT-INITIATED STREAM START SEQUENCE	101
FIGURE 29: BS-INITIATED STREAM START SEQUENCE	101
FIGURE 30: AM RMU DEFINITION	104
FIGURE 31: STREAM RELEASE SEQUENCE	107
FIGURE 32: GRACEFUL STREAM CLOSE SEQUENCE	108
FIGURE 33: KILL FROM L3 RM (CLOSE . CMD UM RECEIVED)	109
FIGURE 34: KILL FROM L3 RM (CLOSE . CMD UM LOST)	109
FIGURE 35: FORCED SHUTDOWN FROM L2 RLC	110
FIGURE 36: L2 RLC OVERVIEW DIAGRAM	115
FIGURE 37: L2 RLC MODEL	116
FIGURE 38: AM TRANSMIT SEQUENCE NUMBERS	119
FIGURE 39: AM RECEIVE SEQUENCE NUMBERS	120
FIGURE 40: L3 COMPONENTS AND INTERACTIONS WITH OTHER LAYERS	130
FIGURE 41: UT L3 CM STATE TRANSITION DIAGRAM	132
FIGURE 42: BS L3 CM STATE TRANSITION DIAGRAM	135
FIGURE 43: UT L3 RM STATE TRANSITION DIAGRAM	141
FIGURE 44: BS L3 RM STATE TRANSITION DIAGRAM	144
FIGURE 45: UPLINK POWER CONTROL	152
FIGURE 46: DOWNLINK POWER CONTROL	155
FIGURE 47: L3 MMC TIMING ALIGNMENT STATE TRANSITION DIAGRAM	158
FIGURE 48: L3 MMC ACQUISITION	160
FIGURE 49: BCH ACQUISITION AND BS SELECTION	163
FIGURE 50: PREMATURE SESSION START REJECTION	164
FIGURE 51: SESSION START (PROACTIVE REGISTRATION)	166
FIGURE 52: SESSION START (REACTIVE REGISTRATION)	167
FIGURE 53: BS-INITIATED STREAM START	168
FIGURE 54: BS-INITIATED EXTENDED STREAM START	169
FIGURE 55: STREAM AGGREGATION (DOWNLINK QUEUE, STREAM START)	171
FIGURE 56: STREAM AGGREGATION (DOWNLINK QUEUE, NEW PACKET ARRIVAL)	172
FIGURE 57: IN-BAND PAGE	173
FIGURE 58: UT-INITIATED STREAM START	174
FIGURE 59: UT-INITIATED EXTENDED STREAM START	175
FIGURE 60: UT L3 RM STREAM RETRY BACKOFF	176
FIGURE 61: UT L3 RM STREAM RETRY LIMITS	177
FIGURE 62: RA LATENCY MESSAGES	178
FIGURE 63: STREAM AGGREGATION (UPLINK QUEUE)	179
FIGURE 64: RESOURCE MAP MESSAGE	180
FIGURE 65: INTRA-UT STREAM CONFLICT RESOLUTION	181
FIGURE 66: GRACEFUL STREAM CLOSURE	182
FIGURE 67: ABRUPT STREAM CLOSURE	183
FIGURE 68: REGISTRATION RENEWAL/PING	184
FIGURE 69: PING TO TCH PROMOTION	185
FIGURE 70: DEREGISTRATION COMMAND	187

FIGURE 71: DEREGISTRATION SUGGESTION	188
FIGURE 72: UT RID REUSE DETECTION	189
FIGURE 73: HANDOVER (PROACTIVE REGISTRATION)	191
FIGURE 74: HANDOVER (REACTIVE REGISTRATION)	193
FIGURE 75: INTER-BS HANDOVER	194
FIGURE 76: INTRA-BS HANDOVER	195
FIGURE 77: SESSION STOP (NO TOP-LEVEL HANDSHAKE)	196
FIGURE 78: SESSION STOP (WITH TOP-LEVEL HANDSHAKE)	197
FIGURE 79: REGISTRATION FAILURE	198
FIGURE 80: SESSION REJECTION	199
FIGURE 81: OUTGOING DATA FLOW	200
FIGURE 82: INCOMING DATA FLOW	202
FIGURE 83: AM MESSAGE STRUCTURE	208
FIGURE 84: VARIABLE-LENGTH FIELD PADDING CASE 1	210
FIGURE 85: VARIABLE-LENGTH FIELD PADDING CASE 2	211
FIGURE 86: VARIABLE-LENGTH FIELD PADDING CASE 3	211
FIGURE 87: BS PARAMS MESSAGE	214
FIGURE 88: UT CAPABILITY MESSAGE	216
FIGURE 89: UT PARAMS MESSAGE	222
FIGURE 90: REG PARAMS MESSAGE	228
FIGURE 91: NET SESS START MESSAGE	229
FIGURE 92: NET SESS STOP MESSAGE	229
FIGURE 93: IN BAND PAGE MESSAGE	230
FIGURE 94: RSRG MAP MESSAGE	233
FIGURE 95: RA LATENCY REQUEST MESSAGE	233
FIGURE 96: RA LATENCY REPORT MESSAGE	234
FIGURE 97: DEREG MESSAGE	235
FIGURE 98: QoS MESSAGE	236
FIGURE 99: AFN VERIFY MESSAGE	237
FIGURE 100: CO REG QUEUE REPORT MESSAGE (EXISTING REGISTRATION)	238
FIGURE 101: CO REG QUEUE REPORT MESSAGE (NEW REGISTRATION)	238
FIGURE 102: RACH RETRY PARAMS MESSAGE	239
FIGURE 103: PROTOCOL LAYER INTERACTIONS	249
FIGURE 104: GENERAL DESCRIPTION OF PRIVACY AND AUTHENTICATION PROTOCOLS	302
FIGURE 105: BASELINE <i>I-HAP</i> EVENT DIAGRAM	303
FIGURE 106: BCH SEARCH AND REGISTRATION	304
FIGURE 107: <i>HC-SDMA</i> BS CERTIFICATE MESSAGE FIELDS. MANDATORY FIELD VALUES ARE IDENTIFIED IN SQUARE BRACKETS. NOT TO SCALE.	307
FIGURE 108: BASELINE <i>I-TAP</i> EVENT DIAGRAM	311
FIGURE 109: UT AUTHENTICATOR MESSAGE FORMATTING (NOT TO SCALE)	317
FIGURE 110: AFN VERIFICATION PROCEDURE	320
FIGURE 111: <i>I-TAP</i> UT PARAMETERS BEFORE SCRAMBLING	321
FIGURE 112: ENCRYPTION IN <i>HC-SDMA</i> SECURE COMMUNICATION PROTOCOL, <i>I-SEC</i>	325
FIGURE 113: <i>I-TAP</i> LINEAR FEEDBACK SHIFT REGISTER	339
FIGURE 114: UT AUTOMATIC CERTIFICATE DOWNLOAD	341
FIGURE 115: USER TERMINAL FREQUENCY SYNCHRONIZATION CHANNEL RECEIVER TEST SETUP FOR THE AWGN CONDITION	344
FIGURE 116: USER TERMINAL FREQUENCY SYNCHRONIZATION CHANNEL RECEIVER TEST SETUP FOR FADING CONDITION	345
FIGURE 117: USER TERMINAL AND BASE STATION TEST SETUP FOR THE AWGN CONDITION	347
FIGURE 118: USER TERMINAL AND BASE STATION TEST SETUP FOR FADING CONDITION	348
FIGURE 119: USER TERMINAL TEST SETUP	356
FIGURE 120: BASE STATION TEST SETUP	359

TABLE OF TABLES

TABLE 1: PROTOCOL LAYERS IN THE <i>HC-SDMA</i> REFERENCE MODEL	8
TABLE 2: BURST TYPES	15
TABLE 3: FREQUENCY SYNCHRONIZATION BURST FIELDS	17
TABLE 4: TIMING SYNCHRONIZATION BURST FIELDS	18
TABLE 5: BROADCAST BURST FIELDS	19
TABLE 6: PAGE BURST IN BOTH SUBSLOTS	19
TABLE 7: PAGE BURST IN FIRST SUBSLOT	20
TABLE 8: PAGE BURST IN SECOND SUBSLOT	20

TABLE 9: CONFIGURATION REQUEST CORE SEQUENCES	22
TABLE 10: CONFIGURATION REQUEST BURST FIELDS	22
TABLE 11: STANDARD UPLINK BURST FIELDS	24
TABLE 12: CORE SEQUENCES FOR STANDARD UPLINK BURST	25
TABLE 13: CYCLIC SHIFT FOR UPLINK TRAINING.....	26
TABLE 14: SCRAMBLING SEQUENCES FOR UPLINK FACCH SYMBOLS.....	28
TABLE 15: STANDARD DOWNLINK BURST FIELDS	29
TABLE 16: CYCLIC SHIFT FOR DOWNLINK TRAINING	29
TABLE 17: CORE SEQUENCES FOR STANDARD DOWNLINK BURST.....	31
TABLE 18: CYCLIC SHIFT FOR DOWNLINK TAIL TRAINING SEQUENCES.....	32
TABLE 19: TAIL TRAINING SEQUENCES FOR STANDARD DOWNLINK BURST.....	33
TABLE 20: SCRAMBLING SEQUENCES FOR DOWNLINK FACCH SYMBOLS.....	34
TABLE 21: USER DATA THROUGHPUT AT VARIOUS MODCLASSES.....	35
TABLE 22: MODULATION AND CODING RATES	36
TABLE 23: MODCLASS VERSUS BURST TYPE.....	37
TABLE 24: BLOCK LENGTHS IN DOWNLINK TRAFFIC BURST	37
TABLE 25: BLOCK LENGTHS IN UPLINK TRAFFIC BURST.....	37
TABLE 26: BLOCK LENGTHS IN CONFIGURATION REQUEST BURST	38
TABLE 27: BLOCK LENGTHS IN CONFIGURATION MESSAGE	38
TABLE 28: BLOCK LENGTHS IN REQUEST ACCESS MESSAGE	38
TABLE 29: BLOCK LENGTHS IN ACCESS ASSIGNMENT MESSAGE	38
TABLE 30: CODE BIT LOCATIONS FUNCTION	40
TABLE 31: MAPPING FOR RATE-3/4 BLOCK SHAPER	42
TABLE 32: MAPPING FOR RATE-4/4 BLOCK SHAPER	42
TABLE 33: MAPPING FOR RATE-5/4 BLOCK SHAPER	43
TABLE 34: SYMBOL MAPPER FOR MODULATION CLASSES 0-3	43
TABLE 35: SYMBOL MAPPER FOR MODULATION CLASS 4	43
TABLE 36: MAPPER FOR MODULATION CLASS 5	44
TABLE 37: SYMBOL MAPPER FOR MODULATION CLASSES 6 - 8.....	44
TABLE 38: INTERLEAVER ROW PERMUTATION.....	44
TABLE 39: SCRAMBLING INITIALIZATION FOR VARIOUS BURST TYPES.....	46
TABLE 40: MODULATION SCALING.....	47
TABLE 41: NOMINAL UT TRANSMIT POWER PER CARRIER FOR VARIOUS MODULATION FORMATS	56
TABLE 42: MODULATION ACCURACY FOR VARIOUS MODULATION FORMATS	59
TABLE 43: MAXIMUM ACPR	59
TABLE 44: REFERENCE SENSITIVITY AT FER = 10 ⁻²	60
TABLE 45: ADJACENT CHANNEL SELECTIVITY	61
TABLE 46: OUT-OF-BAND BLOCKING CHARACTERISTICS	61
TABLE 47: TABULATED RANGE OF ACCEPTABLE DSSI REPORT VALUES	62
TABLE 48: TABULATED RANGE OF ACCEPTABLE SINR REPORT VALUES.....	63
TABLE 49: OUT-OF-BAND SPURIOUS EMISSIONS LIMITS	67
TABLE 50: REFERENCE SENSITIVITY FOR FER = 10 ⁻²	68
TABLE 51: ACS	69
TABLE 52: TABULATED RANGE OF ACCEPTABLE DSSI REPORT VALUES	69
TABLE 53: TABULATED RANGE OF ACCEPTABLE SINR REPORT VALUES.....	70
TABLE 54: FIGURE OF MERIT REQUIREMENTS	72
TABLE 55: LOGICAL CHANNELS, MESSAGES, AND BURST TYPES	75
TABLE 56: CONVENTION EXAMPLE	76
TABLE 57: BCH MESSAGE FIELDS	78
TABLE 58: CONFIGURATION REQUEST MESSAGE FIELDS	82
TABLE 59: CM FIELDS	82
TABLE 60: CM MINIMUM PROTOCOL FIELD	83
TABLE 61: NETWORK OPERATOR ID FIELD	84
TABLE 62: RA FIELD DEFINITIONS	85
TABLE 63: UT CLASS MAPPING	85
TABLE 64: RA TYPE FIELD VALUES	86
TABLE 65: AA MESSAGE FORMAT FOR COMMON FIELDS	87
TABLE 66: aaSUBTYPE FIELD VALUES FOR aaTYPE = 1	88
TABLE 67: aaSUBTYPE FIELD VALUES FOR aaTYPE = 0	88
TABLE 68: MESSAGE FORMAT FOR SUBTYPES AA-CTS AND AA-SHORT.....	89
TABLE 69: MESSAGE FORMAT FOR SUBTYPE AA-RREG	89
TABLE 70: MESSAGE FORMAT FOR SUBTYPE AA-PENDING.....	90
TABLE 71: MESSAGE FORMAT FOR SUBTYPE AA-REJECT	90
TABLE 72: MESSAGE FORMAT FOR SUBTYPE AA-INVALID-ID.....	90

TABLE 73: MESSAGE FORMAT FOR SUBTYPES WITH NO ADDITIONAL FIELDS.....	91
TABLE 74: FRAMEDEC FIELD VALUES	91
TABLE 75: MAPPING OF FREQUENCY AND TIMESLOTS TO CONVENTIONAL CHANNELS	92
TABLE 76: RA TYPE TO AA SUBTYPE CORRESPONDENCE	93
TABLE 77: RA TYPE SENT AS A FUNCTION OF STIMULUS.....	93
TABLE 78: UT START TYPE AS A FUNCTION OF STIMULUS.....	94
TABLE 79: BS START TYPE AS A FUNCTION OF STIMULUS	94
TABLE 80: PARAMETERS OF THE MAC RETRY STRATEGY	95
TABLE 81: FACCH TO MODCLASS MAPPING	99
TABLE 82: FACCH TO AVAILABLE POWER MAPPING	100
TABLE 83: RMU HEADER FIELDS	102
TABLE 84: UM MESSAGE DEFINITIONS	111
TABLE 85: BSLOADEST TO BSLOAD MAPPING	113
TABLE 86: MODULO ARITHMETIC OPERANDS FOR SEQUENCE NUMBERS.....	114
TABLE 87: VALUES OF $N_{TURNAROUND}$ FOR THE BS.....	121
TABLE 88: HC-SDMA LOGICAL CHANNELS	128
TABLE 89: L3 RM STREAM CONTROL.....	146
TABLE 90: NEWLY OPENED STREAMS	146
TABLE 91: DOWNLINK SINR TARGETS	153
TABLE 92: L3 PACKET TAG BYTES.....	200
TABLE 93: HEADER FORMATS.....	203
TABLE 94: PACKET DATA CONVERSION ALGORITHMS.....	204
TABLE 95: QUEUE DEPTH.....	207
TABLE 96: EXAMPLE MESSAGE CODING.....	210
TABLE 97: L3 AM MESSAGES.....	211
TABLE 98: PROTOCOL VERSION FIELD SETTINGS.....	213
TABLE 99: MAPPING OF FREQUENCY AND PAGE RESOURCE TO PCH LOGICAL MASK (0 IS LSB).....	225
TABLE 100: L2 MAC RA RETRY POLICY PARAMETER SETTINGS	226
TABLE 101: CONNECTION REJECTION VALUES	227
TABLE 102: RA LATENCY ENCODING	234
TABLE 103: RA ATTEMPTS ENCODING.....	234
TABLE 104: L3 UM MESSAGES.....	240
TABLE 105: FUNDAMENTAL DATA TYPES	242
TABLE 106: ENUMERATED DATA TYPES.....	243
TABLE 107: SYNONYM DATA TYPES	245
TABLE 108: F BURST FREQUENCY OFFSET SPECIFICATION IN AN AWGN CHANNEL	343
TABLE 109: F BURST TIMING ALIGNMENT SPECIFICATION IN AN AWGN CHANNEL	343
TABLE 110: F BURST FREQUENCY OFFSET SPECIFICATION IN A FADING CHANNEL	344
TABLE 111: F BURST TIMING ALIGNMENT SPECIFICATION IN A FADING CHANNEL.....	344
TABLE 112: T BURST DECODE RATE SPECIFICATION IN AN AWGN CHANNEL.....	346
TABLE 113: T BURST TIMING ALIGNMENT SPECIFICATION IN AN AWGN CHANNEL	346
TABLE 114: T BURST DECODE RATE SPECIFICATION IN A FADING CHANNEL	347
TABLE 115: T BURST TIMING ALIGNMENT SPECIFICATION IN A FADING CHANNEL	348
TABLE 116: B BURST DECODE RATE SPECIFICATION IN AN AWGN CHANNEL.....	349
TABLE 117: B BURST TIMING ALIGNMENT SPECIFICATION IN AN AWGN CHANNEL.....	349
TABLE 118: B BURST FREQUENCY OFFSET SPECIFICATION IN AN AWGN CHANNEL.....	349
TABLE 119: B BURST DECODE RATE SPECIFICATION IN A FADING CHANNEL.....	350
TABLE 120: B BURST TIMING ALIGNMENT SPECIFICATION IN A FADING CHANNEL	350
TABLE 121: B BURST FREQUENCY OFFSET SPECIFICATION IN A FADING CHANNEL.....	351
TABLE 122: BCH ACQUISITION SPECIFICATIONS IN AN AWGN CHANNEL.....	354
TABLE 123: BCH ACQUISITION SPECIFICATIONS IN A FADING CHANNEL	354
TABLE 124: PAGING SPECIFICATIONS IN AN AWGN CHANNEL.....	355
TABLE 125: PAGING SPECIFICATIONS IN A FADING CHANNEL	356
TABLE 126: CR SENSITIVITY SPECIFICATIONS	358
TABLE 127: CM SENSITIVITY SPECIFICATIONS	360
TABLE 128: RA SENSITIVITY SPECIFICATIONS.....	361
TABLE 129: AA SENSITIVITY SPECIFICATIONS	361
TABLE 130: UPLINK TCH SENSITIVITY SPECIFICATIONS.....	362
TABLE 131: DOWNLINK TCH SENSITIVITY SPECIFICATIONS	363
TABLE 132: UPLINK FACCH SENSITIVITY SPECIFICATIONS	364
TABLE 133: DOWNLINK FACCH SENSITIVITY SPECIFICATIONS.....	365

American National Standard for Telecommunications –

High Capacity – Spatial Division Multiple Access (HC-SDMA) Radio Interface Standard

1 SYSTEM OVERVIEW

1.1 Overview of HC-SDMA

1.1.1 Scope of This Document

This document defines the radio (RF), Physical Layer (PHY), Medium Access Control (MAC), and Layer 3 (L3) specifications for the *HC-SDMA* (High Capacity - Spatial Division Multiple Access) protocol.

This specification does not address functionality at the service and application layers. Typical deployments are expected to use a standardized data networking access paradigm, such as L2TP and PPP.

1.1.2 Purpose

HC-SDMA provides wide-area broadband wireless data-connectivity for fixed, portable, and mobile computing devices and appliances that may or may not be handheld. The protocol is designed to be implemented with antenna array techniques to substantially improve the RF coverage, capacity, and performance of the system.

Specifically, the *HC-SDMA* protocol:

- ◆ Specifies base station (BS) radio frequency characteristics such as output power levels, transmit frequencies and timing error, pulse shaping, in-band and out-of-band spurious emissions, receiver sensitivity and selectivity.
- ◆ Specifies user terminal (UT) radio frequency characteristics such as output power levels, transmit frequency and timing error, pulse shaping, in-band and out-of-band spurious emissions, receiver sensitivity and selectivity.
- ◆ Defines associated frame structures for the various burst types including standard uplink and downlink traffic, paging and broadcast burst types.
- ◆ Specifies the modulation, forward error correction, interleaving and scrambling for various burst types.
- ◆ Describes the various logical channels (broadcast, paging, random access, configuration and traffic channels) and their roles in establishing communication over the radio link.
- ◆ Specifies procedures for error recovery and retry.
- ◆ Describes Layer 3 (L3) mechanisms for creating and controlling logical connections (sessions) between user terminal and base station including registration, stream start, power control, handover, link adaptation, stream closure.
- ◆ Describes L3 mechanisms for user terminal authentication and secure transmission of on the uplink and downlink data links.

1.2 Acronyms

AA	Access Assignment
ACLPR	Adjacent Channel Leakage Power Ratio
ACS	Adjacent Channel Selectivity
AFC	Automatic Frequency Control
AFN	Absolute Frame Number
AM	Acknowledged Mode
API	Application Programming Interface
ARQ	Automatic Repeat Request
ASD	Acceleration Spectral Density (mechanical vibration)
BCH	Broadcast Channel
BS	Base Station
BSCC	Base Station Color Code
CA	Certificate Authority
CCB	Connection Control Block
CCH	Configuration Channel
CM	Configuration Message
CoS	Class of Service
CR	Configuration Request
CRC	Cyclic Redundancy Check
CW	Continuous Wave (unmodulated sine wave)
EC	Elliptic Curve
ECDP	Elliptic Curve Domain Parameters
EUD	End User Device
FACCH	Fast Associated Control Channel
FEC	Forward Error Control
FER	Frame Error Rate
FIPS	Federal Information Processing Standard
GPS	Global Positioning System
HC-SDMA	High Capacity Spatial Division Multiple Access
ICRP	Incoming Call Reply
i-HAP	<i>HC-SDMA</i> Handshake and Authentication Protocol
IMSI	International Mobile Station Identifier
IPPR	Intermodulation Product Power Ratio
i-SEC	<i>HC-SDMA</i> Secure Communications Protocol
i-TAP	<i>HC-SDMA</i> Terminal Authentication Protocol
IV	Initialization Vector
IWAN	Interconnection Wide Area Network
K-163	Koblitz Elliptic Curve on $GF(2^{163})$
L2	Layer 2
L2TP	Layer 2 Tunneling Protocol
L3	Layer 3
L3 CM	L3 Connection Management

L3 MMC	L3 Mobility Management and Control
L3 RM	L3 Registration Management
LDAP	Lightweight Directory Access Protocol
LFSR	Linear Feedback Shift Register
LNA	Low Noise Amplifier
LNS	L2TP Network Server
LSB	Least Significant Bit
MAC	Medium Access Control
MSB	Most Significant Bit
PA	Power Amplifier
PCH	Paging Channel
PDCL	Packet Data Conversion Layer
PHY	Physical Layer
PID	Paging Identity
PPM	Parts Per Million
PPP	Point to Point Protocol
PPPoE	PPP over Ethernet
PSS	Packet Services Switch
QoS	Quality of Service
RA	Request Access
RACH	Random Access Channel
RCN	Relative Carrier Number
RSA	Rivest, Shamir, Adleman
RF	Radio Frequency
RFN	Relative Frame Number
RID	Registration Identity
RLC	Radio Link Control
RM	Registration Management
RMU	RLC Message Unit
RRC	Radio Resource Control or Root Raised Cosine
RSSI	Received Signal Strength Indicator
SDU	Service Data Unit
SHA	Secure Hash Algorithm
SINR	Signal-to-Interference plus Noise Ratio
SN	Slot Number
SNR	Signal to Noise Ratio
TCH	Traffic Channel
TDD	Time Division Duplex
TDMA	Time Division Multiple Access
TLVG	Tag Length Value Group
TWAN	Transport Wide Area Network
UM	Unacknowledged Mode
USB	Universal Serial Bus
UT	User Terminal

1.3 Conventions

This section explains the terminology and conventions used in this document.

1.3.1 Measurement Uncertainty

The requirements given in this document are *absolute*. Compliance with these specifications is determined by comparing the measured values to the specified limits without making allowance for measurement uncertainty.

1.3.2 Terminology

The word **shall** or **must** in this document means that compliance with a particular specification is *mandatory*.

The word **should** means that compliance with a particular specification is *recommended*.

1.3.3 Mathematical Expressions

The expression $\lceil x \rceil$ denotes the smallest integer greater than or equal to x .

The expression $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x .

The expression $a \oplus b$ denotes the binary exclusive-or (XOR) of the bits a and b .

The binary expansion a_1, \dots, a_L of a positive integer A is defined so that

$$A = \sum_{l=1}^L a_l 2^{l-1}$$

where $a_l \in \{0,1\}$ and where L is sufficiently large to represent the maximum value of A . The variable a_1 is the *least significant* bit of A , while a_L is the *most significant* bit. In this document the least significant bit (LSB) and most significant bit (MSB) in a binary expansion occur respectively at the left most and right most bit of a binary sequence.

The symbol j is equal to $\sqrt{-1}$.

1.3.4 Message Sequence Chart (MSC) notation

The chapters describing Layers 2 and 3 of the *HC-SDMA* protocol stack frequently use MSCs, also known as *ladder diagrams* or *sequence diagrams*, to illustrate the communication between entities in the BS and UT *HC-SDMA* protocol stacks. This document's use of the MSC notations is modeled after but does not strictly follow the formal MSC language specification of ITU-T Recommendation Z.120. A brief explanation of the graphical form of the MSC language is provided in Figure 1 - MSC Notation.

1.3.5 Reserved Fields

All the reserved fields in the protocol specification must be set to zero, unless otherwise stated.

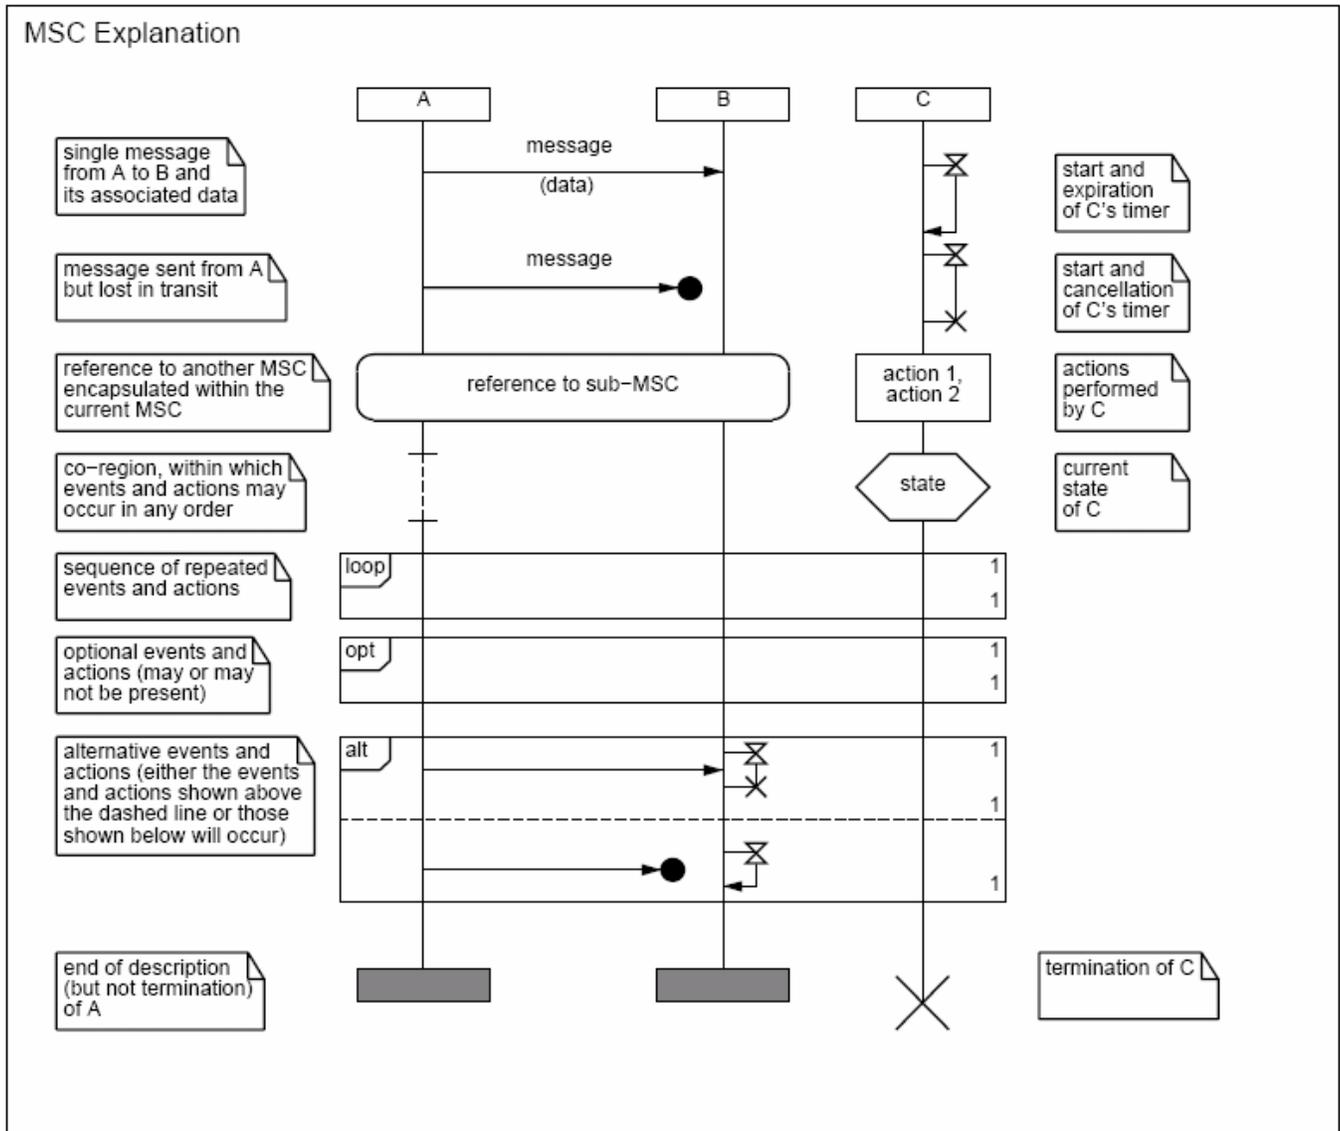


Figure 1: MSC Notation

1.4 Application Overview

This section provides an overview of the intended application and typical deployment scenarios of *HC-SDMA*. This section is an overview of the *HC-SDMA* system and provides context to the protocol specification described in the rest of this document.

1.4.1 Services and Applications

The intended application of *HC-SDMA* is the transport of broadband packet data that provides users with:

- ◆ *Broadband Internet access service:* Always-on Internet access service comparable to DSL and cable.
- ◆ *Mobility:* Anywhere, anytime access with the freedom to move.
- ◆ *High-speed connectivity:* Individual connection speeds of up to 1 Mbps and beyond.
- ◆ *Access through standard devices:* Connections thru standard IP-enabled devices like access routers, laptop and desktop PCs and PDAs

From the end user’s perspective, an *HC-SDMA* system provides high-speed, untethered access to the Internet, virtual private networks (VPNs) and other IP networks from the widest possible range of devices, including laptop computers and PDAs. An *HC-SDMA* system provides a mobile broadband Internet computing and communications experience with peak data rates in excess of 1 Mbps per user.

From an operator’s perspective, an *HC-SDMA* system offers a spectrally efficient, broadband mobile access network. Use of adaptive antenna systems utilizing multiples antennas at the base station (and optionally at the subscriber stations) is integral to the *HC-SDMA* system, resulting in higher coverage, higher data rates, and higher capacity. The wired backhaul and core transport networks utilize open data networking standards and equipment, providing the operator with flexibility in networking technology and vendor and device selection.

1.4.2 Typical Deployment

Figure 2 - Elements in a Typical *HC-SDMA* Deployment depicts the elements directly involved in the transport of user data in a typical end-to-end deployment incorporating the *HC-SDMA* radio access system.

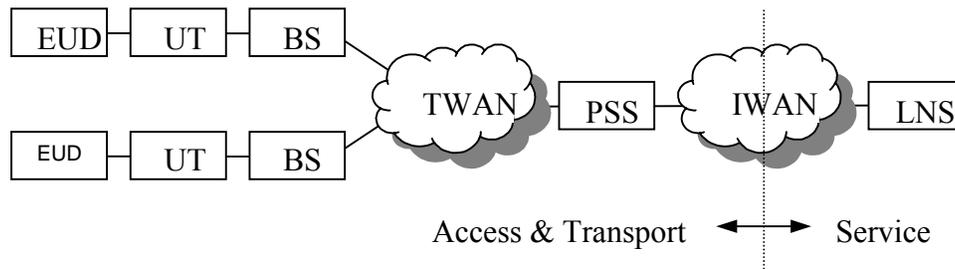


Figure 2: Elements in a Typical *HC-SDMA* Deployment

From left to right, the elements are as follows:

- ◆ An *end-user device (EUD)* such as a laptop computer or personal digital assistant.
- ◆ A *user terminal (UT)*, the end-user *HC-SDMA* modem.
- ◆ A *base station (BS)* that communicates with one or more UTs via the *HC-SDMA* air interface.
- ◆ A *transport wide-area network (TWAN)* connecting one or more base stations to a wired transport network.
- ◆ A *packet services switch (PSS)* that acts as a switch to forward end-user sessions in a particular PSS domain (geographic area) to an appropriate service domain. It also plays a role in providing seamless network-level handovers as UTs experience radio-level handovers from

one BS to another. The PSS can be an L2TP tunnel-switch or a 3GPP2 PDSN; both are supported.

- ◆ An *interconnection wide-area network (IWAN)* interconnecting one or more PSSs with end-user service provider equipment.
- ◆ An *L2TP network server (LNS)* that terminates the end-user PPP sessions originating on the EUDs, providing authentication, authorization, accounting data collection, and COS enforcement.

Figure 2 - Elements in a Typical *HC-SDMA* Deployment only includes those elements of a deployment relating to the flow of end-to-end user data. It does not include the many additional elements present in a complete deployment including platforms for network management, billing, provisioning, content delivery, and so forth.

1.4.3 User Data Transport

End-user IP traffic is conveyed across an *HC-SDMA* deployment through a series of encapsulation/decapsulation steps as depicted in Figure 3 - End-user Traffic Encapsulation. An end-user PPP session bearing an IP header is created between an EUD and an LNS. Each peer-to-peer communication in the diagram is bi-directional and symmetric.

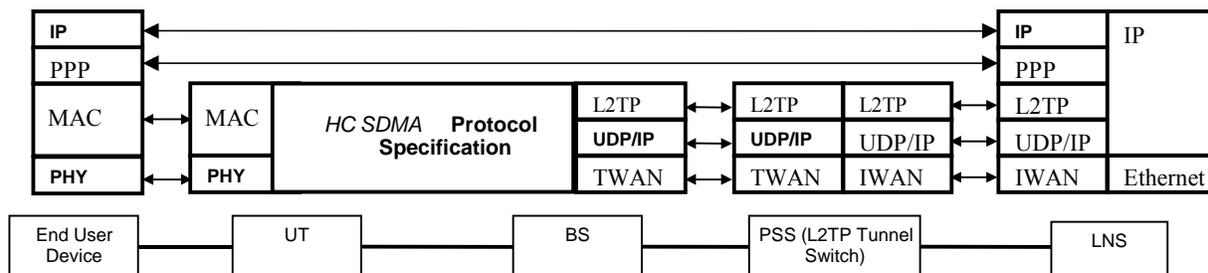


Figure 3: End-user Traffic Encapsulation

Data flows between the UT and the EUD using an encapsulating protocol appropriate to the physical connection between the two entities, generically referred to as the “MAC” and “PHY” layers in the figure. In the case of a USB interface, these are in fact the USB MAC and PHY. Another example is an Ethernet connection, in which case the Ethernet MAC and PHY are used.

Data flows between the UT and the BS over the *HC-SDMA* radio air interface. The UT decapsulates UT-originated traffic from the MAC/PHY encapsulation and then encapsulates it using the *HC-SDMA* air interface encapsulation prior to transmitting it to the base station. The analogous operation is performed at the BS side for network-originated traffic bound for the UT. The *HC-SDMA* protocol specification provides details for the air interface which itself follows a layered protocol model.

1.5 Protocol Overview

1.5.1 Protocol Features

HC-SDMA is designed for the efficient transport of broadband packet data to a large number of users. Key features of the air interface are as follows:

- ◆ TDMA/TDD channel structure.
- ◆ Efficient allocation of traffic resources in response to demand and link conditions.
- ◆ Power control, tiered modulation and forward error control (FEC) to address different link conditions and UT capabilities.
- ◆ ARQ for reliable data delivery.
- ◆ Link-level encryption, BS and UT authentication to guarantee the confidentiality of both user and system control data.
- ◆ Link and network layer handover procedures to ensure seamless continuity of end-user sessions as UTs traverse the access network.
- ◆ Ability to aggregate multiple 625 KHz carriers.
- ◆ Quality of Service (QOS) support.
- ◆ Intrinsic support for adaptive antenna (spatial) processing to achieve high data rates for many users within a cell and network.

1.6 Protocol Reference Model and Interfaces

The HC-SDMA protocol follows a layered structure. For each layer the protocol specifies requirements for actions to be taken in response to and in service of adjoining layers. This layered structure is summarized in Table 1 - Protocol Layers in the HC-SDMA Reference Model.

Table 1: Protocol Layers in the HC-SDMA Reference Model

Layer	Description
L3	Specifications for creating and maintaining logical sessions including registration management, session management, resource control, mobility control, packet fragmentation, slot aggregation, in-band messaging, and security.
L2	Specifications for reliable transmission (RLC), medium access control (MAC), logical channel structure (BCH, PCH, TCH).
L1	Specifications for channelization, burst structure, training data, modulation and FEC, timing.
L0 (RF)	RF specifications for communication over the air link including output power levels, transmit frequency and timing error, pulse shaping, in band and out of band spurious emissions, receiver sensitivity and selectivity, dynamic range.

The HC-SDMA protocol does not provide specifications for applications layers beyond Layer 3.

The messaging and communication interfaces in this layered structure are summarized in Figure 4 - Sublayers and Interfaces in the *HC-SDMA* Protocol Reference Model.

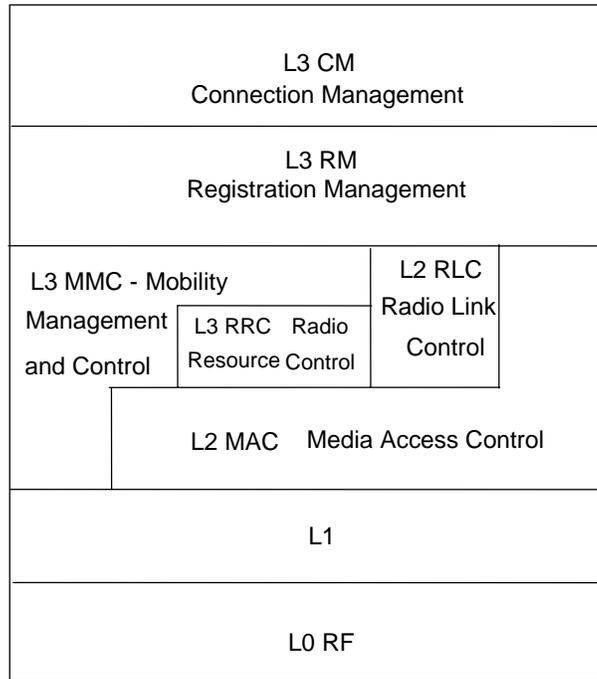


Figure 4: Sublayers and Interfaces in the HC-SDMA Protocol Reference Model

1.6.1 Layer 3 Protocol

The Layer 3 protocol has a number of components with distinct roles in providing transport across the air interface. The components are named as follows:

- ◆ Layer 3 Connection Management (L3 CM);
- ◆ Layer 3 Registration Management (L3 RM);
- ◆ Layer 3 Mobility Management and Control (L3 MMC); and
- ◆ Layer 3 Radio Resource Control (L3 RRC).

1.6.1.1 L3 CM (Layer 3 Connection Management)

This sublayer defines the top-level boundary of the *HC-SDMA* protocol. It provides management of the logical connection between the UT and BS across the air interface. L3 CM receives and sends control and data information to L3 RM.

1.6.1.2 L3 RM (Layer 3 Registration Management)

This sublayer provides management between the virtual connections at L3 CM to physical connections (also known as *streams*) provided by L2 RLC. L3 RM sends (receives) acknowledged mode and unacknowledged mode control and data messages to (from) L2 RLC. L3 RM interfaces to L3 MMC to coordinate registration and deregistration during handovers. In addition, certain control action primitives flow directly between L3 RM and L2 MAC.

1.6.1.3 L3 MMC (Layer 3 Mobility Management and Control)

This sublayer, which exists only at the UT, provides management and control for mobility functionality in the protocol stack. It interfaces to L1 to gather requisite information such as signal levels from various base stations in order to provide the mobility control function to L3 RM. This sublayer is also responsible for tracking frame timing and frequency offset of the user terminals.

1.6.1.4 L3 RRC (Layer 3 Radio Resource Control)

This sublayer interfaces to L1 and to L2 MAC in order to coordinate power control and link adaptation necessary to maintain an RF link.

1.6.2 Layer 2 Protocol

The Layer 2 protocol has two components, L2 RLC and L2 MAC, which control and manage the presentation of reliable streams to L3.

1.6.2.1 L2 RLC (Layer 2 Radio Link Control)

This sublayer provides reliable acknowledge mode (AM) and unreliable unacknowledged mode (UM) peer to peer message delivery services to L3 RM over the L2 MAC sublayer.

1.6.2.2 L2 MAC (Layer 2 Medium Access Control)

This sublayer provides highly dynamic access management and control functions to map and transport logical control and traffic data over various L1 burst types. L2 MAC dynamically controls access to the air interface resources and maps data and control onto these resources available at any time.

1.6.3 Layer 1 Protocol

The Layer 1 protocol enforces the slot and frame structure, modulation and coding, and training sequences necessary for transport of bits over the L0 (RF) interface.

1.6.4 Layer 0 Protocol

The L0 (RF) protocol enforces the RF characteristics of receivers and transmitters in the system including output power levels, transmit frequency and timing error, pulse shaping, in band and out of band spurious emissions, receiver sensitivity, selectivity, and dynamic range.

1.7 Normative References

The following normative references are used within this standard:

- ◆ ITU-T Recommendation Z.120, *Message Sequence Charts (MSC)*; available from the International Telecommunications Union at < <http://www.itu.int/ITU-T/> >.
- ◆ RFC 1662, *PPP in HDLC-like Framing*; available at < <http://www.freesoft.org/CIE/RFC/index.htm> >.

See also clause 11A for references concerning Privacy and Authentication

2 SPECTRAL LAYOUT TERMINOLOGY AND REQUIREMENTS

2.1 Overview

This section defines concepts and specifies requirements for spectral planning in a *HC-SDMA* network.

2.1.1 Definitions for Basic *HC-SDMA* Networks

2.1.1.1 Network Related Definitions

2.1.1.1.1 Basic *HC-SDMA* Networks

A *basic HC-SDMA network* is a wireless communications system operating in accordance with the *HC-SDMA* protocol in a contiguous set of TDD spectrum (up to 10 MHz wide) with an *HC-SDMA* control channel. Basic networks may be combined into *compound HC-SDMA networks* (see 2.1.1.1.4, Compound Networks).

2.1.1.1.2 Carrier Allocation

The *carrier allocation* is the specific set of RF carrier frequencies on which the basic *HC-SDMA* network operates. (A BS in a basic *HC-SDMA* network is not required to support every carrier in the basic network carrier allocation.)

2.1.1.1.3 Frequency Origin and Carrier Numbering

The lowest-frequency carrier for a basic *HC-SDMA* network is called the *frequency origin* (f_{org}). Carriers in the carrier allocation are numbered sequentially from zero, starting with the frequency origin. The RF carrier frequency numbering is the same for every base station (BS) and user terminal (UT) operating in a basic *HC-SDMA* network.

2.1.1.1.4 Compound Networks

A *compound HC-SDMA network* is the combination of several basic *HC-SDMA* networks with disjoint carrier allocations. A particular BS may support more than one basic network within a compound network.

2.1.1.2 Base Station Related Definitions

2.1.1.2.1 Carrier Set

A *carrier set* is a contiguous subset of the basic network carrier allocation that a particular BS may use to support traffic in the basic network. Carrier sets may vary from BS to BS while the basic network carrier allocation remains fixed.

2.1.1.2.2 Relative Carrier Number and Base Carrier

Carriers in a carrier set can be identified by their *relative carrier number*. The relative carriers are numbered sequentially starting from 0.

The *base carrier* is defined as the carrier in the carrier set with relative carrier number zero. A relative carrier is implicitly associated with its carrier set. A particular carrier may therefore have a different relative carrier number from BS to BS and consecutive relative carriers need not correspond to physically adjacent physical carriers. If a carrier is not in the carrier set of a particular BS, then that BS has no relative carrier number for that carrier.

2.1.1.3 UT Related Definitions

2.1.1.3.1 UT Addressable Space

The UT *addressable space* is the set of carriers on which a UT can receive and transmit in accordance to the HC-SDMA specification.

2.1.1.3.2 UT Active Carrier Set

When a UT operates in a basic network, the UT *active carrier set* in the basic network is the set of carriers that the UT uses for reception and transmission with basestations in the basic network.

2.1.2 Requirements

2.1.2.1 Carrier Allocation

The RF carrier frequencies in the carrier allocation shall be a consecutive set of frequencies separated by 625 kHz.

2.1.2.2 Frequency Origin and Preferred Frequencies

The frequency origin should be a *preferred frequency* such that $f_{org} = 312.5\text{kHz} + N \times 625 \text{ kHz}$, where N is a positive integer. For example, an HC-SDMA network operating in a frequency band beginning at 1920 MHz should have a frequency origin of 1920.3125 MHz, corresponding to $N = 3072$. When it is not practical to use such a preferred frequency for f_{org} , the *alternate frequency* $f_{org} = N \times 625 \text{ kHz}$ where N is a positive integer should be used. When it is not practical to use such an alternate frequency for f_{org} , an alternate frequency such that $f_{org} = f_m + N \times 625 \text{ kHz}$ and f_m is selected from among 437.5 kHz, 562.5 kHz, 687.5 kHz, or 812.5 kHz should be used.

2.1.2.3 Carrier Set and Supported Frequencies

A BS providing service in a basic HC-SDMA network shall support only one carrier set within the basic network carrier allocation. An individual BS need not support *all* carriers in the carrier set. The `BCH offset` field in the CM message is the starting index of the BS carrier set.

2.1.2.4 Broadcast Carriers

A broadcast carrier provides the basic network *control channels*, the broadcast and configuration logical channels (BCH and CCH) of the HC-SDMA protocol.

A basic HC-SDMA network shall have at least one *broadcast carrier* within every 10 MHz of the carrier allocation. An HC-SDMA deployment should have a broadcast carrier within every 5 MHz of its carrier

allocation. A broadcast carrier shall not be the lowest or the highest frequency carrier within its carrier allocation.

A broadcast carrier shall support BCH on timeslot 0. A broadcast carrier may support traffic logical channels (RACH, TCH, and PCH) on timeslots not used for control.

2.1.2.5 UT Operation in a Basic Network

In order for a UT to operate in a basic network, the UT addressable space shall contain at least one BCH carrier of the basic network. When a UT operates in a basic network, the UT active carrier set should contain more than one carrier. A UT may operate in more than one basic network within a compound network.

A UT shall not attempt to register on a BS in a basic *HC-SDMA* network unless it is able to transmit and receive on at least one BCH carrier and unless it has received the base station's base carrier. (To receive the base carrier, the UT issues a configuration request and uses the information from the configuration message to initiate registration. See 7.3.4, Configuration Channel (CCH), for details of the CCH channel.)

The UT shall communicate its range of supported carriers to the BS during registration. The BS shall determine whether the UT is suitable for operation within its network and may accept or reject the UT's registration request on this basis. The BS may permit the UT to register and operate even though the UT supports only a subset of the BS carrier set.

Once a UT has registered itself on a basic *HC-SDMA* network, its data traffic is restricted to that basic network until it reregisters itself on another. This means that a UT may only aggregate carriers within the carrier allocation of a single basic *HC-SDMA* network. A UT operating in a basic network is permitted to register itself on another in the same compound network.

3 SLOT AND FRAME STRUCTURE

3.1 Overview

HC-SDMA is a TDD system with 625 kHz allocated to each RF channel. Each RF channel consists of three uplink/downlink time-slot pairs, which together form a *frame*. Each frame has a duration of 5 ms. A transmission on a slot is called a *burst*. Seven different burst types are used to transport control and data messages. The frequency synchronization, timing synchronization, broadcast, page, and standard downlink bursts are transmitted in the downlink, while the configuration request and standard uplink bursts are transmitted in the uplink. Each burst is partitioned into ramp-up, training, payload, ramp-down, and interburst guard.

3.2 RF Channel and Frame Structure

3.2.1 RF Channels

An *HC-SDMA* spectral allocation is divided into one or more RF channels of bandwidth 625 kHz each. Figure 5 - RF Channelization for a 5 MHz Frequency Allocation depicts an allocation of 5 MHz divided into eight RF channels. The symbol rate is 500×10^3 symbols per second. The symbol duration is 2 μ s. See clause 5, User Terminal Radio Transmission and Reception, and clause 6, Base Station Radio Transmission and Reception for specifications on pulse shapes and spectral masks.

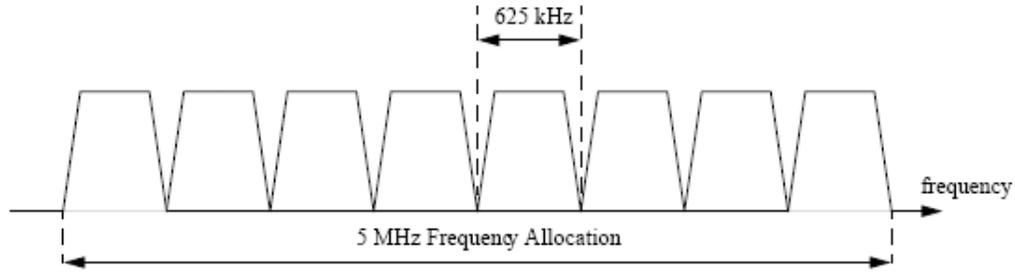


Figure 5: RF Channelization for a 5 MHz Frequency Allocation

3.2.2 Frames and Slots

RF channels are divided into 5 ms frames with the structure shown in Figure 6 – The HC-SDMA Frame Structure (not to scale). A frame consists of:

- ◆ Three 545 μ s uplink slots for communication from UT to BS,
- ◆ A transition guard time of 10 μ s between uplink and downlink,
- ◆ Three 1090 μ s downlink slots for communication from BS to UT,
- ◆ An extended transition guard time of 85 μ s between downlink and uplink for range extension.

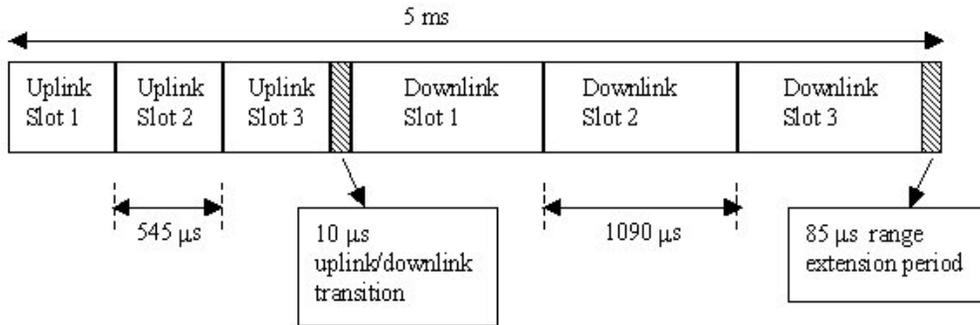


Figure 6: The HC-SDMA Frame Structure (not to scale)

For the purpose of paging, downlink slots are further divided into two equal sized subslots. Every slot includes a ramp-up and ramp-down period of 10 μ s. Uplink slots end with 15 μ s of interburst guard time. Downlink slots end with 14 μ s of interburst guard time. Thus an uplink slot allocates a maximum of 510 μ s for training and information symbols while a downlink slot allocates a maximum of 1056 μ s for training and information symbols.

3.2.3 Burst Types

A transmission on a slot or subslot is referred to as a *burst*. A burst type is defined by the partitioning of a slot into ramp-up, training, payload, ramp-down, and interburst guard intervals. Each burst type also has an associated modulation scheme for transmission in an uplink or downlink slot as appropriate. The physical layer provides seven burst types to transport control and data messages. The burst types

and the corresponding logical channels (defined in clause 7, L2 MAC Protocol Sublayer Specification) are listed in Table 2 – Burst Types.

Table 2: Burst Types

Burst Type	Symbol for Burst Type	Logical Channel
<i>Downlink Bursts:</i>		
Frequency synchronization	F	BCH
Timing synchronization	T	BCH
Broadcast	B	BCH
Page	P	PCH
Standard downlink	D	FACCH, RACH, CCH, TCH
<i>Uplink Bursts:</i>		
Configuration request	C	CCH
Standard uplink	U	FACCH, RACH, TCH

TCH and RACH resources are always allocated in uplink/downlink burst pairs. If the UT sends a message on the RACH logical channel on a particular uplink time slot, the response from the BS on the RACH logical channel occurs on the corresponding downlink burst in the same frame. Analogously, a TCH logical channel is realized as a particular pair of corresponding uplink/downlink bursts in a consecutive sequence of frames on a particular conventional channel.

3.2.4 BCH/CCH Superframe

At least one timeslot at each *HC-SDMA* base station is reserved for the BCH and CCH logical channels. Frames on this timeslot are organized according to the superframe structure (with a period of 20 frames) shown in Figure 7 – *HC-SDMA* BCH/CCH Carrier Superframe Structure. The downlink D-bursts in the superframe carry the downlink portion of the CCH. All the other carriers and timeslots are used for user traffic and associated control messages. Each BS should have at least one BCH carrier in every 5 MHz of spectrum and the carrier frequency should be the same on all BSs in a network.

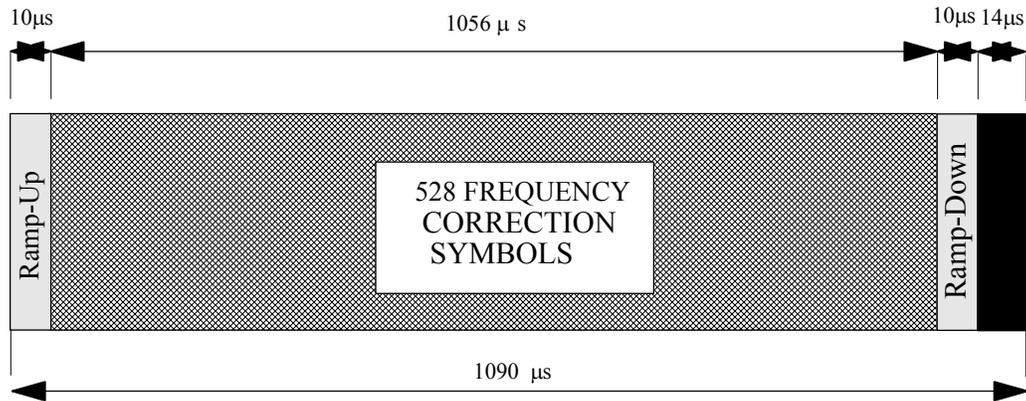


Figure 8: Frequency Synchronization Burst Structure

Table 3: Frequency Synchronization Burst Fields

Duration	Contents
10 μs	ramp-up
1056 μs	frequency correction symbols f_1, f_2, \dots, f_{528}
10 μs	ramp-down
14 μs	interburst guard time

The frequency correction symbols are given by:

$$f_k = \beta_1 \cdot \exp(j\pi k \times 11/16) + \beta_2 \cdot \exp(-j\pi k \times 11/16)$$

for $k = 1, 2, \dots, 528$, where β_1 and β_2 are complex constants selected by the BS.

3.3.2 Timing Synchronization

The timing synchronization burst structure is shown in Figure 9. The burst fields are described in Table 4.

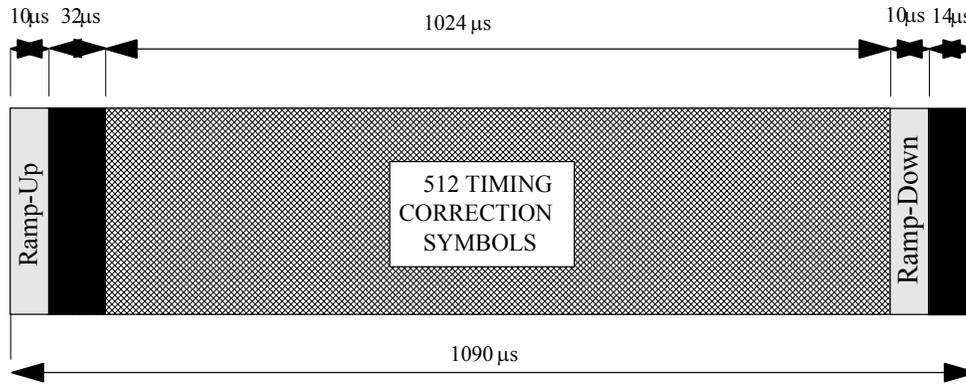


Figure 9: Timing Synchronization Burst Structure

Table 4: Timing Synchronization Burst Fields

Duration	Contents
10 µs	ramp-up
32 µs	timing preamble r_1, r_2, \dots, r_{16}
1024 µs	timing correction symbols t_1, t_2, \dots, t_{512}
10 µs	ramp-down
14 µs	interburst guard time

The timing preamble is the sequence below:

$$r_1 = \exp(j1\pi/2), r_2 = \exp(j2\pi/2), r_3 = \exp(j3\pi/2), r_4 = \exp(j0\pi/2)$$

$$r_5 = \exp(j1\pi/2), r_6 = \exp(j2\pi/2), r_7 = \exp(j3\pi/2), r_8 = \exp(j0\pi/2)$$

$$r_9 = \exp(j3\pi/2), r_{10} = \exp(j2\pi/2), r_{11} = \exp(j1\pi/2), r_{12} = \exp(j0\pi/2)$$

$$r_{13} = \exp(j3\pi/2), r_{14} = \exp(j2\pi/2), r_{15} = \exp(j1\pi/2), r_{16} = \exp(j0\pi/2)$$

The timing correction symbols are constructed from the base station color code. The encoding is defined in clause 4.

3.3.3 Broadcast Burst

The broadcast burst structure is shown in Figure 10. The burst fields are described in Table 5.

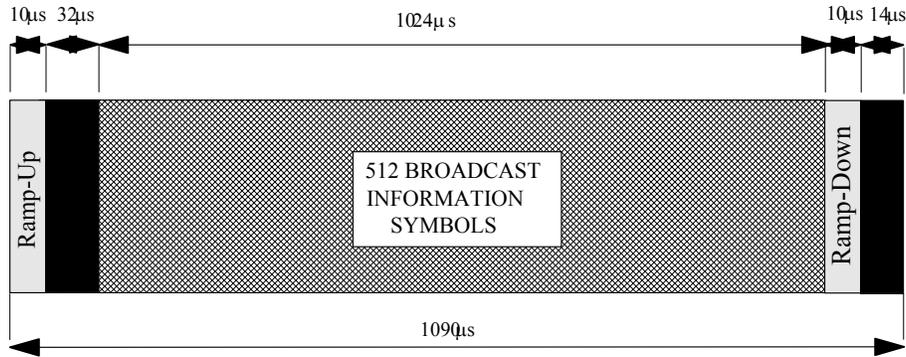


Figure 10: Broadcast Burst Structure

Table 5: Broadcast Burst Fields

Duration	Contents
10 μs	ramp-up
32 μs	broadcast preamble r_1, r_2, \dots, r_{16}
1024 μs	broadcast information symbols b_1, b_2, \dots, b_{512}
10 μs	ramp-down
14 μs	interburst guard time

The broadcast preamble is the sequence below:

$$r_1 = \exp(j1\pi/2), r_2 = \exp(j2\pi/2), r_3 = \exp(j3\pi/2), r_4 = \exp(j0\pi/2)$$

$$r_5 = \exp(j3\pi/2), r_6 = \exp(j2\pi/2), r_7 = \exp(j1\pi/2), r_8 = \exp(j0\pi/2)$$

$$r_9 = \exp(j1\pi/2), r_{10} = \exp(j2\pi/2), r_{11} = \exp(j3\pi/2), r_{12} = \exp(j0\pi/2)$$

$$r_{13} = \exp(j3\pi/2), r_{14} = \exp(j2\pi/2), r_{15} = \exp(j1\pi/2), r_{16} = \exp(j0\pi/2)$$

The broadcast information symbols are constructed from the 6-bit broadcast message using the modulation and coding method described in 4.4.3.

3.3.4 Page Burst

The page burst structure is shown in Figure 11. Page bursts may be transmitted in both subslots, only the first subslot, or only the second subslot of a downlink slot. The burst fields for these three cases are described in Table 6, Table 7, and Table 8, respectively.

Table 6: Page Burst in Both Subslots

Duration	Contents
10 μs	ramp-up
512 μs	page information symbols h_1, h_2, \dots, h_{256}

Duration	Contents
10 μ s	ramp-down
12 μ s	sub-slot guard time
10 μ s	ramp-up
512 μ s	page information symbols $h'_1, h'_2, \dots, h'_{256}$
10 μ s	ramp-down
14 μ s	interburst guard time

Table 7: Page Burst in First Subslot

Duration	Contents
10 μ s	ramp-up
512 μ s	page information symbols h_1, h_2, \dots, h_{256}
10 μ s	ramp-down
544 μ s	extra guard time
14 μ s	interburst guard time

Table 8: Page Burst in Second Subslot

Duration	Contents
544 μ s	extra guard time
10 μ s	ramp-up
512 μ s	page information symbols h_1, h_2, \dots, h_{256}
10 μ s	ramp-down
14 μ s	interburst guard time

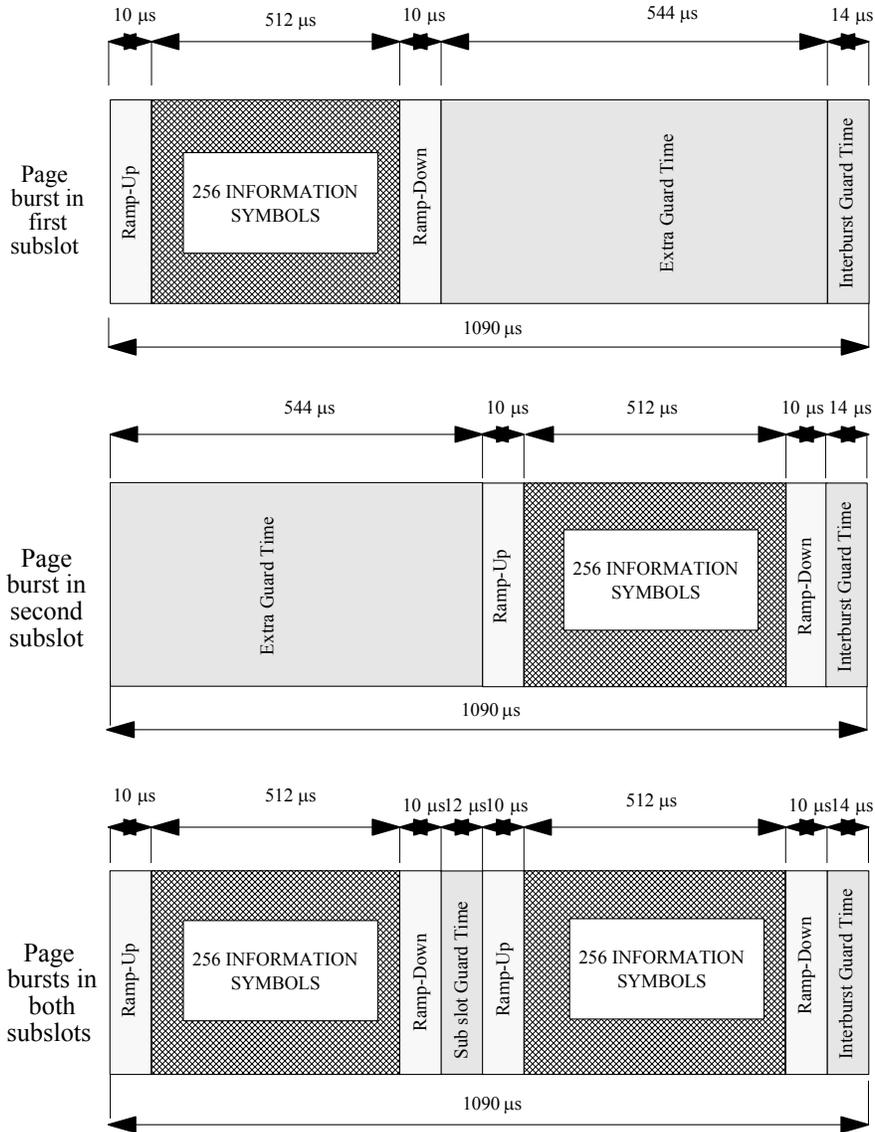


Figure 11: Division of a Downlink Slot into Two Page Bursts

The page information symbols are constructed from the 15-bit page message using the modulation and coding method described in 4.3.4.

3.3.5 Configuration Request Burst

The configuration request burst structure is shown in Figure 12. The burst fields are described in Table 10.

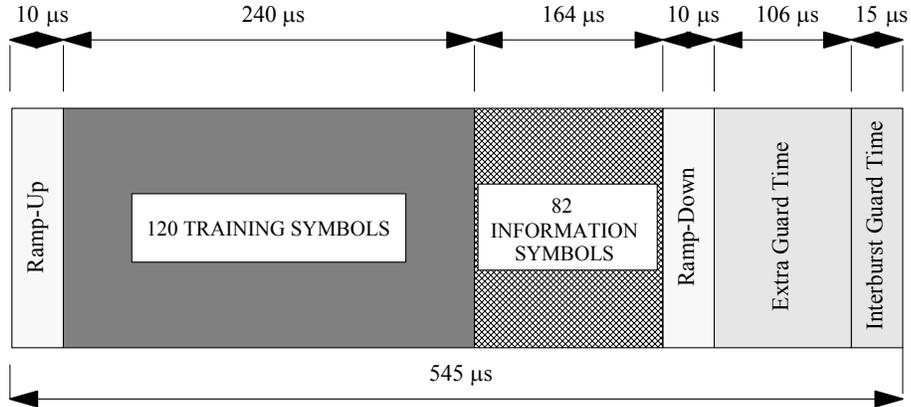


Figure 12: Configuration Request Burst Structure

The training symbols for the configuration request burst are based on either a “primary” or “secondary” core sequence, defined in Table 9. Base stations with an even color code use the primary sequence, while the other base stations use the secondary sequence.

Table 9: Configuration Request Core Sequences

core sequence index	1	2	3	4	5	6	7	8	9	10	11	12
primary	1	1	1	1	1	0	0	1	0	0	0	0
secondary	1	1	1	0	0	0	1	1	0	0	0	1

Training sequence symbols (\tilde{a}_k 's) are formed by repeating 10 times a core sequence (c_k) of period 12 symbols, taking one's-complement of the last repetition, and transmitting the resulting sequence by $\pi/2$ - BPSK modulation.

For $\{c_1, c_2, \dots, c_{12}\}$, the chosen core sequence, $l = 0, 1, \dots, 9$ and $k = 1, 2, \dots, 12$,

$$\tilde{a}_{12l+k} = \begin{cases} c_k & \text{if } l < 9, \\ 1 \oplus c_k & \text{if } l = 9. \end{cases}$$

Table 10: Configuration Request Burst Fields

Duration	Contents
10 μs	ramp-up
240 μs	training symbols a_1, a_2, \dots, a_{120}
164 μs	information symbols p_1, p_2, \dots, p_{82}
10 μs	ramp-down
106 μs	extra guard time
15 μs	interburst guard time

The 120 training symbols are constructed by applying $\pi/2$ - BPSK modulation to $\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_{120}$:

$$a_k = \begin{cases} \exp(j0\pi/2) & \text{if } \tilde{a}_k = 1 \text{ and } k \text{ even,} \\ \exp(j2\pi/2) & \text{if } \tilde{a}_k = 0 \text{ and } k \text{ even,} \\ \exp(j1\pi/2) & \text{if } \tilde{a}_k = 1 \text{ and } k \text{ odd,} \\ \exp(j3\pi/2) & \text{if } \tilde{a}_k = 0 \text{ and } k \text{ odd.} \end{cases}$$

for $k = 1, 2, \dots, 120$.

The 82 information symbols p_1, p_2, \dots, p_{82} are constructed from the configuration request message using modulation class 0, defined in *clause 4 Modulation and Channel Coding*.

A *pseudo-random* time delay is applied to the configuration request burst, where the timing is delayed, not advanced. This is done to avoid repeated collisions of CR bursts from UTs at the same range. For each configuration burst, the user terminal selects a timing delay pseudo-randomly with uniform probability from the set:

$$\{0, 2, 4, 6, 8, 10, 12, 14, 16, 18\} \mu\text{s}.$$

3.3.6 Standard Uplink Burst (FAACH, RACH, & TCH)

The standard uplink burst structure is shown in Figure 13. The uplink RACH and the uplink traffic channel (TCH) adhere to the standard uplink burst format. The burst fields are described in Table 11.

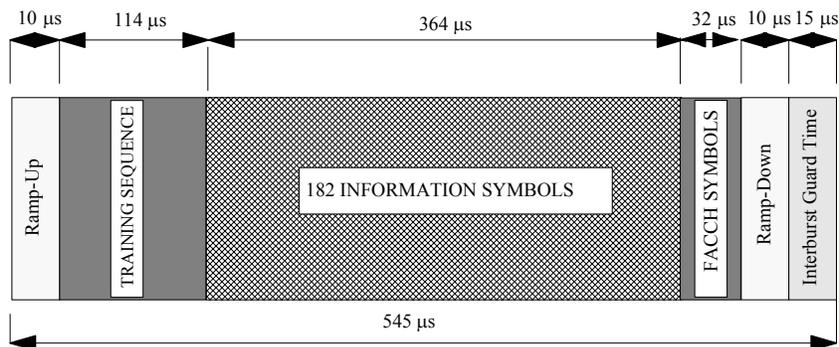


Figure 13: Standard Uplink Burst Structure

Table 11: Standard Uplink Burst Fields

Duration	Contents
10 μ s	ramp-up
114 μ s	training sequence a_1, a_2, \dots, a_{57}
364 μ s	information symbols p_1, p_2, \dots, p_{182}
32 μ s	FACCH symbols b_1, b_2, \dots, b_{16}
10 μ s	ramp-down
15 μ s	interburst guard time

The standard uplink burst carries a 57-symbol long training sequence at the beginning of the burst. The symbols for the training sequence are selected from a family of 160 sequences indexed by the five least significant bits $BSCC_{0-4}$ of the base station color code (32 values) and by a training offset t_{Offset} (5 values), both provided by higher layers.

The training sequence is constructed as follows:

1. *Select Core Sequence:* Use the five least significant bits of the base station color code $BSCC_{0-4}$ to select one of the 32 candidate core sequences from Table 12, denoted c_1, c_2, \dots, c_{51}
2. *Shift and Append:* Cyclically shift the core sequence by an amount l determined from t_{Offset} according to Table 13. (Uplink RACH bursts use t_{Offset} equal to zero, while TCHs on the spatial channels use other time shifts provided by higher layers), and extend it by 6 symbols to obtain the appended core sequence \tilde{c}_k of length 57 symbols:

$$\tilde{c}_k = c_{(k+l-3) \bmod 51}$$

for $k = 1, 2, \dots, 57$ and $x \bmod y = x - y \cdot (\lceil x/y \rceil - 1)$.

Table 12: Core Sequences for Standard Uplink Burst

BSCC0-4	c1...c10	c11...c20	c21...c30	c31...c40	c41...c51
0	0000110010	1101111010	1000100111	0100000100	11000001010
1	0110110011	1101001010	1110001000	0111101111	11100011011
2	0011100011	1011000000	1111010110	0111100101	10010001000
3	1101101011	0110000101	0100111111	1101110111	10011100000
4	0101111100	1110101111	0100001110	1111001100	01001010110
5	0101001100	1000000000	0111001110	0011010101	11100000011
6	1011110011	1110010100	0011001111	1101001011	10101010101
7	0010000101	1001000100	1010101010	0111110011	11001100111
8	1000000100	1101011101	1111110000	0011000010	10100001001
9	0010100111	0110100100	1100110010	1010000010	10000011110
10	1011010101	1011100111	1000000001	0111010111	00110101100
11	0111011101	1001001101	1110001101	1110110101	11000010101
12	0111001100	0111101111	1010101011	0110100000	01110110111
13	0000111100	0100100011	1110100011	1011010110	01000000001
14	1110111101	1101000001	1111100110	0100111101	01001011001
15	0110100001	0100000101	1011010000	0100111010	001111100110
16	0001111100	0011010100	0000010001	0100110010	10100100111
17	0100011010	1100101011	0001110010	1111111110	01000010011
18	1110011010	1010010110	1011111101	0011110111	00010000111
19	1001100010	1111011001	1010001100	0001111101	01011111111
20	0000001100	0010001100	1101100011	0100110100	00001011111
21	0100100010	1100000111	0111001011	1111001011	00111011111
22	0111010110	1000100011	1010001010	1111101111	10011001000
23	1010111100	0010011110	0100101100	0000001000	10010101011
24	1101001011	1100101100	1100001100	0100010111	10100000000
25	0000100100	1110100110	0000101001	0001111011	00101001101
26	0001010001	0111001001	1111011000	1010010010	10010001100
27	1011010100	1001000011	0111010100	1111100110	10100001111
28	0010110100	1101001111	0110111111	0110001101	01100111010
29	1001101110	0011001001	0010100011	1101111110	11001011111
30	1011111000	1111111010	0111001000	1001011101	10101000101
31	0000101110	1101101100	1011111011	1101111000	01001100111

Table 13: Cyclic Shift for Uplink Training

tOffset	<i>l</i>
0	0
1	30
2	20
3	40
4	10

1. *Sequence Modification:* Modify the appended core sequence \tilde{c}_k , i.e., for $k = 1, \dots, 57$, by performing the following toggle operation based on the values of k and l :

$$\tilde{a}_k = \begin{cases} \tilde{c}_k & \text{if } (k+l) \text{ umod } 4 < 3 \\ 1 - \tilde{c}_k & \text{otherwise} \end{cases}$$

where *umod* is defined in Step 2.

2. $\pi/2$ - BPSK modulation: Apply $\pi/2$ - BPSK modulation ($j^2 = -1$):

$$a_k = \begin{cases} \exp(j0\pi/2) & \text{if } \tilde{a}_k = 1 \text{ and } k \text{ is odd,} \\ \exp(j2\pi/2) & \text{if } \tilde{a}_k = 0 \text{ and } k \text{ is odd,} \\ \exp(j1\pi/2) & \text{if } \tilde{a}_k = 1 \text{ and } k \text{ is even,} \\ \exp(j3\pi/2) & \text{if } \tilde{a}_k = 0 \text{ and } k \text{ is even} \end{cases}$$

for $k = 1, 2, \dots, 57$.

The Fast Associated Control Channel (FACCH) symbols in Figure 13 – Standard Uplink Burst Structure carry a four bit FACCH message v_1, v_2, v_3, v_4 . The FACCH symbols are constructed as follows:

1. *Select Walsh-Hadamard Sequence:* $v = v_1 + 2 \times v_2 + 4 \times v_3 + 8 \times v_4$ is the integer representation of the FACCH message and l_1, l_2, \dots, l_{16} is the $(v + 1)$ -th row of the Hadamard matrix H_4 defined in 4.3.2.
2. *Select Scrambling Sequence:* The scrambling sequence s_1, s_2, \dots, s_{16} is selected from a family of 160 sequences, shown in Table 14, indexed by the five least significant bits $BSCC_{0-4}$ of the base station color code (32 values) and by a training offset *tOffset* (5 values), both provided by higher layers. The scrambling sequence corresponding to *tOffset* equal to 0 is used only by the RA. For $k = 1, 2, \dots, 16$, compute:

$$b_k = l_k \oplus s_k$$

3. $\pi/2$ - BPSK modulation: For $k = 1, 2, \dots, 16$ and $j^2 = -1$:

$$b_k = \begin{cases} \exp(j0\pi/2) & \text{if } \check{b}_k = 1 \text{ and } k \text{ even,} \\ \exp(j2\pi/2) & \text{if } \check{b}_k = 0 \text{ and } k \text{ even,} \\ \exp(j1\pi/2) & \text{if } \check{b}_k = 1 \text{ and } k \text{ odd,} \\ \exp(j3\pi/2) & \text{if } \check{b}_k = 0 \text{ and } k \text{ odd.} \end{cases}$$

Table 14: Scrambling Sequences for Uplink FACCH symbols

BSCC0-4	tOffset = 0	tOffset = 1	tOffset = 2	tOffset = 3	tOffset = 4
0	0001001001000111	0001111011100001	0001100011011011	0000010101011111	0000000001101001
1	0000110001100101	0000101001010000	0011010111001010	0011001100001111	0010001010001000
2	0011010100110101	0011011001101100	0011000010100110	0011110011111111	0010001001110111
3	0010110111011101	0010111010000100	0101010101010101	0101011000111111	0101000010010011
4	0010110100101101	0101011011111100	0101000010100000	0101001100111010	0101101010011001
5	1001010111111100	1001011001010101	1001000001101111	1001110011000110	1000011111010010
6	1001110001101100	1001111100001001	1001100101011010	1000011110000111	1000000110110010
7	1101110111010010	1101111001110100	1100000001010110	1100001111000011	1100100111110101
8	1101110100101101	1101101111100111	1100000011110011	1100001110010110	1100100110100000
9	1101110101111000	1101111000011101	1100001100111100	1100111110011010	1100100101011111
10	1101110101110111	1101111000010010	1100001100000000	1100111110100110	1100100101101100
11	1111011000001001	1111000011001100	1111001101100101	1111101010100000	1110000101001011
12	1111011001100000	1111000000110011	1111110001010110	1111101011110101	1110000111100001
13	1010111111000101	1010100110010110	1010010111110011	1011111010001110	1011100001000100
14	1010101010010101	1010011000110011	1010000000001001	1011110110001101	1011100010111011
15	1010010101100101	1010011011001100	1010000011110110	1011101100100001	1011011101001011
16	1011101100011101	1011011110001000	1011000100100100	1011001010110001	1001110000110101
17	1011010001110100	1011000111101000	1000111001110010	1000101100011110	1001100100110000
18	1011010010110111	1011001010000010	1000110100011000	1000101111011101	1001101001011010
19	1000110111100111	1000111010111110	1000100001110100	1000010000101101	1001101010100101
20	1001100110100110	1001101011110000	1001001111000101	1110110111101110	1110111001110100
21	1001010100001111	1001011001010110	1110110110000111	1110111011101101	1110100001000001
22	1001010111111111	1110111000101110	1110100001110010	1110101111101000	1110001001001011
23	1111001110011001	1100110000001100	1100100101100000	1101110110111000	1101111000100010
24	1100110001100101	1100111111110000	1100100101011100	1101111011011101	1101100011101000
25	1100011010100011	1100001110010101	1101010011010111	1101000110111011	1101001011101101
26	1101001011010001	0010111001000100	0010100011100111	0010101101111101	0011110011111100
27	0010110100101110	0010111010000111	0010100010111101	0010010000010100	0011111100000000
28	0010000100100010	0010001001000111	0011111110101010	0011100101101111	0011001100001100
29	0010000111011101	0010001010000100	0011100110101100	0011000000001111	0011001101100101
30	0011100110010000	0011000001011010	0000111100110000	0001110100011110	0001111001111011
31	0011011010011111	0011000010100101	0000000000111111	0001110111100001	0001111010000100

3.3.7 Standard Downlink Burst (CM, AA, & TCH)

The standard downlink burst structure is shown in Figure 14. The configuration message (CM), access assignment (AA), and the downlink traffic channel (TCH) bursts adhere to the standard downlink burst format. The burst fields are described in Table 15

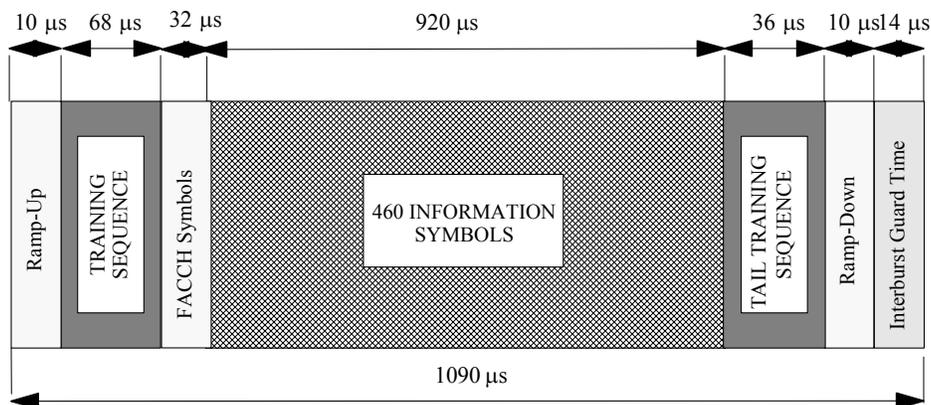


Figure 14: Standard Downlink Burst Structure

Table 15: Standard Downlink Burst Fields

Duration	Contents
10 μs	ramp-up
68 μs	training sequence a_1, a_2, \dots, a_{34}
32 μs	FACCH symbols g_1, g_2, \dots, g_{16}
920 μs	information symbols p_1, p_2, \dots, p_{460}
36 μs	tail training sequence h_1, h_2, \dots, h_{18}
10 μs	ramp-down
14 μs	interburst guard time

Table 16: Cyclic Shift for Downlink Training

tOffset	l
0	0
1	16
2	11
3	6
4	21

The standard downlink burst carries two training sequences, one in the beginning of the burst and one in the end. The training sequence at the beginning is 34 symbols long and the tail training sequence is 18 symbols long. The symbols for each of the training sequences are selected from two families of 160 sequences indexed by the five least significant bits $BSCC_{0-4}$ of the base station color code (32 values) and by a training offset $tOffset$ (5 values), both provided by higher layers.

The training symbols (a_k 's) at the beginning of the burst are constructed as follows:

1. Use the five least significant bits $BSCC_{0-4}$ of the color code to select the core sequence c_1, c_2, \dots, c_{26} from Table 17.

2. Cyclically shift the core sequence by an amount l determined from $tOffset$ according to Table 16. (The training sequence with $tOffset$ equal to 0 is used for CM and AA. The other delays are used for TCH), and extend it by 8 symbols:

$$\tilde{a}_k = c_{(k+l-4) \bmod 26}$$

for $k = 1, 2, \dots, 34$, where $x \bmod y = x - y \cdot (\lceil x/y \rceil - 1)$.

3. Apply QPSK modulation:

$$a_k = \exp(j(2\tilde{a}_k + 1)\pi/4)$$

for $k = 1, 2, \dots, 34$.

Table 17: Core Sequences for Standard Downlink Burst

BSCC₀₋₄	c_{1...c₈}	c_{9...c₁₆}	c_{17...c₂₆}
0	11322230	21101012	0031210200
1	12211210	33112003	3020202001
2	11302212	23123232	0311103310
3	11332003	22332302	3331020201
4	22200020	02202021	2010331112
5	22230211	01011131	3210021230
6	00010213	01231332	1121203320
7	30333122	21023301	3021300320
8	33220300	31212333	0223120310
9	01323211	11220102	1013302113
10	22102330	30333121	1302032001
11	21213310	22322131	2301000330
12	03121233	33220302	2313311131
13	22332301	13320032	3131010013
14	20012211	21033112	2020100020
15	22020222	20002003	2332130101
16	33010311	10203201	2231211101
17	11230313	10023003	1300103222
18	31120012	21121030	3310131300
19	23221321	21331023	1233230000
20	23231130	22122310	2200010323
21	31110203	20330100	0120301332
22	23133000	10213010	2210023221
23	00010213	01231333	0303211002
24	21111220	10013231	3200220100
25	00030231	03213112	2312232231
26	12031012	11312021	1001133331
27	10131300	21003323	3020011133
28	03233132	02320133	0220000110
29	01123031	31002301	0313321111
30	11000302	31032130	2011323300
31	32023201	30323312	2330031100

The tail training sequence symbols (h_k 's) are constructed as follows:

1. Use the five least significant bits BSCC₀₋₄ of the color code to select the core sequence $\{d_1, d_2, \dots, d_{18}\}$ from Table 19.

2. Cyclically shift the sequence $\{d_1, d_2, \dots, d_{18}\}$ by an amount l determined from t_{Offset} according to Table 18. (The training sequence with t_{Offset} equal to 0 is used for CM and AA. The other delays are used for TCH):

$$f_k = \begin{cases} d_{k+l}, & \text{if } k+l \leq 18 \\ d_{k+l-16}, & \text{if } k+l > 18 \end{cases}$$

for $k = 1, 2, \dots, 18$, and $l \in \{0, 10, 4, 7, 13\}$.

3. Apply QPSK modulation:

$$h_k = \exp(j\pi(2f_k + 1)/4)$$

for $k = 1, 2, \dots, 18$.

Table 18: Cyclic Shift for Downlink Tail Training Sequences

tOffset	l
0	0
1	10
2	4
3	7
4	13

Table 19: Tail Training Sequences for Standard Downlink Burst

BSCC₀₋₄	<i>d</i>_{1...d₆}	<i>d</i>_{7...d₁₂}	<i>d</i>_{13...d₁₈}
0	120022	020222	200012
1	330001	023103	231133
2	110003	021301	213311
3	121021	103231	112012
4	112003	021323	233311
5	302200	000020	222030
6	122010	011030	111312
7	321223	101211	132232
8	123221	303233	312212
9	322312	103313	002132
10	311010	131100	232031
11	102332	301313	000310
12	331332	023221	121233
13	113112	021223	323211
14	100232	031232	133310
15	311010	331122	210231
16	133030	113322	230213
17	322030	033010	333132
18	311210	133300	030231
19	133230	311100	010213
20	320020	000220	020232
21	332001	023121	211133
22	323023	301213	332032
23	303221	301211	112030
24	303233	332001	223130
25	133030	313300	212013
26	122132	301131	002312
27	322232	231012	331132
28	202130	101113	002120
29	211230	313100	230021
30	220132	123111	002122
31	012223	021321	211101

The FACCH symbols in Figure 14 carry a four-bit message v_1, v_2, v_3, v_4 . The FACCH symbols are constructed as follows:

1. *Select Walsh-Hadamard Sequence:* $v = v_1 + 2 \times v_2 + 4 \times v_3 + 8 \times v_4$ is the integer representation of the

FACCH message and l_1, l_2, \dots, l_{16} is the $(v + 1)$ -th row of the Hadamard matrix H_4 defined in 4.3.2.

2. *Select Scrambling Sequence:* The scrambling sequence s_1, s_2, \dots, s_{16} is selected from a family of 160 sequences, shown in Table 20 indexed by the five least significant bits $BSCC_{0-4}$ of the base station color code (32 values) and by a training offset $tOffset$ (5 values), both provided by higher layers. The scrambling sequence corresponding to $tOffset$ equal to zero is used only by the CM and the AA.
3. *QPSK modulation:* For $k = 1, 2, \dots, 16$:

$$g_k = (2l_k - 1) \cdot \exp(j\pi(2s_k + 1)/4)$$

Table 20: Scrambling Sequences for Downlink FACCH Symbols

BSCC0-4	tOffset = 0	tOffset = 1	tOffset = 2	tOffset = 3	tOffset = 4
0	1212102113331221	1212100333311023	1212122313113223	1210100133113221	1210122133333023
1	2023311221012000	2023331021232220	2021313223232200	2021333021030000	2021331223010220
2	1300003222212013	1300023022032233	1302001220032213	1302021022230013	1302023220210233
3	1012030213212212	1012032033010210	1012010033232230	1012230013030232	1010032213010032
4	3132210131010012	3132212311212010	3132230311030030	3132010331232032	3130212131212232
5	3030211212331121	3030213010313123	3032011210111123	3032211030311101	3032213232113121
6	2013323310203210	2013301310021230	2031301330003210	2033121112021030	2033123310221232
7	3200103211302112	3200123031320312	3222301211302332	3222321013100132	3220323013302312
8	1312233312221333	1312213132203133	1330031312221113	1330011110023313	1332013110221133
9	3121133310011001	3121131112211203	3121113110231221	3123331132213023	3123313132033221
10	0333300320000303	0333302122200101	0333320120220123	0331102102202321	0331120102022123
11	2130131233012221	2130133031212023	2132313231232003	2132133211032223	2132111233030203
12	0203211112121030	0021013330101032	0021031332123232	0021033112101212	0021233332301230
13	0102111221101213	0102131021321033	0100113223321013	0100133021123213	0100131223103033
14	1300300331213321	1300320131033101	1302302333033121	1302322131231321	1302320333211101
15	2010113331000010	2010111133200212	2012331333220232	2012111313020012	2012133331022032
16	3033221320212322	3033203302212120	3033201122010122	3033001300010320	3031203122212302
17	3332121110200003	3332301130000223	3332323130222203	3330103310222003	3330123110002223
18	0002332202232131	0002330000210133	022013222230131	0220130020032111	0220112000230311
19	3102002102123233	3102020100103211	3102200102303013	3102222122103031	3100020322123231
20	3012032202320301	3012212222300103	3012230202102303	3012232022122101	3032010002122301
21	1012330313313011	1030112133333213	1030132331131013	1032112313113013	1032110133131033
22	0023121312012220	0023123132210222	0023101132032202	0023321112230200	0021123312210000
23	310322233101203	3103220031123201	3101022231321201	3101222011121223	3101220213323203
24	3231200301122000	3231202103100002	3233000303302002	3233200123102020	3233202321300000
25	0022013100301203	0022011302101001	0020231102121021	0020011122321201	0020033100323221
26	3322311122302111	3322333120322133	3322113122122331	3322131102322313	3320333302302113
27	1303210320232100	1303232320012302	1321010100030322	1321032122030120	1323032302030302
28	1020033030013133	1020031232031131	1022233032233131	1022033212033113	1022031010231133
29	2013012102221312	2013030120223332	2011232320203132	2011230122003330	2011032122223110
30	3223302112123201	3223300310101203	3223322332103223	3223102330121221	3221320310123221
31	3311031013322211	3333213233302013	3333233031100213	3331213013122213	3331211233100233

4 MODULATION AND CHANNEL CODING

4.1 Overview

A symbol rate of 500 kSymbols/second is used in all burst types.

The standard uplink and downlink bursts employ coding and modulation schemes to provide different data rates and minimize ARQ retransmissions. *HC-SDMA* defines nine such modulation classes (or ModClasses). The adaptation mechanism between ModClasses (link adaptation) operates in concert with power control algorithms. Downlink and uplink ModClass selections are performed independently.

A rate-1/2 convolutional encoder provides the channel coding. In addition, some of the coding schemes employ puncturing for increased data rates and some coding schemes employ block coding in addition to convolutional coding. The modulation types used are $\pi/2$ - BPSK, QPSK, 8-PSK, and QAM. Each burst payload is encrypted, interleaved, and a cyclic redundancy check of the encrypted message is included for error detection. Each burst is scrambled before transmission.

End-user data throughput varies depending on both the ModClass and the number of aggregated streams. Table 21 shows the user data throughput (i.e., delivered end-user data rate after removing overhead) as a function of ModClass and stream aggregation. Note that ModClass 8 is not used on the uplink.

Table 21: User Data Throughput at various ModClasses

ModClass	Single Stream Downlink Throughput (kbps)	Aggregated 3 Stream Downlink Throughput (kbps)	Single Stream Uplink Throughput (kbps)	Aggregated 3 Stream Uplink Throughput (kbps)
0	35.2	105.6	6.4	19.2
1	49.6	148.8	12.8	38.4
2	81.6	244.8	25.6	76.8
3	126.4	379.2	43.2	129.6
4	161.6	484.8	57.6	172.8
5	198.4	595.2	72.0	216.0
6	262.4	787.2	97.6	292.8
7	307.2	921.6	115.2	345.6
8	353.6	1060.8		

Finally, RACH and CCH logical channels are transmitted at ModClass 0 for reliability. BCH and PCH are transmitted using $\pi/2$ -BPSK, but with more robust coding than that used in ModClass 0. The timing correction and broadcast bursts are encoded using Walsh-Hadamard codes.

4.2 Standard Modulation and Coding

HC-SDMA physical layer supports different data rates by selecting among various coding and modulation schemes. Figure 15 illustrates the coded modulation system that achieves rates from approximately 1/2 to 4 bits/symbol.

Table 22 lists the modulation classes and signal sets, together with the associated parameters for puncturing, shaping, and block coding. Table 23 indicates the allowed modulation class (modClass) values for different burst types.

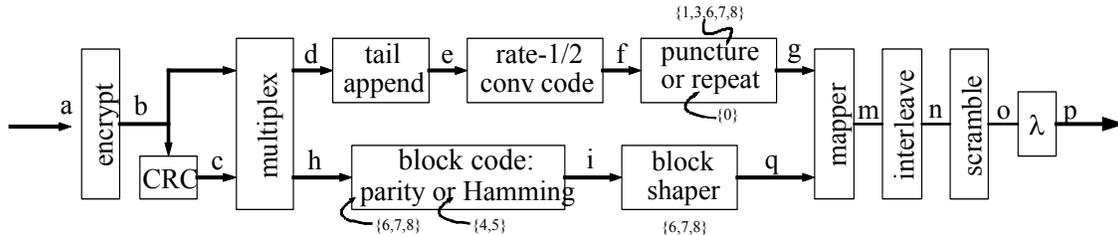


Figure 15: Block Diagram of Error Control Coding Scheme

The notation $\{i, j\}$ indicates that a block is active only for modulation classes i and j .

Coding is provided by a rate 1/2 convolutional code with 256 states combined in some cases with a block code. In some modClasses, periodic puncturing is applied to increase the rate of the convolutional code to 2/3 or 3/4. Repetition is used in modClass 0 to construct RA, AA and CM bursts (and is not used in modClass 0 to construct CR and TCH bursts). ModClasses 6, 7, and 8 employ four-dimensional block shaping to generate 12, 16, and 24QAM signal sets, respectively.

The operation of each block in Figure 15 is detailed in the subsequent sections. In the descriptions that follow, $\{a_1, a_2, \dots, a_{N_a}\}$ denotes the sequence of bits N_a input to the “encrypt” block, $\{b_1, b_2, \dots, b_{N_b}\}$ denotes the N_b bits input to the “CRC” block. The block sizes $\{N_a, N_b, \dots, N_q\}$ are listed in Table 24, Table 25, Table 26, Table 27, Table 28, and Table 29 as a function of message type, modulation class, and uplink/downlink direction.

Table 22: Modulation and Coding Rates

ModClass	Bits/Sym	Signal Set	Puncture	Shaper	Block Code
0	0.5	$\pi/2$ -BPSK	repeat	-	-
1	0.67	$\pi/2$ -BPSK	1 of 4	-	-
2	1.0	QPSK	-	-	-
3	1.5	QPSK	2 of 6	-	-
4	2.0	8-PSK	-	-	(64,57)
5	2.5	8-PSK	-	-	(64,57)
6	3	12-QAM	2 of 6	3/4	(48,47)
7	3.5	16-QAM	2 of 6	4/4	(64,63)
8	4	24-QAM	2 of 6	5/4	(80,79)
9 - 15	RESERVED				

Table 23: ModClass versus Burst Type

Logical Channel	Burst Type	ModClass
TCHuplink	standard uplink	0 - 7
TCHdownlink	Standard downlink	0 - 8
RACHuplink	standard uplink	0
RACHdownlink	Standard downlink	0
CCHuplink	configuration request	0
CCHdownlink	Standard downlink	0
BCH and PCH	broadcast and page	See clause 4.3.

Table 24: Block Lengths in Downlink Traffic Burst

ModClass	$N_{a,b}$	N_c	N_d	N_e	N_f	N_g	N_h	N_i	N_q	N_m	$N_{m,n,o,p}$
0	206	16	222	230	460	460	0	0	0	460	460
1	282	16	298	306	613	460	0	0	0	460	460
2	436	16	452	460	920	920	0	0	0	920	460
3	666	16	682	690	1380	920	0	0	0	920	460
4	840	16	452	460	920	920	404	460	460	460	460
5	1021	16	222	230	460	460	815	920	920	460	460
6	1341	16	682	690	1380	920	675	690	920	920	460
7	1571	16	682	690	1380	920	905	920	920	920	460
8	1801	16	682	690	1380	920	1135	1150	920	920	460

Table 25: Block Lengths in Uplink Traffic Burst

ModClass	$N_{a,b}$	N_c	N_d	N_e	N_f	N_g	N_h	N_i	N_q	N_m	$N_{m,n,o,p}$
0	67	16	83	91	182	182	0	0	0	182	182
1	97	16	113	121	242	182	0	0	0	182	182
2	158	16	174	182	364	364	0	0	0	364	182
3	249	16	265	273	546	364	0	0	0	364	182
4	319	16	174	182	364	364	161	182	182	182	182
5	389	16	83	91	182	182	322	364	364	182	182
6	516	16	265	273	546	364	267	273	364	364	182
7	607	16	265	273	546	364	358	364	364	364	182

Table 26: Block Lengths in Configuration Request Burst

ModClass	$N_{a,b}$	N_c	N_d	N_e	N_f	N_g	N_h	N_i	N_q	N_m	$N_{m,n,o,p}$
0	17	16	33	41	82	82	0	0	0	82	82

Table 27: Block Lengths in Configuration Message

ModClass	$N_{a,b}$	N_c	N_d	N_e	N_f	N_g	N_h	N_i	N_q	N_m	$N_{m,n,o,p}$
0	105	16	121	129	258	460	0	0	0	460	460

Table 28: Block Lengths in Request Access Message

ModClass	$N_{a,b}$	N_c	N_d	N_e	N_f	N_g	N_h	N_i	N_q	N_m	$N_{m,n,o,p}$
0	23	16	39	47	94	182	0	0	0	182	182

Table 29: Block Lengths in Access Assignment Message

ModClass	$N_{a,b}$	N_c	N_d	N_e	N_f	N_g	N_h	N_i	N_q	N_m	$N_{m,n,o,p}$
0	105	16	121	129	258	460	0	0	0	460	460

4.2.1 Encryption

TCH messages are encrypted by exclusiveoring the message with a stream cipher. The stream cipher is initialized using a part of the shared secret exchanged during registration, and an initialization vector (IV). The IV is derived from the AFN, slot number (SN), and relative carrier number (RCN). Details of the encryption process are presented in 11.4.

4.2.2 Cyclic Redundancy Check

Error detection is provided by appending a 16-bit cyclic redundancy check based on the polynomial $g_{CRC}(x) = x^{16} + x^{12} + x^5 + 1$. The output $\{c_1, c_2, \dots, c_{16}\}$ of the CRC block is selected so that the polynomial:

$$b_1x^{N_b+15} + b_2x^{N_b+14} + \dots + b_{N_b}x^{16} + c_1x^{15} + c_2x^{14} + \dots + c_{16}$$

over GF(2) yields a remainder of 0 when divided by $g_{CRC}(x)$.

4.2.3 Multiplexing

The multiplex block (refer to Figure 15 - Block Diagram of Error Control Coding Scheme) divides the input sequence $\{b_1, b_2, \dots, b_{N_b}\}$ into subsequences $\{d_1, d_2, \dots, d_{N_d}\}$ and $\{h_1, h_2, \dots, h_{N_h}\}$, and appends the CRC to the former. Specifically noting that $N_b = N_h + N_d - 16$:

$$d_k = \begin{cases} b_k & \text{if } 1 \leq k \leq N_d - 16, \\ c_{k-N_d+16} & \text{if } N_d - 15 \leq k \leq N_d \end{cases}$$

Furthermore, $h_m = b_{m+N_d-16}$ for $m = 1, 2, \dots, N_k$.

4.2.4 Tail Append

The tail append block (refer to Figure 15) appends 8 zero bits:

$$e_k = \begin{cases} d_k & \text{for } k = 1, \dots, N_d \\ 0 & \text{for } k = N_d + 1, \dots, N_d + 8. \end{cases}$$

4.2.5 Convolutional Encoding

The 256-state convolutional encoder is shown in Figure 16. The constraint length is 9 and the coding rate is 1/2. The initial value of the shift register is the vector $[0,0,0,0,0,0,0,0]$. The input is taken serially in the order $\{e_1, e_2, \dots\}$. The output sequence $\{f_1, f_2, \dots, f_N\}$ is formed by concatenating the two shift register outputs in the order $\{\text{Output 0}, \text{Output 1}, \text{Output 0}, \text{Output 1}, \dots\}$. An exception to this rule occurs when N_f is odd; for such cases the final output bit f_{N_f} is set equal to the second-to-last output bit f_{N_f-1} .

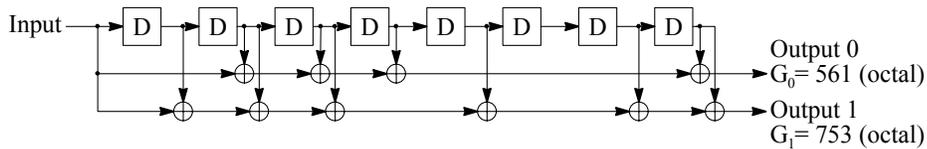


Figure 16: Convolutional Encoder

4.2.6 Puncturing and Repeating

The coded outputs $\{f_1, f_2, \dots\}$ are punctured or repeated to create the sequence $\{g_1, g_2, \dots\}$. Repetition is used only for modulation class 0.

- ◆ For modulation class 0, the coded outputs are repeated sufficiently many times to fill all information symbols in the selected burst type. $P = N_g - N_f$ is the number of bits to be repeated. The pseudocode below defines the repetition rule.

```

A := 0
l := 1
for k = 1, . . . , N_f
    g_l := f_k (copy bit)
    l := l + 1
    A := A + P
    while A ≥ N_g then
        g_l := f_k (repeat bit)
        l := l + 1
        A := A - N_f
    end if
end for
    
```

- ◆ For modulation class 1, the coded outputs are punctured by a periodic puncturing pattern that deletes one bit from every block of four:

The code bit locations function p is defined in Table 30 - Code Bit Locations Function. For $l = 0, 1, \dots, R - 2$ the output is determined by first setting some of the output bits according to the following pseudocode:

```

N := 0
for k := 1, . . . , B
    if (k + N) == P (N + 1) then
        N := N + 1
    end if
     $i_{k+N+C \cdot l} = h_{k+B \cdot l}$ 
end for
    
```

and then selecting the remaining output bits $i_{P(1)+C \cdot b}, \dots, i_{P(7)+C \cdot l}$ so that

$$H_{HAM} \cdot [i_{1+C \cdot b}, i_{2+C \cdot b}, \dots, i_{64+C \cdot l}]^T = 0$$

over GF(2).

For $l = R - 1$, the output is determined by first setting some of the output bits according to the following pseudocode:

```

N := 0
for k := 1, . . . ,  $N_h - B \cdot l$ 
    if (k + N) == P (N + 1) then
         $Q (N + 1) := k + N$ 
        N := N + 1
    end if
     $i_{k+N+C \cdot l} = h_{k+B \cdot l}$ 
end for
    
```

$$H'_{HAM} = [h_{HAM,1} \cdot \begin{matrix} \bullet \\ \bullet \\ \bullet \end{matrix} \cdot h_{HAM,2} \cdot \dots \cdot h_{HAM,N+N_h-B \cdot l}]$$

for $r := (N + 1), \dots, 7$

$$Q(r) := r + N_h - B \cdot l$$

$$H'_{HAM} = [H_{HAM} \cdot \begin{matrix} \bullet \\ \bullet \\ \bullet \end{matrix} \cdot h_{HAM, P(r)}]$$

end for

and then selecting the remaining output bits $i_{Q(1)+C \cdot b}, \dots, i_{Q(7)+C \cdot l}$ so that

$$H'_{HAM} \cdot [i_{1+C \cdot b}, i_{2+C \cdot b}, \dots, i_{Q(7)+C \cdot l}]^T = 0$$

over GF(2).

4.2.7.2 Parity Check Code

A single parity check code is employed for modulation classes 6, 7, and 8 with respective input block lengths $B = 47$, $B = 63$, and $B = 79$.

The output sequence i_1, i_2, \dots, i_{N_i} is computed by inserting one parity bit every B input bits. Specifically, for $l = 0, 1, \dots, R - 2$ where $R = \lceil N_h / B \rceil$, the output of the block code is:

$$i_{k+(B+1)l} = \begin{cases} h_{k+Bl} & \text{for } k = 1, \dots, B \\ h_{1+Bl} \oplus \dots \oplus h_{B+Bl} & \text{for } k = B+1 \end{cases}$$

while for $l = R - 1$ the output of the block code is:

$$i_{k+(B+1)l} = \begin{cases} h_{k+Bl} & \text{for } k = 1, \dots, N_h - Bl, \\ h_{1+Bl} \oplus \dots \oplus h_{N_h} & \text{for } k = N_h + 1 - Bl. \end{cases}$$

4.2.8 Block Shaper

Modulation classes 6, 7, and 8 employ block shaping. For other modulation classes, no block shaping is performed and the output of the block shaper is equal to its input: $q_k = i_k$ for $k = 1, 2, \dots, N_i$.

The block shaper processes a binary-valued input sequence into a ternary-valued output sequence. The input sequence is divided into blocks of size $B = 3$, $B = 4$, and $B = 5$ bits for modulation classes 6, 7, and 8, respectively. Each input block is converted into an output block of 4 trits.

The input-output mappings for the block shaper for different rates are given in Table 31 through Table 33. The block of input bits is denoted by x_1, x_2, \dots, x_k where k is the input block size. The variable l ranges from 0 to $\lceil N_i / B \rceil - 1$. In Table 33, the notation \bar{x}_4 denotes the binary complement of x_4 .

Table 31: Mapping for Rate-3/4 Block Shaper

$i_{1+3l} \ i_{2+3l} \ i_{3+3l}$	$q_{1+4l} \ q_{2+4l} \ q_{3+4l} \ q_{4+4l}$
000	0000
001	0001
010	0010
011	0101
100	0100
101	0110
110	1001
111	1000

Table 32: Mapping for Rate-4/4 Block Shaper

$i_{1+4l} \ i_{2+4l} \ i_{3+4l} \ i_{4+4l}$	$q_{1+4l} \ q_{2+4l} \ q_{3+4l} \ q_{4+4l}$
x_1, x_2, x_3, x_4	x_1, x_2, x_3, x_4

Table 33: Mapping for Rate-5/4 Block Shaper

i_{1+5l}	i_{2+5l}	i_{3+5l}	i_{4+5l}	i_{5+5l}	q_{1+4l}	q_{2+4l}	q_{3+4l}	q_{4+4l}
0	x_2	x_3	x_4	x_5	x_2	x_3	x_4	x_5
1	0	0	x_4	x_5	2	0	x_4	x_5
1	0	1	x_4	x_5	0	2	\bar{x}_4	x_5
1	1	0	x_4	x_5	\bar{x}_4	x_5	2	0
1	1	1	x_4	x_5	x_4	x_5	0	2

4.2.9 Symbol Mapping

The mapper (refer to Figure 15) combines the binary sequence $\{g_1, g_2, \dots\}$ with the ternary sequence $\{q_1, q_2, \dots\}$ to produce a sequence of complex numbers. The mapper first performs a table look-up on the input sequence to produce a sequence $\{\tilde{m}_1, \tilde{m}_2, \dots\}$. The look-up tables for the different modulation classes are given in Table 34, Table 35, Table 36, and Table 37 where $k = 1, 2, \dots, N_{\tilde{m}}$.

For modulation classes 2, 3, 6, 7, and 8, pairs of outputs from the lookup table are multiplexed into a single complex output symbol:

$$m_k = \tilde{m}_{2k-1} + j \times \tilde{m}_{2k}$$

for $k = 1, 2, \dots, N_m$.

The multiplexing step is omitted for modulation classes 0, 1, 4, and 5: $m_k = \tilde{m}_k$ for $k = 1, 2, \dots, N_m$.

Table 34: Symbol Mapper for Modulation Classes 0-3

g_k	\tilde{m}_k
0	-1
1	+1

Table 35: Symbol Mapper for Modulation Class 4

g_{2k-1}	g_{2k}	q_k	\tilde{m}_k
0	0	0	$\exp(j0\pi/4)$
1	0	0	$\exp(j1\pi/4)$
1	1	0	$\exp(j2\pi/4)$
0	1	0	$\exp(j3\pi/4)$
0	0	1	$\exp(j4\pi/4)$
1	0	1	$\exp(j5\pi/4)$
1	1	1	$\exp(j6\pi/4)$
0	1	1	$\exp(j7\pi/4)$

Table 36: Mapper for Modulation Class 5

g_k	q_{2k-1}	q_{2k}	\tilde{m}_k
0	0	0	$\exp(j0\pi/4)$
1	0	0	$\exp(j1\pi/4)$
0	1	0	$\exp(j2\pi/4)$
1	0	1	$\exp(j3\pi/4)$
0	1	1	$\exp(j4\pi/4)$
1	1	1	$\exp(j5\pi/4)$
0	0	1	$\exp(j6\pi/4)$
1	1	0	$\exp(j7\pi/4)$

Table 37: Symbol Mapper for Modulation Classes 6 - 8

g_k	q_k	\tilde{m}_k
0	2	-5
1	1	-3
0	0	-1
1	0	+1
0	1	+3
1	2	+5

4.2.10 Interleaving

The interleave block (refer to Figure 15) employs depth-8 rectangular interleaving with row permutation.

Table 38: Interleaver Row Permutation

$P(1)$	$P(2)$	$P(3)$	$P(4)$	$P(5)$	$P(6)$	$P(7)$	$P(8)$
1	5	3	7	2	6	4	8

The permutation function P is defined in Table 38. The operation of the interleaver is defined by the following pseudocode:

```

row := 1
l := P (row)
for k := 1, . . . , Nm
    nk := ml
    l := l + 8
    if l > Nm then
        row := row + 1
        l := P (row)
    end if
end for
    
```

4.2.11 Scrambling

The scramble block (refer to Figure 15 - Block Diagram of Error Control Coding Scheme) applies complex rotations of 0° , 90° , 180° , and 270° to the input sequence $\{n_1, n_2, \dots\}$. The rotations are determined by a binary scrambling sequence $\{s_1, s_2, \dots\}$. For modulation classes 0 and 1, the output of the scramble block is:

$$o_k = \begin{cases} n_k & \text{if } s_{2k} = 0 \\ -n_k & \text{if } s_{2k} = 1 \end{cases}$$

for $k = 1, 2, \dots, N_o$.

For modulation classes 2 - 8, the output of the scramble block is:

$$o_k = \begin{cases} n_k, & \text{if } s_{2k-1} = 0 \text{ and } s_{2k} = 0 \\ j \cdot n_k, & \text{if } s_{2k-1} = 1 \text{ and } s_{2k} = 0 \\ -n_k, & \text{if } s_{2k-1} = 0 \text{ and } s_{2k} = 1 \\ -j \cdot n_k, & \text{if } s_{2k-1} = 1 \text{ and } s_{2k} = 1 \end{cases}$$

for $k = 1, 2, \dots, N_o$.

The scrambling sequence is the output of a linear feedback shift register (LFSR) defined by the generator polynomial $g_{SCR}(x) = x^{31} + x^3 + 1$, as shown in Figure 17. If the vector $[u_{0,k}, u_{1,k}, \dots, u_{30,k}]$ denotes the state of the shift register at time k , then:

$$\begin{aligned} s_k &= u_{0,k} \oplus u_{28,k} \\ u_{30,k+1} &= s_k \\ u_{29,k+1} &= u_{30,k} \\ u_{28,k+1} &= u_{29,k} \\ &\vdots \\ u_{0,k+1} &= u_{1,k} \end{aligned}$$

The initial value $[u_{0,1}, u_{1,1}, \dots, u_{30,1}]$ of the shift register is set to a 31-bit value shown in Table 39. Here $c1(\text{LSB})\dots c6(\text{MSB})$ is the binary expansion of BSCC and $a1(\text{LSB})\dots a10(\text{MSB})$ are the 10 least significant bits of AFN. $r1(\text{LSB})\dots r15(\text{MSB})$ is the binary expansion of the combination of the random ID field

($r_1...r_{10}$) and UT class field ($r_{11}...r_{15}$) in RA-rreg. For all other RA bursts, $r_1...r_{15}$ are obtained from the RID or PID.

Table 39: Scrambling Initialization for Various Burst Types

	CR	CM	RA	AA	TCH (Up)	TCH (Dn)
$u_{0,1}$	c1	c1	c1	c1	c1	c1
$u_{1,1}$	c2	1	c2	c2	c2	c2
$u_{2,1}$	1	c2	c3	c3	c3	c3
$u_{3,1}$	c3	1	c4	c4	c4	c4
$u_{4,1}$	c4	0	c5	c5	c5	c5
$u_{5,1}$	c5	1	c6	c6	c6	c6
$u_{6,1}$	c6	1	a1	a1	a1	a1
$u_{7,1}$	1	0	a2	a2	a2	a2
$u_{8,1}$	1	c3	a3	a3	a3	a3
$u_{9,1}$	0	0	a4	a4	a4	a4
$u_{10,1}$	1	0	a5	a5	a5	a5
$u_{11,1}$	1	1	a6	a6	a6	a6
$u_{12,1}$	0	1	a7	a7	a7	a7
$u_{13,1}$	0	1	a8	a8	a8	a8
$u_{14,1}$	0	c4	a9	a9	a9	a9
$u_{15,1}$	0	0	a10	a10	a10	a10
$u_{16,1}$	0	0	1	r1	r1	r1
$u_{17,1}$	0	0	0	r2	r2	r2
$u_{18,1}$	0	1	1	r3	r3	r3
$u_{19,1}$	1	1	0	r4	r4	r4
$u_{20,1}$	0	c5	1	r5	r5	r5
$u_{21,1}$	1	0	0	r6	r6	r6
$u_{22,1}$	0	1	1	r7	r7	r7
$u_{23,1}$	0	0	0	r8	r8	r8
$u_{24,1}$	0	1	1	r9	r9	r9
$u_{25,1}$	1	1	0	r10	r10	r10
$u_{26,1}$	1	c6	1	r11	r11	r11
$u_{27,1}$	1	1	0	r12	r12	r12
$u_{28,1}$	0	1	1	r13	r13	r13
$u_{29,1}$	0	0	0	r14	r14	r14
$u_{30,1}$	0	1	1	r15	r15	r15

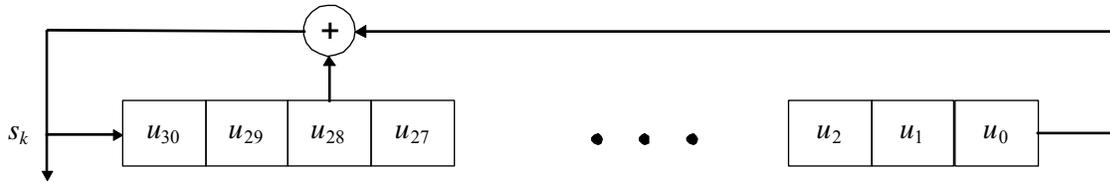


Figure 17: Scrambling Sequence Generator

4.2.12 $\pi/2$ Rotation and Scaling

The “ λ ” block in Figure 15 supplies $\pi/2$ rotation for the BPSK modulation classes and rescales all of the symbols to provide an average symbol energy equal to 1. For modulation classes 0 and 1 the output $\{p_1, p_2, \dots\}$ is formed by rotating every other symbol by 90° in the complex plane:

$$p_k = \begin{cases} o_k, & \text{if } k \text{ is even,} \\ j \cdot o_k, & \text{if } k \text{ is odd} \end{cases}$$

for $k = 1, 2, \dots, N_o$.

For modulation classes 2 - 8, the “ λ ” block provides scaling: $p_k = \lambda o_k$ for $k = 1, 2, \dots, N_o$, where the value of λ is a function of modulation class (refer to Table 40).

Table 40: Modulation Scaling

ModClass	λ
2	$1/\sqrt{2}$
3	$1/\sqrt{2}$
4	1
5	1
6	$1/\sqrt{7}$
7	$1/\sqrt{10}$
8	$1/\sqrt{14}$

4.3 Broadcast Channel Modulation and Coding

The broadcast logical channel has four burst types: frequency correction, timing correction, broadcast message, and paging. The frequency correction burst is defined in clause 3. The information-bearing portions of the other three burst types are defined below.

4.3.1 General Definitions

The timing and broadcast bursts use Hadamard matrices. Define the Hadamard family of matrices H_0, \dots, H_6 recursively as follows:

$$H_{k+1} = \begin{bmatrix} H_k & H_k \\ H_k & \bar{H}_k \end{bmatrix}$$

for $k = 0, 1, \dots, 5$, where $H_0 = [1]$ and \bar{H}_k is the binary complement of H_k .

4.3.2 Timing Correction Burst

The timing correction symbols $\{t_1, t_2, \dots, t_{512}\}$ in a timing synchronization burst are an encoded version of the base station color code. $b \in \{0, \dots, 63\}$ is the BSCC, and $\{h_1, h_2, \dots, h_{64}\}$ is the $(b+1)$ -th row of the Hadamard matrix H_6 defined in 4.3.1, General Definitions. The timing correction symbols consist of 8 scrambled repetitions of $\{h_1, h_2, \dots, h_{64}\}$ with QPSK modulation. Each repetition $l = 0, \dots, 7$ may have a different complex gain λ_l , which in addition may vary from burst to burst. The values λ_l are not specified by the *HC-SDMA* protocol.

For $l = 0, \dots, 7$ and $k = 1, \dots, 64$,

$$t_{64 \cdot l + k} = \begin{cases} \lambda_l \cdot \exp(j(1 + 2s_{64 \cdot l + k})\pi/4), & \text{if } h_k = 1 \\ -\lambda_l \cdot \exp(j(1 + 2s_{64 \cdot l + k})\pi/4), & \text{if } h_k = 0 \end{cases}$$

where $\{s_1, s_2, \dots, s_{512}\}$ is defined by a linear congruential generator. The initial state of the generator is $g_1 = 124543770$, and the outputs $\{g_2, g_3, \dots, g_{512}\}$ are defined recursively by:

$$g_{k+1} = (69069 \times g_k) \bmod 2^{32}$$

for $k = 1, \dots, 511$. The scrambling sequence is taken from the two most significant bits of g_k :

$$s_k = g_{31,k} + 2 \times g_{32,k}$$

4.3.3 Broadcast Burst

The broadcast burst carries the six-bit broadcast message $\{p_1, p_2, \dots, p_6\}$. $p = p_1 + 2 \cdot p_2 + 4 \cdot p_3 + 8 \cdot p_4 + 16 \cdot p_5 + 32 \cdot p_6$ is the integer representation of the broadcast message and $\{h_1, h_2, \dots, h_{64}\}$ is the $(p+1)$ -th row of the Hadamard matrix H_6 defined in clause 4.3.1 General Definitions.

The broadcast information symbols $\{b_1, b_2, \dots, b_{512}\}$ consist of 8 scrambled repetitions of $\{h_1, h_2, \dots, h_{64}\}$ with QPSK modulation. Each repetition $l = 0, \dots, 7$ may have a different complex gain λ_l , subject to

the constraint that $\lambda_l = \lambda_{l+4}$ for $l = 0, \dots, 3$. The coefficients λ_l may vary from burst to burst. The values λ_l are not specified by the *HC-SDMA* protocol.

For $l = 0, \dots, 7$ and $k = 1, \dots, 64$,

$$b_{64 \cdot l + k} = \begin{cases} \lambda_l \cdot \exp(j(1 + 2s_{64 \cdot l + k})\pi/4), & \text{if } h_k = 1 \\ -\lambda_l \cdot \exp(j(1 + 2s_{64 \cdot l + k})\pi/4), & \text{if } h_k = 0 \end{cases}$$

where the scrambling sequence $\{s_1, s_2, \dots, s_{512}\}$ is defined by a linear congruential generator (refer to clause 4.3.2, Timing Correction Burst).

The generator is initialized using the base station color code $\text{BSCC} \in \{0, \dots, 63\}$. The initial state of the generator is $g_1 = (\text{BSCC} + 1)^2$, and the outputs $\{g_2, g_3, \dots, g_{512}\}$ are defined recursively by:

$$g_{k+1} = (69069 \times g_k) \bmod 2^{32}$$

for $k = 1, \dots, 511$. The scrambling sequence is taken from the two most significant bits of g_k :

$$s_k = g_{31, k} + 2 \times g_{32, k}$$

4.3.4 Paging Burst

The page information symbols h_1, \dots, h_{256} are constructed by QPSK modulating the output $\{s_1, s_2, \dots, s_{256}\}$ of a quaternary sequence generator. The QPSK modulation is defined by

$$h_k = \lambda \cdot \exp(j(1 + 2s_k)\pi/4)$$

$k = 1, 2, \dots, 256$, where λ is a complex gain not specified by the *HC-SDMA* protocol.

The quaternary sequence generator consists of one quaternary and two binary shift registers (refer to Figure 18 - Quaternary Sequence Generator). All of the shift registers have 8 registers. The quaternary shift register is defined by the polynomial $g_{SCR,0}(x) = x^8 + x^5 + 3 \cdot x^3 + x^2 + 2 \cdot x + 1$. The first binary shift register is defined by the polynomial $g_{SCR,1}(x) = x^8 + x^7 + x^5 + x + 1$. The second binary shift register is defined by the polynomial $g_{SCR,2}(x) = x^8 + x^7 + x^5 + x^4 + 1$.

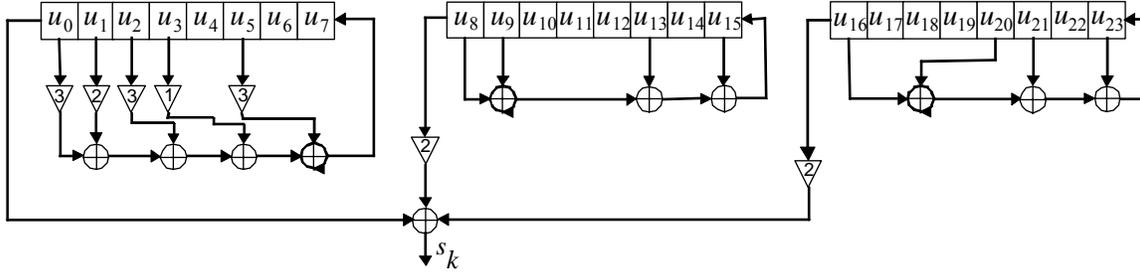


Figure 18: Quaternary Sequence Generator

The state of the quaternary shift register at time k , $[u_{0,k}, u_{1,k}, \dots, u_{7,k}]$, is given by:

$$\begin{aligned} u_{7,k+1} &= (3 \times u_{5,k} + u_{3,k} + 3 \times u_{2,k} + 2 \times u_{1,k} + 3 \times u_{0,k}) \bmod 4 \\ u_{6,k+1} &= u_{7,k} \\ u_{5,k+1} &= u_{6,k} \\ &\vdots \\ u_{0,k+1} &= u_{1,k} \end{aligned}$$

The state of the first binary shift register at time k , $[u_{8,k}, u_{9,k}, \dots, u_{15,k}]$, is given by:

$$\begin{aligned} u_{15,k+1} &= u_{15,k} \oplus u_{13,k} \oplus u_{9,k} \oplus u_{8,k} \\ u_{14,k+1} &= u_{15,k} \\ u_{13,k+1} &= u_{14,k} \\ &\vdots \\ u_{8,k+1} &= u_{9,k} \end{aligned}$$

The state of the second binary shift register at time k , $[u_{16,k}, u_{17,k}, \dots, u_{23,k}]$, is given by:

$$\begin{aligned} u_{23,k+1} &= u_{23,k} \oplus u_{21,k} \oplus u_{20,k} \oplus u_{16,k} \\ u_{22,k+1} &= u_{23,k} \\ u_{21,k+1} &= u_{22,k} \\ &\vdots \\ u_{16,k+1} &= u_{17,k} \end{aligned}$$

The output of the quaternary sequence generator at time k is:

$$s_k = (u_{0,k} + 2 \times u_{8,k} + 2 \times u_{16,k}) \bmod 4$$

The sequence generator is initialized using the binary expansion $[c_1, c_2, \dots, c_6]$ of the base station color code (BSCC), the binary expansion $[p_1, p_2, \dots, p_{15}]$ of the paging identity (PID) of the target user terminal, and the 10 least significant bits $[a_1, a_2, \dots, a_{10}]$ of the absolute frame number (AFN). The initial state for the sequence generator is

$$\begin{array}{lll} u_{0,1} = 2p_1 + 1 & u_{8,1} = p_9 \oplus a_4 & u_{16,1} = c_2 \\ u_{1,1} = 2p_2 & u_{9,1} = p_{10} \oplus a_5 & u_{17,1} = c_3 \\ u_{2,1} = 2p_3 & u_{10,1} = p_{11} \oplus a_6 & u_{18,1} = c_4 \\ u_{3,1} = 2p_4 & u_{11,1} = p_{12} \oplus a_7 & u_{19,1} = c_5 \\ u_{4,1} = 2p_5 & u_{12,1} = p_{13} \oplus a_8 & u_{20,1} = c_6 \\ u_{5,1} = 2p_6 & u_{13,1} = p_{14} \oplus a_9 & u_{21,1} = a_1 \\ u_{6,1} = 2p_7 & u_{14,1} = p_{15} \oplus a_{10} & u_{22,1} = a_2 \\ u_{7,1} = 2p_8 & u_{15,1} = c_1 & u_{23,1} = a_3 \end{array}$$

5 USER TERMINAL RADIO TRANSMISSION AND RECEPTION

5.1 HC-SDMA Introduction

This chapter describes the radio frequency (RF) performance characteristics of an *HC-SDMA* user terminal, and serves as the document from which manufacturers must generate compliance specifications.

It is the UT manufacturer's responsibility to ensure compliance (defined below) with this specification and applicable local regulations. The following sections define the terminology used for both base station and user terminal specification chapters.

5.1.1 Terminology

5.1.1.1 Carrier aggregation: All definitions and specifications for UT transmission describe transmission and reception on a single modulated carrier. The UT may initiate data streams on multiple carriers to increase its data rate; this is termed *carrier aggregation*. In this case, reception and transmission specifications apply independently to each of the carriers. For example, a UT transmitting on two carriers simultaneously may transmit twice the composite power as a UT transmitting on a single carrier.

5.1.1.2 Compliance: Means that for any given performance metric in this document, the equipment shall conform to whichever specification is more restrictive, that of this document or the applicable

local regulations. Emissions behaviors specified in this document may or may not satisfy the regulatory requirements of any particular country or region.

5.1.1.3 Transmitter or Receiver: Refers in this text to a single transmitter or receiver. In the case of a base station that supports multiple antennas, these terms refer individually to the transmitters and receivers for each of the antennas.

5.1.1.4 Useful symbols: Are all the symbols in a time slot with the exception of the ramp-up, ramp-down, and guard symbols.

5.1.2 Measurements

5.1.2.1 Cabled environment: Means that the radio paths among explicitly described test equipment, a base station and one or more user terminals are fixed in a way to exclude other RF signals. This may be accomplished by connecting the equipment antenna ports directly to a distribution network of cables and couplers, it may be accomplished by placing the equipment in a sealed RF enclosure such as a screen room, or it may be accomplished by a combination of the two techniques. In any case, it is understood that no interfering signals or pathloss variabilities are present other than those explicitly called out.

5.1.2.2 Operating Temperature Range: The operating temperature range as specified by the base station or user terminal manufacturer.

5.1.2.3 Peak Transmit Power: Is the instantaneous output power of the radio transmitter that is not exceeded 99.9% of the time. Output power is measured at the antenna connector, and the measurement shall be within the channel bandwidth centered around the specified carrier frequency. Unless otherwise specified, the measurement procedure requires a 5-pole synchronously tuned measurement filter with bandwidth 625 kHz.

5.1.2.4 Receive Power: Is the average input power to the radio receiver. Input power is measured at the antenna connector, and the measurement shall be within the channel bandwidth centered around the specified carrier frequency. Averaging takes place between the start of the first useful symbol and the end of the last useful symbol of a downlink time slot (ramp-up, ramp-down, and guard symbols excluded). Unless otherwise specified, the measurement procedure requires a root-raised cosine filter with rolloff factor $\alpha = 0.25$ and symbol period $T_s = 2 \mu\text{sec}$.

5.1.2.5 SRRC Filtered signal on a carrier: Is measured over a 625 kHz bandwidth centered at the named carrier frequency, when a square root raised cosine filter with roll-off factor $\alpha = 0.25$ and symbol period $T_s = 2 \mu\text{s}$ is applied. When a power measurement is required and a root-raised cosine filter is unavailable or impractical for measurement, it is permissible to measure SRRC-filtered power by acquiring the power over 500 kHz bandwidth centered at the named carrier frequency in 10 kHz sub-bands, and summing the measured power.

5.1.2.6 Temperature Dependent Adjustments to Noise Related Specifications: Noise-Related Specifications are given with an assumed reference temperature of 300 K (27°C). Conformance to the specification may be tested at any temperature within the Operating Temperature Range by adjusting the noise levels by $10 \times \log_{10}(T/300)$, where T is the measurement temperature in Kelvins. For example, when testing at 50°C, add 0.32 dB to noise levels.

5.1.2.7 Transmit Power: Is measured at the antenna connector, and the measurement shall be within the 625 kHz carrier bandwidth centered around the specified carrier frequency. Averaging takes place

between the start of the first useful symbol (defined in clause 1.3, Conventions) and the end of the last useful symbol of an uplink time slot (ramp-up, ramp-down, and guard symbols excluded).

5.1.3 Symbols

Symbols which are common to all protocol chapters are defined in clause 1.3, Conventions.

N_f	The number of frequency carriers supported by a given <i>HC-SDMA</i> system is designated N_f and depends on the allocated spectrum.
f_{org}	The FREQUENCY ORIGIN is the lowest-frequency HC-SDMA carrier's center frequency for a given installation. See clause 2, Spectral Layout Terminology and Requirements for a discussion of preferred values for f_{org} .
f_n	The carrier frequencies are given by $f_n = f_{org} + n \times 625 \text{ kHz}$ where n is the carrier number and f_{org} is the frequency origin. A single carrier's RF power is spectrally distributed over 625 kHz centered on the carrier frequency. See the discussion of preferred carrier frequencies in clause 2, Spectral Layout Terminology and Requirements.
N_s	The number of spatial channels supported by a given <i>HC-SDMA</i> installation is called N_s .
M	NUMBER OF ANTENNAS is the number of independent transmitters and receivers supported by a given <i>HC-SDMA</i> installation. The term "number of antennas" is used even though it is recognized that each transmitter or receiver might be coupled to more than one antenna structure through multiple feed points. Diversity in polarization is not distinguished from diversity in beam patterns for purposes of this document.
P_R	Average SRRC filtered input power for a given carrier to a radio receiver. INPUT POWER is measured at the antenna, and is not reduced to account for cable losses. Averaging takes place between the start of the first useful symbol and the end of the last useful symbol of an uplink or downlink time slot. Ramp-up, ramp-down, and guard symbols are excluded.
T_s	The symbol period for <i>HC-SDMA</i> modulation formats, always $2 \mu s$.

MA	<p>The modulation accuracy (MA) for a transmitted burst is defined by:</p> $MA = \min_{c_0, \omega_0, \tau_0} \sqrt{\frac{\sum_k d(kT_s) - c_0 \cdot e^{j\omega_0 kT_s} \cdot x(kT_s - \tau_0) ^2}{\sum_k d(kT_s) ^2}}$ <p>where:</p> <ul style="list-style-type: none"> t is the time, with the burst's first useful symbol occurring at $t = 0$, $x(t)$ is the complex transmitted voltage waveform at time t, after SRRC filtering, $d(t)$ is the complex desired (ideal) waveform at time t, after SRRC filtering, c_0 is the complex amplitude value which minimizes least-squares error for the training sequence(s) ω_0 is the least-squares optimized real-valued frequency offset, τ_0 is the least-squares optimized positive or negative timing offset, k is an integer index variable that runs over the useful symbols of the burst, including the training sequences and payload but not the ramp-up or ramp-down symbols, <p>Note that all bursts have a training sequence or a preamble consisting of symbols with unity amplitude (see clause 3.3, Burst Formats) in <i>HC-SDMA</i>. The value of c_0 for a measured burst should be estimated from the mean signal phase and amplitude at the symbol points in its training sequence or preamble. Note that the same reference amplitude (the denominator in the above equation) applies to all symbols in a burst, whether they are part of a preamble, training sequence or payload. Scaling between constellation points of different modclasses are described in Table 40 – Modulation Scaling.</p> <p>Most instruments designed to report “Error Vector Magnitude,” or EVM use the maximum constellation point amplitude as the reference amplitude, rather than c_0. Measurements on QAM payloads obtained from such instruments shall be adjusted for this difference to obtain MA.</p>
-----------	---

5.2 Time and Frequency Characteristics

This section describes the UT equipment requirements for timing and frequency accuracy. The UT is responsible to maintain frequency and timing (as per the following sections) during any transmit or receive interval, and will need the same accuracy to demodulate TCH bursts. The base station broadcast burst described in clause 3.3.3, Broadcast Burst is designed to assist the UT in aligning its frequency and time sources with those of the base station, beginning with a much less accurate UT-internal standard. Typical UT implementations use internal oscillators with significantly less accuracy, and are fine-tunable for alignment with the base station time and frequency. The UT-internal clock reference should be sufficiently accurate to support clause 5.2.1.2, Carrier Frequency Error.

5.2.1 RF Synchronization Characteristics

This section defines the performance of the local oscillators in the UT transmitter.

5.2.1.1 Initial Carrier Frequency Error

Definition: The difference between the commanded and actual UT carrier frequency during BCH burst reception. Initial carrier frequency error is specified as an absolute frequency.

Specification: The initial carrier frequency error shall be within ± 10 kHz.

5.2.1.2 Carrier Frequency Error

Definition: The difference between the commanded and actual UT carrier frequency during any active uplink burst, using the received base station BCH frequency as a reference.

Specification: The carrier frequency error of the UT shall be within ± 100 Hz.

5.2.1.3 Transmit Burst Timing Error

Definition: The accuracy with which the UT controls its timing advance.

Specification: The burst timing error shall be less than ± 0.25 symbol. Timing shall be changeable on a frame-by-frame basis.

5.2.2 Radio Tuning Characteristics

This section defines the required radio tuning capabilities of the UT.

5.2.2.1 Carrier Tuning Frequencies

Definition: The carrier tuning frequencies are the carriers on which the UT supports reception and transmission of *HC-SDMA* bursts.

Specification: The UT shall be able to tune to carriers on the carrier raster, $f_{org} = 312.5 \text{ kHz} + N \times 625 \text{ kHz}$ where N is a positive integer. The UT shall be able to tune to carriers on the carrier raster, $f_{org} = N \times 625 \text{ kHz}$ where N is a positive integer. For example, a UT operating in a frequency band beginning at 1920 MHz shall be able to support a carriers of center frequency 1920.3125 MHz and 1920.625 MHz. The UT should be able to tune to carriers on the carrier rasters, $f_{org} = f_m + N \times 625 \text{ kHz}$ and f_m is selected from among 437.5 kHz, 562.5 kHz, 687.5 kHz, or 812.5 kHz.

5.2.2.2 Carrier Tuning Speed

Definition: The carrier tuning speed is the speed with which the UT tunes between any two supported carrier frequencies as indicated by the “lowest freq” and “highest freq” fields, which the UT transmits to the BS as part of the UT Capability Message (see 9.5.1.3.2, UT Capability Message).

Specification: The UT *shall* be capable of tuning to any carrier within its active carrier set on a frame-by-frame basis. That is, the carrier tuning speed *shall* be fast enough to support operation on any given time slot, using different carrier frequencies on subsequent frames.

UT *should* be capable of tuning to any carrier within its active carrier set on a slot-by-slot basis.

5.3 Transmitter Characteristics

This section describes the transmitter performance of the UT radio equipment.

5.3.1 Output RF Power Characteristics

This section defines the output power and the transient response characteristics of the UT transmitter under various conditions.

5.3.1.1 Absolute Maximum Transmit Power

Definition: Absolute Maximum Transmit Power is the highest SRRC-filtered transmit power per carrier, averaged over all useful symbols of an active uplink burst, of UT radio emissions in the HC-SDMA band, whether originating in the transmitter or receiver.

Specification: Absolute maximum transmit power of the UT shall be no greater than 33 dBm.

5.3.1.2 Nominal Output Power

Definition: Nominal output power is the SRRC-filtered transmit power that the UT supports, while meeting all HC-SDMA protocol specifications. The nominal output power depends on the UT's power class.

Specification: Table 41 - Nominal UT transmit power per carrier for various modulation formats defines the nominal output power by class that the UT shall support. The UT transmit power shall not be less than 3 dB below the nominal power stated in Table 41 - Nominal UT transmit power per carrier for various modulation formats. A user terminal may restrict its transmit power to 6 dB less than the tabulated value when operating on carriers 0 (lowest carrier) or $N_f - 1$ (highest carrier) if this is needed to meet out-of-band emission requirements.

Table 41: Nominal UT transmit power per carrier for various modulation formats

Modulation Format	Nominal Output Power		
	Power class 1	Power class 2	Power class 3
16QAM	30 dBm	25 dBm	20 dBm
12QAM	30 dBm	25 dBm	20 dBm
8PSK	31 dBm	26 dBm	21 dBm
QPSK	31 dBm	26 dBm	21 dBm
$\pi/2$ BPSK	32 dBm	27 dBm	22 dBm

5.3.1.3 Transmit Leakage Power

Definition: The transmitter leakage power is the highest transmit power of the UT, measured with at least 1 MHz bandwidth, such that the measurement time interval is *not* within an active uplink time slot or guard times adjacent to an active time slot.

Specification: Transmitter leakage power of the UT shall not exceed -65 dBm.

5.3.2 Ramp-Up and Ramp-Down Transmit Power

For each burst type, the transmitter ramp-up and ramp-down shall occur respectively during the ramp-up and ramp-down periods defined in clause 3.3, Burst Formats.

Definition: Ramp-up and ramp-down transmit power is the highest power transmitted during a ramp-up or ramp-down period. Ramp-up and Ramp-Down transmit power should be measured with instrument bandwidth of at least 1 MHz.

Specification: Ramp-up and ramp-down transmit power shall not exceed 6dB relative to the average power transmitted over the useful symbols of the corresponding burst. See Figure 19 – Ramp-up and Ramp-down Transmit Power.

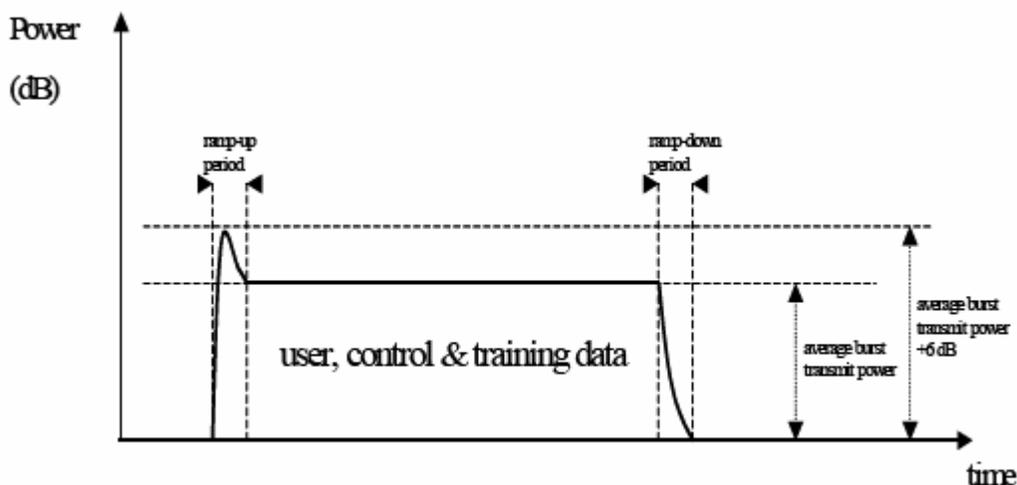


Figure 19: Ramp-up and Ramp-down Transmit Power

5.3.3 Uplink Power Control

The UT shall support power control for its transmitter. The UT's transmit power is initially based on the UT's estimate of received base station power (open loop control). The UT then responds appropriately to base station commands to implement closed loop control. This method is described in the following sections.

5.3.3.1 Open Loop Power Control Absolute Error

Definition: The open loop power control absolute error is defined as the ratio between the actual and the commanded transmit power, expressed in dB.

Specification: The open loop power control absolute error shall not exceed ± 6 dB.

5.3.4 Minimum Transmit Power

Definition: The minimum transmit power is SRRC-filtered transmit power of the UT when the radio transmitter operates at the minimum power setting.

Specification: Minimum transmit power of the UT shall not exceed -20 dBm.

5.3.4.1 Power Control Step Size

Definition: The power control step size is the absolute change in the transmit power of the UT in response to a power control step command. The transmit power is increased by this amount in response to a power increment step command, and decreased by this amount in response to a power decrement step command.

Specification: The power control step size of the UT shall be 1 dB (tolerance: ± 0.5 dB).

5.3.4.2 Power Control Maximum Rate

Definition: The power control maximum rate is the frequency at which the UT adjusts its transmit power.

Specification: The UT shall independently control transmit power for all active uplink time slots, and shall update the transmitted power for any active time slot on a frame-by-frame basis.

5.3.5 Transmit Modulation

This section defines the performance of the transmit modulator. Transmit modulation is the process of converting a digital symbol sequence into an analog waveform.

5.3.5.1 Transmit Pulse Shaping Filter

Definition: Digital pulses are bandlimited by using the transmit the pulse shaping filter.

Specification: The pulse shaping filter on the UT transmitter shall be a root raised cosine filter with rolloff factor $\alpha = 0.25$ and symbol period $T_s = 2\mu s$.

5.3.5.2 Transmitter Modulation Classes

Specification: The UT transmitter shall support at least *HC-SDMA* uplink modulation class 0 as described in clause 4.2, Standard Modulation and Coding.

5.3.6 Modulation Accuracy

Definition: The modulation accuracy is the ratio of the root mean square error vector magnitude to the reference amplitude, averaged over the useful symbols of an uplink time slot. The error vector is the difference between the theoretically optimal reference waveform and the transmitted waveform measured at the UT antenna connector (or a suitable receiver antenna if the UT antenna has no connector). The transmitted waveform is to be SRRC filtered as described in clause 3.5.1.2, and the initial phase, amplitude, frequency offset, and timing offset are to be identified and removed by a leastsquares search. Modulation accuracy (*MA*) is to be calculated as described in clause 3.5.1.3.

Specification: The modulation accuracy of the transmitter shall be in accordance with the specifications given in Table 42 - Modulation Accuracy For Various Modulation Formats.

Table 42: Modulation Accuracy For Various Modulation Formats

Modulation Format	Modulation Accuracy
16QAM	< 6%
2QAM	< 7%
8PSK	< 9%
QPSK	< 10%
$\pi/2$ BPSK	< 10%

5.3.7 Output RF Spectrum Characteristics

This section defines restrictions on the output RF spectrum of the UT radio transmitter – i.e., spectral masks.

5.3.7.1 Occupied Bandwidth

Definition: The occupied bandwidth is the frequency range that contains 99% of the total power of the transmitted spectrum with 0.5% of the total transmitted power below this range and 0.5% of the total transmitted power above this range.

Specification: The occupied bandwidth of the UT shall not exceed 625 kHz.

5.3.7.2 Adjacent Channel Power Ratio (ACPR)

Definition: Adjacent Channel Power Ratio (ACPR) is expressed as a decibel ratio of undesired SRRC-filtered power transmitted by the UT on adjacent channels relative to the desired transmitted signal. The desired transmit signal power is averaged over the useful symbols of an uplink burst. Both the undesired and desired signals are measured as SRRC-filtered power.

Specification: The ACPR for any carrier frequencies within the carrier allocation shall not exceed than the values in Table 43 – Maximum ACPR_ when the transmit power is greater than +10 dBm. If the ACPR limit in the table, together with the transmit power results in an ACPR limit less than -40 dBm, -40 dBm is applied as the limit instead of the tabulated value.

Table 43: Maximum ACPR

Carrier	Frequency Offset (Δf)	ACPR
First Adjacent Carrier	625 kHz	-35 dBc
econd Adjacent Carrier	1250 kHz	-45 dBc
Other Inband Carrier	> 1250 kHz, < 5000 kHz	-50 dBc

5.3.7.3 Out-of-Band Spurious Emissions

Definition: Out-of-band spurious emission performance is evaluated by measuring the peak transmit power over all the useful symbols of a burst, in which UT transmits at maximum power.

Specification:

- ◆ Out-of-band spurious emission of the UT shall be within local regulatory limits.

- ◆ UT out-of-band emissions at frequency offsets more than 4687.5 kHz from the edge of the nominal carrier bandwidth shall be less than -30 dBm, measured within a 1 MHz bandwidth.

5.4 Receiver Characteristics

This section characterizes the receiver performance of the UT radio equipment.

5.4.1 Input RF Power Characteristics

This section defines the input power requirements of the UT radio receiver.

5.4.1.1 Input Modulation Class

Specification: The UT receiver shall at a minimum support *HC-SDMA* downlink modulation class 0 as described in clause 4.2, Standard Modulation and Coding.

5.4.1.2 Reference Sensitivity Level

Definition: The reference sensitivity level is that minimum SRRC-filtered receive power at the UT antenna port such that the frame error rate (FER) does not exceed a specific value. See clause 5.1.2, Measurements for temperature-dependent adjustments to the measurements.

Specification: Reference sensitivity level of the UT shall be no more than 1dB higher than the nominal limits in Table 44 - Reference Sensitivity at $FER = 10^{-2}$.

Table 44: Reference Sensitivity at $FER = 10^{-2}$

Mod Class	Nominal Reference Sensitivity (dBm)
0	-108.5
1	-106.7
2	-105.2
3	-102.3
4	-100.1
5	-97.9
6	-95.8
7	-94.5
8	-92.6

5.4.1.3 Maximum Non-Distorting Input Power

Definition: The maximum receive power at the UT antenna port such that the frame error rate (FER) does not exceed 10^{-2} .

Specification: The maximum input power of the UT shall be greater than -35 dBm.

5.4.2 Receiver Selectivity

This section defines the ability of the UT radio to receive a desired signal in the presence of interfering signals.

5.4.2.1 Adjacent Channel Selectivity

Definition: Adjacent Channel Selectivity (ACS) measures the receiver’s ability to receive a desired signal on its assigned carrier $n:0 \leq n < N_f$ in the presence of a modulated interfering signal on an adjacent carrier. The ACS is the ratio (in dB) of the interfering signal receive power at the UT antenna connector and desired signal receive power at the UT antenna connector when the desired signal receive power is at 3dB above the reference sensitivity values in Table 44 - Reference Sensitivity at FER = 10^{-2} and the interfering signal power is such that the desired signal FER reaches 10^{-2} .

Specification: When the desired signal uses any modclass between 0-6, ACS shall be at least 30 dB for an interfering signal on the first adjacent carrier. When the desired signal uses modclass 7 or 8, ACS shall be at least 27 dB for an interfering signal on the first adjacent carrier. For all modclasses (of the desired signal), ACS shall be at least 47 dB for interfering carriers within the UT tuning range which are beyond the first adjacent carrier. The interfering signal shall use modclass 8. See Table 45 - Adjacent Channel Selectivity.

Table 45: Adjacent Channel Selectivity

Interfering Carrier (Modulation Class 8)	Desired Signal Modulation Class	ACS
First Adjacent Carrier	0-6	30 dB
First Adjacent Carrier	7 or 8	27 dB
Other Inband Carrier	0-8	47 dB

5.4.2.2 Out-of-Band Blocking Characteristics

Definition: Out-of-Band Blocking measures the receiver’s ability to receive a desired signal on its assigned carrier in the presence of a CW interfering signal in the vicinity of its assigned carrier. The out-of-band blocking performance is the power of the CW signal, expressed (in dBm) measured at the UT antenna connector, when the desired signal power at the UT antenna connector is fixed at 3dB above the reference sensitivity values in Table 44 - Reference Sensitivity at FER = 10^{-2} and when the CW signal power is such that the desired signal FER is 10^{-2} .

Specification: The out-of-band blocking shall be as specified in Table 46 - Out-of-Band Blocking Characteristics.

Table 46: Out-of-Band Blocking Characteristics

Parameter	Value	
Desired Signal Power	Reference Sensitivity + 3 dB	
Interference Signal Frequency	0.1 to (X - 15) MHz	(Y+15) to 12750 MHz

Parameter	Value	
Interference Signal Power	≤ -40 dBm	≤ -40 dBm

Where:

- X – lower end of spectrum allocation.
- Y – upper end of spectrum allocation.

5.4.3 Receiver Leakage Power

Definition: The Receive Leakage Power is the power of all emissions generated or amplified by the radio receiver during any reception interval.

Specification: The receiver leakage power of the UT shall be less than -65 dBm measured with a minimum of 1 MHz bandwidth within the UT receiver tuning range.

5.4.4 RF Measurement Characteristics

This section defines the ability of the UT radio receiver to measure signal parameters required for radio resource management.

5.4.4.1 DSSI Estimator Accuracy

Definition: The DSSI estimator is required to support open loop TX gain control. The difference between the output value of the Desired Signal Strength Indicator (DSSI) estimator and the RF input level of the UT receiver P_R expressed in dB. The DSSI estimator reports a value of SRRC-filtered RF power, at the antenna connector.

Specification: DSSI Estimator accuracy shall be within ± 4 dB for signals having P_R greater between -105 dBm and -45 dBm. DSSI Estimator accuracy shall be within ± 6 dB for signals having P_R between -110 dBm and -105 dBm. Refer to Table 47 - Tabulated Range of Acceptable DSSI Report Values.

Table 47: Tabulated Range of Acceptable DSSI Report Values

Input power P_R (dBm)	Min. DSSI report (dBm)	Max. DSSI report (dBm)
$-45 < P_R$	-49	$P_R + 4$
$-105 < P_R \leq -45$	$P_R - 4$	$P_R + 4$
$-110 < P_R \leq -105$	$P_R - 6$	$P_R + 6$
$P_R \leq -110$	no minimum	-104

5.4.4.2 SINR Estimator Accuracy

Definition: The SINR Estimator is required for closed loop power control. The SINR Estimator Accuracy is the difference between the output value of the SINR estimator and the received

SINR at the antenna connector. For bursts with training sequences, SINR should be calculated from the training sequences alone.

Specification: Refer to Table 48 - Tabulated Range of Acceptable SINR Report Values.

Table 48: Tabulated Range of Acceptable SINR Report Values

Input SINR, S (dB)	5th Percentile (dB)	95th Percentile (dB)
$S < -3$ dB	no minimum	0 dB
$-3 \leq S < 25$ dB	$S - 3$ dB	$S + 3$ dB
$S \geq 25$ dB	22 dB	$S + 3$ dB

6 BASE STATION RADIO TRANSMISSION AND RECEPTION

6.1 Introduction

This chapter describes the radio frequency (RF) performance characteristics of the *HC-SDMA* base station, and serves as the reference from which manufacturers must generate compliance specifications.

Emissions behaviors specified in this document may or may not satisfy the regulatory requirements of any particular country or region. Equipment shall conform to whichever specification is more restrictive, that of this document or the applicable local regulations.

6.2 Carrier And Spatial Channel Support

The base station shall support all three HC-SDMA timeslots. The base station shall at a minimum support $N_f = 1$ carriers and $N_s = 1$ spatial channels. A spatial channel is one of possibly multiple data communication channels that share the same carrier and timeslot are prevented from interfering with one another through the use of multi-antenna spatial processing.

6.3 Transmitter Characteristics

The base station supports communication with multiple user terminals, each potentially transferring data on multiple data streams (see Glossary definition in *Annex A: Glossary*).

The base station shall support power control on a per-stream and per-frame basis. To accomplish this, output power levels shall be independently controllable for each of the streams on the three time slots, N_f carriers and N_s spatial channels. All performance specifications shall be met over the entire range of power levels and combinations of power levels for streams, and it shall be possible to vary each stream's power independently of the others. In an installation with multiple antennas, it is likely that due to a combination of power control and antenna weighting, one or more carriers on one or more antennas will be commanded to power levels lower than that described in clause 6.3.1.2, Minimum Stream Power. The base station shall support per carrier and per antenna power level commands down to zero power, but the distortion and spurious performance requirements do not apply to signals at these low power levels.

6.3.1 Output RF Power Characteristics

6.3.1.1 Rated Power Per Data Stream

Definition: The rated power per data stream P_{RAT} is defined as the highest SRRC-filtered power level such that when the base station opens a data stream with a user terminal, the power available to the new stream is at least P_{RAT} , while meeting all HC-SDMA specifications. For the case of a multi-antenna base station, P_{RAT} is the incoherently summed power of signal for the new data stream from all antennas.

6.3.1.2 Minimum Stream Power

Definition: The minimum stream power is defined as a SRRC-filtered power level, summed over all antennas. It is expressed in dBm:

$$P_{min} = P_{RAT} - 15\text{db}$$

Specification: The base station power control shall support transmit powers per data stream between at least P_{min} and P_{RAT} .

6.3.1.3 Transmit Power Control Step Size

Definition: The Transmit Power Control Step Size is the smallest interval that can be commanded between any two values of data stream power between P_{min} and P_{RAT} , expressed in decibel. The power levels are measured as SRRC-filtered power.

Specification: The transmit power control step size shall be no greater than 1 dB.

6.3.1.4 Downlink Power Control Maximum Rate

Definition: The Downlink Power Control Maximum Rate specifies the frequency with which the BS is able to update the transmit power of every active data stream.

Specification: The BS shall update the downlink power independently for every individual carrier, time slot and spatial channel combination.

6.3.1.5 Transmit Power Control Nonlinearity

Definition: The transmit power control nonlinearity (in dB) is defined as the absolute value of the error in output power step size when the transmitter is commanded to change power for a given data stream between any two valid power levels. That is, if the transmitter is commanded to change power by P_1 dB, and the resulting change in power is P_2 dB, the transmit power control nonlinearity is $|P_1 - P_2|$.

Specification: The transmit power control nonlinearity shall be less than 1 dB.

6.3.1.6 Transmit Leakage Power

Definition: Transmit Leakage Power is the sum of power radiated by all antennas of a base station, during any uplink time slot. Transmit leakage power is measured on a given carrier using a minimum 1 MHz bandwidth.

Specification: Transmit leakage power shall be less than -60 dBm for all frequencies within the carrier allocation.

6.3.1.7 Guard Time Transmit Power

Definition: Guard time transmit power is the sum of powers radiated by all antennas of a base station, during any interval designated as “Guard Time” in clause 3.2.2, Frames and Slots. Guard time transmit power is measured on a given carrier using a minimum 1 MHz measurement bandwidth

Specification: Guard time transmit power shall be less than $P_{RAT} - 40$ dBm for all carrier frequencies in the base station’s carrier sets.

6.3.1.8 Transmitter Power Absolute Error

Definition: The transmitter power error is defined as the ratio between commanded output power for any data stream and the actual power transmitted for that data stream. The ratio is expressed in decibels.

Specification: The transmitter power error shall be between -3 dB and $+3$ dB.

6.3.1.9 Carrier Frequency Error

Definition: Carrier frequency error is the difference between the programmed and actual transmitted base station carrier frequency, measured in parts per million (PPM).

Specification: Carrier frequency error shall not exceed 0.05 PPM.

6.3.2 Transmit Modulation

This section defines the performance of the transmit modulator. Transmit modulation is the process of converting a digital symbol sequence into an analog waveform.

6.3.2.1 Transmit Symbol Clock

Definition: The transmit symbol clock governs the timing of symbols within a burst, as well as the beginning of each burst.

Specification: The transmit symbol clock has a rate of $T_s = 2\mu\text{s}$ and is synchronized for all carriers, all spatial channels and all transmitters of a given base station.

6.3.2.2 Transmit Burst Timing Error

Definition: Transmit burst timing error is the absolute accuracy of the base station's symbol clock.

Specification: Transmit burst timing error shall be less than $2\mu\text{s}$.

6.3.2.3 Transmit Pulse Shaping Filter

Specification: The pulse shaping filter on the base station transmitter shall be a root-raised cosine filter with roll-off factor $\alpha = 0.25$ and symbol period $T_s = 2 \mu\text{s}$.

6.3.2.4 Phase Polarity

Definition: A transmitter is said to have positive (rather than negative) phase polarity if tones with faster advancing phase at the base-band input produce a higher frequency signal at the antenna.

A receiver is said to have positive (rather than negative) phase polarity if higher frequency tones incident at the antenna produce signals with faster-advancing phase at the base-band output.

Specification: *HC-SDMA* radios shall have positive phase polarity both for transmitters and receivers. It is permissible to use radios with negative phase polarity only if the base-band signal phase is correspondingly reversed.

6.3.2.5 Ramp-Up and Ramp-Down Transmit Power

For each burst type, the transmitter ramp-up and ramp-down shall occur respectively during the ramp-up and ramp-down periods defined in clause 3.3, Burst Formats.

Definition: Ramp-up and ramp-down transmit power is the highest power transmitted by a given antenna on a given carrier during a ramp-up or ramp-down period. Ramp-up and ramp-down transmit power should be measured within at least a 1 MHz bandwidth.

Specification: Ramp-up and ramp-down transmit power shall not exceed 6dB relative to the average power transmitted over the useful symbols of the corresponding burst.

6.3.2.6 Modulation Accuracy

The modulation accuracy calculation is described in clause 5.3.6, Modulation Accuracy. The error vector is the difference between the theoretically optimal desired waveform and the transmitted waveform at the symbol points, after receive SRRC filtering is applied to both waveforms and the initial phase, amplitude, frequency offset, and timing offset have been identified by a least-squares search.

Definition: Let a single stream be active on frequency carrier n , with transmitted power level P_{RAT} for the entire array. The *MA* for the array shall be the highest *MA* for the individual transmitters in that array.

Specification: The *MA* for the array shall not exceed 3.5% for all modulation classes with equal weighting over all N antennas and total transmitted power P_{RAT} .

See the definition of modulation accuracy in clause 5.3.6, Modulation Accuracy. See also the EVM measurement discussion in clause 5.3.6, Modulation Accuracy.

6.3.3 Output RF Spectrum Characteristics

This section defines restrictions on the output RF spectrum of the BS radio transmitter i.e., spectral masks.

6.3.3.1 Occupied Bandwidth

Definition: Occupied bandwidth is the frequency range that contains 99 percent of the total power of a single carrier's transmission with 0.5 percent of the total transmitted power below this range and 0.5 percent of the total transmitted power above this range. For base stations with $N > 1$ antennas, Occupied bandwidth of the base station is the highest Occupied bandwidth of the signals at the individual antennas. The measurement is to be made while one data stream is present, total power radiated by all the antennas is P_{RAT} , and equal power is radiated by each of the antennas.

Specification: The occupied bandwidth of any transmitter and any carrier for a base station shall not exceed 625 kHz.

6.3.3.2 Adjacent Carrier Power

Definition: Adjacent carrier power (ACP) is the SRRC filtered power radiated from all antennas on any carrier adjacent to carrier n , averaged over the entire downlink time slot s . The result is expressed in dBm. This is to be measured on adjacent carriers to one active data stream of modclass 8 at power P_{RAT} within the carrier allocation

Specification: ACP shall be less than $(P_{RAT} - 43)$ dBm in the adjacent carrier within the carrier allocation, and less than $(P_{RAT} - 50)$ dBm for carriers with center frequency more than 625 kHz away from f_n .

6.3.3.3 Multi-carrier Inter-modulation Products

Definition: Given any unoccupied carrier, the multi-carrier inter-modulation product (MCIP) is defined as the highest SRRC filtered output power on that unoccupied carrier, summed over all antennas, with equal power on all other carriers and equal composite power on all antennas. The measurement is expressed in dBm.

Specification: MCIP shall be less than $(P_{RAT} - 40)$ dBm with one unoccupied carrier, equal power on all occupied carriers, and equal composite power on all antennas.

6.3.3.4 Out-of-Band Spurious Emissions

Definition: Out-of-band spurious performance is defined as any radio emanation outside the HC-SDMA band allocated to the base station.

Specification: The base station shall meet all regulatory requirements in the jurisdiction within which it is installed. Emissions shall not exceed the limits in Table 49 - Out-of-Band Spurious Emissions Limits.

Table 49: Out-of-Band Spurious Emissions Limits

Offset from nearest HC-SDMA band edge	Emissions limit
0 kHz to 500 kHz	-3 dBm/100kHz

Offset from nearest HC-SDMA band edge	Emissions limit
500 kHz to 5 MHz	-16 dBm/100kHz
Beyond 5 MHz	-20 dBm/100kHz

6.4 Receiver Characteristics

6.4.1 Input RF Power Characteristics

This section defines the input power requirements of the base station radio receivers. Performance is defined for single receivers rather than for the base station as a whole.

Noise is to be measured in the 500 kHz wide nominal channel, with an idealized “brick wall” filter. Noise power present in the 62.5 kHz to either side of the nominal channel is not included in the specification.

6.4.1.1 Reference Sensitivity Level

Reference sensitivity level requirements for the base station (or “uplink”) receiver are based on frame error rate (FER) in the presence of additive Gaussian white noise, and are described in clause 12.9.1.2, Uplink Performance Specifications. Signal power measurements are to be made on SRRC-filtered waveforms. See clause 5.1.1, Terminology for temperature-dependent adjustments to the measurements.

Table 50: Reference Sensitivity for FER = 10⁻²

Mod Class	Nominal Reference Sensitivity (dBm)
0	-109.8
1	-108.2
2	-106.5
3	-103.6
4	-101.4
5	-99.1
6	-97.1
7	-95.8

Specification: The reference sensitivity level of the base station receiver shall be no greater than 1.2dB above the nominal values tabulated in Table 50 - Reference Sensitivity for FER = 10⁻².

6.4.2 Adjacent Channel Selectivity

Adjacent channel selectivity (ACS) measures the receiver's ability to receive a desired signal on its assigned carrier in the presence of a modulated interfering signal on an adjacent carrier.

Definition: Given a single data stream active on carrier $n: 0 \leq n < N_f$, with 3 dB more received power than the tabulated “maximum” value of REFERENCE SENSITIVITY for 10⁻² FER and a

second stream of uncorrelated data on carrier $m:m \neq n, 0 \leq m < N_f$, the ACS is defined as the ratio of INPUT POWERS (expressed in dB) of stream m relative to stream n when the power of stream m is increased so that the FER for stream n is 10^{-2} .

Specification: The ACS shall be at least 30 dB adjacent carriers 625 kHz apart, and at least 46 dB for streams 1250 kHz or more apart.

Table 51: ACS

Interfering Carrier	ACS
First Adjacent Carrier	30 dB
Other Inband Carrier	46 dB

6.4.2.1 Non-Distorting Input Power

Definition: Non-distorting input power is defined as the maximum SRRC-filtered receive power at any antenna port such that the frame error rate (FER) does not exceed 10^{-2} .

Specification: The non-distorting input power shall be greater than - 45 dBm.

6.4.3 RF Measurement Characteristics

This section defines the ability of the base station to measure signal parameters required for radio resource management.

6.4.3.1 DSSI Estimator Accuracy

Definition: The Desired Signal Strength Indicator (DSSI) is required to support open loop power control. The DSSI is an estimate of SRRC-filtered input power P_R for a given active data stream. The DSSI Estimator accuracy is expressed as a decibel ratio between the actual value of P_R and the estimated value.

Specification: DSSI Estimator Accuracy shall be within the permitted range of the template shown in Table 52 - Tabulated Range of Acceptable DSSI Report Values.

Table 52: Tabulated Range of Acceptable DSSI Report Values

Input power P_R (dBm)	Min. DSSI report	Max. DSSI report
$-45 < P_R$	-49	$P_R + 4$
$-105 < P_R \leq -45$	$P_R - 4$	$P_R + 4$
$-110 < P_R \leq -105$	$P_R - 6$	$P_R + 6$
$P_R \leq -110$	no minimum	-104

6.4.3.2 SINR Estimator Accuracy

Definition: The SINR estimator is used for closed loop power control. SINR estimator accuracy is defined as the difference between the output value of the SINR estimator and the received SINR at the antenna connector. TCH bursts from an established stream shall be present at the

antenna (for testing purposes, the stream may or may not be communicating with the base station under test). The SRRC-filtered input power of the bursts and the SRRC-filtered input power of added gaussian noise are measured independently of the base station. Then the SINR estimator accuracy is the decibel ratio of the externally measured burst to noise power and the base station SINR estimator output. SINR should be calculated from the training sequence portions of the bursts.

The SINR estimator error is the difference between the output value of the SINR estimator and the SINR present at the antenna.

Specification: Refer to Table 53 - Tabulated Range of Acceptable SINR Report Values.

Table 53: Tabulated Range of Acceptable SINR Report Values

Input SINR, S (dB)	5th Percentile (dB)	95th Percentile (dB)
$S < -5$ dB	no minimum	-2 dB
$-5 \leq S < 25$ dB	$S - 3$ dB	$S + 3$ dB
$S \geq 25$ dB	22 dB	$S + 3$ dB

6.5 Multi-antenna Performance

6.5.1 Multi-antenna Architecture

Wide area *HC-SDMA* base stations shall employ a multi-antenna architecture that simultaneously utilizes multiple antennas for sending and receiving of data to a particular subscriber. Limited coverage base stations may use single antenna architecture. Figure 20 - Example of Base Station Multi-antenna Architecture shows an example of the block diagram for such an architecture.

Multi-antenna architectures make use of ordinary, off-the-shelf antennas. Received radio signals are digitized and accumulated by the receiver bank. This received data is then packaged and processed in the spatial temporal processing block. This block involves the use of high-performance digital signal processing to extract and demodulate the various signals of interest and appropriately package the results for transport through the network interface and on into the network.

At the same time, data is being received from the network bound for subscribers. The spatial temporal processing block communicates with the transmitter bank to indicate how the data is to be weighted across the different antennas. Finally, the modulated data is routed through power amplifiers (PAs), one for each antenna, and transmitted across the array.

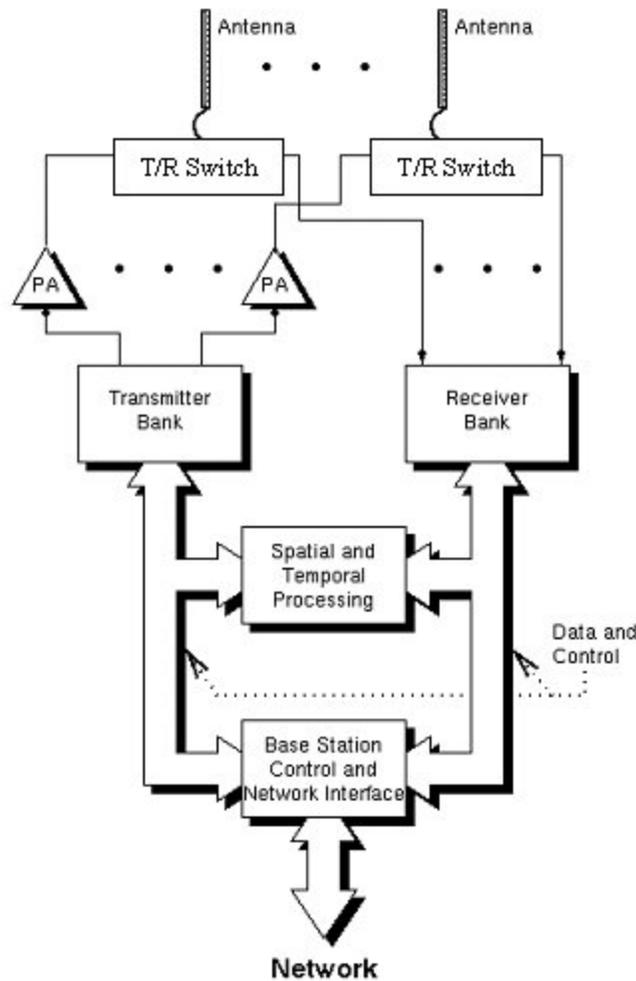


Figure 20: Example of Base Station Multi-antenna Architecture

6.5.2 Uplink and Downlink Multi-antenna Operation

On the uplink, typically, the received signal from each of the spatially distributed antenna elements is multiplied by spatial temporal weights, a bank of filters with complex delay tap adjustments of amplitude and phase for each of the array elements. These signals are combined to yield the array output. An adaptive algorithm controls the weights according to predefined objectives such as “tuning in” to a particular user while “tuning out” interference and noise. This processing should be performed independently and simultaneously for each of the users being served by the base station. Uplink bursts in *HC-SDMA* contain spatial training data for the express purpose of aiding the process of weight calculation.

A similar process should occur on the downlink. The signals to be transmitted are multiplied by filter banks with complex delay tap adjustments of amplitude and phase for each of the array elements. The weighting factors are chosen dynamically to ensure that the transmitted signals constructively combine and add at the user of interest while at the same time presenting minimal interference to other co-channel users. The weight factors are again chosen dynamically based on predefined objectives. Because, *HC-SDMA* is a TDD system, reciprocal knowledge from the uplink should be used in adapting downlink weights.

Finally, downlink broadcast and paging channels contain additional coding over and above that normally used on traffic channels in order to equalize the link budgets for these less directive channels against the link budgets for directive channels such as the traffic channels.

These dynamic adjustments should enable the base station to tune itself so that equivalent received signal level on the uplink and downlink are improved while at the same time co-channel interference received from other user terminals is rejected and co-channel interference transmitted towards other co-channel user terminals is minimized.

6.5.3 Calibration Figure of Merit

Definition: The calibration figure of merit (FOM) is defined as the decibel ratio of the maximum power the base station can potentially deliver to a UT to the minimum power the base station can potentially deliver to the same UT while maintaining the total downlink transmitted power constant at P_{RAT} . See Figure 21 – FOM Definition.

Specification: FOM shall be in accordance with Table 54 - Figure of Merit Requirements.

Table 54: Figure of Merit Requirements

Number of antennas	FOM (dB)
1	No requirement
2 - 3	20
4 - 7	25
More than 7	35

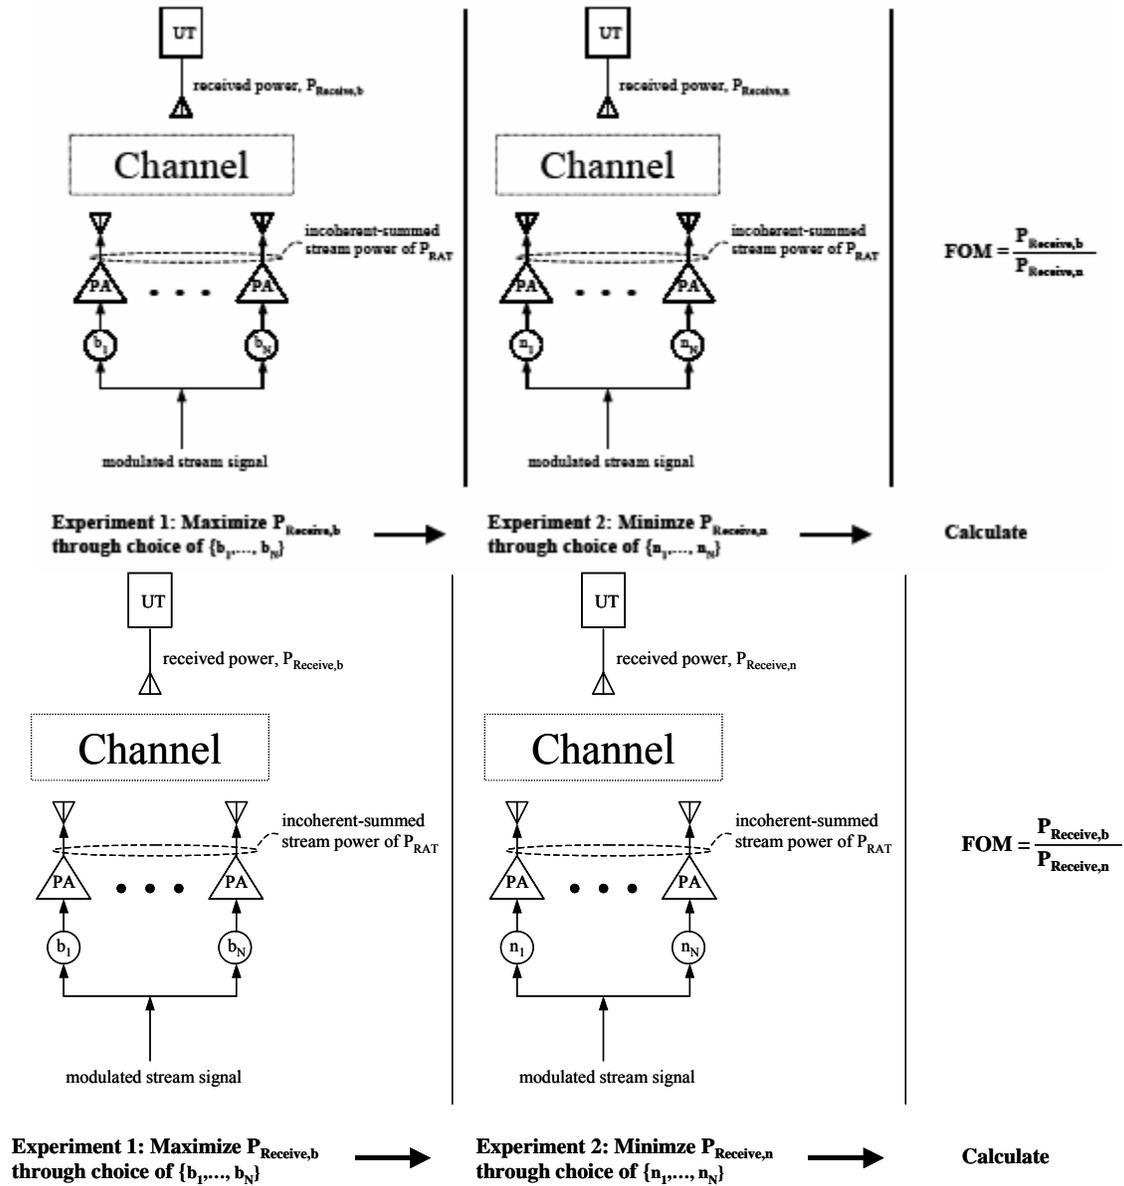


Figure 21: FOM Definition

7 L2 MAC PROTOCOL SUBLAYER SPECIFICATION

7.1 L2 MAC Overview

The L2 MAC layer provides the following services:

- ◆ Access management and control functions between *HC-SDMA* user terminals (UT) and base stations (BS).
- ◆ Mapping of logical to physical channels.
- ◆ Transfer services for control and traffic data by way of logical channels.

7.1.1 Access Overview

An *HC-SDMA* UT forms a relationship with a base station called a *registration*. A registration enables a given UT and a given BS to securely exchange control and user datagrams over *streams*. The UT may have multiple simultaneous registrations to one or more BSs. For example, the registrations to two BSs can overlap during a handover (a make-before-break handover), or there can be multiple modems on a UT. Additionally, multiple streams that share a common registration can be opened simultaneously and transfer packets in parallel. This is called *stream aggregation*. Though the multiple streams share a common registration, each stream is managed independently.

During a typical registration, many streams start, transfer data for a short time (a few seconds or less), and end. If no streams are open, the UT periodically checks for downlink pages. After a long idle period, the registration times out and ends. The UT can postpone deregistration by polling the BS.

The registration machinery is transparent to highlayer connectivity. In particular, the Layer 4 *network session*, used to transport IP datagrams between the UT and the wireline network, persists across registrations and handovers.

7.1.2 L2 MAC Logical Channels Overview

The L2 MAC layer provides transfer services for data and control messages by way of logical channels. The types of logical channels and the structure of messages carried by these channels are defined below.

There are both shared and dedicated logical channels. For shared channels, an address is included in each message header to identify the sender and intended recipient (except if it is broadcast). In the case of dedicated channels, the sender and intended recipient are identified implicitly by the physical channel and the training data used in the corresponding bursts.

Traffic channels (TCH) are dedicated. Control channels are either shared or dedicated, depending on the type.

One type of dedicated control channel is the Fast Associated Control Channel (FACCH). The FACCH is a control channel associated with all standard uplink and standard downlink bursts. It provides the UT with the current and recommended mod class information. It provides the BS with the current mod class and power information. Other control channels of the dedicated type are multiplexed with traffic. Multiplexing procedures are defined in clause 7.3.7.2, RMU Mux Function (UM control) and in clause 9.4.2, Outgoing Data Flow (AM control).

The other class of control channels, shared control channels, is further classified into uplink, downlink, and bidirectional channels:

- ◆ The Broadcast Channel (BCH) is a shared downlink control channel that provides synchronization and system access parameters.
- ◆ The Paging Channel (PCH) is a shared downlink control channel used by the BS to initiate data streams.
- ◆ The Configuration Channel (CCH) is a shared bidirectional control channel that enables the UT to learn information needed to initiate a data stream and register with the BS.
- ◆ The Random Access Channel (RACH) is a shared bidirectional control channel used by the UT to initiate data streams and by the BS to confirm stream requests and assign resources.

The relationship between logical channels, messages, and burst types is described in Table 55 – Logical Channels, Messages, and Burst Types. See clause 3.3.3, Broadcast Burst for burst definitions.

Table 55: Logical Channels, Messages, and Burst Types

Logical Channel	Messages	Burst Types
BCH	frequency synch message timing synch message broadcast message	frequency synchronization timing synchronization broadcast
PCH	page message	page
RACH	request access message access assignment message	standard uplink standard downlink
CCH	configuration request message configuration message	configuration request standard downlink
TCH	uplink traffic downlink traffic	standard uplink standard downlink
FACCH	current modulation class (odd RFN's) ^a additional available power (even RFN's) recommended modulation class (odd RFN's) current modulation class (even RFN's)	standard uplink standard uplink standard downlink standard downlink

a. See 7.3.5.1.2.3.7, Field definitions for subtype-specific fields for the definition of the relative frame number (RFN).

7.2 Layer to Layer Communication

7.2.1 L3 RM Communication

Refer to clause 10.3.2.11, L3 RM to L2 MAC Interface Primitives for a list of the communication primitives from L3 RM to L2 MAC.

Refer to clause 10.3.2.10, L2 MAC to L3 RM Interface Primitives for a list of the communication primitives from L2 MAC to L3 RM.

7.2.2 L3 RRC Communication

Refer to clause 10.3.2.7, L3 RRC to L2 MAC Interface Primitives for a list of the communication primitives from L3 RRC to L2 MAC.

Refer to clause 10.3.2.6, L2 MAC to L3 RRC Interface Primitives for a list of the communication primitives from L2 MAC to L3 RRC.

7.2.3 L3 MMC Communication

Refer to clause 10.3.2.9, L3 MMC to L2 MAC Interface Primitives for a list of the communication primitives from L3 MMC to L2 MAC.

Refer to clause 10.3.2.8, L2 MAC to L3 MMC Interface Primitives for a list of communication primitives from L2 MAC to L3 MMC.

7.2.4 L2 RLC Communication

Refer to clause 10.3.2.5, L2 RLC to L2 MAC Interface Primitives for a list of communication primitives from L2 RLC to L2 MAC.

Refer to clause 10.3.2.4, L2 MAC to L2 RLC Interface Primitives for a list of communication primitives from L2 MAC to L2 RLC.

7.2.5 L1 Communication

Refer to clause 10.3.2.1, L1 to L2 MAC Interface Primitives for a list of communication primitives from L1 to L2 MAC.

Refer to clause 10.3.2.2, L2 MAC to L1 Interface Primitives for a list of communication primitives from L2 MAC to L1.

7.3 Logical Channels

7.3.1 Conventions

7.3.1.1 Ordering of bits in message fields

Unless stated otherwise, in this chapter, the bits in the message fields are ordered from least significant bit (LSB) to most significant bit (MSB). For example, assume an L2 MAC message is composed of 2 fields, `field 1` and `field 2`. The field sizes and order are specified as:

Table 56: Convention Example

Example Message		
Field	# of Bits	Bit Positions
<code>field 1</code>	4	0:3
<code>field 2</code>	2	4:5
Total	6	

In this example, `field 1` is composed of 4 bits, while `field 2` is composed of 2 bits. The LSB of `field 1` is in message bit position 0, while the MSB of `field 1` is in message bit position 3. Similarly, the LSB of `field 2` is in message bit position 4, while the MSB of `field 2` is in message bit position 5.

Additionally, L1 receives the bits of an L2 message in bit position order with bit position 0 delivered first. In this example, L2 MAC delivers the LSB of `field 1` first.

7.3.1.2 Message Field to Value Conversion

Unless otherwise defined, when a message field `msgfield` corresponds to a scalar value, the message field is mapped to an unsigned integer `uint(msgfield)` in the natural way. For example, `uint(10110) = 22`. A subsequent formula can be applied to this value to complete the mapping.

7.3.2 Broadcast Channel (BCH)

The broadcast logical channel (BCH) provides the information needed for a UT to open a configuration logical channel (CCH) to the BS. It also provides information to guide UT initial cell selection and handover decisions.

The broadcast logical channel transports broadcast messages. The payload structures of the Frequency Synchronization (F), Timing Synchronization (T), and Broadcast (B) bursts shall be as given in clause 3.3, Burst Formats.

All BCH's for every BS of the *HC-SDMA* network shall be located on timeslot 0. Any BS may use any RF carrier for its BCH. A conventional channel used for BCH, corresponding to timeslot 0 and the RF carrier used for that BCH, shall not be used for access requests or traffic.

7.3.2.1 Basic Access Description

Figure 22 - BCH Basic Access Diagram shows the basic access flow for BCH scan. The UT may initiate a BCH scan for several reasons. Examples include finding candidate base stations in order to register (see clause 9 - L3 Protocol Specification for details on registration) or looking for candidate base stations for handover. In these cases, the BCH scan enables the UT to send a CR -- see clause 7.3.4, Configuration Channel (CCH) -- to these candidate base stations. Additionally, the UT may use a BCH scan to maintain timing/frequency synchronization to the base station during long idle periods.

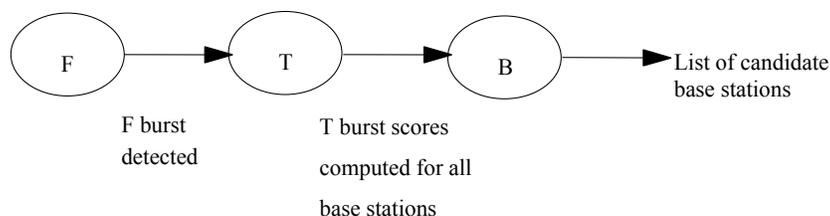


Figure 22: BCH Basic Access Diagram

7.3.2.2 Superframe Structure

Clause 3, Slot and Frame Structure specifies the F, T, and B bursts. These bursts shall be placed in a BCH super frame structure as shown in Figure 23 - BCH Superframe. The F burst contains no information and is used only to give the UT a coarse estimate of frame timing and an accurate estimate of the frequency offset and to synchronize to the BCH superframe given by Figure 23 - BCH Superframe. The F burst shall be transmitted on frames where $AFN \bmod 20 = 0$. All BSs shall transmit the F burst simultaneously.

The T burst is a Walsh-Hadamard code that allows the UT to estimate which base stations are the strongest and gives the UT an estimate of an accurate frame timing to these base stations. T burst is distinct for each BS since it is constructed based on the color code, thus allowing timing and received signal strength to be determined for each BS. All BSs shall transmit their T bursts simultaneously.

The B bursts allow the UT to determine a list of candidate base stations from the set of strongest base stations given by the T burst. Also, the UT improves its timing offset estimate for each of these base

stations from the B burst. The UT also obtains information on the traffic load for each of the base stations on the candidate list from the B-burst. There shall be eight B burst frames ($B_0, B_1, B_2, \dots, B_7$). A base station whose color code BSCC satisfies $BSCC \bmod 8 = n$ shall transmit its B burst during frame B_n . This is illustrated in the superframe structure of Figure 23 – BCH Superframe.

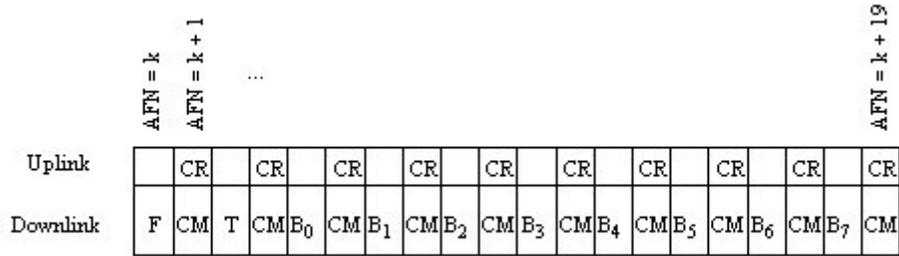


Figure 23: BCH Superframe

7.3.2.3 BCH Message Fields

The broadcast (B) burst shall convey six bits of information p_1, \dots, p_6 as described in clause 3.3.3, Broadcast Burst and clause 4.3.3, Broadcast Burst. This shall be mapped to BCH messages as given by Table 57 – BCH Message Fields with p_6 as the MSB.

Table 57: BCH Message Fields

BCH Message		
Field	# of Bits	Bit Positions
bsTxPwr	4	0:3
bsLoad	2	4:5
Total	6	

- ◆ bsTxPwr: Shall be the mean transmitted power per antenna element of the B burst = $3 * \text{uint}(\text{bsTxPwr}) \text{dBm}$.
- ◆ bsLoad: Shall conform to Appendix 7A. Setting bsLoad on BCH.

7.3.3 Paging Channel (PCH)

The paging logical channel prompts the UT to initiate a data stream with the BS. The paging logical channel transports paging messages. Paging messages shall be mapped onto page bursts by the physical layer.

The paging identifier (PID) and registration identifier (RID) shall be assigned during registration. See clause 9, L3 Protocol Specification for information. Each UT shall have one PID for all of its RIDs.

The information contained in the page burst shall be the following: BSCC, PID, and the 10 least significant bits of the AFN.

A page burst may be transmitted simultaneously with the maximum number of TCH channels on a given conventional channel (i.e., a frequency/timeslot pair).

7.3.3.1 Frequency and Timeslot Hopping for Paging Messages

The BS and a particular UT synchronize on the frequency and timeslot on which the particular UT will be shall paged using a *page hopping function*.

The page hopping function shall have the following form:

$$\text{page resource} = \text{PageHop}(\text{AFN}, \text{PCH logical set}, \text{PID}) \quad (7-1)$$

The PCH logical set shall be set by the PCH logical mask of the `reg params AM` message sent during registration. The UT shall scan the PCH logical set for paging messages. The paging resources in the PCH logical set shall be ordered in the order of their conventional channel ID (see Table 75 - Mapping of Frequency and Timeslots to Conventional Channels). See clause 7.5, Page Hopping Sequences for details on the `PageHop` function.

7.3.3.2 Base Station Page Procedures

To open a stream on the BS, L3 RM shall send an `L2MacBsPageTx.req` request to BS L2 MAC. BS L2 MAC shall then schedule paging messages for this UT as described in this section.

The paging messages shall be mapped onto page bursts. Up to two subslots may contain page bursts with identical payloads (and with possibly different spatial transmit weights). Each paging resource in the PCH logical set shall be mapped to a pair of subslots on the same carrier frequency, and this mapping is defined by the PCH conventional mask in the `reg params AM` message sent during registration. For every UT, the PCH logical set shall be configured to be a subset of the conventional channels marked “on” in the PCH conventional mask. Different settings of the PCH conventional mask on a given carrier frequency shall correspond to a subslot structure as shown in Figure 24 - Paging Subslot Diversity.

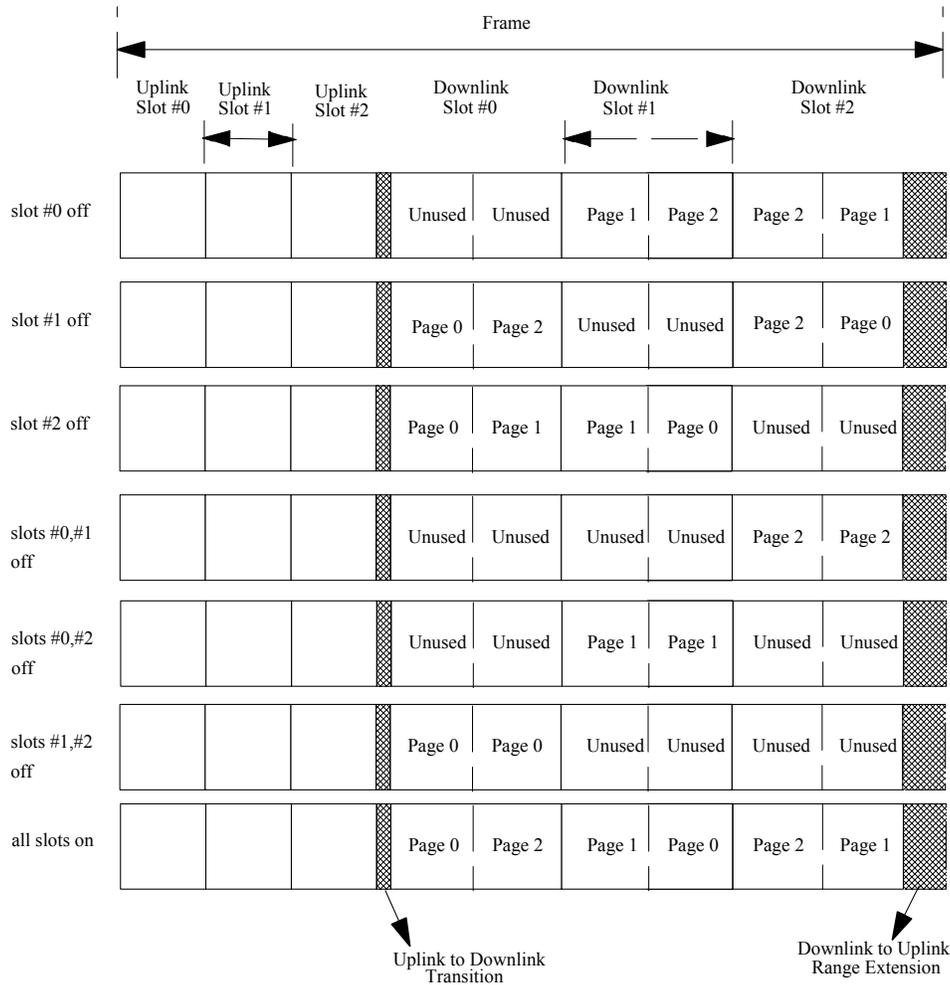


Figure 24: Paging Subslot Diversity

A page burst has an implicit correspondence with a paging resource on the carrier, called the implied resource of the page. In Figure 24 - Paging Subslot Diversity, a subslot labeled “Page k ” shall corresponds to the timeslot k paging resource on that carrier, $k = 0, 1, 2$.

For example, assume for a given carrier that all slots in the PCH conventional mask are on. Then if the UT detects a paging message using the second subslot from timeslot #1, the implied resource is the paging resource corresponding to timeslot #0. As a second example, assume for a given carrier that timeslot #0 is off while timeslots #1 and #2 are on. Then if the UT detects a paging message using the second subslot from timeslot #1, the implied resource is the paging resource corresponding to timeslot #2.

The BS should introduce no inconsistencies in the PCH conventional masks given to different UT’s, allowing all UT’s that listen to the same carrier frequency for page bursts to share the same paging subslot structure on that carrier.

The BS can send page bursts out of subslots of the conventional channels that are turned “on” in the union of all its UT’s PCH conventional masks. For every paging resource, the BS shall compute a list of the PID’s to which it could send a page message on that paging resource during that frame. Specifically:

- ◆ For a given paging resource c , the base station shall compute a list of PID's whose outcome from Equation 7-1 is c .
- ◆ Of these, the base station shall include only PID's that have paging enabled (using the `L2MacBsPageTx.req` request).
- ◆ Of these, the base station shall include only PID's that satisfy $(AFN \bmod PL) = (x \bmod PL)$, where PL is the current paging activity level for the PID as described in clause 9.3.4.1.2, Paging States and Transitions, and x is the UT's PID.

For each paging resource, the base station may choose from among these PID's to send one or more paging messages. The base station should avoid introducing `RA-page-response` collisions on the RACH resource.

The paging activity level for each PID shall be managed by L3 RM. Each time it changes, a new `L2MacBsPageTx.req` request or `L2MacUtPageRx.req` request shall be issued by L3 RM.

The *HC-SDMA* specification does not specify the arbitration method when there are multiple eligible PID's listening to a paging resource. The method used should ensure fairness and should give preferential treatment to PID's that are less frequently eligible to be paged (based on the paging activity levels).

If the base station receives an `RA-ping` with a PID that is enabled for paging, the BS MAC shall treat this `RA-ping` as a `RA-page-response` as described in clause 9.3.5.2, Ping Promotion to TCH.

A BS shall not page a UT that has any TCH streams open. When UT has TCH streams open, BS L3 RM may use an existing open stream to send an `aggregate.cmd` UM control message if the desired stream is for the same RID that owns the stream, or it sends an `in band page` AM control message if the desired stream is for a different RID.

7.3.3.3 User Terminal Page Procedures

If a UT has no TCH streams open, the UT shall listen for page bursts when $(AFN \bmod PL) = (x \bmod PL)$, where PL is its current paging level and x is the UT's PID. The paging level shall be sent to UT L2 MAC by the `L2MacUtPageRx.req` request from L3 RM. The UT shall listen on the paging resource as given by Equation 7-1, and the subslots of the carrier frequency where it listens for page bursts shall be deduced from its PCH conventional mask.

If the (BSCC, PID) of the decoded page message matches the (BSCC, PID) of one of the UT's current registrations, the UT shall respond with an `RA-page-response` using the access request algorithm described in clause 7.3.5, Random Access Channel (RACH). Additionally, the UT shall use the directed RACH retry policy, as configured by the `reg params` AM message sent during registration and as defined in clause 7.3.5.2.2, Directed RACH retry policy. If the UT detects another page before receiving an AA for its `RA-page-response` message, the UT shall reinitialize the access request algorithm as if the latest page were the initial page.

7.3.4 Configuration Channel (CCH)

The configuration channel allows the UT to obtain the information such as uplink timing alignment necessary to open an initial stream with the BS. The CCH shall be located on the same frequency/timeslot as the BCH (see Figure 23 - BCH Superframe).

The configuration channel shall transport configuration request (CR) messages on the uplink and configuration messages (CM) on the downlink. A configuration request message shall be mapped onto

a configuration request burst by the physical layer. A configuration message shall be mapped onto a standard downlink burst by the physical layer. A CR/CM exchange shall follow a BCH acquisition during a UT's initial contact with a base station (for example, when a UT powers on, or when a UT performs a handover). Additionally, a CR/CM exchange may occur when a UT realigns its uplink timing, such as when constant RA failures might point to such a problem. Only the UT may initiate a CR/CM exchange. On the UT, L3 MMC shall control CR/CM exchange initiation. The CR burst power level shall be controlled by L3 RRC, and the CR retry strategy shall be controlled by L2 MAC. The CR/CM shall exchange last 1 frame.

7.3.4.1 Configuration Request (CR)

The CR message shall conform to Table 58 - Configuration Request Message Fields.

Table 58: Configuration Request Message Fields

Configuration Request Message		
Field	# of Bits	Bit Positions
identity	8	0:7
adjTxPwr	5	8:12
capability	1	13
Reserved	3	14:16
Total	17	

- ◆ **identity**: Shall be random bits that disambiguate simultaneous CR messages from multiple UTs.
- ◆ **adjTxPwr**: Shall be the power used by the UT to transmit the configuration request burst, adjusted for the UT noise floor relative to the UT reference noise floor: $\text{adjusted UT power} = (3 \cdot \text{uint}(\text{adjTxPwr}) - 56) \text{ dBm}$, where $\text{adjusted UT power} = \text{UT transmit power} + (\text{UT noise floor} - \text{UT reference noise floor})$, and where $\text{UT reference noise floor} = -110 \text{ dBm}$.
- ◆ **capability**: Vendor specific definition. The capability field shall be set equal to 1 during normal operation.

7.3.4.2 Configuration Message (CM)

The CM message shall conform to Table 59 - CM Fields.

Table 59: CM Fields

Configuration Message		
Field	# of Bits	Bit Positions

Minimum protocol	4	0:3
identity	8	4:11
BCH offset	4	12:15
rachCarrierMask	16	16:31
rachSlotMask	3	32:34
pwrAdjust	4	35:38
pwrCtrl	1	39
AFN	10	40:49
toa	7	50:56
Network Operator ID	8	57:64
Reserved	40	65:104
Total	105	

- ◆ **Minimum Protocol:** Shall specify the lowest version of the *HC-SDMA* protocol that the BS can support per Table 60 - CM Minimum Protocol Field. If the UT protocol version is less than the lowest supported version inferred from the CM minimum protocol field, the UT shall disconnect from the BS and may try another BS, since the BS's protocol version is too recent to interoperate with the UT's. Similarly, if the UT protocol version is greater than the highest supported version inferred from the CM minimum protocol field, the UT shall reject the BS. In this case, the BS's protocol version is too old to interoperate with the UT's. The value of the minimum protocol field shall be determined based on compatibility of a protocol release with earlier versions.

Table 60: CM Minimum Protocol Field

minimum protocol field value	lowest <i>HC-SDMA</i> protocol version supported
0	reserved
1	1.0
2-15	reserved

- ◆ **identity:** Shall be the identity of the requester, learned from the configuration request message.
- ◆ **BCH offset:** Shall indicate the base carrier that the BS will be using for this UT, specified relative to the BCH. For example, if the offset field is equal to 0010, then the base carrier is $2 \times 625 = 1250$ kHz down from the BCH. The base carrier, denoted f_0 , is used to number the resources on the BS (see Table 75 - Mapping of Frequency and Timeslots to Conventional Channels). This, together with the RACH carrier set parameter, tells the UT which carriers it can use for RACH.
- ◆ **rachCarrierMask:** Shall indicate which carriers the UT is allowed to use for RACH, namely, the RACH carrier set. For example, suppose the BCH offset is equal to 0001 and rachCarrierMask is equal to 0000000011111111. In this case, the UT sends a registration RA, RA-rreg, on one of the 8 carriers, f_0, f_1, \dots, f_7 where f_1 is the BCH. The RACH carrier set only applies to RA-rreg attempts. During the registration stream, the BS shall inform the registration of the RACH resource set in the reg params AM message, which supersedes the RACH carrier set (and RACH slot mask).

- ◆ **rachSlotMask**: Shall indicate which timeslots the UT may use for RACH. The LSB of the RACH slot mask corresponds to timeslot 0. The RACH slot mask only applies to RA-rreg attempts. During the registration stream, the BS shall inform the registration of the RACH resource set in the reg params AM message, which supersedes the RACH slot mask (and RACH carrier set).
- ◆ **pwrAdjust**: Shall be the power adjustment the UT should apply to future configuration request bursts and standard uplink bursts: $\text{adjustment} = 2^{\text{uint}(\text{pwrAdjust})} - 26$ dB. The pwrAdjust field shall be set from the L2MacBsCm.req primitive from L3 RRC.
- ◆ **pwrCtrl**: Shall indicate whether the 1-bit pwrCtrl field in the RMU header of the downlink TCH burst is active on every burst. Setting pwrCtrl to 0 indicates that the field will be inactive. However, when a pwrCtrl.cmd UM control message is received, the 1-bit pwrCtrl field of the RMU header of the downlink TCH burst shall always used (see clause 9.2.3.4.2 Closed-loop Power Control).
- ◆ **AFN**: Shall be the 10 least significant bits of the Absolute Frame Number on which the CM was transmitted.
- ◆ **toa**: Shall be mapped to the actual time of arrival of the CR burst, T_A , by the equation $T_A = \text{uint}(\text{toa}) - 5$, in units of microseconds. This field shall be used by the UT to compute the timing advance adjustment it must apply to RAs that follow this burst. This toa field has a wider range and granularity than the toa field in the AA burst (see clause 7.3.5.1.2, AA field definitions). Once the UT has an initial timing estimate, the toa field in the AA burst shall be used to refine the timing advance adjustment. The toa field shall be set from the timing error measurement in the L1RxDetect.ind primitive from L1.
- ◆ **Network Operator ID**: Shall be the identification of the network operator that owns the BS. The UT may use this network operator identity to assist its decision of which BS to choose.

Table 61: Network Operator ID Field

Network Operator ID	Interpretation
0-15	reserved
16	reserved
17	reserved
18-255	reserved

7.3.5 Random Access Channel (RACH)

The UT uses the uplink portion of the random access channel to open a data stream with the BS. The BS uses the downlink portion of the random access channel to grant access requests and to assign resources to the associated data stream.

The random access logical channel transports request access (RA) messages on the uplink and access assignment (AA) messages on the downlink. A request access message shall be mapped onto a standard uplink burst by the physical layer. An access assignment message shall be mapped onto a standard downlink burst by the physical layer.

When the BS grants an access request, it shall assign a resource for the associated data stream. The resource shall be specified by the resource field of the AA message. A traffic logical channel (TCH) shall be created on the specified resource for the associated data stream.

7.3.5.1 RACH Message Definitions

7.3.5.1.1 RA Field Definitions

The RA field definitions shall conform to Table 62 – RA Field Definitions.

Table 62: RA Field Definitions

Request Access (except RA-rreg)			Request Access (RA-rreg)		
Field	# of Bits	Bit Positions	Field	# of Bits	Bit Positions
raType	3	0:2	raType	3	0:2
RID or PID	15	3:17	rreg extended ID	(15)	(3:17)
			ID	10	3:12
			UT class	5	13:17
UTtxPwr	5	18:22	UTtxPwr	5	18:22
Total	23		Total	23	

- ◆ raType: Shall be the type of the request access message. See Table 64 – raType Field Values.
- ◆ RID: Shall be the registration identifier, present in all but RA-rtts, RA-rtts-short, and RA-rtts-directed.
- ◆ PID: Shall be the page identifier, present in RA-page-response and RA-ping.
- ◆ rreg extended ID: Shall be the combination of the ID field followed by the UT class field of the RA-rreg message.
- ◆ ID: Shall be a random identifier, present in RA-rreg. This is used to distinguish multiple simultaneous RA-rreg messages.
- ◆ UT class: Shall give the noise figure capabilities of the UT according to the class identifier as in Table 63 – UT Class Mapping. If the BS receives a reserved UT class value, it shall respond with AA-reject (see clause 7.3.5.1.2.2, AA subtype definitions).

Table 63: UT Class Mapping

UT class	Noise Figure (dB)
0	5
1	7
2	9
3-31	reserved

- ◆ UTtxPwr: Shall be the power used by the UT to transmit the request access message on the standard uplink burst: $UT\ power = (3 * \text{uint}(UTtxPwr) - 56) \text{ dBm}$.

The RID field shall be set by the L2MacUtStart.req request from L3 RM. Request access messages include a 3 bit raType field allowing for a variety of requests. The request types shall conform to Table

64 - raType Field Values. In response to a reserved request type, the BS shall respond with AA-reject (see clause 7.3.5.1.2.2, AA subtype definitions).

Table 64: raType Field Values

Notation	Value (MSB..LSB)	Meaning
RA-rts	000	Stream request
RA-ping	001	Keepalive poll request
RA-rts-short	010	Short stream request
RA-rts-directed	011	Directed stream request
RA-page-response	100	Paged stream request
	101	Reserved
	110	Reserved
RA-rreg	111	Registration request

- ◆ **RA-rts:** The UT shall send this message to open a TCH data stream that is initiated on the UT side or that is in response to an in-band control message received on an ongoing stream, where the in-band control message does not indicate a directed resource.
- ◆ **RA-page-response:** The UT shall send this message to open a TCH data stream in a directed response to a downlink page.
- ◆ **RA-rts-directed:** The UT shall send this message to open a new data stream in response to an in-band control message received on an ongoing stream, where the in-band control message indicates a directed resource.
- ◆ **RA-rts-short:** The UT shall send this message to open a twoframe stream.
- ◆ **RA-ping:** The UT should send this message to notify the BS that the UT remains active and within communication range. An acknowledgment from the BS indicates that the RID expiration timers have been reset. This acknowledgment allows further adjustment to timing advance and power and possibly notifies the terminal that downlink data is waiting. This message is also sent to open a TCH stream in an undirected response to a downlink page.
- ◆ **RA-rreg:** The UT shall send this message to inform the BS that a UT wants to open a new registration.

7.3.5.1.2 AA field definitions

7.3.5.1.2.1 Common message format and field definitions

The AA field format shall conform to Table 65 - AA Message Format for Common Fields. There are some fields that vary depending on the subtype given by the aaSubType field. The field formats for each of these subtypes are given in clause 7.3.5.1.2.3, Subtype-dependent message formats and field definitions.

Table 65: AA Message Format for Common Fields

Access Assignment Common Fields		
Field	# of Bits	Bit Positions
ID	15	0:14
toa	6	15:20
pwrAdjust	4	21:24
aaType	1	25
aaSubType (subtype dependent fields)	4 variable	26:29 30:104
Total	105	

- ◆ ID: Shall indicate the intended recipient of the message. For AA subtypes AA-cts and AA-cts-short, this field shall contain the RID. For all other AA subtypes, this field shall contain the value of the PID, RID, or rreg extended ID contained in the corresponding RA.
- ◆ toa: Shall be the time of arrival of corresponding RA relative to BS frame timing. $T_a = (\text{uint}(\text{toa})/2 - 9)$, in units of microseconds (see Figure 25 - RA Time of Arrival Relative to Base Station Frame Timing). It shall be used by the UT to compute the timing advance adjustment it must apply during the subsequent stream.

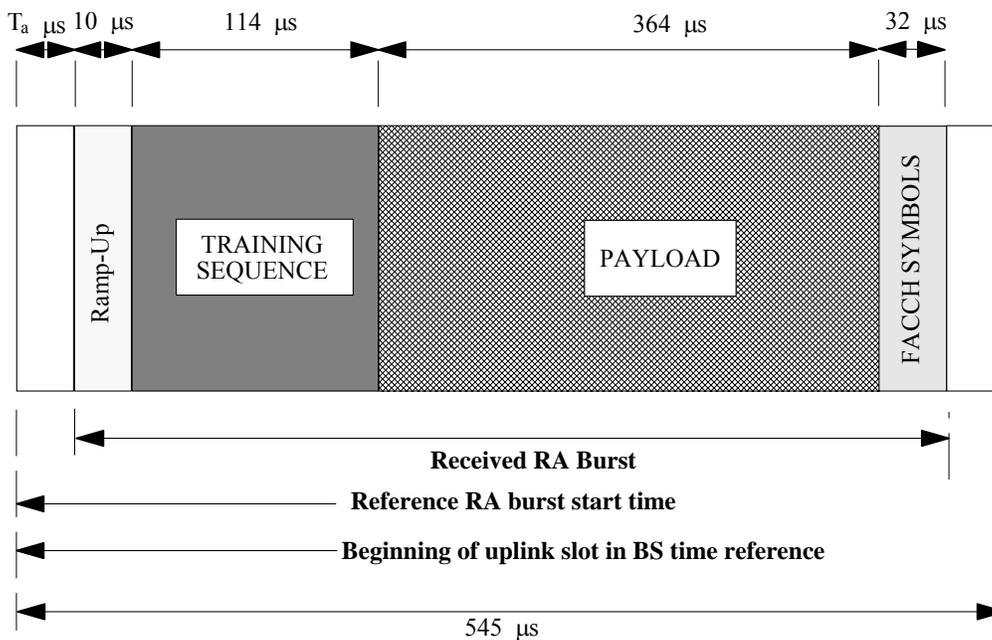


Figure 25: RA Time of Arrival Relative to Base Station Frame Timing

- ◆ pwrAdjust: Shall be the power control adjustment the UT should apply to the initial burst of subsequent stream: $\text{adjustment} = 2 * \text{uint}(\text{pwrAdjust}) - 26 \text{ dB}$.
- ◆ aaType: This field shall indicate whether a stream is granted with the AA, and also acts as a type bit that separates the AA messages into two sets of subtypes as indicated by the

aaSubType field. A value of 1 means that a stream is granted, and a value of 0 means that a stream is not granted.

- ◆ aaSubType: This field further specifies the meaning of the AA message.

7.3.5.1.2.2 AA subtype definitions

The AA subtypes corresponding to aaType = 1 and aaType = 0 shall respectively conform to Table 66 - aaSubType Field Values for aaType = 1 and Table 67 - aaSubType Field Values for aaType = 0. The last column in these tables indicate whether the paging activity levels (see clause 9.3.4.1.2, Paging States and Transitions) are reset upon sending the message at the BS and receiving the message at the UT.

Table 66: aaSubType Field Values for aaType = 1

Value (MSB..LSB)	Notation	Meaning	Resets Page Level
0000	AA-cts	Stream grant	Y
0011	AA-cts-short	Short uplink grant	Y
0111	AA-rreg	Registration grant	NA

- ◆ **AA-cts:** The BS shall send this message in order to open a stream in response to an RA-*rts*, RA-*rts-directed*, and RA-*page-response*. The BS shall send this message in order to open a stream in response to an RA-*ping* or RA-*rts-short* when there is downlink data queued for one of the registrations on the UT.
- ◆ **AA-cts-short:** The BS shall send this message in order to begin a two-frame stream in response to a short stream request. The downlink contains an acknowledgment of the uplink burst, then shuts down.
- ◆ **AA-rreg:** The BS shall send this message in order to begin an unencrypted stream for connection parameter exchange in response to an RA-*rreg*.

Table 67: aaSubType Field Values for aaType = 0

Value (MSB..LSB)	Notation	Meaning	Resets Page Level
0001	AA-reject	Stream reject	Y
0010	AA-ping-ack	Keepalive poll ack	N
0100	AA-cancel	Cancel false positive page	N
0110	AA-pending	Pending channel assignment	N
1001	AA-invalid-id	ID field in RA not valid	N

- ◆ **AA-reject:** The BS shall send this message in order to reject the UT's request and tells the UT when it is next allowed to request a stream. For this AA subtype, the ID field in the common fields portion shall be the same as the ID field contained in the corresponding RA.
- ◆ **AA-ping-ack:** The BS shall send this message in order to acknowledge an RA-*ping* when there is no downlink data queued for any connection on the UT. The BS shall concurrently reset the RID expiration timers for the registrations that have the ping renew bit set (see clause 9.5.1.3.3, UT Params Message).

- ◆ **AA-cancel:** If a paged stream request is received by the BS but no page is sent, the BS shall reply with the AA-cancel message to instruct the UT to return to a prepaged state.
- ◆ **AA-pending:** The BS shall send this message in order to request the sender to repeat the same RA message on the same conventional channel at a time specified by the AA-pending message.
- ◆ **AA-invalid-id:** The BS shall send this message in order to indicate the ID field in the corresponding RA (either a RID or PID) is not recognized as valid by the BS.

7.3.5.1.2.3 Subtype-dependent message formats and field definitions

In addition to the fields given in clause 7.3.5.1.2, AA field definitions, certain subtypes have other fields appended to this set.

7.3.5.1.2.3.1 Message Format for AA-cts and AA-short

The message format for AA subtypes AA-cts and AA-short shall conform to Table 68 - Message Format for Subtypes AA-cts and AA-short. The AA common fields consist of the first 5 fields of Table 65 - AA Message Format for Common Fields. For field definitions see clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.

Table 68: Message Format for Subtypes AA-cts and AA-short

AA-cts and AA-cts-short		
Field	# of Bits	Bit Positions
(AA common fields)	30	0:29
frameDec	4	30:33
tchIndex	3	34:36
modClassUp	4	37:40
modClassDown	4	41:44
resource	6	45:50
rmuHeaderType	1	51
Reserved	53	52:104
Total	105	

7.3.5.1.2.3.2 Message Format for AA-rreg

The message format for subtype AA-rreg shall conform to Table 69. The AA common fields consist of the first 5 fields of Table 65. For field definitions see clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.

Table 69: Message Format for Subtype AA-rreg

AA-rreg		
Field	# of Bits	Bit Positions
(AA common fields)	30	0:29
frameDec	4	30:33
tchIndex	3	34:36
Reserved	68	37:104
Total	105	

7.3.5.1.2.3.3 Message Format for AA-pending

The message format for subtype AA-pending shall conform to Table 70 – Message Format for Subtype AA-pending. The ID field in the AA common field portion shall be the same ID type as in the corresponding RA. The AA common fields consist of the first 5 fields of Table 65 - AA Message Format for Common Fields. For field definitions see clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.

Table 70: Message Format for Subtype AA-pending

AA-pending		
Field	# of Bits	Bit Positions
(AA common fields)	30	0:29
raIdType	2	30:31
raDelay	4	32:35
Reserved	69	36:104
Total	105	

7.3.5.1.2.3.4 Message Format for AA-reject

The message format for subtype AA-reject shall conform to Table 71 – Message Format for Subtype AA-reject. The AA common fields consist of the first 5 fields of Table 65 - AA Message Format for Common Fields. For field definitions see clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.

Table 71: Message Format for Subtype AA-reject

AA-reject		
Field	# of Bits	Bit Positions
(AA common fields)	30	0:29
raIdType	2	30:31
holdPage	1	32
holdTimer	5	33:37
Reserved	67	38:104
Total	105	

7.3.5.1.2.3.5 Message Format for AA-invalid-id

The message format for subtype AA-invalid-id. shall conform to Table 72 - Message Format for Subtype AA-invalid-id. The AA common fields consist of the first 5 fields of Table 65 - AA Message Format for Common Fields. For field definitions see clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.

Table 72: Message Format for Subtype AA-invalid-id

AA-invalid-id		
Field	# of Bits	Bit Positions
(AA common fields)	30	0:29
ridOrPid	1	30
Reserved	74	31:104
Total	105	

7.3.5.1.2.3.6 Message Format for other subtypes

The message format for subtypes not given in clauses 7.3.5.1.2.3.1, Message Format for AA-cts and AA-short to 7.3.5.1.2.3.5, Message Format for AA-invalid-id shall conform to Table 73 - Message Format for Subtypes with No Additional Fields. These other subtypes have no fields other than the common AA fields in clause 7.3.5.1.2, AA field definitions.

Table 73: Message Format for Subtypes with No Additional Fields

AA-ping-ack and AA-cancel		
Field	# of Bits	Bit Positions
(AA common fields)	30	0:29
Reserved	75	30:104
Total	105	

7.3.5.1.2.3.7 Field definitions for subtype-specific fields

- ◆ `frameDec`: Shall define the fractional rate (in frames) of the granted stream, from full-rate progressively down to fifth-rate. The BS grants a stream that occupies a subset of AFN's, as specified in Table 74 - `frameDec` Field Values. The RA/AA exchange that results in the stream grant takes place on any AFN.

Table 74: `frameDec` Field Values

Stream Rate	Value (MSB..LSB)	AFN's granted for stream
Full	0000	Every AFN
Half	0001	(AFN mod 2) = 0
Half	0010	(AFN mod 2) = 1
Third	0011	(AFN mod 3) = 0
Third	0100	(AFN mod 3) = 1
Third	0101	(AFN mod 3) = 2
Quarter	0110	(AFN mod 4) = 0
Quarter	0111	(AFN mod 4) = 1
Quarter	1000	(AFN mod 4) = 2
Quarter	1001	(AFN mod 4) = 3
Fifth	1010	(AFN mod 5) = 0
Fifth	1011	(AFN mod 5) = 1
Fifth	1100	(AFN mod 5) = 2
Fifth	1101	(AFN mod 5) = 3
Fifth	1110	(AFN mod 5) = 4
Reserved	1111	

- ◆ `tchIndex`: `tchIndex` shall indicate which training sequence to use for the uplink and the downlink. See clause 3.3.6, Standard Uplink Burst (FAACH, RACH, & TCH) and clause 3.3.7, Standard Downlink Burst (CM, AA, & TCH) for a detailed description.

- ◆ **modClassUp**: Shall be the modulation and coding used for uplink. The modulation class (as defined in clause 4.2, Standard Modulation and Coding) is `uint(modClassUp)`.
- ◆ **modClassDown**: Shall be the modulation and coding used for downlink. The modulation class (see clause 4.2, Standard Modulation and Coding) is `uint(modClassDown)`.
- ◆ **resource**: Shall indicate the conventional channel ID of the assigned resource. The assigned conventional channel ID = `uint(resource)`. The conventional channels shall be numbered based on the base carrier indicated in the BCH offset field of the CM burst, denoted f_0 . Table 75 - Mapping of Frequency and Timeslots to Conventional Channels defines the conventional channel numbering based on f_0 .

Table 75: Mapping of Frequency and Timeslots to Conventional Channels

Timeslot	Frequency				
	f_0	f_1	f_2	...	f_{N_f-1}
0	0	3	6	...	$3N_f-3$
1	1	4	7	...	$3N_f-2$
2	2	5	8	...	$3N_f-1$

- ◆ **rmuHeaderType**: Shall indicate whether the basic RMU header or minimized RMU header shall be used for the TCH bursts of the stream opened by the AA message. If `rmuHeaderType = 0`, the basic RMU header shall be used. Else, the minimized RMU header shall be used. The BS shall set the RMU header type field to 0 unless the BS has selected in the Reg Params registration message use of a protocol version that defines the minimized RMU header.
- ◆ **raIdType**: Shall indicate whether the ID field in the AA common field portion (see clause 7.3.5.1.2.1, Common message format and field definitions) is a RID, a PID, or an rreg extended ID. If `raIdType = 0`, the ID is a PID. If `raIdType = 1`, the ID is a RID. If `raIdType = 2`, the ID is an rreg extended ID. The setting `raIdType = 3` is currently reserved.
- ◆ **holdPage**: Shall indicate whether the BS may page a UT during the hold time defined by the `holdTimer` parameter. If this bit is set, the BS may page the UT during the hold time. Otherwise, the BS shall not page the UT during the hold time. If this bit is set, the BS should schedule the UT to be paged according to its (implementation-specific) page scheduling algorithm.
- ◆ **holdTimer**: Shall indicate the number of frames to not send an RA to the BS. For the case of an AA-reject following an RA-page-response, the restriction shall apply to the UT as a whole. In all other cases, the restriction shall apply to the registration that requested the stream. The encoding of this field shall be the same as that given in the `hold.cmd` UM message given in clause 7.4, Communication Elements.
- ◆ **ridOrPid**: Shall indicate whether the ID field in the AA common field portion (see clause 7.3.5.1.2.1, Common message format and field definitions) is a RID or PID. If `ridOrPid = 0`, the ID is a PID. Otherwise, it is a RID.
- ◆ **raDelay**: This value shall inform the UT to send an RA on the same conventional channel (the channel that received the AA) after a delay of `max[1 ; uint(raDelay)]` frames. Thus, if `uint(raDelay) = 0` or 1, the UT sends the RA on the next frame. In response to RA-rreg, the minvalue that the BS sets for `raDelay` shall equal or greater than 2. Additionally, the minimum value for `raDelay` shall be 2 if the RACH decode bit for the corresponding timeslot is set to 1 in

the UT Params message that is sent during registration (see clause 9.5.1.3.3, UT Params Message).

The modulation classes (modClassUp, modClassDown) shall be set by the

L2MacBsAaInfoGet.resp primitive from L3 RRC. The pwrAdjust field shall be set from the L2MacBsAaInfoGet.resp primitive from L3 RRC. toa shall be set by the timing error measurement in the L1RxDetect.ind primitive from L1.

7.3.5.1.2.4 RA to AA Mapping

Table 76 - RA Type to AA Subtype Correspondence shows the AA subtypes that may be sent on the downlink in response to each RA type.

Table 76: RA Type to AA Subtype Correspondence

AA subtype	RA type					
	RA-rts	RA-ping	RA-rts-short	RA-directed	RA-page-response	RA-rreg
AA-cts	Y	Y	Y	Y	Y	
AA-cts-short			Y		Y	
AA-rreg						Y
AA-reject	Y		Y	Y	Y	Y
AA-ping-ack		Y				
AA-cancel					Y	
AA-invalid-id	Y	Y	Y	Y	Y	
AA-pending	Y	Y	Y	Y	Y	Y

7.3.5.2 RACH Policy and Procedures

A UT shall send a different type of RA depending on stimulus from higher layers or the BS MAC as described in Table 77 - RA Type Sent as a Function of Stimulus.

Table 77: RA Type Sent as a Function of Stimulus

Stimulus	Source	RA type
Start registration	UT L3 RM	RA-rreg
Start TCH	UT L3 RM	RA-rts (no existing stream for this RID)
Start short TCH	UT L3 RM	RA-rts-short
Start ping	UT L3 RM	RA-ping
Start TCH (existing stream for this RID)	UT L3 RM	RA-rts or RA-rts-directed
Downlink page	BS L2 MAC	RA-page-response

If the BS L2 MAC receives an RA, the BS L2 MAC shall send an L2MacL3RmBsStart.ind indication to BS L3 RM. The BS L3 MAC sublayer within BS L3 RM chooses to grant or deny stream resources and shall inform BS L2 MAC of its decision with the L2MacL3RmBsStart.resp primitive. BS L2 MAC shall send an AA to the UT in the same frame. The AA-type and resource fields shall be set to reflect BS L3 MAC’s stream resource grant or denial. Similarly, if the UT receives an AA, it shall send an

L2MacUtStart.conf confirmation (for streams started by an UT L3 RM request) or L2MacL3RmUtStart.ind (for streams started by a page rather than from UT L3 RM) to UT L3 RM. The Start type shall be as given by Table 78 - UT Start Type as a Function of Stimulus and Table 79 - BS Start Type as a Function of Stimulus.

Table 78: UT Start Type as a Function of Stimulus

Stimulus	Start Type
AA-cts	TCH
AA-ping-ack	RID renewal
AA-cts-short	Short TCH
AA-rreg	Registration

Table 79: BS Start Type as a Function of Stimulus

Stimulus	Start Type
RA-rts	TCH
RA-rts-directed	TCH
RA-page-response	TCH
RA-ping (no page outstanding for any RID with this PID)	RID renewal
RA-ping (page outstanding for a RID with this PID)	TCH
RA-rts-short	Short TCH
RA-rreg	registration

7.3.5.2.1 Undirected RACH Retry Policy

The UT MAC RA retry strategy deals with when (in frame time) and where (which channel and timeslot) the UT retries an RA for a given RID. In the following, a RACH resource refers to a carrier and timeslot on which the UT may send an RA. There are N_RACH resources for sending RAs. These resources are the conventional channels of the RACH resource set. The RACH resource set is initially determined by the rachCarrierMask, the rachSlotMask, and the BCH offset fields received during the CR/CM exchange -- see clause 7.3.4.2, Configuration Message (CM). It is later overwritten using the reg params AM message received during registration.

Initially, N_RACH is determined as

$$N_{RACH} = A \times B - C,$$

where A equals the number of carriers in the RACH carrier set, B equals the number of timeslots allowed for sending RA's (rachSlotMask), and C equals the number of BS control channels whose physical carrier number is a member of the RACH carrier set and whose timeslot is allowed for RAs, as indicated by rachSlotMask. That is, C accounts for the fact that UTs are never allowed to send RAs on control channels.

In response to UT L3 RM's L2MacUtStart.req to UT L2 MAC, the UT shall use the following retry policy. The UT L2 MAC RA retry strategy is called an exponential sub-scan retry policy. The UT sends RAs to a fixed number of RACH resources, scanning a subset of the RACH resources. This is referred to

as a sub-scan. A UT shall exit the undirected RACH retry policy upon receiving an AA other than AA-pending. The UT shall wait a small, random amount of time, called an intra-sub-scan delay, between RA attempts within a sub-scan. If the UT does not exit the retry policy during a sub-scan, it shall wait for an inter-sub-scan delay before beginning another sub-scan. The inter-sub-scan delay shall grow exponentially with the number of failed sub-scans, to an upper limit. After a time limit from its first RA attempt, the UT shall exit the retry policy, and UT L2 MAC shall return an `L2MacUtStart.conf(fail)` to UT L3 RM.

The 6 parameters used in the retry policy shall have default values as specified in Table 80 - Parameters of the MAC Retry Strategy. These default values shall be replaced by the BS during registration in the `reg params` message (Section 9.5.1.3.4, Reg Params Message). The values for these 6 parameters may be replaced by the BS after registration by sending the `RACH retry params` message (see clause 9.5.1.3.15, RACH Retry Params Message).

Table 80: Parameters of the MAC Retry Strategy

Parameter	What it Parameterizes	Default Value
N_a	Number of resources per sub-scan	$\min[N_{RACH}; 3]$
N_b	Intra-sub-scan wait	5
N_c	Inter-sub-scan wait -- exponential growth (1 of 2)	4
N_d	Inter-sub-scan wait -- exponential growth (1 of 2)	3
N_e	Maximum backoff window	120
N_f	Maximum delay (in frames) before notifying UT L3 RM of the failure to open a stream (using <code>L2MacUtStart.conf(fail)</code>)	400

The intra-sub-scan wait shall be as given by Equation 7-2 in units of 5 millisecond frames. In Equation 7-2, ρ is the output of a pseudo-random number generator whose output lies between 0 and 1, but never takes on the value 0.0. The pseudo-random number generator is not specified by the *HC-SDMA* specification; it is implementation-specific. A new ρ is generated and W_{intra} is recalculated and used for every intra-sub-scan RA attempt within a sub-scan.

$$W_{intra} = \lceil N_b \times \rho \rceil \tag{7-2}$$

Note that it is possible in Equation 7-2 that $W_{intra} = 1$. In such cases, the UT may increment the intra-sub-scan wait by 1, resulting in $W_{intra} = 2$.

The inter-sub-scan wait shall be as given by Equation 7-3 and Equation 7-4 in units of 5 millisecond frames, where v_1 is the number of failed sub-scans and ρ is the output of a pseudo-random number generator whose output lies between 0 and 1. A new ρ is generated and used for every inter-sub-scan wait.

$$v_2 = \min[(N_c \times N_d^{v_1}); N_e] \tag{7-3}$$

$$W_{inter} = \left\lceil \frac{1}{2} \times (1 + \rho) \times v_2 \right\rceil \tag{7-4}$$

Note that it is possible in 7-4 that $W_{inter} = 1$. In such cases, the UT may increment the inter-sub-scan wait by 1, resulting in $W_{inter} = 2$.

W_{inter} is randomly recalculated before every inter-sub-scan RA attempt.

The UT's undirected RACH retry policy shall conform to Figure 26 - UT L2 MAC Undirected RACH Retry Policy for `L2MacUtStart.req`, wherein AFN equals the current frame number.

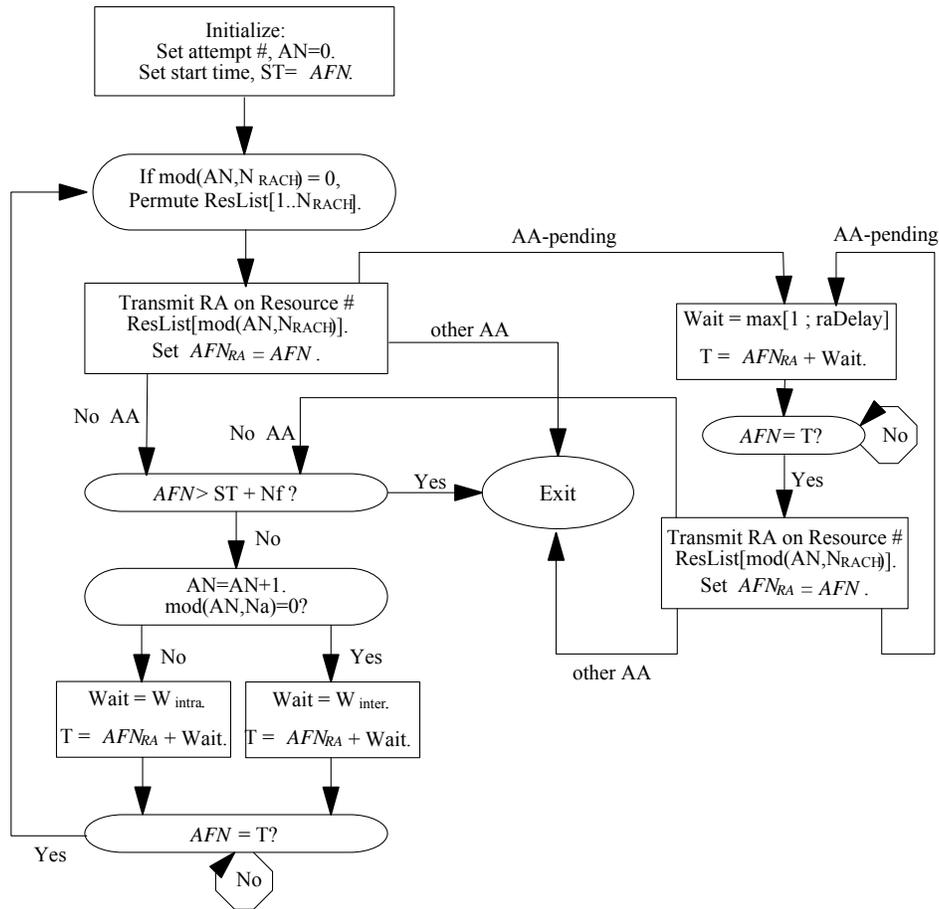


Figure 26: UT L2 MAC Undirected RACH Retry Policy for `L2MacUtStart.req`

In the figure, `ResList[]` is a randomly permuted list of RACH resources numbered (and indexed) 1 through N_{RACH} . As described in the figure, this list is randomly permuted every time the UT finishes a complete scan of all RACH resources. The righthand side of the figure illustrates the assignment of TCH using AA-pending. This is used when the BS wants the flexibility to assign a TCH resource other than the resource on which the RA is sent, but the BS needs additional time to make the assignment. As the figure indicates, an RA is sent after waiting $\max[1 ; raDelay]$ frames (see clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields for the definition of `raDelay`). The `raType` field of this RA shall be the same as the previously sent RA.

L3 RM does not have visibility into the L2 MAC RACH retry policy state diagram given above. From L3 RM's perspective, the L2 MAC RACH retry policy is one access request attempt. As shown in the

state diagram, an `L2MacUtStart.conf(fail)` result is sent to UT L3 RM if the retry policy finishes without UT L2 MAC receiving a correctly decoded AA corresponding to a sent RA.

7.3.5.2.2 Directed RACH retry policy

The UT directed RACH retry policy shall conform to Figure 27 - Directed RACH Retry Policy (for $X > 0$).

When a BS pages a UT or sends an aggregation command, the UT typically has an implied resource. There are 2 methods the BS may use to page the UT -- a PCH page and an in-band page. In the case of a PCH page, the implied resource is the carrier on which the UT heard the page, along with the timeslot associated with the page's subslot (see clause 7.3.3.2, Base Station Page Procedures). In the case of an aggregation or in-band page message, the message itself indicates whether it is a directed or undirected message, i.e., whether to use the directed RACH retry policy or the undirected RACH retry policy². When the aggregate or in-band page message is directed, the implied resource is indicated in the `aggregate.cmd` UM control message or the `in-band page` L3 AM message, respectively. For any of these forms of directed access, the conventional channel corresponding to the implied resource might not be a member of a RACH resource set. Even if the implied resource is not a member of a RACH resource set, the UT may send an RA-rtts-directed or RA-page-response on the implied resource.

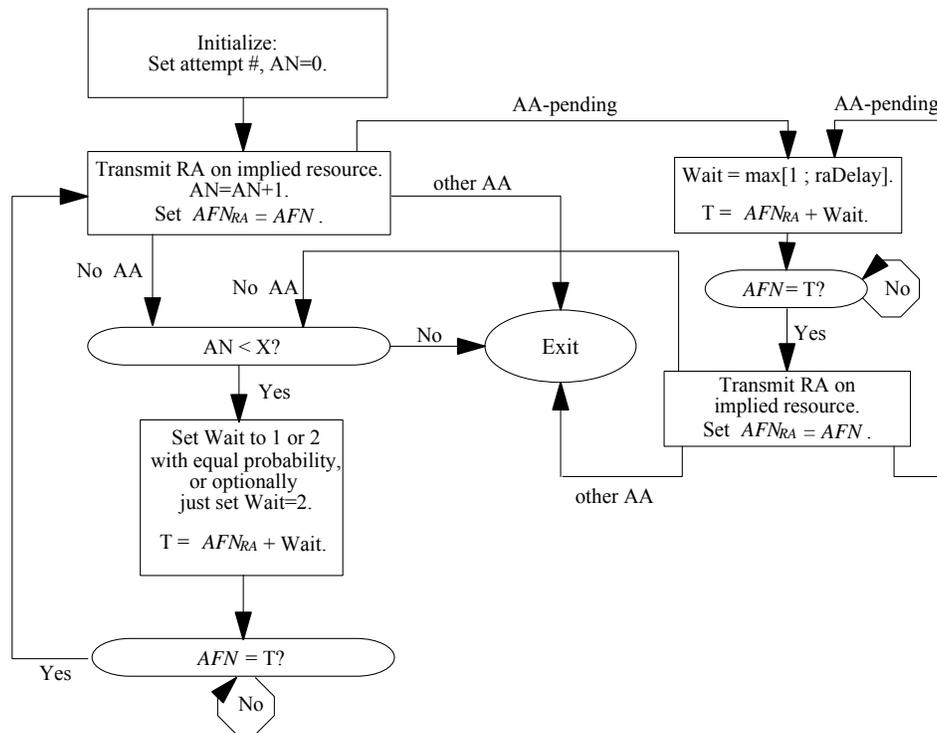


Figure 27: Directed RACH Retry Policy (for $X > 0$)

² See clause 9.5.2, Unacknowledged Mode (UM) Messages, and clause 9.5.1.3.7, In Band Page Message, for the interpretation of the resource field in the aggregate and in-band page messages, respectively. For an undirected aggregation or in-band page message, the UT uses the undirected RACH retry policy.

The RACH retry algorithm for responding to PCH pages, directed aggregation, and directed in-band page messages is called the directed RACH retry policy. It is parametrized by 2 parameters, X and Y. For a PCH page response, these are $X=N_{\text{priirp}}$ and $Y=\text{RA page rsrc}$ (taking values of 0 or 1). For a directed aggregation or directed in-band page message, these are $X=(N_{\text{rddrp}}+1)$ and $Y=0$. The parameters N_{priirp} , RA page rsrc, and N_{rddrp} are passed to the UT in the `reg params` AM message sent during registration (see clause 9.5.1.3.4, Reg Params Message). A new UT L3 RM instance trying to register with a BS shall use the default value 2 for N_{rddrp} when responding to an in-band page.

For the directed RACH retry policy, if $X=0$, the UT shall use the undirected RACH retry policy defined in clause 7.3.5.2.1, Undirected RACH Retry Policy. Specifically, if responding to a PCH page when $N_{\text{priirp}}=0$, the UT L2 MAC shall immediately send `L2MacUtPageRx.ind` to UT L3 RM. UT L3 RM then initiates the undirected RACH retry policy by issuing an `L2MacUtPing.req`. Similarly, if responding to an undirected aggregation or undirected in-band page message, UT L3 RM shall send `L2MacUtStart.req` to UT L2 MAC, configuring UT L2 MAC to send RA-rtts messages using the undirected RACH retry policy.

Conversely, if $X>0$, the UT shall begin to send RA's on the implied resource. If responding to a PCH page, the UT shall send RA-page-response. If responding to a directed aggregation or directed in-band page message, the UT shall send RA-rtts-directed.

For PCH pages, after X RA's have been sent unsuccessfully, UT L2 MAC shall send `L2MacUtPageRx.ind` to UT L3 RM. If $Y=0$, UT L3 RM halts stream initiation. If $Y=1$, then UT L3 RM shall fall back to the undirected RACH retry policy by issuing an `L2MacUtPing.req`.

For directed aggregation and directed in-band page messages, after X RA's have been sent unsuccessfully, UT L2 MAC shall stop sending RAs and shall send `L2MacUtStart.conf(fail)` to UT L3 RM. UT L3 RM shall then halt stream initiation.

7.3.6 Fast Associated Control Channel (FACCH)

The FACCH is used to carry power control and link adaptation information when physically transported on TCH bursts. FACCH is also transported on RA, AA and CM bursts. In these latter cases, FACCH carries information about the mode of operation. Specifically, it conveys the information whether the current mode of operation is "normal mode" (compliant with the protocol) or "vendor specific mode" (usually done for testing purposes).

7.3.6.1 FACCH on CM, RA, and AA Bursts

FACCH carries 4 bits of information, b3 (MSB), b2, b1, and b0 (LSB). The LSB (b0) is the allocated "capability bit" for CM, RA, and AA. b3 through b0 are reserved bits and their value shall be 0000.

7.3.6.2 FACCH on TCH Bursts

The FACCH carries power control and link adaptation information and is associated with either a standard uplink burst or a standard downlink burst. The information carried in the FACCH is a function of the Relative Frame Number (RFN) of the stream. See clause 7.3.5.1.2, AA field definitions for the definition of RFN. The uplink FACCH shall carry the current mod class (CMC) during the RFN frames labeled 1,3,5, etc. and shall carry the additional available power during the RFN frames labeled 2,4,6 etc. The downlink FACCH shall carry the recommended mod class (RMC) during the RFN frames labeled 1,3,5, etc. and shall carry the CMC during the RFN frames labeled 2,4,6, etc.

See Figure 13 – Standard Uplink Burst Structure and Figure 14 – Standard Downlink Burst Structure for more information on where the FACCH lies within the uplink and downlink bursts.

7.3.6.2.1 Modulation Class

HC-SDMA supports adaptive modulation. The physical layer supports different modulation/coding scheme pairs called modulation classes. The CMC and RMC are encoded in the FACCH. The FACCH, associated with the TCH burst, contains 4 bits (b_3, \dots, b_0) of information encoded as a Walsh-Hadamard code (see clause 4, Modulation and Channel Coding, for details), where b_0 is the LSB. The CMC or RMC of the payload of that burst is $8b_3 + 4b_2 + 2b_1 + b_0$. The mapping from the FACCH message to modclass shall be as described in Table 81.

Table 81: FACCH to Modclass Mapping

FACCH Message (b_3, b_2, b_1, b_0)	ModClass
0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	8
1001-1111	reserved

7.3.6.2.2 Additional Available Power

The additional power available to the UT shall be transmitted in the FACCH in the uplink during the even frames. The reported available power is assumed to be in dB and is a number between 0 and 15 dB. The reported available power is equal to $8b_3 + 4b_2 + 2b_1 + b_0$. The mapping from the FACCH message to available power shall be as described in Table 82.

Table 82: FACCH to Available Power Mapping

FACCH Message (b3,b2,b1,b0)	Available Power (dB)
0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	8
1001	9
1010	10
1011	11
1100	12
1101	13
1110	14
1111	≥ 15

7.3.7 Traffic Channel (TCH)

The TCH logical channel transports both UM (unacknowledged mode) and AM (acknowledged mode) control and data in an RMU (RLC message unit).

The L2 RLC sublayer is responsible for reliable delivery and packaging the AM portion of the RMU. The L2 MAC sublayer is responsible for the packaging of the entire RMU, which is primarily the multiplexing of UM and AM portions of the RMU.

7.3.7.1 Stream Start

A stream is initiated either by the UT or by the BS. The UT-initiated stream start sequence shall conform to Figure 28 - UT-initiated Stream Start Sequence. The BS-initiated stream start sequence shall conform to Figure 29 - BS-initiated Stream Start Sequence.

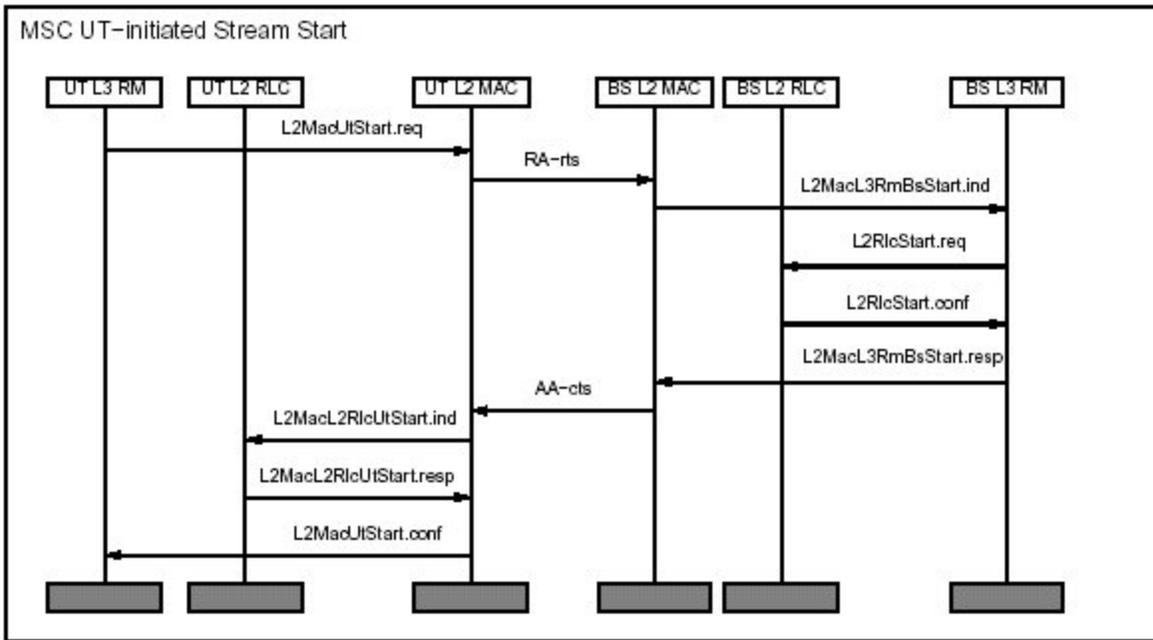


Figure 28: UT-initiated Stream Start Sequence

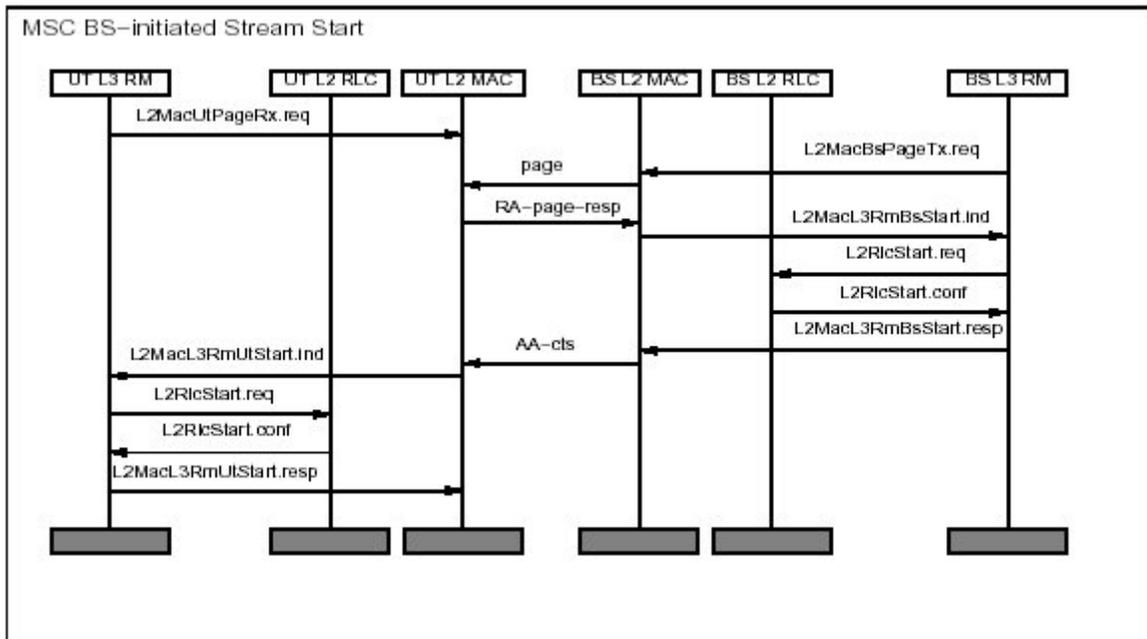


Figure 29: BS-initiated Stream Start Sequence

As shown in these figures, L2 RLC is initialized differently for the UT and BS, and the message sequence depends on whether the stream is UT-initiated or BS-initiated. For a UT-initiated stream start, on the UT, before UT L2 MAC issues an `L2MacUtStart.conf` confirmation to UT L3 RM to open a TCH stream, UT L2 MAC sends to UT L2 RLC an `L2MacL2RlcUtStart.ind` indication that instructs UT L2 RLC to initialize the ARQ state machine for this TCH stream (see clause 8.2.2, State Variables for

information). On the BS, BS L3 RM sends to BS L2 RLC an `L2RlcStart.req` request to initialize its ARQ state machine for this TCH stream.

For a BS-initiated stream start, on the UT, UT L2 MAC sends an `L2MacL3RmUtStart.ind` indication to UT L3 RM to open a TCH stream. UT L3 RM then sends an `L2RlcStart.req` request to UT L2 RLC to initialize the ARQ state machine for this TCH stream. On the BS, BS L3 RM sends to BS L2 RLC an `L2RlcStart.req` request to initialize its ARQ state machine for this TCH stream.

7.3.7.2 RMU Mux Function

7.3.7.2.1 RMU Field Definition

The structure of AM RMU shall conform to Figure 30 - AM RMU Definition. The gray portion is the RMU Header. The basic RMU header is depicted in Figure 30 - AM RMU Definition. The minimized RMU header may be used instead, at the discretion of the BS, if both the BS and UT support it. The header shall be present in every AM RMU. The white portion is the RMU Payload. The payload may be empty; that is, there might be no UM data and no AM data to transmit in a given RMU. The maximum size of the RMU depends on the modulation class the physical layer uses to transport RMUs from the RLC entity. The modulation class may change from frame to frame.

Note that the payload size can change when octets are retransmitted. For example, suppose that the AM octets of RMU 1 are lost, and suppose further that RMU 2 is to retransmit as many of these octets as possible. Since the UM portions of RMU 1 and RMU 2 can be different sizes, it is possible that the set of AM octets in RMU 1 will not fit into RMU 2.

7.3.7.2.2 Header Field Insertion

The RMU header fields shall conform to either the basic RMU header or the minimized RMU header shown in Table 83 - RMU Header Fields. The `pwrCtrl` bit in a downlink burst defines the relative adjustment to apply to transmit power on the uplink. The adjustment is 1 dB higher if the bit is 1 and 1 dB lower if the bit is 0. On an uplink burst, the `pwrCtrl` bit indicates the UT's measurement of SINR relative to the target SINR for a received downlink burst (which is a function of the modulation class of the downlink burst). The measured SINR is lower than the target if the bit is a 1 and higher if the bit is a 0. The `seqNum` and `ack` fields are used for the retransmission scheme as defined in clause 8.5.1, Transmit Procedure and clause 8.5.2 Receive Procedure.

Table 83: RMU Header Fields

Basic RMU Header Fields			
Field	# of Bits	Bit Positions	Interpretation
<code>pwrCtrl</code>	1	0	Power Control (downlink) / SINR (uplink) bit
<code>type</code>	2	1:2	00: payload is pure AM data 01: payload is mixture of AM and UM data 10, 11: reserved
<code>seqNum</code>	13	3:15	sequence number of the first AM octet in the payload. If the AM payload size is 0 octets, this value is ignored.
<code>ack</code>	12	16:27	RLC acknowledgment
Minimized RMU Header Fields			

Field	# of Bits	Bit Positions	Interpretation
pwrCtrl	1	0	Power Control (downlink) / SINR (uplink) bit
type	1	1	0: payload is pure AM data 1: payload is mixture of AM and UM data
arqFormat	1 or 3	2	Coded indication of the format of the ARQ seq num and ack fields. 0: 8 bit sequence number, ack field is not present. 001: 8bit sequence number, ack field is present. 111: 13 bit seq num, ack field is not present. 011: 13 bit seq num, ack field is present. 101: reserved
seqNum	8 or 13		Sequence number of the first AM octet in the payload. If the AM payload size is 0 octets, this value is ignored. If the arqFormat field is 0 or 001, this field contains the 8 LSBs of the full sequence number.
ack	0 or 12		RLC acknowledgment. This field is only present when the arqFormat is 001 or 011.
Total:	11, 18 , 25, or 30		

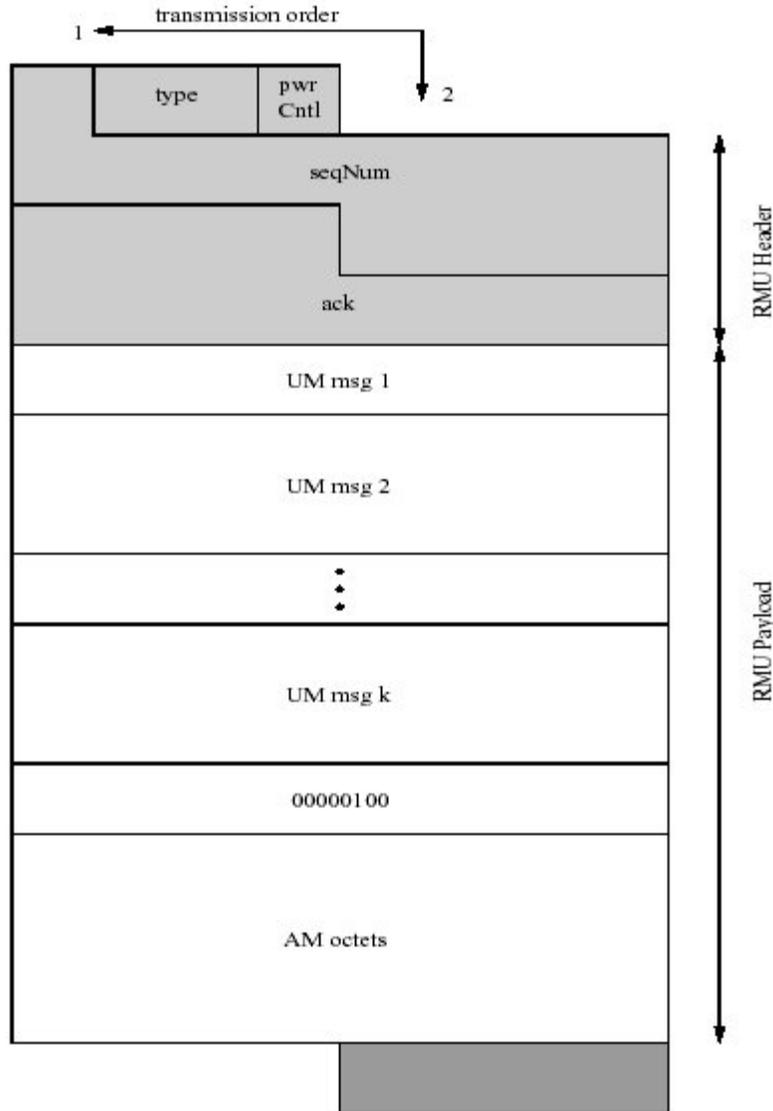


Figure 30: AM RMU Definition

During each frame, L2 MAC shall set the `pwrCtrl` bit with the value indicated by an `L2MacTchTxInfoGet.resp` primitive from L3 RRC. The `seqNum` and `ack` fields are filled by L2 RLC.

7.3.7.2.3 UM Message Insertion

The L2 MAC has two priority levels for AM data (see clause 8, L2 RLC Protocol Specification). The L2 MAC may prioritize UM messages above the high priority AM data, or below the high priority AM data and above the low priority AM data, or below the low priority AM data. These respective priority levels for UM messages will hereafter be referred to as high, medium, and low.

The method that L2 MAC uses to do this prioritization shall be as follows:

- ◆ Let N_{tot} be the total size of the RMU in bits (which is a function of modulation class) and N_{hdr}

the size of the RLC header in bits (see Figure 30 – AM RMU Definition). If the basic RMU header is used, then $N_{hdr} = 28$ (see Table 83 - RMU Header Fields). If the minimized RMU header is used, then L2 MAC must ask L2 RLC on each frame whether the RMU header will be 11, 18, 25, or 30 bits (see Table 83 - RMU Header Fields). The maximum number of bytes available for UM and AM messages is $N_{rmu} = \lfloor (N_{tot} - N_{hdr}) / 8 \rfloor$ bytes, where the floor function $\lfloor x \rfloor$ equals the largest integer less than or equal to x .

- ◆ Determine N_{high} , the total number of bytes of high priority UM messages that L2 MAC wishes to send on this frame, and N_{med} , the total number of bytes of medium priority UM messages that L2 MAC wishes to send on this frame.
- ◆ Compute N_{max} , the maximum AM data payload size available to RLC for this frame. If $N_{high} > 0$, $N_{max} = N_{rmu} - N_{high} - 1$; else, $N_{max} = N_{rmu}$.
- ◆ Compute $N_{reserve}$, the number of payload bytes L2 RLC must leave available for L2 MAC medium priority UM control data unless needed for high priority AM data (specifically, of the N_{max} payload bytes available for L2 RLC, high priority AM data consumes the unreserved payload bytes before consuming the reserved payload bytes). If $N_{high} = 0$ and $N_{med} > 0$, $N_{reserve} = N_{med} + 1$; else, $N_{reserve} = N_{med}$.
- ◆ Notify L2 RLC of N_{max} and $N_{reserve}$.
- ◆ L2 RLC notifies L2 MAC of the actual AM payload size that it will use in this frame, N_{actual} bytes.
- ◆ If $N_{actual} < N_{rmu}$ L2 MAC sets type=01 in the RMU header (mixture of AM and UM data). Otherwise L2 MAC sets type=00 in the RMU header (pure AM data).
- ◆ If type=01, L2 MAC fills in the first bytes of the burst payload with UM messages, beginning with high priority, then medium, then low, up to a maximum of $N_{rmu} - N_{actual} - 1$ bytes. Let N_{UMdata} be the total number of UM data bytes³ that L2 MAC uses, and define N_{gap} as the number of remaining bytes: $N_{gap} = N_{rmu} - N_{actual} - N_{UMdata}$. The remaining N_{gap} bytes are filled as described by the procedure below. Note that the UM messages and N_{gap} filler bytes occur before the AM data, as shown in Figure 30 – AM RMU Definition.

The manner in which the L2 MAC buffers and prioritizes UM messages is not specified by the HC-SDMA specification. High priority message should be sent immediately, as indicated in the above procedure. Low priority messages should each have a timer, and expiry of the timer should elevate the message to high priority.

The ‘fractional byte’ portion at the end of the RMU (the heavily shaded portion of Figure 30 – AM RMU Definition) shall be unused. Its size is $(N_{tot} - N_{hdr}) - 8 \lfloor (N_{tot} - N_{hdr}) / 8 \rfloor$. Of the remainder, the N_{actual} bytes of AM data payload fill contiguously the last part of the RMU (the last byte of the AM data payload is immediately ahead of the heavily shaded portion of Figure 30 – AM RMU Definition). When type = 00, there are no UM control message bytes -- $N_{UMdata} = 0$ and $N_{gap} = 0$. When type = 01, the N_{UMdata} UM control message bytes shall contiguously fill the RMU immediately after the L2 RLC header.

³ N_{UMdata} includes bytes used for UM control messages.

The N_{gap} unused bytes between the last UM control message byte and the first AM data byte shall be filled according to the following procedure. When $N_{gap} = 1$, the only unused byte shall be filled with the AM data tag value as given in Table 84 - UM Message Definitions. When $N_{gap} > 1$, a NOP structure shall be used. The NOP structure shall have 3 components: a NOP tag byte (using the NOP tag value as given in Table 84 - UM Message Definitions), an extensible length field, and a series of random bytes. The extensible length field shall be formatted as an arbitrary number of bytes in which the MSB of each byte indicates whether the length field is continued in the subsequent byte. If the MSB is set, then the subsequent byte's 7 LSBs shall be appended as the MSBs of the length. The length describes how many bytes there are after the first byte of the extensible length field until the start of the AM data bytes; specifically, the length shall equal $N_{gap} - 2$. Thus, the length includes the series of random bytes and all bytes of the length field except the first one. Note that when the NOP structure is used, the AM data begins immediately after the random bytes; there is no AM data tag following the random bytes.

For example, given a TCH burst with 221 payload bytes ($N_{rmu} = 221$), and given L2 MAC with no UM control messages of any priority ($N_{high} = N_{med} = 0$), and given L2 RLC with 40 bytes of AM data ($N_{actual} = 40$), then since $N_{gap} = 181$, the NOP structure is used. Since there is no UM control message data, payload byte 0 is the NOP tag byte. The extensible length field is comprised of 2 bytes, payload bytes 1 and 2. Using hexadecimal notation, payload byte 1 equals 0xb3 (MSB set, 7 LSBs are length LSBs -- 0110011), while payload byte 2 equals 0x01 (MSB cleared, 7 LSBs are length MSBs -- 0000001). Payload bytes 3 through 180 are random. Payload bytes 181 through 220 are the AM data bytes, where AM data byte 0 fills payload byte 181.

7.3.7.3 RMU Demux Function

The RMU Demux shall take the AM portion of the RMU and shall deliver it to the L2 RLC sublayer. Since there are UM messages that the L2 MAC protocol layer consumes, it shall parse the UM payload of the RMU and route UM messages destined for higher layers. The L2 MAC layer shall issue a `L2MacTchRxPwrCtrlInfo.ind` indication to L3 RRC when it receives a correctly decoded frame from L1. For the uplink, this is an SINR bit and possibly an `SinrErr.rpt` UM control message. For the downlink, it is a power control bit and possibly a `pwrCtrl.cmd` UM control message.

7.3.7.4 Stream Closure

There are several different reasons for a stream to close. L2 MAC sees all of them as two basic types of stream closure message sequences between UT and BS: *graceful* and *abrupt*. Following these stream closure sequences, the standard stream release sequence shall then be followed on the UT and BS. These message sequences are described below. Note that a stream closure does not imply that the registration also closes.

7.3.7.4.1 Stream Release Sequence

The stream release sequence shall follow Figure 31 - Stream Release Sequence. L2 MAC initiates the stream release sequence by sending an `L2MacL3RmStop.ind` indication to L3 RM. At this point the stream is released so that L2 MAC can now use the physical resources (for example, the conventional channel) of this stream for a new TCH stream. L3 RM then sends L2 RLC an `L2RlcStop.req` request. This request prompts L2 RLC to deliver any outstanding incoming data to L3 RM with the `L2RlcData.ind` indication, and to then inform L3 RM how much unsent or unacknowledged outgoing data is left via the `L2RlcStop.conf` confirmation.

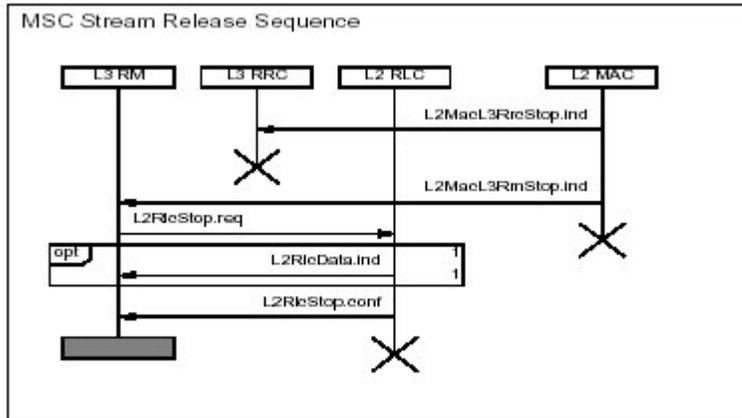


Figure 31: Stream Release Sequence

7.3.7.4.2 Graceful Stream Closure

The graceful stream closure sequence shall follow Figure 32 - Graceful Stream Close Sequence. The portion in the box labeled "opt" sometimes does not occur. For example, both sides might send all pending data, causing the graceful shutdown timers in both BS and UT L2 RLC to expire. The graceful shutdown time is the same as the $T_{shutDown}$ described in clause 8.2.2.3, Timers. When `L2MacTxDataInfoGet.resp` primitive indicates that L2 RLC is idle, BS L2 MAC sends `done.rpt` messages to UT L2 MAC repeatedly, every successive frame. BS L2 MAC will start the stream release sequence of clause 7.3.7.4.1, Stream Release Sequence as soon as either of two events occur: it receives a `done.rpt` message from the peer, or $N_{implicitAck}$ frames go by with neither a `done.rpt` nor a non-zero AM payload. UT L2 MAC will similarly start its stream release sequence using the same conditions.

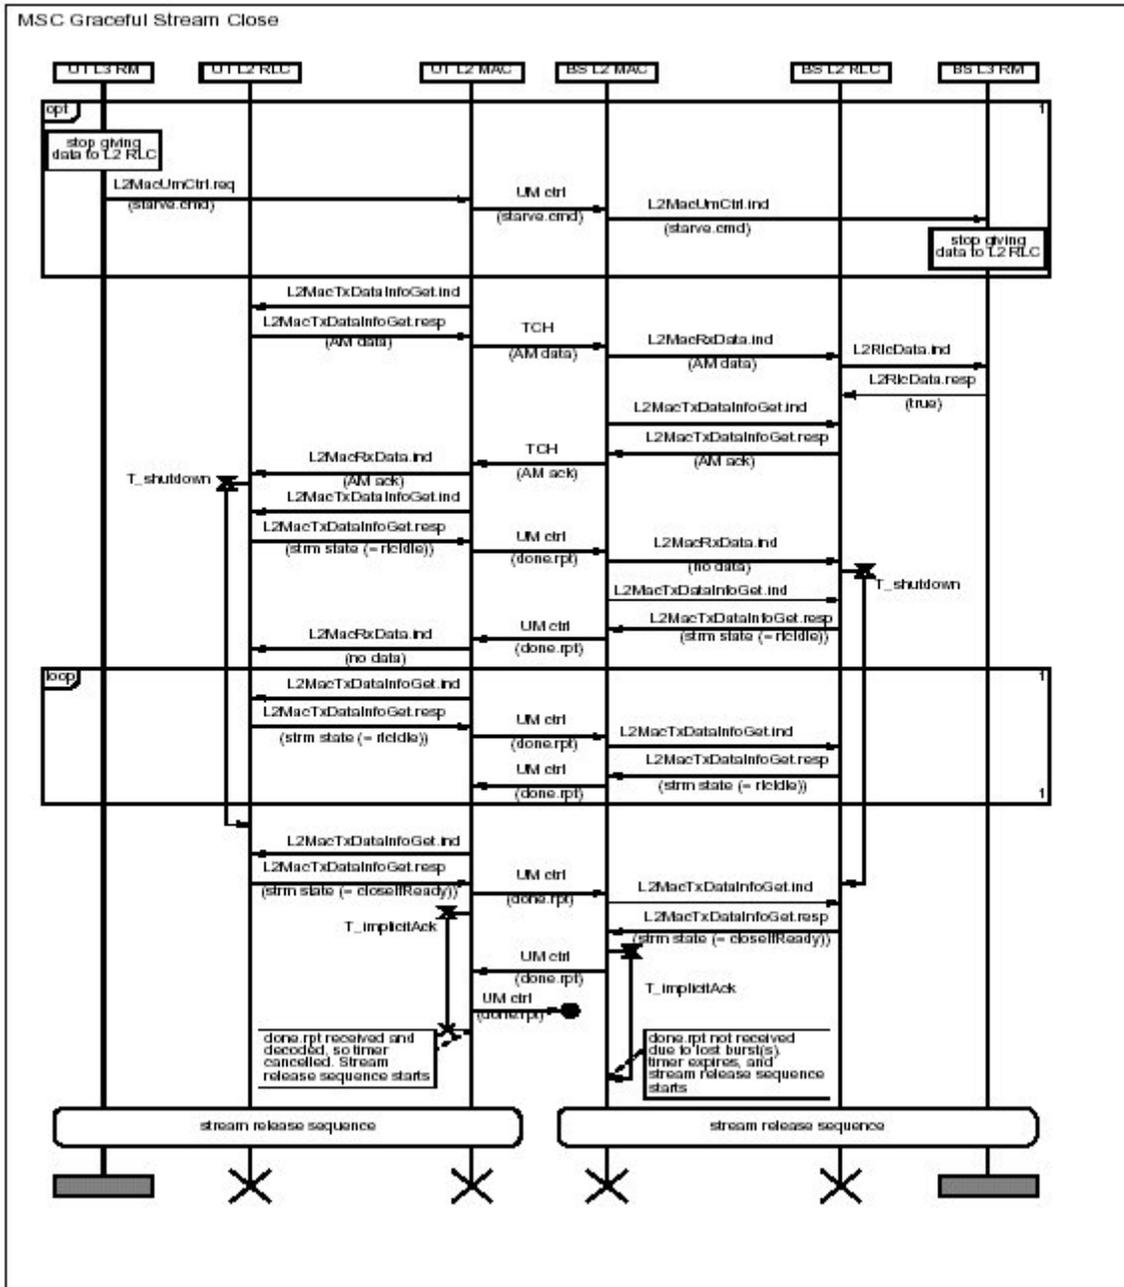


Figure 32: Graceful Stream Close Sequence

7.3.7.4.3 Abrupt stream Closure

The abrupt stream closure sequences shall conform to Figure 33 - Kill from L3 RM (`close.cmd` UM Received), Figure 34 - Kill from L3 RM (`close.cmd` UM Lost), and Figure 35 - Forced Shutdown from L2 RLC. L2 MAC performs an abrupt stream closure when it receives an `L2MacKill.req` request from L3 RM or from L2 RLC. In either case, L2 MAC sends a `close.cmd` message to the peer L2 MAC and immediately begins the stream release sequence. The `L2MacKill.req` can also come from the BS side, in which case the arrows are reversed.

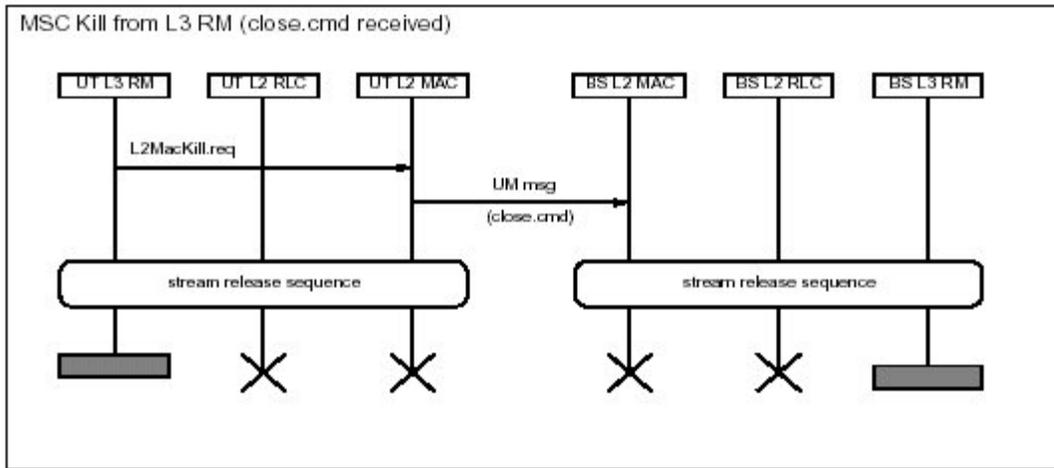


Figure 33: Kill from L3 RM (close.cmd UM Received)

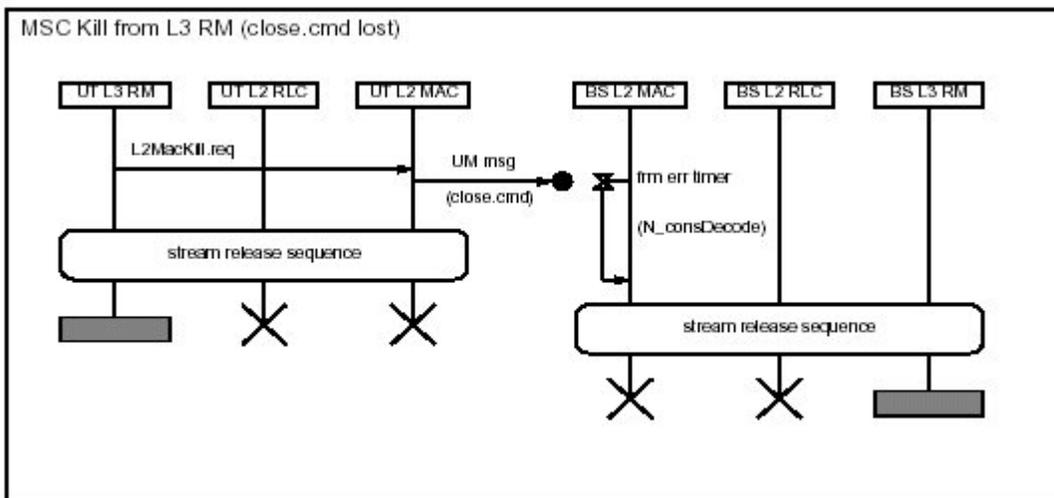


Figure 34: Kill from L3 RM (close.cmd UM Lost)

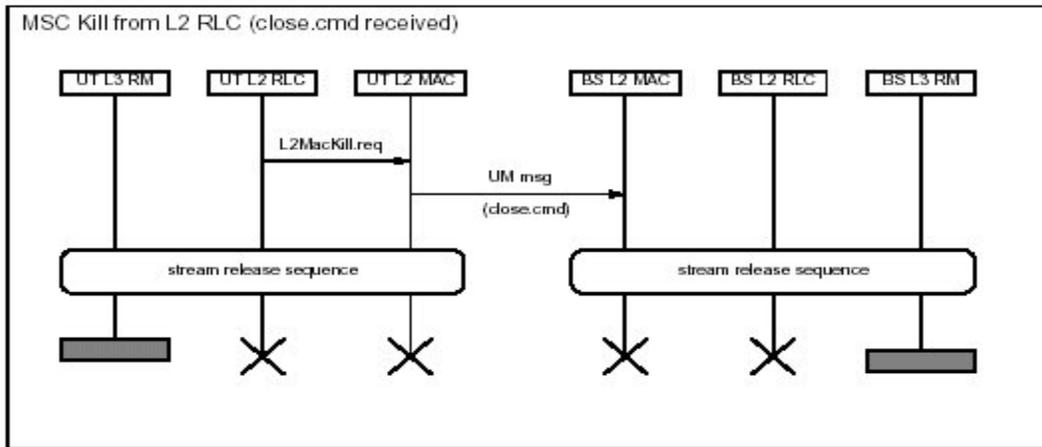


Figure 35: Forced Shutdown from L2 RLC

7.3.7.4.4 Stream closure parameters

This section lists the parameters and timers associated with stream closure, including parameters used by L2 RLC.

- ◆ $N_{consDecode}$ - The stream shall be closed after $N_{consDecode}$ consecutive RFNs of decode failures. $N_{consDecode}$ shall be set to 10.
- ◆ $N_{consDecodeInit}$ - The stream will close if the first $N_{consDecodeInit}$ consecutive RFNs of a stream fail decode. $N_{consDecodeInit}$ shall be set to 5.
- ◆ $T_{shutDown}$ - Used by L2 RLC. Specified in clause 8.2.2.3, Timers.
- ◆ $T_{transmit}$ - Used by L2 RLC. Specified in clause 8.2.2.3, Timers.
- ◆ $T_{receive}$ - Used by L2 RLC. Specified in clause 8.2.2.3, Timers.
- ◆ $N_{implicitAck}$ - The number of frames after the local L2 RLC becomes ready to close the stream (i.e., $T_{shutDown}$ has expired) when L2 MAC will conclude that the peer L2 RLC is also ready to close, in the absence of explicit information from the peer (e.g., the peer sends more AM data, or the peer sends a done.rpt message). The value for this timer shall be set to 1.
- ◆ $streamThreshUT$ - Timer used for security. See clause 11.4.2.6, Burst Counter Test for a description of how this timer is used.
- ◆ $streamThreshBS$ - Timer used for security. See clause 11.4.2.6, Burst Counter Test for a description of how this timer is used.

7.4 Communication Elements

The UM message headers and payloads shall conform to Table 84 - UM Message Definitions.

Table 84: UM Message Definitions

Header Value (MSB..LSB)	Direction	Label	Interpretation
L2 RLC Control Messages			
1xxxxxxx	both ^a	UM data	UM data of length xxxxxx octets follows.
00000100	both	AM data	AM data follows until end of payload.
L3 RM Control Messages			
00000011 001xxxxx	downlink	hold.cmd	Request for the UT L3 RM peer to not start any streams in the time interval specified by xxxxb after the last current stream closes.
0000011x	both	starve.cmd	Request the L3 RM peer to not send new SDUs on the stream that carries the starve.cmd message, but to finish sending outstanding SDUs. The field x states whether data transfer is being halted bidirectionally (x = 1) or unidirectionally (x = 0)
01xxxxxx	downlink ^a	aggregate.cmd	Request the UT L3 RM peer to open a stream on the resource identified by xxxxxx ^c .
010xxxxx	uplink ^a	queue.rpt	Queue depth state of UT (see 9.3.4.2.4, Stream Aggregation (Uplink Backlog) for details).
L3 RRC Control Messages			
001xxxxx	downlink	pwrCtrl.cmd	Command the UT L3 RRC peer to adjust uplink power. See 9.2.3 Radio Resource Control for details.
001xxxxx	uplink	sinrErr.rpt	Report from UT L3 RRC of the download signal quality. See 9.2.3 Radio Resource Control for details.
L2 MAC Control Messages			
00000001	both	close.cmd	Command L2 MAC peer to immediately shut down.
00000010	both	done.rpt	Request L2 MAC peer to gracefully shut down.
L3 MMC Control Messages			
00000011 01xxxxxx	downlink	timing.rpt	Time of arrival of corresponding uplink TCH relative to BS frame timing. $T_a = ((\text{uint}(\text{toa})/2 - 9) \mu\text{s})$.
Other			
00000000	both	NOP	Padding to extend RMU to a given size.
00000011	both	(extension)	This allows for the insertion of additional UM messages.
00000101	both available	Unused	
00001xxx	both available	Unused	
0001xxxx	both available	Unused	
011xxxxx	uplink available	Unused	

TABLE FOOTNOTES:

- a. Does not apply to registration streams.
- b. Define value = uint(xxxxx). Then value is mapped to a hold time as follows. If value is less than 25, then hold time = 8 * value frames. Otherwise, hold time = 192*(2^(val-ue - 24)) frames. See clause 9.5.2, Unacknowledged Mode (UM) Messages for detail on how the hold.cmd message and the hold time are interpreted.
- c. The data field is interpreted as an unsigned integer, indicating a conventional channel ID between 0 and 63. See clause 9.5.2, Unacknowledged Mode (UM) Messages, for detail on how this message is interpreted by UT L3 RM.

7.5 Page Hopping Sequences

7.5.1 Properties

The page hopping sequence for sending (from the BS) and listening (from the UT) to pages has the following properties:

- ◆ The BS knows when (frequency, time slots) the UT will be listening to pages during each frame.
- ◆ The users are spread evenly, i.e., the number of UTs listening to each resource should be approximately the same.
- ◆ The users are grouped differently on every frame. This means that, as far as possible, no pairs of users listen on the same resources for many consecutive frames.
- ◆ Each user scans through all the resources in a reasonable amount of time.
- ◆ The algorithm is easily adaptable to any number of resources.

7.5.2 Page Hop Function

The page hop function is used by the UT to determine the paging resource to scan for a page from the BS. The page hopping function shall be as defined by the following pseudocode:

```

c = pagehop(AFN, resource list, pid)
/* identify the resources in an array */
nChannels = length of resource list;

for j= 0,...,nChannels-1
    r[j] = jth resource on resource list;
end

Set nHop to be smallest prime larger than or equal to nChannels;
Set nDelta to be the next smaller prime, so nDelta is the largest prime less than
nHop;

/* mod(x,y), y > 0, is the modulo operator returning a value in the range
{0,...,y-1} */
init = mod(pid, nChannels);
/* [x] returns the smallest integer j greater than or equal to x */
ug = [(pid)/nChannels] /* ug is the user group */

t = mod(AFN, 1024);
block = [t/(nHop)] ;
index_init = mod(ug, nDelta);
delta = mod(index_init + block* [(ug)/(nDelta)] , nDelta) + 1;
user_pos = hopcore(block, init, floor(ug/nDelta), nChannels, nHop);

j = hopcore(t, user_pos, delta, nChannels, nHop);
c = r[j];

```

The hopcore function is as follows:

```

c = hopcore(t,r,delta,nR,nP)
/* generic core hopping function. The inverse of this function is
* implicitly applied by setting delta negative.
* Inputs:

```

```

* t - time index.
* r - the "row" index of the hopping function.
* delta - increment on each successive hop.
* nR - number of resources.
* nP - hop period. Must be a prime number.
* Outputs:
* c - the numerical value of the hop table. */

c = mod(r + t*delta, nP);

k = c - r;
while c >= nR
    c = c + k;
    if c >= nP
        c = c - nP;
    end
end

```

Appendix 7A. Setting bsLoad on BCH

The 2-bit bsLoad field shall be set as follows:

The BS maintains an estimate of its load, bsLoadEst, a number between 0 and 1. There are three parameters that define the mapping from bsLoadEst to the bsLoad field: bl_1 , bl_2 , and bl_3 . These parameters satisfy $0 < bl_1 < bl_2 < bl_3 < 1$ in particular, $bl_1 = 0.2$, $bl_2 = 0.4$, $bl_3 = 0.7$. The mapping is defined as in Table 85 - bsLoadEst to bsLoad Mapping.

Table 85: bsLoadEst to bsLoad Mapping

bsLoad value (MSB..LSB)	bsLoadEst
11	$Bl_3 \leq bsLoadEst \leq 1$
10	$bl_2 \leq bsLoadEst < bl_3$
01	$bl_1 \leq bsLoadEst < bl_2$
00	$0 \leq bsLoadEst < bl_1$

The bsLoadEst is updated once per BCH superframe (see clause 7.3.2.2, Superframe Structure). bsLoadEst is a time-varying estimate of the fraction of total uplink frame resources that are not idle, hereafter referred to as busy uplink frame resources.

A frame resource corresponds to a particular RF frequency/timeslot pair. For a particular frame resource, the BS knows the number of active TCH channels and can estimate whether there was any RA activity, and it knows whether that frame resource corresponds to the BCH/CCH resource. With this information, the BS estimates whether the resource was busy as follows. If this frame resource corresponds to the BCH/CCH resource, if the resource could not support an additional TCH, or if the BS estimates that there was at least one RA attempt, the BS declares that the frame resource was busy and sets its estimate to 1. Otherwise, the BS declares that the frame resource was idle and sets its estimate to 0.

Having computed this estimate for every frame resource in the superframe, the BS averages the estimates, yielding bsCurrentLoadEst. Finally, the BS updates bsLoadEst as follows:

$$bsLoadEst = df_{bs} * bsLoadEst + (1 - df_{bs}) * bsCurrentLoadEst. (7-5)$$

The discount factor parameter, df_{bs} , satisfies $0 \leq df_{bs} < 1$. In particular, we recommend $df_{bs} = 0.95$.

8 L2 RLC PROTOCOL SPECIFICATION

8.1 General

8.1.1 Definitions

8.1.1.1 Sequence Number Operands

The operations on sequence numbers are often done in modulo 2^S arithmetic, where S is the number of bits in the sequence number space. S is 13 bits in this section. Relational operands are defined as follows: A sequence number N is greater than $N - k$ for $k \in [1, 2^{s-1}]$ and less than $N + l$ for $l \in [1, 2^{s-1} - 1]$. Addition and subtraction are done in the usual modulo arithmetic. Special symbols as shown in Table 86 - Modulo Arithmetic Operands For Sequence Numbers will be used throughout this section for these modulo operations.

Table 86: Modulo Arithmetic Operands For Sequence Numbers

Symbol	Meaning
$=_m$	Equal, modulo 2^S
$<_m$	Less than as defined in 8.1.1.1 Sequence Number Operands.
$>_m$	Greater than as defined in 8.1.1.1 Sequence Number Operands.
\leq_m	Less than or equal to.
\geq_m	Greater than or equal to.
$+_m$	Addition, modulo 2^S .
$-_m$	Subtraction, modulo 2^S .
$:=_m$	$x :=_m y$ means $x =_m y$ and $0 \leq x < 2^S$.
$[k, l]_m$	Defined for $0 \leq k, l < 2^S$. If $k \leq l$, it is the usual interval $[k, l]$. Otherwise, it is $[k, 2^S - 1] \cup [0, l]$.

8.1.2 Sublayer Overview

The Radio Link Control (RLC) sublayer of the *HC-SDMA* Layer 2 Protocol provides a set of data link services for L3 peertopeer messages. Figure 36 - L2 RLC Overview Diagram gives an overview of the L2 RLC sublayer. The L2 RLC sublayer provides both Acknowledged Mode (AM) and Unacknowledged Mode (UM) methods of message delivery for both data and control message units from L3. Each L2 RLC entity combines the messages from these modes of delivery as directed by the L2 MAC sublayer and provides this RLC message unit (RMU) payload to the *HC-SDMA* L2 MAC sublayer.

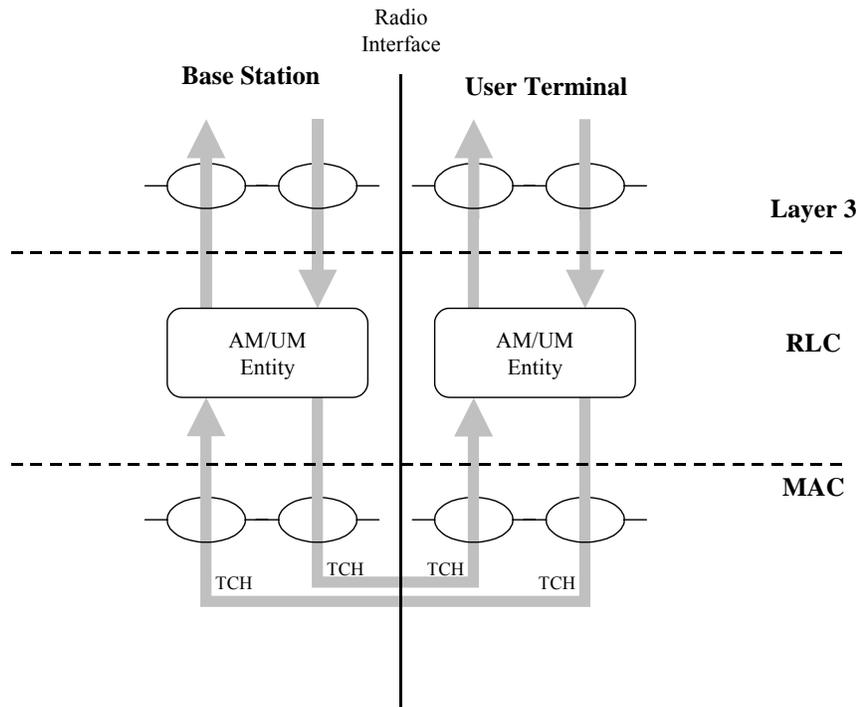


Figure 36: L2 RLC Overview Diagram

8.1.3 L2 RLC Model

HC-SDMA L2 RLC entities shall conform to the model in Figure 37 – L2 RLC Model. The left side of the vertical dashed line is the transmitting side and the right side of the line is the receiving side of the same L2 RLC entity (not the peer entity). The two sides will be referred to as the reverse entities of each other.

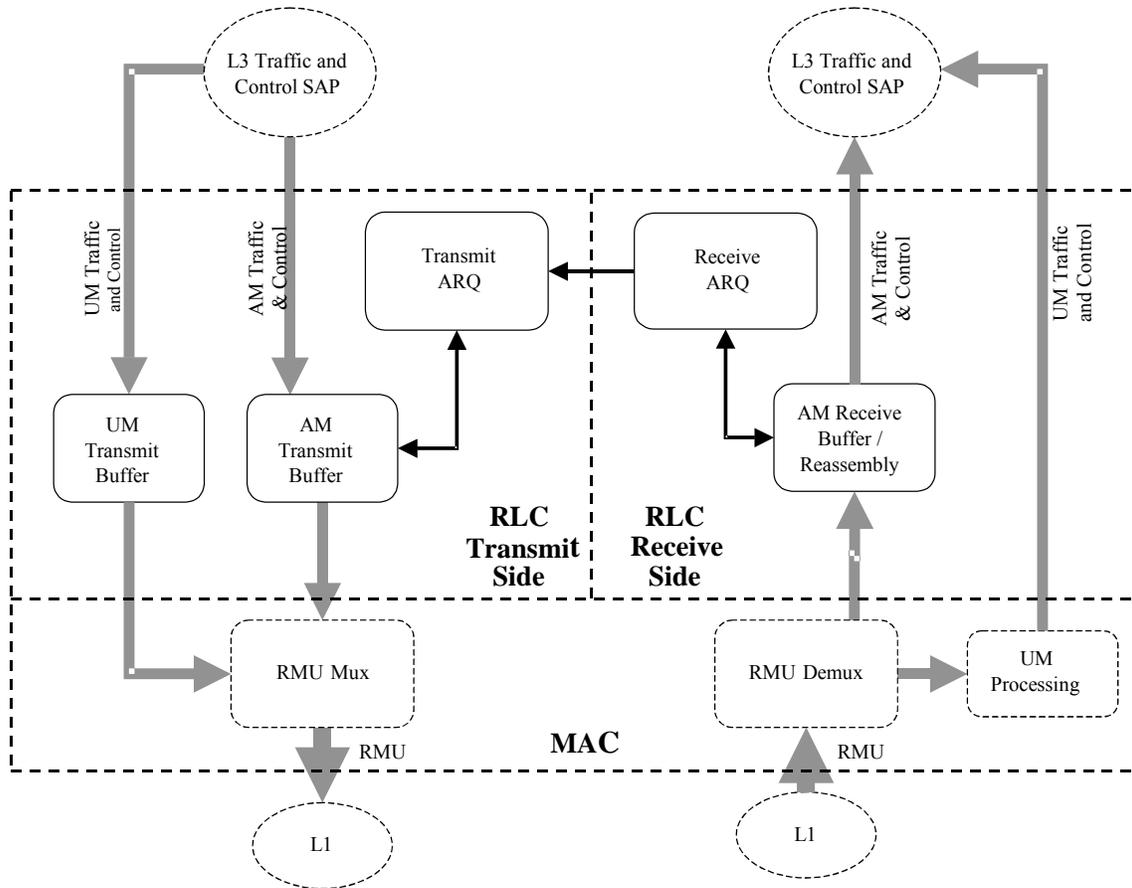


Figure 37: L2 RLC Model

8.1.3.1 AM Transmit Buffer

The AM Transmit Buffer shall receive L3 acknowledged mode SDUs and shall buffer them for transmission. Each octet shall be held in the buffer until it has been successfully acknowledged by the peer Receive entity (that is, a given octet in the buffer is possibly transmitted multiple times) or until the L2 RLC is shut down. Upon request by the transmit ARQ, a contiguous segment of octets in the AM Transmit Buffer shall be sent to the RMU Mux for transmission.

8.1.3.2 RMU Mux

The RMU Mux shall take AM data and control from the AM Transmit Buffer and UM data and control from the L2 MAC sublayer, and shall multiplex them into one RMU in the format described in clause 8.2.1.1, AM RMU.

8.1.3.3 Transmit ARQ

The transmit ARQ shall manage the transmission of new AM octets and the retransmission of octets known or suspected to be lost. High priority octets shall consist of new AM octets or AM octets known to be lost. Low priority octets shall consist of AM octets previously transmitted but not yet known to

be lost. In order to accommodate medium priority UM messages, L2 MAC shall inform L2 RLC of two values pertaining to the number of octets that can be sent on the current frame, n_{max} and $n_{reserve}$. See clause 8.5.1, Transmit Procedure for a description of how these values are used in constructing the RMU payload. The transmit ARQ may send less than the number of octets implied by these limits, for example, if it has no unsent or unacknowledged data queued for transmission.

8.1.3.4 RMU Demux

The RMU Demux shall parse the RMU.

8.1.3.5 Receive ARQ

The Receive ARQ manages the reliable reception of the AM data stream for the peer Transmit entity. The inbound Receive ARQ shall acknowledge transmissions from the peer Transmit ARQ by piggybacking acks (see clause 8.2.1.1, AM RMU) onto the outbound Transmit ARQ data stream. The inbound Receive ARQ shall also perform a dual role by forwarding piggybacked acks sent by the peer Transmit ARQ data stream to the outbound Transmit ARQ.

8.1.3.6 AM Receive Buffer/Reassembly

The AM Receive Buffer shall queue the AM octets which have been successfully received. Octets will arrive out of order due to lost RMUs. The data shall be buffered for ordered delivery to L3.

8.2 Peer-to-Peer Communication Elements

This section describes the elements used to enable communication between L2 RLC AM peer entities.

8.2.1 RMU Fields and Formats

8.2.1.1 AM RMU

The AM RMU shall send L3 SDUs and some control messages generated by L2 MAC using the L2 RLC entity described in clause 8.1.3, L2 RLC Model. Figure 30 - AM RMU Definition shows the format of the RMU payload.

8.2.1.1.1 RMU Header

The RMU header fields are defined in Table 83 - RMU Header Fields. For the basic RMU header and for the minimized RMU headers which include an explicit acknowledgement, the ack field shall contain the 12 least significant bits of a sequence number.

8.2.1.1.2 RMU Payload

The RMU payload contains a variable number of octets. The size of the UM portion of the RMU payload (and by inference the maximum number of AM octets that can be sent in a given RMU) shall be controlled by the L2 MAC sublayer. The first portion of the payload shall consist of zero or more UM tag/field pairs. The tag/field definitions are specified in Table 84 - UM Message Definitions. Certain messages, e.g., RMU control message `close.cmd`, consist only of a bit tag. The corresponding field size is then 0.

The second portion of the payload shall consist of exactly one contiguous segment from the AM octet stream preceded by the AM data tag (see Figure 30 – AM RMU Definition). If the L2 RMU payload contains no UM octets, the RMU type field shall be set to 00 and the AM data tag field shall be omitted.

8.2.2 State Variables

The octets transmitted in the AM portion of a AM RMU shall be numbered sequentially in the range 0 to $2^S - 1$. The octet after $2^S - 1$ shall be numbered 0.

8.2.2.1 Transmit State Variables

- ◆ **RFN (Relative Frame Number):** Shall be the frame counter used by the stream. (See 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.) The fractional rate stream grant defines how the Relative Frame Number (RFN) increments for the stream. The RFN shall be set equal to 0 during the RA/AA exchange frame. The RFN shall thereafter be incremented on the AFN frames granted for TCH bursts, as specified by the frameDec field of the AA message (note that frames punctured for BCH when a TCH burst would otherwise have been scheduled increment the RFN). Therefore for full-rate streams, the RFN advances at the same rate as the AFN. For half-rate streams, the RFN advances at half the rate as the AFN. The RFN is used by the FACCH (see clause 7.3.6, Fast Associated Control Channel (FACCH)), by the L2 RLC sublayer (see clause 8, L2 RLC Protocol Specification), and by encryption (see clause 11, Privacy and Authentication).

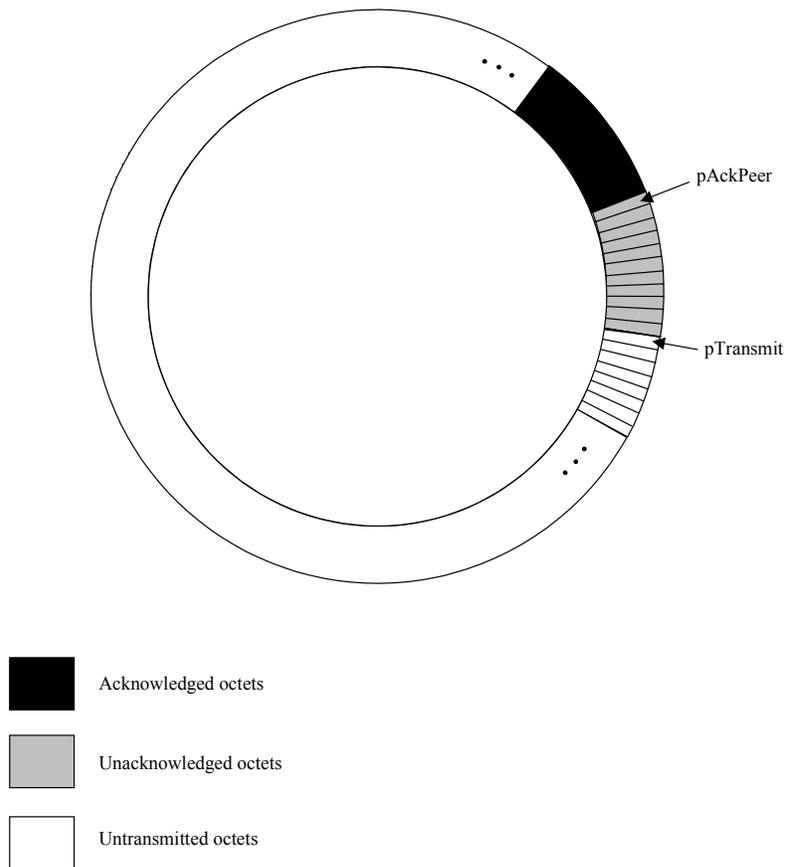


Figure 38: AM Transmit Sequence Numbers

- ◆ **pTransmit:** Shall be the sequence number of the first untransmitted octet. The octets preceding pTransmit shall have been transmitted at least once (see Figure 38: AM Transmit Sequence Numbers). The initial value of this variable shall be 0.
- ◆ **nTransmit:** Shall be the number of new octets still untransmitted. These shall be the octets after pTransmit. The initial value of this variable shall be 0.
- ◆ **pAckPeer:** Shall be the sequence number of the first unacknowledged octet. The Transmit ARQ knows that the octets preceding pAckPeer have been contiguously received by the peer Receive ARQ (see Figure 39 - AM Receive Sequence Numbers). The initial value of this variable shall be 0.
- ◆ **pRetransmit:** Shall be the sequence number of the first octet to be retransmitted in Retransmit state (defined in clause 8.5.1.1, Transmit State Execution). The initial value of this variable shall be 0.
- ◆ **nRetransmit:** Shall be the number of octets to be retransmitted in Retransmit state. The initial value of this variable shall be 0.
- ◆ **pSweep:** Shall be the sequence number of the first octet to be retransmitted in Sweep state (defined in clause 8.5.1.1, Transmit State Execution). The initial value of this variable shall be 0.

- ◆ **mapTx:** Shall be a vector indexed by sequence number. The values in this vector shall correspond to the RFN in which the octet was last sent. All entries shall be initialized to -1.

8.2.2.2 Receive State Variables

- ◆ **pReceive:** Shall be $pReceive = R_{last} +_m 1$, where R_{last} is the largest sequence number received; that is, octet R_{last} is received and if octet k is received and $k \geq_m pAck$, then $R_{last} \geq_m k$. The initial value of this variable shall be 0.
- ◆ **pAck:** pAck is defined recursively. Let p be the current value of pAck. After newly received octets are processed by the Receive ARQ, the octets in $[p, q]_m$ are all received, and octet $q +_m 1$ is not received. The new value of pAck shall be $q +_m 1$ (see Figure 39 - AM Receive Sequence Numbers). The initial value of this variable shall be 0.

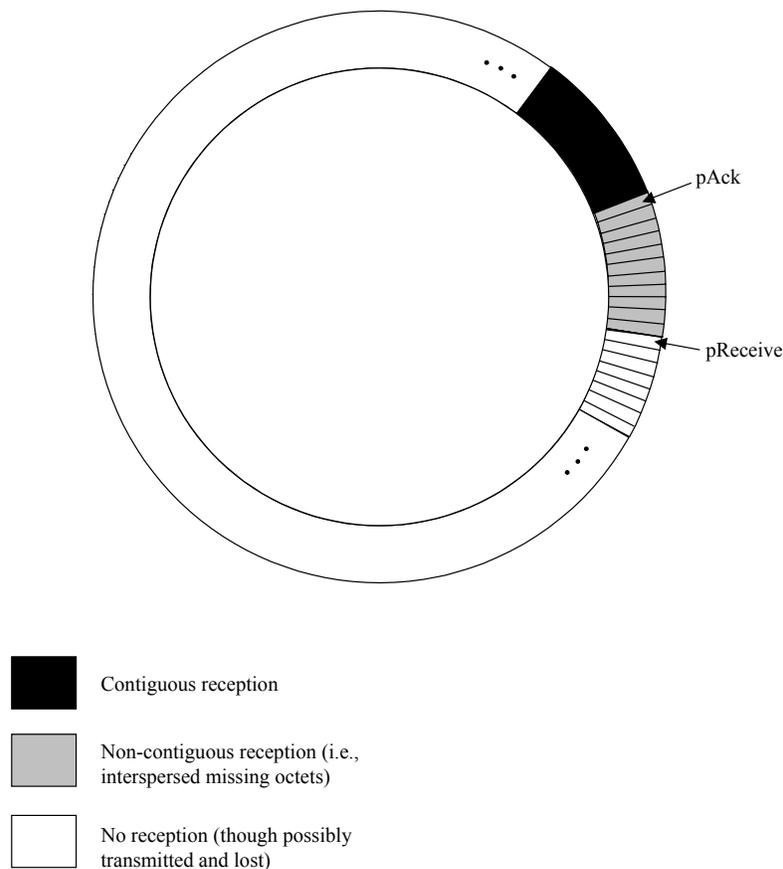


Figure 39: AM Receive Sequence Numbers

- ◆ **mapRx:** Shall be a vector of one bit for each sequence number. The 1s in this map shall identify the received octets that the Receive ARQ has not acknowledged (for example, because of non-contiguous reception). All entries shall be initialized to 0.

8.2.2.3 Timers

The following timers are defined for the L2 RLC entity. All of the timers below are in units of RFN (i.e., the timers increment when RFN changes).

- ◆ **N_{turnAround}**: For the UT, $N_{turnAround}$ shall be equal to the value of the TCH decode bit of the BS. For the BS, $N_{turnAround}$ depends on the value of the TCH decode bit (see clause 9.5.1.3.1, BS Params Message and clause 9.5.1.3.3, UT Params Message) of the UT and also whether the channel is full or fractional rate (see clause 7.3.5.1.2.3.1, Message Format for AA-cts and AA-short). It shall conform to the following table:

Table 87: Values of $N_{turnAround}$ for the BS

TCH decode bit (UT)	channel rate	
	full	fractional
0	1	1
1	2	1

For octets sent at $RFN = k$, if the L2 RLC entity receives an acknowledgement of these octets from the peer entity at $RFN \leq k + N_{turnAround}$, these octets shall not be retransmitted. Therefore, this value is the expected turnaround time (in units of RFN) for the peer L2 RLC entry to transmit an acknowledgement of the octets in question.

- ◆ **T_{shutDown}**: Shall be the timer for graceful shut down when transmit and receive buffers are empty. The timer shall start when the following conditions hold: $nTransmit = 0$, $pAckPeer = pTransmit$, $pAck = pReceive$, a complete packet has been transmitted, a complete packet has been received, $pAck$ has not changed for at least 1 frame, and the timer is not already running. If any of these state variables change while the timer is running but the condition still holds (for example, $pAckPeer$ and $pTransmit$ change by the same amount simultaneously), the timer shall be reset and continue running. If $nTransmit \neq 0$ or $pAckPeer \neq pTransmit$ or $pAck \neq pReceive$, the timer shall be reset and stopped. The configuration of $T_{shutdown}$ shall be done via the `L2RlcConfig.req` request by L3 RM.
- ◆ **T_{transmit}**: Shall be the timer for forced shutdown. A forced shutdown shall be done when there is no transmission progress. The timer shall be reset and stopped when $pAckPeer = pTransmit$. The timer shall run when $pAckPeer \neq pTransmit$. The timer shall be reset when $pAckPeer$ changes. Upon expiry, the L2 RLC shall send a `L2MacKill.req(L2RlcTxStuck)` primitive to L2 MAC. The value for this timer shall be 30.
- ◆ **T_{receive}**: Shall be the timer for forced shutdown. The timer shall be reset and stopped when $pAck = pReceive$. The timer shall run when $pAck \neq pReceive$. The timer shall be reset when $pAck$ changes. Upon expiry, the L2 RLC shall send a `L2MacKill.req(L2RlcRxStuck)` primitive to L2 MAC. The value for this timer shall be 30.

8.3 Error Handling

The list of error cases is shown below.

- ◆ E_{seqNum_OB} : Shall occur when at least one received octet is out of bounds in that the Receive ARQ can deduce that the Transmit ARQ is transmitting beyond its allowable set of sequence numbers (see clause 8.5.2.1, Receive Task Execution for a precise criterion). This shall cause `L2MacKill.req(L2RlcBadSeq)` to be sent to L2 MAC.
- ◆ E_{ack_OB} : Shall occur when Received ack is out of bounds (Received ack $>_m pTransmit$). This shall cause `L2MacKill.req(L2RlcBadAck)` to be sent to L2 MAC.
- ◆ E_{dup} : Shall occur when octet(s) with the same sequence number(s) have data that is inconsistent with previously received octet(s). The new inconsistent octets shall be ignored.

8.4 Layer-to-Layer Communication

This section describes the communication interface primitives of the L2 RLC sublayer to other layers/sublayers.

8.4.1 L3 RM

Refer to clause 10.3.2.13, L3 RM to L2 RLC Interface Primitives for a listing of the communication primitives from L3 RM to L2 RLC.

Refer to clause 10.3.2.12, L2 RLC to L3 RM Interface Primitives for a listing of the communication primitives from L2 RLC to L3 RM.

8.4.2 L2 MAC

Refer to clause 10.3.2.4, L2 MAC to L2 RLC Interface Primitives for a listing of the communication primitives from L2 MAC to L2 RLC.

Refer to clause 10.3.2.5, L2 RLC to L2 MAC Interface Primitives for a listing of the communication primitives from L2 RLC to L2 MAC.

8.5 L2 RLC Procedures and Process States

8.5.1 Transmit Procedure

During each frame, L2 MAC shall compute the maximum number of octets n_{max} available for transmitting high priority AM octets (see clause 8.1.3.3, Transmit ARQ), by considering the size of the transmit burst's RMU payload, how many bits are needed for the RMU header, and how many octets are needed for high priority UM messages. The number of bits needed for the RMU header is constant (28) if the basic RMU header is used, and variable if the minimized RMU header is used. Similarly, L2 MAC shall compute $n_{reserve}$ the number of octets desired for medium priority UM messages if L2 RLC has less than n_{max} high priority AM octets. If the number of high priority AM octets is less than $n_{max} - n_{reserve}$ L2 RLC may use the remaining space to transmit any low priority AM octets contiguous to the sequence of high priority AM octets. The Transmit ARQ shall also set `seqNum` and `ack` in the RMU header. If the basic RMU header is used, `seqNum` shall be set to the sequence number of the first octet written to the payload. `ack` shall be set to the $S - 1$ least significant bits of the value of `pAck`. If the minimized RMU header is used, L2 RLC shall determine whether to put the full sequence number or only the LSBs into the RMU header, and whether to include or exclude the ack. The L2 RLC shall set `arqFormat` to indicate whether the RMU header contains the full sequence number or the LSBs, and whether the ack is present or implicit (see Table 83 - RMU Header Fields).

If the minimized RMU header is used, the RMU header shall contain only the 8 LSBs of the sequence number if $pTransmit - m pAckPeer < 128$ and $pSweep = pAckPeer$. Otherwise, the full sequence number shall be used.

If the minimized RMU header is used, then the RMU header shall not include the ack if the following three conditions are met:

1. $pAck = m pReceive$;
2. The last received RMU was decoded successfully; and
3. The AM data length of the last received RMU payload is zero, or the sequence number of the last byte of the last received RMU payload = $pAck - m - 1$.

Otherwise, the ack shall be included in the RMU header.

The ack from the peer receive entity shall be processed and the transmit state variables shall be updated as described in clause 8.5.2.1, Receive Task Execution prior to the execution of the transmit states in 8.5.1.1, Transmit State Execution below.

8.5.1.1 Transmit State Execution

The execution of the Transmit ARQ shall follow the following steps:

- ◆ If $nRetransmit > 0$ (**Retransmit state**)
 1. Set $n := \min(nRetransmit, n_{max} - n_{reserve})$.
 2. Send the segment of octets in $[pRetransmit, pRetransmit + m (n - 1)]_m$.
 3. Update mapTx as described in clause 8.5.1.2, mapTx Update Procedure.
 4. Set $pRetransmit := m pRetransmit + n$.
 5. Set $nRetransmit := nRetransmit - n$.
- ◆ Else if $nTransmit > 0$ and $pTransmit - m pAckPeer < 2^{S-1} - 1$ (Transmit state).
 1. Set $n := \min(nTransmit, n_{max} - n_{reserve}, 2^{S-1} - (pTransmit - m pAckPeer) - 1)$.
 2. Send the segment of octets in $[pTransmit, pTransmit + m (n - 1)]_m$.
 3. Update mapTx as described in clause 8.5.1.2, mapTx Update Procedure.
 4. Set $pTransmit := m pTransmit + n$.
 5. Set $nTransmit := nTransmit - n$.
 6. Cancel timer Tshutdown if it is running.
- ◆ Else if $pAckPeer \neq pTransmit$ (Sweep state)
 1. If $pSweep < m pAckPeer$, set $pSweep := pAckPeer$.
 2. If $pSweep \geq m pTransmit$, set $pSweep := pAckPeer$.
 3. Set $n := \min(pTransmit - m pSweep, n_{max} - n_{reserve})$.
 4. Send the segment of octets in $[pSweep, pSweep + m (n - 1)]_m$.

5. If the basic RMU header is used, update mapTx as described in clause 8.5.1.2 mapTx Update Procedure.
 6. Set $pSweep := m pSweep + n$.
- ◆ Else nothing to transmit (Idle State)

8.5.1.2 mapTx Update Procedure

The following procedure shall be used to update the mapTx state variable introduced in clause 8.2.2.1, Transmit State Variables.

- ◆ **Update mapTx:** In any state, for every segment of octets transmitted at *RFN* whose sequence numbers are in $[p, q]_m$, set $mapTx[k] := RFN$ for $k \in [p, q]_m$.

8.5.2 Receive Procedure

8.5.2.1 Receive Task Execution

During each frame, the RMU Demux shall deliver to the Receive ARQ the AM payload (possibly 0 octets) as well as the RMU header fields seqNum and ack unless there is a CRC error. (In case of a CRC error, none of the above shall be delivered to the AM unit, and the following tasks are not executed.)

If the minimized RMU header is used, then the receiver shall check the arqFormat RMU header field to determine whether the sequence number and ack fields need to be expanded to full length. If arqFormat indicates that the RMU header contains only the 8 LSBs of the sequence number, then the receiver shall recreate the 13 bit sequence number with the following formulas. In these formulas, the & operator indicates a bit-wise AND operation, and the ~ operator indicates bit-wise negation.

```

if seqNum8bits >= ((pAck -m 128) & 0xff)
    seqNum13bits = ((pAck -m 128) & ~0xff) + seqNum8bits
else
    seqNum13bits = ((pAck -m 128) & ~0xff) + 0x100 + seqNum8bits
    
```

If arqFormat indicates that the RMU header contains no ack, then the receiver shall recreate the ack through the following steps:

1. Set $AckRFN = RFN - N_{turnaround}$ where RFN is the RFN in which the burst was received.
2. Find the range of mapTx elements = AckRFN
3. Set ack to the 12 LSBs of the sequence number of the next mapTx element beyond the range identified in step 2.

The receiver shall perform the following three tasks.

- ◆ **Task 1: Receive the data in the AM payload**
1. Set $first := seqNum$.
 2. Set $last :=_m seqNum + \text{size of payload}$.

3. If $last \geq_m pAck + 2^{S-1}$ and $first \geq_m pAck$.
 - i) Recommend forced shut down ($E_{seqNum-OB}$ error).
 - ii) Go to task 2.
4. If $last \geq_m pAck$, ignore the AM octets in this RMU and go to task 2.
5. Store data in appropriate location.
6. Cancel timer $T_{shutdown}$ if it is running.
7. If $last >_m pReceive$, set $pReceive :=_m last$.
8. If $first \leq_m pAck$,
 - i) Find p ($0 \leq p < 2^S$), the first sequence number that satisfies $p \geq_m last$ and $mapRx[p] = 0$ (that is, there is not a $q <_m p$ such that $q \geq_m last$ and $mapRx[q] = 0$).
 - ii) Set $mapRx[k] = 0$ for $k \in [pAck, p - m1]_m$.
 - iii) Deliver and record octets in $[pAck, p - m1]_m$ to L3.
 - iv) Set $pAck := p$.
 - v) Cancel timer $T_{receive}$.
 - vi) If $pAck \neq pReceive$, restart timer $T_{receive}$.

Else (that is, $first >_m pAck$) Set $mapRx[k] = 1$ for $k \in [first, last - m1]_m$.

◆ **Task 2: Process ack in RMU header (for transmit side)**

1. ack is $S - 1$ bits only. Convert it to S bits:
If $ack <_m pAckPeer$, set $ack :=_m ack + 2^{S-1}$.
2. If $ack >_m pTransmit$
 - i) Recommend forced shut down (error E_{ack-OB}).
 - ii) Go to task 3.
3. If $ack >_m pAckPeer$
 - i) Set $mapTx[k] = -1$ for $k \in [pAckPeer, ack - m1]_m$.
 - ii) If $pSweep <_m ack$ Set $pSweep :=_m ack$
 - iii) Set $pAckPeer :=_m ack$.
 - iv) Cancel timer $T_{transmit}$.
 - v) If $pAckPeer \neq pTransmit$, restart timer $T_{transmit}$
 - vi) If $nRetransmit > 0$ and $pRetransmit <_m ack$
Perform a mapTx retransmit update as described in clause 8.5.2.2, mapTx Retransmit Update.
4. If $RFN - mapTx[pAckPeer] \geq N_{turnaround}$ (here, RFN is when the relevant octets were transmitted on the air interface)
Perform a mapTx retransmit update as outlined in clause 8.5.2.2, mapTx Retransmit Update.

◆ **Task 3: Check if stream is ready to end**

If timer $T_{shutDown}$ is not running

If $nTransmit = 0$, $pAckPeer = pTransmit$, $pAck = pReceive$, a complete packet has been transmitted, a complete packet has been received, and $pAck$ has not changed for at least 1 frame, start timer

$T_{shutDown}$

8.5.2.2 mapTx Retransmit Update

The following procedure shall be used to update $nRetransmit$:

1. Set $pRetransmit = pAckPeer$

2. If $mapTx[pAckPeer] < 0$

Set $nRetransmit = 0$

Else

Set $nRetransmit = k -_m pAckPeer$ where k is the first sequence number such that $k >_m pAckPeer$ and $mapTx[k] \neq mapTx[pAckPeer]$ (that is, there is not an $l <_m k$ such that $l >_m pAckPeer$ and $mapTx[l] \neq mapTx[pAckPeer]$)

8.5.2.3 Reset Procedure

When the $L2RlcStart.req$ (see clause 10.3.2.13, L3 RM to L2 RLC Interface Primitives) or $L2MacL2RlcUtStart.ind$ (see clause 10.3.2.4, L2 MAC to L2 RLC Interface Primitives) primitive is received, all the state variables in the Transmit and Receive ARQs shall be set to their initial values as indicated in clause 8.2.2, State Variables.

8.5.3 Shutdown Procedure

There are two ways that the L2 RLC entity can shut down. One method is a *forced shutdown* and the other method is a *graceful shutdown*.

8.5.3.1 Forced Shutdown

A forced shutdown initiated by a L2 RLC error condition shall begin with sending the $L2MacKill.req$ request (see clause 10.3.2.11, L3 RM to L2 MAC Interface Primitives). After L2 MAC closes the stream, it shall inform L3 RM, which shall close L2 RLC with the $L2RlcStop.req$ request (see clause 10.3.2.13, L3 RM to L2 RLC Interface Primitives). As the L2 RLC entity completes the termination, it shall send a $L2RlcStop.conf$ confirmation (see clause 10.3.2.12, L2 RLC to L3 RM Interface Primitives).

8.5.3.2 Graceful Shutdown

A graceful shutdown is negotiated at other layers (for example, L2 MAC and L3 RM), and it is implicit from the point of view of the L2 RLC sublayer. The timer $T_{shutdown}$ shall indicate how long the L2 RLC entity has not been sending and receiving any data, and when this timer expires, the L2 RLC shall inform the L2 MAC sublayer that it is ready to close the stream. After L2 MAC ascertains that the peer is also ready, it shall inform L3 RM, which shall close L2 RLC with the $L2RlcStop.req$ request (see clause 10.3.2.13, L3 RM to L2 RLC Interface Primitives). As the L2 RLC entity completes the

termination, it shall send a L2RlcStop.conf confirmation (see clause 10.3.2.12, L2 RLC to L3 RM Interface Primitives).

9 L3 PROTOCOL SPECIFICATION

9.1 Overview

9.1.1 Introduction

HC-SDMA Layer 3 is responsible for presenting and managing a logical connection between the user terminal (UT) and base station (BS) across the air interface.

L3 is responsible for the following functionality:

- ◆ Creating logical sessions to transfer the higher layer's L4 data across the air interface.
- ◆ Giving logical sessions the ability to open physical streams by registering the session.
- ◆ Converting L4 data packets to forms suitable for transport across the air interface (HC-SDMA packets), and undoing the conversion on the remote side.
- ◆ Determining when to open streams for data transport.
- ◆ Performing power control and modulation class selection to keep streams operating efficiently.
- ◆ Handing over a session from one BS to another.

The L3 protocol layer is made up of components with distinct roles in supporting a connection across the air interface. This section fully describes the following components:

- ◆ Connection management (CM)
- ◆ Registration management (RM)
- ◆ Radio resource control (RRC)
- ◆ Mobility management and control (MMC)

The management modules (CM, RM, and MMC) involve logical functionality such as state machines and resource usage maps. The control modules (RRC and MMC) involve controlling physical quantities such as power levels or timing using measurements as inputs to control loops.

9.1.2 General Protocol Concepts

The duration of an end-to-end transport of end-user data between an *HC-SDMA* network and a particular UT is referred to as a session or a logical connection. A session is always initiated by a UT, typically at the request of some higher layer protocol such as PPP running on the EUD. Within the *HC-SDMA* protocol, the initiation of a session invokes the registration process. A registration is the air interface association between a BS and a UT that permits UT-specific information including certain types of control data and end-user data to flow between those devices. Finally, a stream is the association between a BS and a UT that allows end-user data to flow between those devices as well as multiple classes of control data⁴. Streams are generally initiated and destroyed multiple times in the course of a session. Streams may be aggregated to increase throughput for a particular data flow. In

⁴ The remaining classes of control data that flow between those devices within the context of a session are handled via paging and random access channels that will be described subsequently.

summary, end-user data is transported in streams, created as necessary in the context of a registration to support a session.

HC-SDMA is designed for efficient operation with spatial processing technology. Downlink spatial processing works best when there is recent uplink information from the target UT. For this reason, downlink bursts immediately follow uplink bursts on the same physical resource. The broadcast and paging channels are exceptions to this rule because the messages they bear must be available throughout the cell. Downlink spatial processing provides only diversity gain rather than combining gain for these channels; to compensate, their payloads are heavily coded and their information rate is low.

9.1.3 Logical Channels and Typical Ladder Diagram

HC-SDMA provides the logical channels listed in Table 88 - HC-SDMA Logical Channels.

Table 88: HC-SDMA Logical Channels

Channel	Direction	Spatial Processing	Function
Broadcast Channel (BCH)	Downlink	Non-directive	Synchronization, pathloss, and handover measurements, BS color code announcement.
Paging Channel (PCH)	Downlink	Non-directive	Paging to initiate stream for downlink data.
Configuration Channel (CCH)	Bidirectional	Directive	Exchange of basic BS and UT parameters necessary to initiate registration.
Random Access Channel (RACH)	Bidirectional	Directive	Stream request and setup, power control and modulation selection.
Traffic Channel (TCH)	Bidirectional	Directive	Exchange of acknowledged and unacknowledged mode data, including user data and control information for modulation selection, power control, exchange of security parameters, stream aggregation.

Over the lifetime of a session, a typical sequence of use for these channels is shown below.

◆ **Pre-Registration**

- UT acquires BCH, obtains base station color code (BSCC), and synchronizes time and frequency.
- UT sends configuration request on CCH including UT capabilities information; BS responds on CCH with site information including the set of resources on which the UT can send requests (RACH resource set) and power control setting.
- UT issues registration request on RACH; BS responds on RACH with registration grant including power control settings. The RACH sent from the BS confirms the UT's request for stream the initiation. This stream is termed as the registration stream, over which information messages are passed to complete the registration process.
- UT sends a UT capability message including the version of the protocol that the sending party conforms to, and the uplink and downlink modclass capabilities of the UT.

- BS sends a detailed BS parameters message on TCH including security information and absolute frame number. UT responds with a detailed UT parameters message on TCH including hardware identifier, and security information.
- BS sends a registration parameter message on TCH including paging and registration identifiers, and random access guidelines.

(Registration is now complete.)

◆ **Pre-Session**

- UT requests TCH resources on RACH. BS responds on RACH with a TCH resource grant (stream grant).
- UT sends session start message on TCH.

(Session is now established.)

◆ **Post Registration, Post Session Establishment, Pre-Stream**

- BS pages UT on PCH in response to traffic destined for UT from network. This step is skipped when the UT wishes to acquire uplink traffic resources in response to traffic from the EUD destined for the network.
- UT requests TCH resources on RACH. BS responds on RACH with a TCH resource grant.

(Stream is now established.)

- ◆ End-user data and possibly control traffic is exchanged on TCH.
- ◆ Additional streams are potentially aggregated with the initial one. Streams are aggregated following an aggregation message passed as in-band control information on the downlink TCH. While the aggregation messages themselves are always sent by the BS, they are generated in response to both uplink and downlink traffic throughput requirements. A BS/UT RACH exchange, similar to the one that opened the original stream, follows to set up a new TCH resource. Multiplexing of data onto and off of streams is managed at Layer 3 of the protocol.
- ◆ Streams are terminated either by in-band messaging in TCH and/or through timeouts at the MAC layer of the protocol.

(Stream is now terminated.)

Additional streams are created and terminated as necessary.

9.2 Layer 3 Components

Figure 40 - L3 Components and Interactions with Other Layers illustrates how L3 components interact with each other and with other protocol layers.

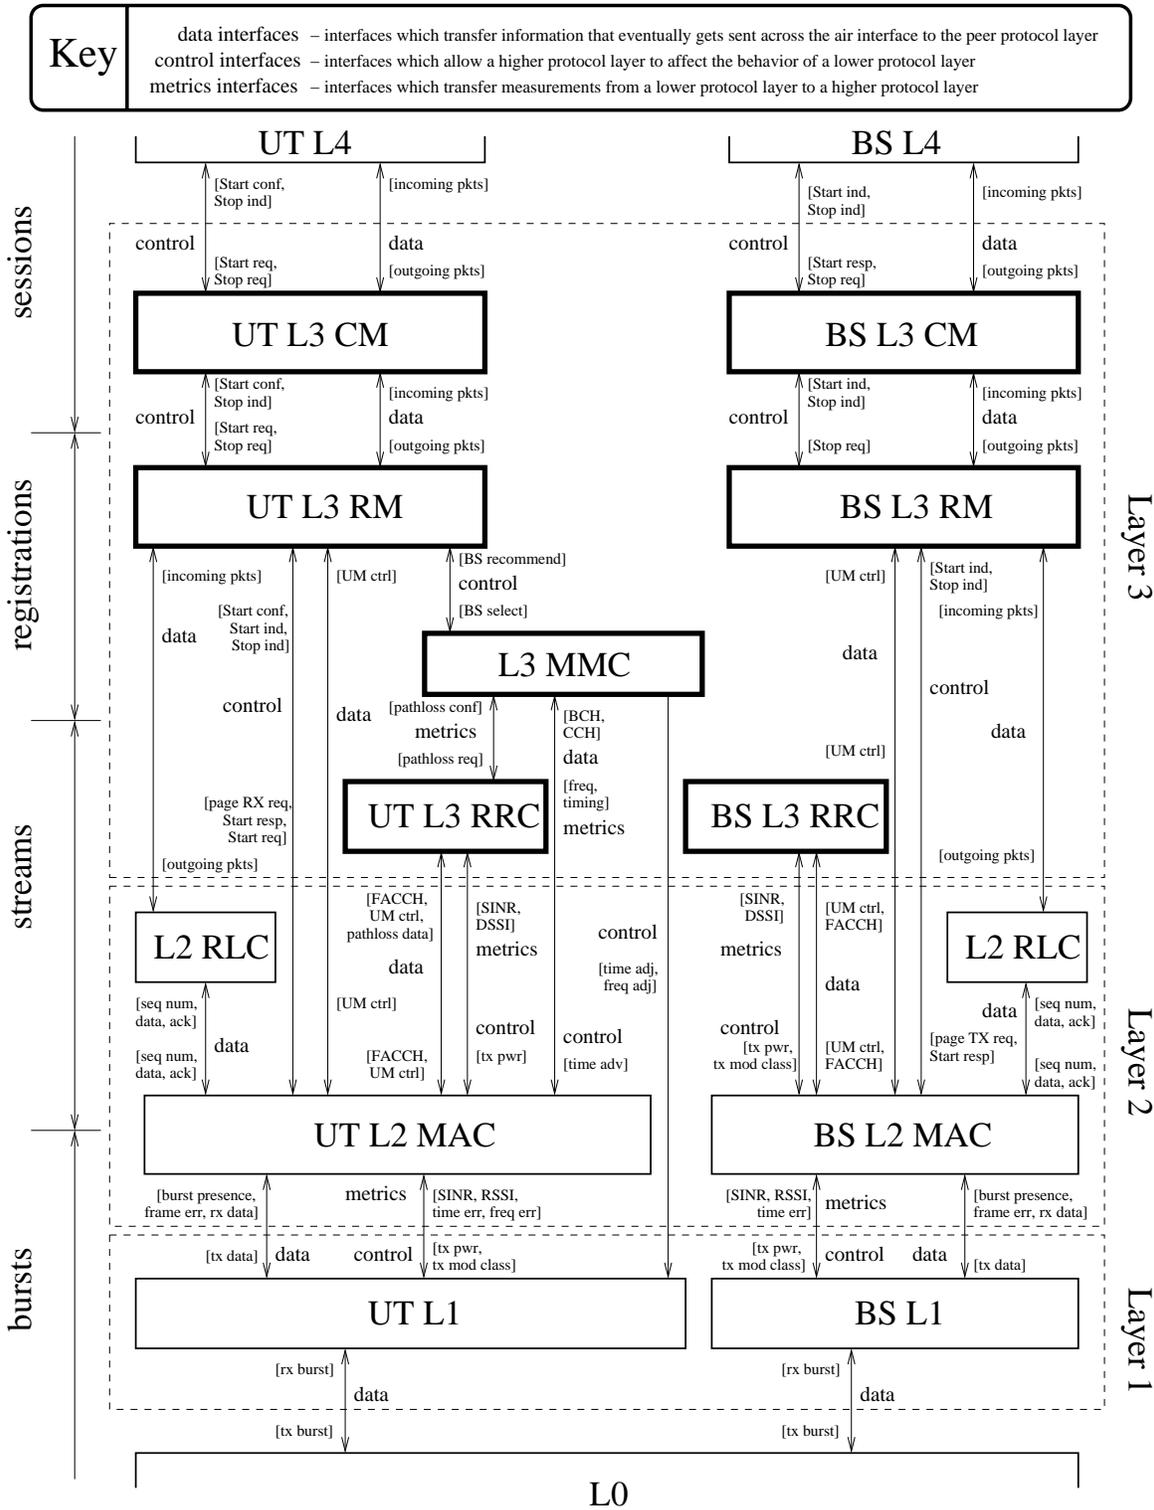


Figure 40: L3 Components and Interactions with Other Layers

9.2.1 Connection Management

The L3 CM component shall be the top sub-level of L3. It shall provide the interface to L4, which is also the boundary of the *HC-SDMA* protocol stack. L3 CM shall provide L4 with a virtual connection between the BS and UT across the air interface.

9.2.1.1 L3 CM Interface with L4

The L3 CM Application Programming Interface (API) to L4 (the *HC-SDMA* Service API) shall consist of an L3 CM service access point that can be used by L4 and an L4 service access point for use by L3 CM. Refer to clause 10.3.2.21, L4 to L3 CM Interface Primitives for the service access point that shall be provided by L3 CM for the use of L4, and clause 10.3.2.20, L3 CM to L4 Interface Primitives for the service access point provided by L4 for the use of L3 CM.

9.2.1.2 L3 CM Interface with L3 RM

Refer to clause 10.3.2.19, L3 CM to L3 RM Interface Primitives regarding the service access point that shall be provided by L3 RM for the use of L3 CM, and clause 10.3.2.18, L3 RM to L3 CM Interface Primitives for the service access point that shall be provided by L3 CM for the use of L3 RM.

9.2.1.3 L3 CM State Machines

The UT L3 CM state machine shall conform to Figure 41 - UT L3 CM State Transition Diagram.

9.2.1.3.1 UT L3 CM State Machine

Figure 41 - UT L3 CM State Transition Diagram illustrates how the UT L3 CM state machine shall respond to requests from UT L4 for virtual connections across the air interface, and requests registrations from UT L3 RM to allow the virtual connections to use physical channels (streams). The following sections describe the actions and events associated with each state.

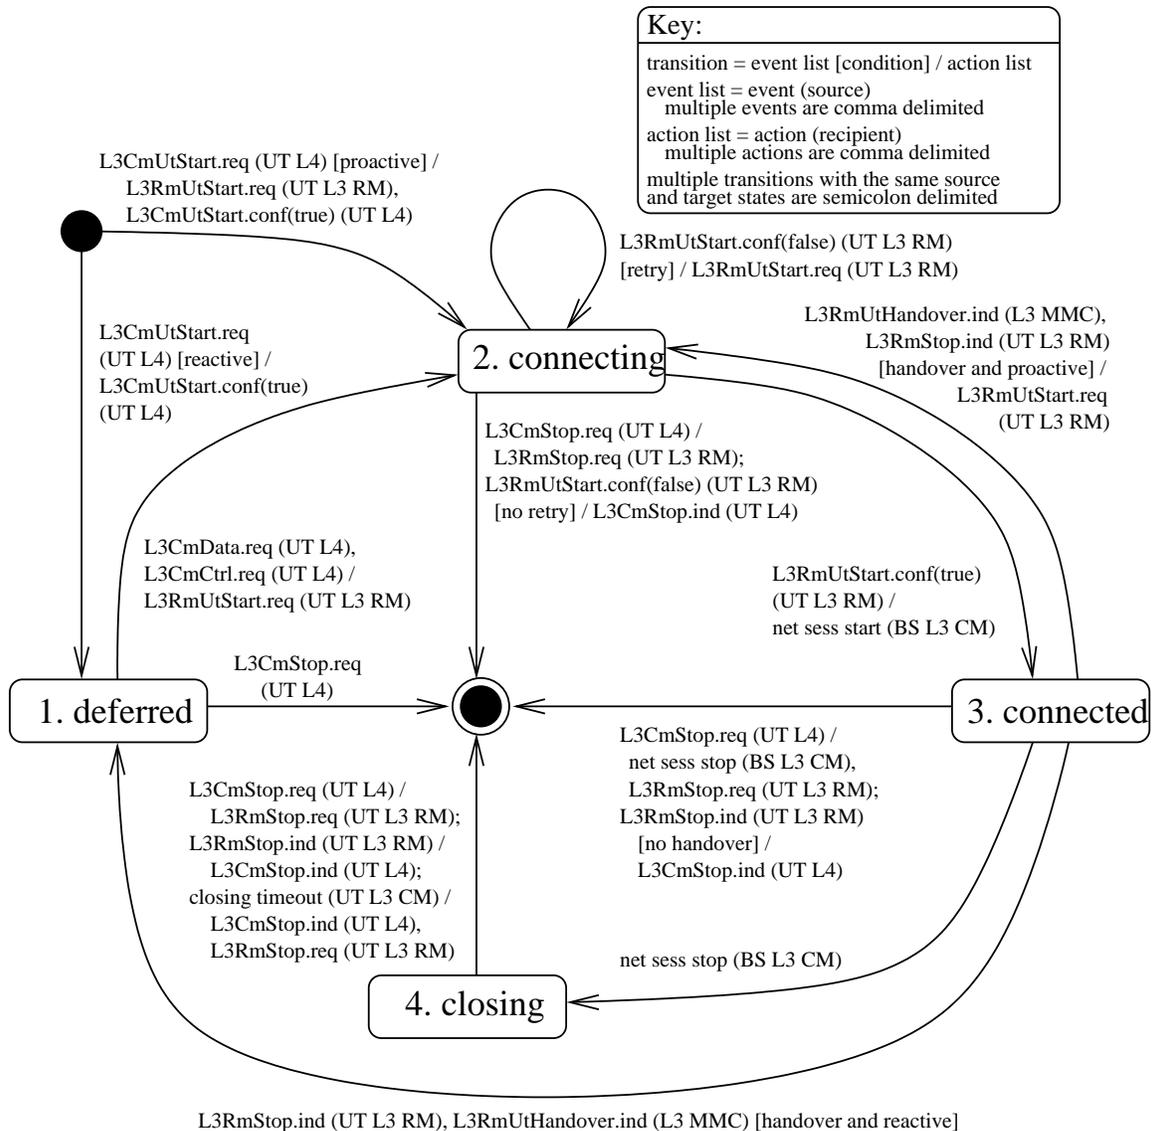


Figure 41: UT L3 CM State Transition Diagram

UT L3 RM shall give different indications to UT L3 CM for a failed registration (`L3RmStop.ind`) versus a recommendation to register with (handover to) a different BS (`L3RmUtHandover.ind`), but UT L3 CM shall respond to these indications in nearly the same way. In particular, UT L3 CM shall give an `L3CmUtHandover.ind` indication to UT L4 in both cases so that UT L4 replies with an `L3CmUtHandover.resp` response.

9.2.1.3.1.1 UT L3 CM Initial State

When UT L3 CM receives an `L3CmUtStart.req` request from UT L4, it shall allocate and initialize a connection control block (CCB) and shall check whether the UT is configured to register proactively or reactively. *Reactive registration* minimizes power and network capacity usage by not registering the session until UT L4 gives UT L3 CM data to send. In contrast, *proactive registration* maximizes performance by registering immediately so that the latency of registration is absorbed at the start of the session, and does not delay the first data sent by UT L4.

If the UT has been configured for proactive registration, UT L3 CM shall send an `L3RmUtStart.req` request to UT L3 RM and transition into the connecting state. Otherwise, if the UT has been configured for reactive registration, UT L3 CM shall enter the deferred state without issuing an `L3RmUtStart.req` request to UT L3 RM. In either case, UT L3 CM shall reply to UT L4 with an `L3CmUtStart.conf(true)`. (At the UT L3 CM level, requests for new connections shall always be honored.)

9.2.1.3.1.2 UT L3 CM Deferred State

When the first data or control packet arrives from UT L4, UT L3 CM shall send an `L3RmUtStart.req` request to UT L3 RM, and transition to the connecting state. The packet received from UT L4 shall be saved until an `L3RmUtStart.conf` confirmation is received from UT L3 RM. If UT L4 issues an `L3CmStop.req` request, UT L3 CM shall transition to the final state, freeing the CCB as it does so.

9.2.1.3.1.3 UT L3 CM Connecting State

UT L3 CM shall remain in the connecting state until UT L3 RM communicates that the air interface connection is complete by giving an `L3RmUtStart.conf` confirmation to UT L3 CM. At that point, UT L3 CM shall send a `net sess start` message to BS L3 CM and shall transition to the connected state. See clause 9.5.1.3.5, Net Sess Start Message and Figure 51 - Session Start (Proactive Registration) for a complete description of the `net sess start` message.

If UT L3 RM is unable to register, it shall inform UT L3 CM. For a limited number of times, UT L3 CM shall wait and shall reissue the `L3RmUtStart.req` request to UT L3 RM. When UT L3 CM times out, it shall send an `L3CmStop.ind` indication to UT L4, deallocate its CCB, and transition to the final state.

If UT L4 issues an `L3CmStop.req` request, UT L3 CM shall pass an `L3RmStop.req` request on to UT L3 RM and then transition to the final state, freeing the CCB as it does so.

9.2.1.3.1.4 UT L3 CM Connected State

UT L3 CM shall remain in the connected state until either UT L4 shuts down the connection with an `L3CmStop.req` request, or the BS peer shuts down the connection with a `net sess stop` message. See clause 9.5.1.3.6, Net Sess Stop Message and Figure 78 - Session Stop (With Top-level Handshake) for a complete description of the `net sess stop` message.

When UT L3 CM receives an `L3CmStop.req` request from UT L4, it shall first send a `net sess stop` message to its BS peer, then forward an `L3RmStop.req` request to UT L3 RM, then free its CCB and transition to the final state.

When UT L3 CM receives a `net sess stop` message from its BS peer, it shall stop sending any outgoing data to its associated UT L3 RM registration, and transition to the closing state to attempt to receive any remaining incoming data.

If UT L3 RM is unable to maintain a registration, for example, due to an extended lack of RF connectivity, UT L3 RM shall give an `L3RmStop.ind` indication to UT L3 CM. In response, UT L3 CM shall attempt to set up a new registration to service the connection. UT L3 CM shall transition to either the deferred or connecting state, depending on whether the UT is configured for reactive or proactive registration respectively. In the proactive case, when UT L3 CM transitions to the connecting state, it shall give a new `L3RmUtStart.req` request to UT L3 RM.

9.2.1.3.1.5 UT L3 CM Closing State

After it has been informed that its BS L3 CM peer is closing the connection, UT L3 CM shall remain in the closing state until either it finds that all outstanding incoming data from the BS has been received, or until the time limit for receiving outstanding incoming data has been exceeded.

UT L3 CM shall use the following procedure to indirectly determine whether it has received all outstanding incoming (downlink) data. Just after having sent the `net sess stop` message to UT L3 CM, BS L3 CM shall issue a `L3RmStop.req` to BS L3 RM, but BS L3 RM shall attempt to deliver outstanding outgoing (downlink) data before closing the registration by sending a `Dereg AM` message to UT L3 RM. Upon receipt of the `Dereg` message, UT L3 RM shall inform UT L3 CM with a `L3RmStop.ind` indication. If UT L3 CM receives an `L3RmStop.ind` from UT L3 RM, it shall send an `L3CmStop.ind` to UT L4, free its CCB, and transition to the final state. Conversely, if UT L3 CM receives an `L3CmStop.req` from UT L4, it shall send an `L3RmStop.req` to UT L3 RM, free its CCB, and transition to the final state. If the time limit for remaining in the closing state is exceeded, UT L3 CM shall send an `L3CmStop.ind` to UT L4 and an `L3RmStop.req` to UT L3 RM, then free its CCB and transition to the final state.

9.2.1.3.2 BS L3 CM State Machine

The BS L3 CM state machine shall conform to Figure 42 - BS L3 CM State Transition Diagram.

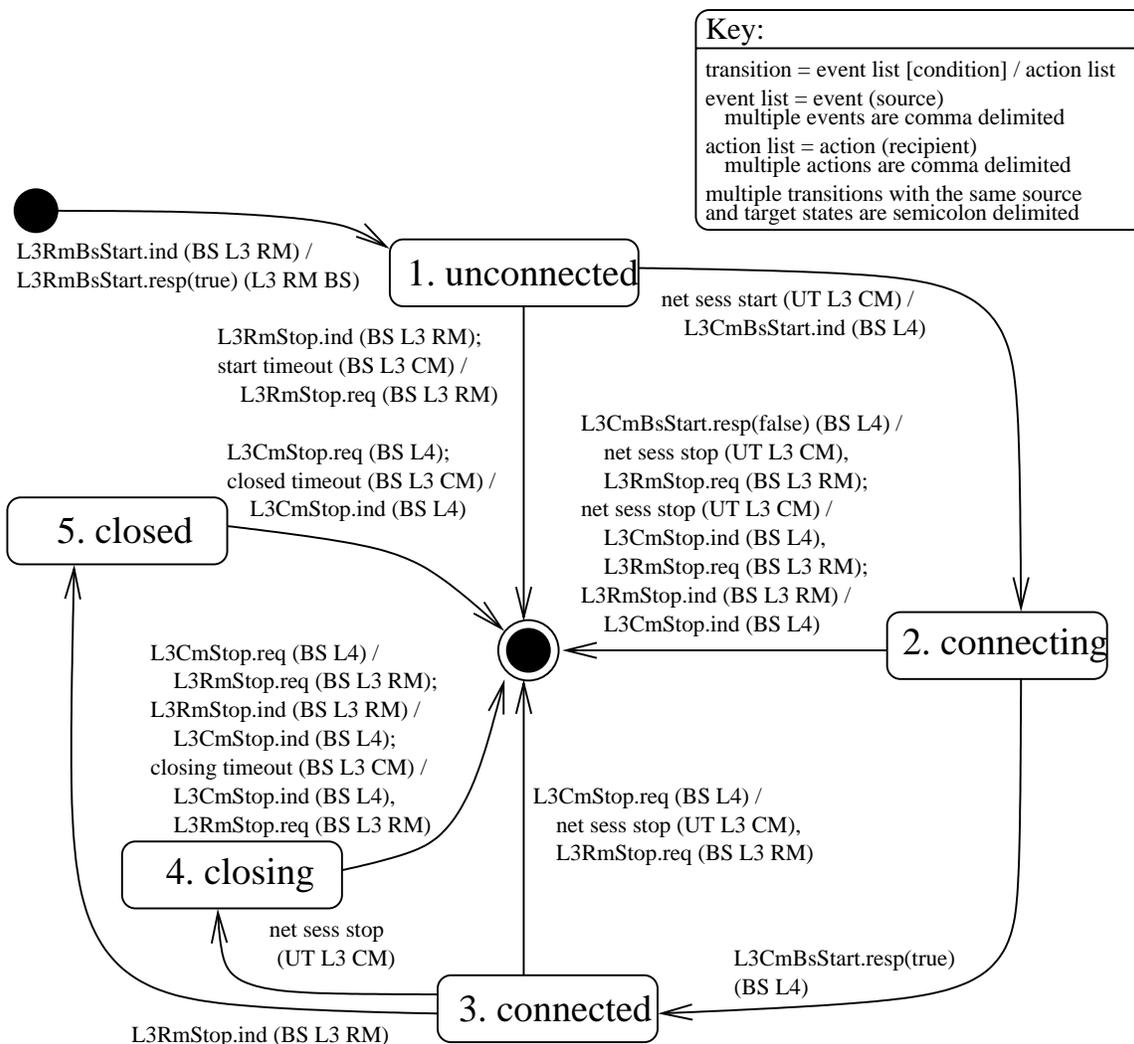


Figure 42: BS L3 CM State Transition Diagram

9.2.1.3.2.1 BS L3 CM Initial State

When BS L3 RM completes a new registration, it shall inform BS L3 CM with an `L3RmBsStart.ind` indication. BS L3 CM shall then allocate and initialize a connection control block (CCB), and transition to the unconnected state.

9.2.1.3.2.2 BS L3 CM Unconnected State

When BS L3 CM receives the `net sess start` message from its UT peer, it shall send an `L3CmBsStart.ind` indication to BS L4, and transition to the connecting state. See clause 9.5.1.3.5, Net Sess Start Message and Figure 51 - Session Start (Proactive Registration) for a complete description of the `net sess start` message. If BS L3 CM times out before it receives the `net sess start` message from its UT peer, it shall send an `L3RmStop.req` request to BS L3 RM, free its CCB, and transition to the final state.

9.2.1.3.2.3 BS L3 CM Connecting State

If BS L4 accepts BS L3 CM's L3CmBsStart.ind indication, BS L3 CM shall move into the connected state and deliver to BS L4 any data it has already received from the UT. If BS L4 rejects the session, BS L3 CM shall send a net sess stop message to its UT L3 CM peer, give an L3RmStop.req request to BS L3 RM, free its CCB, and transitions to the final state. See clause 9.5.1.3.6, Net Sess Stop Message and Figure 80 - Session Rejection for a complete description of the net sess stop message.

If BS L3 CM receives a net sess stop message from its peer, it shall send an L3CmStop.ind indication to BS L4 and an L3RmStop.req request to BS L3 RM, then free its CCB and transition to the final state.

If BS L3 RM sends an L3RmStop.ind indication, BS L3 CM forwards an L3CmStop.ind indication to BS L4 and transitions into the final state, deallocating its CCB in the process.

9.2.1.3.2.4 BS L3 CM Connected State

BS L3 CM shall remain in the connected state until either BS L4 shuts down the connection with an L3CmStop.req request, or the UT peer shuts down the connection with a net sess stop message. When BS L3 CM receives an L3CmStop.req request from BS L4, it shall check to see whether to send a net sess stop message to its UT peer. BS L3 CM shall compare the i-TAP UT certificate from the registration associated with the connection being closed against the i-TAP UT certificates from all other connections' registrations. The presence of an identical certificate indicates that the UT has performed a handover within the same BS. In this case, when the BS infers that the connection is only undergoing a handover rather than a termination, the BS shall not send the net sess stop message. If no matching certificate is found, BS L3 CM shall send the net sess stop message. In either case, BS L3 CM shall then forward an L3RmStop.req request to BS L3 RM, then free its CCB and transition to the final state. When BS L3 CM receives a net sess stop message from its UT peer, it shall stop sending any outgoing data to its associated BS L3 RM registration, and transition to the closing state to attempt to receive any remaining incoming data. If BS L3 RM is unable to maintain a registration, for example, due to an extended lack of RF connectivity, BS L3 RM shall give an L3RmStop.ind indication to BS L3 CM, which shall cause BS L3 CM to transition to the closed state.

9.2.1.3.2.5 BS L3 CM Closing State

After it has been informed that its UT L3 CM peer is closing the connection, BS L3 CM shall remain in the closing state until either it finds that all outstanding incoming data from the UT has been received, or until the time limit for receiving outstanding incoming data has been exceeded.

BS L3 CM shall use the following procedure to indirectly determine whether it has received all outstanding incoming (uplink) data. Just after having sent the net sess stop message to BS L3 CM, UT L3 CM shall issue a L3RmStop.req to UT L3 RM, but UT L3 RM shall attempt to deliver outstanding outgoing (uplink) data before closing the registration by sending a Dereg AM message to BS L3 RM. Upon receipt of the Dereg message, BS L3 RM shall inform BS L3 CM with a L3RmStop.ind indication. If BS L3 CM receives an L3RmStop.ind from BS L3 RM, it shall send an L3CmStop.ind to BS L4, free its CCB, and transition to the final state. Conversely, if BS L3 CM receives an L3CmStop.req from BS L4, it shall send an L3RmStop.req to BS L3 RM, free its CCB, and transition to the final state. If the time limit for remaining in the closing state is exceeded, BS L3 CM shall send an L3CmStop.ind to BS L4 and an L3RmStop.req to BS L3 RM, then free its CCB and transition to the final state.

9.2.1.3.2.6 BS L3 CM Closed State

BS L3 CM remains in the closed state for an extended time, to give the UT a chance to perform a handover. If BS L4 shall send BS L3 CM an `L3CmStop.req` to indicate that the UT has handed over, and the PSS is now tearing down this connection in favor of the connection through the new BS. Upon receipt of the `L3CmStop.req`, BS L3 CM shall free its CCB and transition to the final state. If the time limit for remaining in the closed state is exceeded, BS L3 CM shall send an `L3CmStop.ind` to BS L4, free its CCB, and transition to the final state.

9.2.1.4 Multiple Connections per UT

A UT may maintain several connections concurrently. Thus a UT may have independent concurrent instances of the L3 CM state machines shown in Figure 41 - UT L3 CM State Transition Diagram and Figure 42 - BS L3 CM State Transition Diagram.

9.2.1.4.1 RID versus PID

The protocol provides a paging ID (PID) to identify which UT a page is intended for, but provides a registration ID (RID) to identify which registration a stream is for. This distinction between RIDs and PIDs allows a UT to support multiple registrations without having PCH collisions between the different registrations. See clause 9.2.2.5.1.3, UT L3 RM Negotiating State and clause 9.2.2.5.2.3, BS L3 RM Negotiating State for more information regarding how the BS assigns a PID on behalf of all the registrations on a UT.

When the BS receives the RA-page-resp containing the PID, BS L3 MAC shall decide which of the registrations within the UT the stream should be assigned to, and shall specify the RID for the AA-cts sent in answer to the RA-page-resp. See the `L2MacL3RmBsStart.ind` and `L2MacL3RmBsStart.resp` descriptions in clause 10.3.2.10, L2 MAC to L3 RM Interface Primitives and clause 10.3.2.11, L3 RM to L2 MAC Interface Primitives for more information regarding the BS's choice of which registration receives the stream opened by paging.

There are other parameters related to paging in addition to the PID. All paging-related parameters are shared for all registrations on the UT. For example, all concurrent registrations within the UT share the same paging level durations and the same PCH conventional channel and logical channel masks. See clause 9.5.1.3.3, UT Params Message, clause 9.5.1.3.4, Reg Params Message, and clause 9.5.1.3.8, Rsrc Map Message regarding paging-related parameters that are common among all registrations on the UT.

The UT shall use the same values for paging related parameters across all its concurrent registrations.

9.2.1.4.2 in-band paging

If a stream is required for a registration 1 on a UT, registration 1, and another registration (registration 2) on the same UT already has a stream open, the BS shall send an in-band page message on registration 2's stream to tell registration 1 to open a stream. Moving paging from the PCH to in-band on the TCH avoids the problem of one registration's TCH colliding with another's PCH. See clause 9.3.4.1.5, In-band Page and clause 9.5.1.3.7, In Band Page Message for further description.

9.2.1.4.3 in-band queue depth reporting

If the UT wants to open a stream for registration 1 but is currently unable to do so, because registration 2 is using all the stream resources available to the UT, the UT should send a `co reg queue report` message to the BS in-band on registration 2's stream on behalf of registration 1. Once the BS receives this information, it may choose to break one of registration 2's streams to make room for the stream desired by registration 1. See clause 9.5.1.3.14, Co Reg Queue Report Message and clause 9.3.4.2.6, Intra-UT Stream Conflicts for more information.

9.2.1.4.4 ping renews all selected registrations

When each registration is created, the UT shall choose whether to have it refreshed by a `ping` even if it is not automatically refreshed by data transfer. (See clause 9.3.5.1, Registration Renewal (Ping) for a description of renewing a registration via a `ping`.) The UT shall send a single `ping` on behalf of all registrations that want to be renewed even if they are idle. When the BS receives the `ping`, the BS shall renew all registrations that stated when they were created that they should be renewed by a `ping`. See clause 9.5.1.3.3, UT Params Message for a description of how the UT tells the BS whether a particular registration should get renewed when the BS receives a `ping` from the UT. Because the `ping` represents the UT as a whole rather than a specific registration, the `ping` shall carry the UT's PID rather than a RID.

9.2.1.4.5 TCH renews all selected registrations

When any registration on a UT opens a stream, the BS and UT shall renew that registration along with all the other registrations on the UT which chose to be renewed by `pings`. See clause 9.3.5, Registration Maintenance.

9.2.1.4.6 Network Access UT certificate contains session reference

The i-TAP Network Access UT certificate contains a `session reference ID` as well as a UT ID, so that each connection on the UT is uniquely identified and authenticated. See clause 11.3.1 UT Certificates.

9.2.1.4.7 concerted handover of all connections

If the UT were to have registration 1 to a different BS than registration 2, the PCH of registrations 1 and 2 interfere with each other. Similarly, the TCH of one registration interferes with and is interfered by the PCH of the other registration.

To avoid this, the UT shall make sure that all active connections are registered with the same BS. To accomplish this, when each new registration is fully established, it shall confirm that it is registered with the currently selected BS. See the `L3RmUtCheckBs.req` and `L3RmUtCheckBs.conf` descriptions in clause 10.3.2.19, L3 CM to L3 RM Interface Primitives and clause 10.3.2.18, L3 RM to L3 CM Interface Primitives, respectively.

9.2.1.4.8 connection-based QoS

The HC-SDMA system is optimized for application of Quality of Service (QoS) to connections. Different connections within a UT may have different QoS markings. If so, the BS may apportion the stream resources available to the UT between its registrations based on the appropriate relative priority. See

the description of the `ibQoS` argument of the `L3CmBsData.req` and `L3RmBsData.req` primitives in clause 10.3.2.21, L4 to L3 CM Interface Primitives and clause 10.3.2.19, L3 CM to L3 RM Interface Primitives, respectively.

9.2.2 Registration Management

The L3 RM component maps the virtual connection across the air interface provided by L3 CM to physical connections (streams) provided by Layer 2 Medium Access Control (L2 MAC). Furthermore, it associates each stream from L2 MAC with the corresponding instance of the Layer 2 AM Radio Link Control (L2 RLC), which provides reliable data transport over the stream. L3 RM determines when to open streams, tracks which streams belong to which registrations, and chooses which of a registration's streams to send data on if there is more than one available. The number of concurrent registrations is limited to the range of RID values.

In addition, L3 RM takes the L4 data packets provided by L4 (through L3 CM) and converts them into a form that can be sent over the air interface. On the receiving side, L3 RM converts packets received from the air interface back into network packets before giving them to L3 CM. clause 9.4.6, Packet Data Conversion and clause 9.4.5, Packet Delimiting and Reconstruction for information about network/*HC-SDMA* packet translations.

9.2.2.1 L3 RM Interface with L3 CM

Refer to clause 10.3.2.19, L3 CM to L3 RM Interface Primitives, regarding the service access point provided by L3 RM for the use of L3 CM, and clause 10.3.2.18, L3 RM to L3 CM Interface Primitives, regarding the service access point provided by L3 CM for the use of L3 RM.

9.2.2.2 L3 RM Interface with L2 RLC

Refer to clause 10.3.2.13 L3 RM to L2 RLC Interface Primitives regarding the service access point provided by L2 RLC for the use of L3 RM, and clause 10.3.2.12 L2 RLC to L3 RM Interface Primitives, regarding the service access point provided by L3 RM for the use of L2 RLC.

At the end of a stream, when L2 RLC receives the `L2RlcStop.req` request from L3 RM, L2 RLC shall send any remaining incoming data bytes to L3 RM with the `L2RlcData.ind` primitive, and then inform L3 RM how much outgoing data was acknowledged (and thus implicitly, how much remains unacknowledged) with the `L2RlcStop.conf` primitive.

L2 RLC shall send L3 an `L2RlcHungry.ind` indication to inform L3 that the specified stream might shut down soon due to lack of data unless L3 refills it. (The stream might not shut down anyway if the peer is still sending data on it.) If L3 RM wants to fill multiple streams in parallel, it may use the `L2RlcTxStatus.req` to ask each stream's L2 RLC instance how much unacknowledged data the stream has. Upon receiving this request, L2 RLC shall provide this information to L3 RM with the `L2RlcTxStatus.conf` primitive. This information allows L3 RM to roughly predict how long it will take each stream to run out of outgoing data, and thus make informed decisions about how to route data between streams.

9.2.2.3 L3 RM Interface to L2 MAC

Refer to clause 10.3.2.11, L3 RM to L2 MAC Interface Primitives regarding the service access point provided by L2 MAC for the use of L3 RM, and clause 10.3.2.10, L2 MAC to L3 RM Interface Primitives, regarding the service access point provided by L3 RM for the use of L2 MAC.

BS L3 RM shall start streams with the `L2MacBsPageTx.req` request. With this request, BS L3 RM shall inform L2 MAC of what paging ID (PID) to use for the paging messages and how frequently to send paging messages. BS L3 RM shall use this primitive to disable paging by specifying that the paging frequency is zero. Similarly, UT L3 RM shall use the `L2MacUtPageRx.req` request to tell L2 MAC to start listening for pages. As soon as a registration is set up, UT L3 RM shall give an `L2MacUtPageRx.req` request to UT L2 MAC. When the registration closes, UT L3 RM shall tell UT L2 MAC to stop listening for pages by sending an `L2MacUtPageRx.req` request with a zero paging frequency.

UT L3 RM shall start streams with the `L2MacUtStart.req` request. With this request, UT L3 RM shall inform L2 MAC what type of stream it wants (registration, traffic channel (TCH), or short TCH), which BS to open the stream to (through the BS color code), which conventional channels it is allowed to use for the RA/AA exchange to start the stream, and which RID to send in the stream request (RA) messages to the BS. UT L2 MAC shall respond with either an `L2MacUtStart.conf(success)` or an `L2MacUtStart.conf(failure)` confirmation. When a stream starts at the UT due to receipt of a page from the BS, UT L2 MAC shall inform UT L3 RM with the `L2MacL3RmUtStart.ind` indication. When a stream starts at the BS, BS L2 MAC shall inform BS L3 RM with `L2MacL3RmBsStart.ind`. In response to this, L3 RM shall use the `L2MacIsec.req` request to tell L2 MAC which encryption key to use for the stream, and shall use the `L2RlcConfig.req` request to tell L2 RLC how long to leave the stream open after the stream becomes idle.

When L2 MAC closes a stream, it shall send the `L2MacL3RmStop.ind` indication to L3 RM. L3 RM shall then clean up the stream by giving L2 RLC the `L2RlcStop.req` request, and by waiting for L2 RLC's `L2RlcStop.conf` confirmation.

9.2.2.4 L3 RM Interface to L3 MMC

L3 MMC interacts with UT L3 RM to select which base station to register with, during both session creation and session handover. This is only relevant for the UT. Refer to clause 10.3.2.17, L3 RM to L3 MMC Interface Primitives regarding the service access point provided by L3 MMC for the use of L3 RM, and clause 10.3.2.16, L3 MMC to L3 RM Interface Primitives regarding the service access point provided by L3 RM for the use of L3 MMC.

The `L3MmcUtBestBs.req` request is used to determine which base station to connect a new registration to, both at session creation and during handover. In either case, the `L3MmcUtBestBs.req` request shall occur in response to an `L3RmUtStart.req` request from UT L3 CM to UT L3 RM.

In session creation, the `L3RmUtStart.req` request shall be generated in response to an `L3CmUtStart.req`. In session handover, the `L3RmUtStart.req` request shall be generated in response to an `L3CmUtHandover.resp`, which in turn, shall be generated in response to an `L3CmUtHandover.ind` indication. The `L3CmUtHandover.ind` indication be generated in response to an `L3RmUtHandover.ind` indication, and the `L3RmUtHandover.ind` indication shall in turn be generated in response to an `L3MmcUtNewBestBs.ind` indication. See clause 9.3.6, Handover for related information.

9.2.2.5 L3 RM State Machines

9.2.2.5.1 UT L3 RM State Machine

The UT L3 RM state machine shall conform to Figure 43 - UT L3 RM State Transition Diagram.

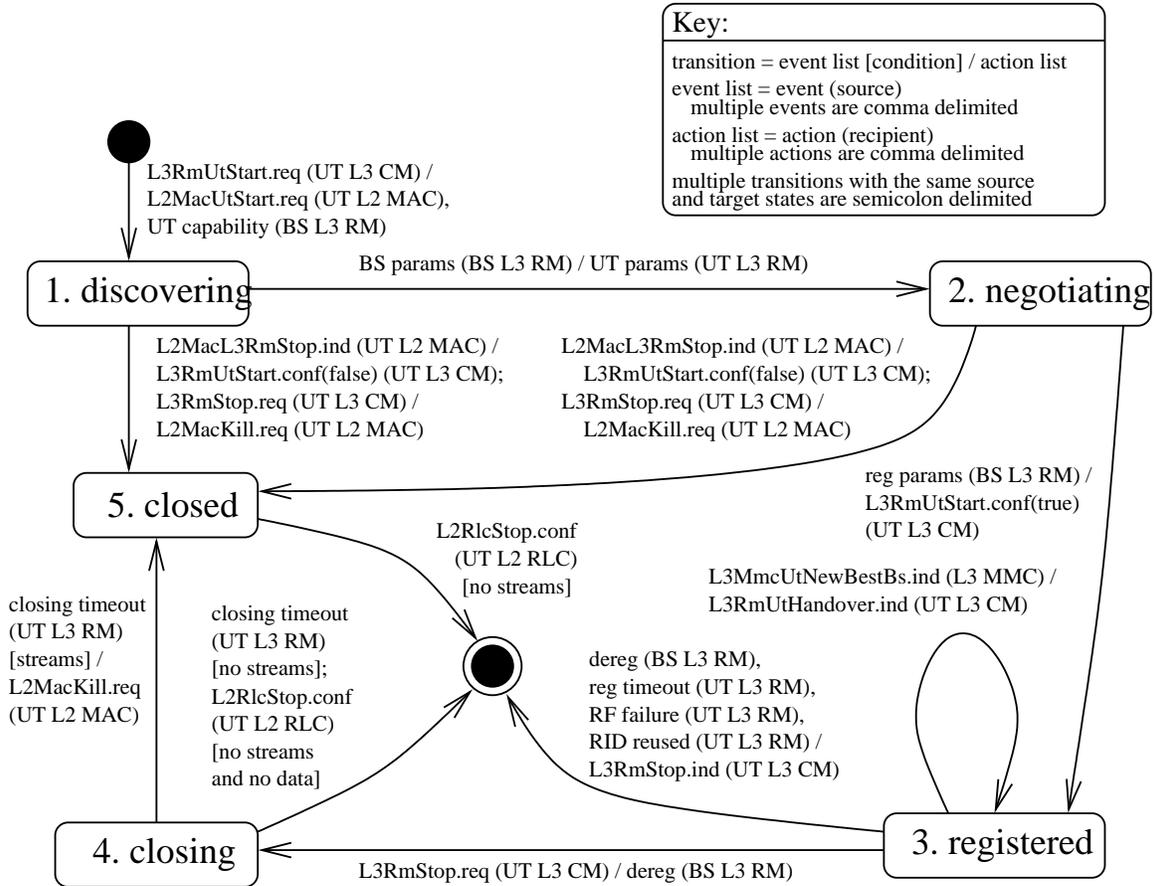


Figure 43: UT L3 RM State Transition Diagram

9.2.2.5.1.1 UT L3 RM Initial State

When UT L3 RM receives an L3UtRmStart.req request from UT L3 CM, it shall allocate a registration control block (RCB), it shall initialize it with the specification of which BSCC to register with and what conventional channels to consider for the registration stream, it shall send UT L2 MAC for a L2MacUtStart.req request, it shall build a UT capability message to inform its BS peer about the hardware characteristics of the UT (in particular, its ability to do power control and link adaptation) and it shall transition to the discovering state. See clause 9.5.1.3.2, UT Capability Message for a complete description of the UT capability message.

9.2.2.5.1.2 UT L3 RM Discovering State

When UT L3 RM receives the BS params message on the registration stream, it shall transition to the negotiating state. In the process, it shall construct a UT params message to tell the BS peer, among other things, whether it already has a valid PID or if it requires a new one, which encryption shared secret to use for the registration, what certificate the UT has to authenticate itself, which Certificate

Authority the certificate is signed by, and which of the optional network/*HC-SDMA* packet data conversion layer algorithms to use. See clause 9.4.6, Packet Data Conversion for information. The encryption shared secret shall be encrypted with the BS public key sent in the BS params message's BS certificate. See clause 9.5.1.3.1, BS Params Message and clause 9.5.1.3.3, UT Params Message for complete descriptions of the BS params and UT params messages, respectively. If UT L3 CM gives UT L3 RM an `L3RmStop.req` request, UT L3 RM shall kill the registration stream and transition to the closed state to wait for the registration stream closure handshake to complete. If UT L2 MAC informs UT L3 RM that the registration stream has closed, UT L3 RM shall give an `L3RmUtStart.conf(fail)` notification to UT L3 CM and shall transition to the closed state to wait for the registration stream closure handshake to complete.

9.2.2.5.1.3 UT L3 RM Negotiating State

If UT L3 RM receives the `reg params` (registration parameters) message on the registration stream, it shall transition to the registered state. In the process, it shall store the registration parameters provided by BS L3 RM in the message. These include the RID, PID, RACH and PCH conventional channel masks, and a set of packet data conversion layer algorithms. Also, UT L3 RM shall notify UT L3 CM that the registration was successful by sending an `L3UtRmStart.conf(success)` confirmation. See 9.5.1.3.4, Reg Params Message for a complete description of the registration parameters message. A `L3RmStop.req` request from UT L3 CM or an `L2MacL3RmStop.ind` indication from UT MAC shall be handled just as in the discovering state.

9.2.2.5.1.4 UT L3 RM Registered State

UT L3 RM shall stay in the registered state until it gets an `L3RmStop.req` request from UT L3 CM, a `dereg` notification message from its BS L3 RM peer, or if it determines that it cannot maintain the registration.

There are three reasons why a particular registration is possibly not maintainable:

1. The registration time-to-live timer expires.
2. BS L3 RM reassigns the RID originally given to this registration to a different registration.
3. UT L3 RM is repeatedly unable to start a stream due to RF connectivity problems.

In each registration maintenance failure case, UT L3 RM shall give an `L3RmStop.ind` indication to UT L3 CM, free the RCB, and transition to the final state. When UT L3 RM receives an `L3RmStop.req` request from UT L3 CM, UT L3 RM shall send a `dereg` AM control message to inform its BS L3 RM peer that the registration is being terminated, and shall also send `starve.cmd` UM control messages to the BS L3 RM peer on all open streams. UT L3 RM shall then transition to the closing state. (See clause 9.5.1.3.11, Dereg Message for a description of the `dereg` message, Table 84 - UM Message Definitions for a definition of the `starve.cmd` UM control message and 9.3.5.3, Deregistration for a more complete discussion of deregistration.)

UT L3 RM shall check a configuration flag to determine whether to prevent the registration time-to-live timer from expiring by resetting it. To reset the timer, it shall perform a "ping" exchange with BS L3 RM.

If BS L3 RM chooses to reassign a RID to a different registration, UT L3 RM detects that its RID is no longer valid as it opens streams. Each stream quickly shuts down due to decryption failures. The decryption failures will result in one of the following problems: invalid L2 RLC sequence numbers,

invalid L2 RLC acknowledgement numbers, or packet delimiting errors. When UT L3 RM sees every stream close due to these errors, it shall infer that its RID has been reused. See clause 9.3.5.4, RID Reuse Detection for details.

9.2.2.5.1.5 UT L3 RM Closing State

UT L3 RM shall stay in the closing state until either all outgoing data has been sent and the last stream closed, or the closing state cleanup timer expires. In the former case, UT L3 RM shall free its RCB and transition to the final state. In the latter, it shall transition to the closed state. In both cases, UT L3 RM shall give an `L2MacUtPageRx.req` request to UT L2 MAC with the paging frequency argument set to zero to tell UT L2 MAC to stop listening for pages (again, unless some other registration still needs to listen). While in the closing state, UT L3 RM shall continue to request streams to send the outgoing data. As each stream starts it shall send a `starve.cmd` UM control message to the BS L3 RM peer. See clause 9.3.5.3, Deregistration for information.

9.2.2.5.1.6 UT L3 RM Closed State

When UT L3 RM enters the closed state, it shall give UT L2 MAC an `L2MacKill.req` request for any open streams it has. While it is in the closed state, UT L3 RM shall immediately kill any streams that start. Due to a potential race condition between UT L3 RM telling UT L2 MAC not to listen for pages and UT L2 MAC telling UT L3 RM that a stream has started due to a page, UT L3 RM shall handle stream starts even though in theory, no streams can start because UT L3 RM is not requesting streams and has told UT L2 MAC to not listen for pages. When the last stream closes, UT L3 RM shall free its RCB and transitions to the final state. See clause 9.3.5.3, Deregistration for information.

9.2.2.5.2 BS L3 RM State Machine

BS L3 RM shall conform to the state transition diagram in Figure 44 - BS L3 RM State Transition Diagram.

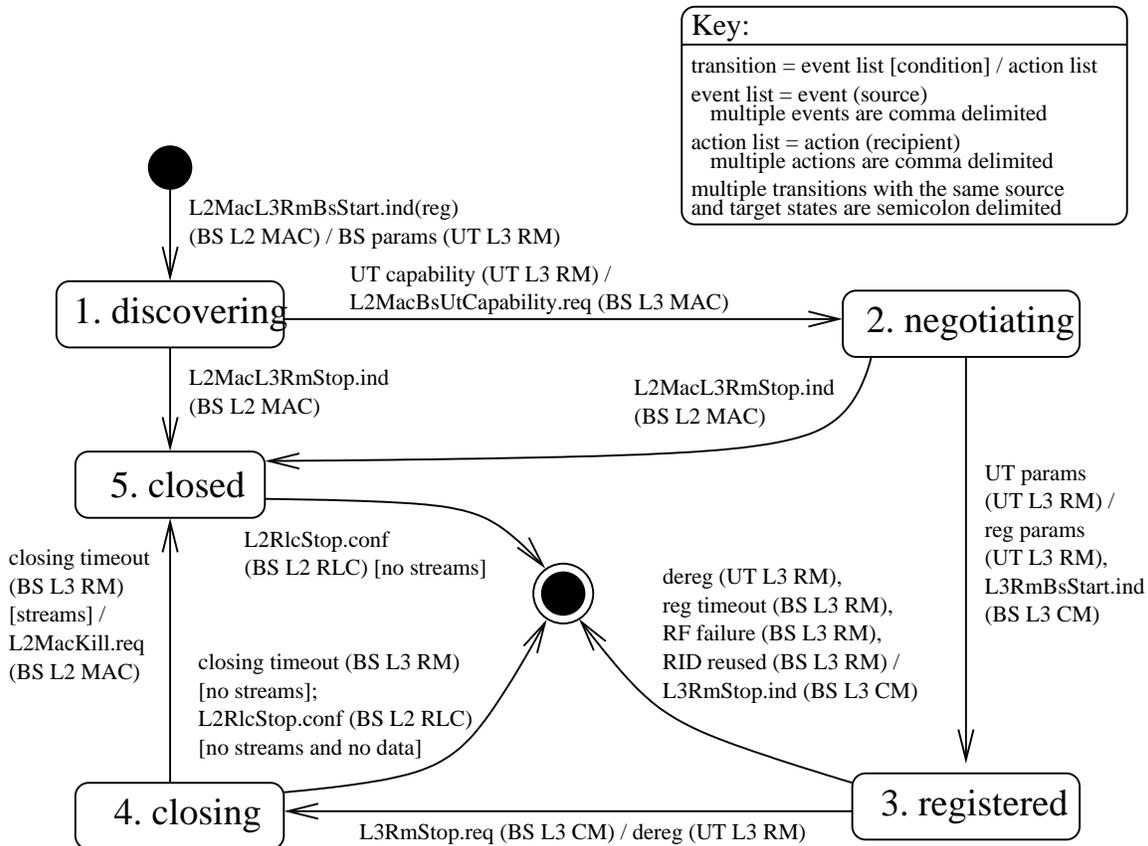


Figure 44: BS L3 RM State Transition Diagram

9.2.2.5.2.1 BS L3 RM Initial State

When BS L3 RM receives an `L2MacL3RmStart.ind` indication from BS MAC for a registration stream, it shall allocate and initialize a RCB, send the `BS params` message on the stream, and transition to the discovering state. See clause 9.5.1.3.1, BS Params Message for a complete description of the BS params message.

9.2.2.5.2.2 BS L3 RM Discovering State

When BS L3 RM receives a `UT capability` message, it shall store the UT capability information in the RCB and also forward it via BS L2 MAC to BS L3 RRC (with the `L2MacBsUtCapability.req`, `L2MacBsUtCapability.ind` primitive sequence) so that BS L3 RRC can do suitable link adaptation on the registration stream. BS L3 RM shall then transition into the negotiating state. See clause 9.5.1.3.2, UT Capability Message for a complete description of the `UT capability` message. If BS L2 MAC informs BS L3 RM with the `L2MacL3RmStop.ind` that the registration stream has closed, BS L3 RM shall temporarily transition to the closed state until the stream closure handshaking with L2 RLC is complete, and then it shall free the RCB and transition to the final state.

9.2.2.5.2.3 BS L3 RM Negotiating State

If BS L3 RM receives a `UT params` message, it shall record the information in the RCB. The encryption shared secret shall be decrypted with the BS private key. See clause 9.5.1.3.3, UT Params Message for a complete description of the `UT params` message. BS L3 RM shall then choose final parameters for the

session, including a RID, PID, RACH, and PCH conventional channel masks, and the set of packet data conversion layer (PDCL) algorithms that is a subset of the ones UT L3 RM requested.

If a UT has only one connection and thus one registration, the UT shall specify in the `UT params` message that it has no PID already assigned. However, if the UT is supporting multiple connections, then after a PID is assigned to the UT during the first registration, subsequent registrations shall inform the BS that the UT already has a PID (by specifying the PID in the `opt PID` field of the `UT params` message). The BS shall then re-use the same PID for the subsequent registrations. The BS shall ensure that it does not re-issue a PID to a new UT if the old UT which was previously assigned the PID still considers the PID valid. If BS L3 RM is not able to allocate an RID or PID, it shall set the reject cause field in the `reg params` message accordingly, and transition to the closing state. Otherwise, BS L3 RM shall send these final parameters in a `reg params` (registration parameters) message to its UT peer, send an `L3RmBsStart.ind` indication BS L3 CM, and transition to the registered state. See clause 9.5.1.3.4, Reg Params Message for a complete description of the registration parameters message. If BS L2 MAC informs BS L3 RM with the `L2MacL3RmStop.ind` that the registration stream has closed, BS L3 RM shall temporarily transition to the closed state until the stream closure handshaking with L2 RLC is complete, and shall then free the RCB and transition to the final state.

9.2.2.5.2.4 BS L3 RM Registered State

BS L3 RM shall stay in the registered state until it gets an `L3RmStop.req` request from BS L3 CM, gets a `dereg` notification message from the UT L3 RM peer, or determines that it cannot maintain the registration.

There are four reasons why a particular registration is possibly not maintainable:

1. The registration time-to-live timer expires.
2. BS L3 RM reassigns the RID originally given to this registration to a different registration.
3. BS L3 RM is repeatedly unable to start a stream due to RF connectivity problems.
4. The registration stream ends without delivering the `reg params` message.

Unlike UT L3 RM, BS L3 RM knows a priori when an RID is reassigned.

BS L3 RM shall determine whether the registration parameters message was delivered after it receives an `L2MacL3RmStop.ind` indication for the registration stream and after it gives L2 RLC an `L2RlcStop.req` request. If L2 RLC returns unacknowledged outgoing data for the registration stream in the associated `L2RlcStop.conf`, the `reg params` message did not go through. When the `reg params` message was not delivered or a registration cannot be maintained, BS L3 RM shall give an `L3RmStop.ind` indication to BS L3 CM, free the RCB, and transition to the final state.

In response to an `L3RmStop.req` request from BS L3 CM, BS L3 RM shall send a `dereg` AM control message to inform its UT L3 RM peer that the registration is being terminated, and also shall send `starve.cmd` UM control messages to the UT L3 RM peer on all open streams. BS L3 RM shall then transition to the closing state. (See clause 9.5.1.3.11, Dereg Message for a description of the `dereg` message, Table 84 - UM Message Definitions for a definition of the `starve.cmd` UM control message and clause 9.3.5.3, Deregistration for a more complete discussion of deregistration.) In all cases where BS L3 RM transitions to the final state, it shall tell BS L2 MAC to not accept RA messages for the registration's RID any more, and not to page the registration's PID, unless there is another registration sharing the PID which is still active.

9.2.2.5.2.5 BS L3 RM Closing State

BS L3 RM shall stay in the closing state until either all outgoing data has been sent and the last stream closed, or the closing state cleanup timer expires. In the former case, BS L3 RM shall free its RCB and transition to the final state. In the latter, it shall transition to the closed state. In both cases, BS L3 RM shall tell BS L2 MAC to not accept RA messages for the registration's RID any more, and not to page the registration's PID, unless there is another registration sharing the PID which is still active. While in the closing state, BS L3 RM may request streams to send the outgoing data, but as each stream starts it shall send a `starve.cmd` UM control message to the UT L3 RM peer. See clause 9.3.5.3, Deregistration for information.

9.2.2.5.2.6 BS L3 RM Closed State

When BS L3 RM enters the closed state, it shall give BS L2 MAC an `L2MacKill.req` request for any open streams it has. While it is in the closed state, BS L3 RM shall immediately kill any streams that start. Due to a potential race condition between BS L3 RM telling BS L2 MAC not to accept RA-rtts and BS L2 MAC telling BS L3 RM that a stream has started, BS L3 RM shall handle stream starts even though no streams can start because BS L3 RM is not paging the UT and has told BS L2 MAC to not accept stream requests (RA-rtts) from the registration's RID. When the last stream closes, BS L3 RM frees its RCB and transitions to the final state. See clause 9.3.5.3, Deregistration for information.

9.2.2.5.3 L3 RM Stream Control Summary

BS L3 RM and UT L3 RM shall respond to protocol events on existing streams as specified in Table 89: L3 RM Stream Control. (In all cases, L3 RM is in the registered state when the listed event occurs.)

Table 89: L3 RM Stream Control

Event	Applies To	Action
L3 CM stop	BS, UT	Send <code>starve.cmd</code> UM ctrl msg (all streams).
L3 RM reg fail	BS, UT	Kill streams with the <code>L2MacKill.req</code> (all streams).
L3 RM strm aggr decision	BS only	Send <code>aggregate.cmd</code> UM ctrl msg (single stream).
L2 RLC hungry	BS, UT	Give data with the <code>L2RlcData.req</code> (single / multiple streams).
L2 RLC data available	BS, UT	Read data from the <code>L2RlcData.ind</code> (single stream).

Table 90 - Newly Opened Streams summarizes operations performed on newly-opened streams.

Table 90: Newly Opened Streams

L3 RM State	Applies To	Action
Registered	BS, UT	Associate the stream with its registration; turn off paging.
Registered	BS only	Send <code>aggregate.cmd</code> UM ctrl msg (if required).
Closing	BS, UT	Send <code>starve.cmd</code> UM ctrl msg; associate the stream with its registration; turn off paging.
Closed	BS, UT	Kill with the <code>L2MacKill.req</code> .

9.2.3 Radio Resource Control

The L3 RRC component shall maintain the RF connectivity of individual streams by performing power control and modulation class selection. Both transmit power and modulation class vary during the stream.

9.2.3.1 L3 RRC Interfaces

L3 RRC interacts with L2 MAC and L3 MMC. Refer to clause 10.3.2.15, L3 MMC to L3 RRC Interface Primitives regarding the service access point provided by L3 RRC for the use of L3 MMC, and clause 10.3.2.14, L3 RRC to L3 MMC Interface Primitives regarding the service access point provided by L3 MMC for the use of L3 RRC. Refer to clause 10.3.2.7, L3 RRC to L2 MAC Interface Primitives regarding the service access point provided by L2 MAC for the use of L3 RRC, and clause 10.3.2.6, L2 MAC to L3 RRC Interface Primitives regarding the service access point provided by L3 RRC for the use of L2 MAC.

9.2.3.2 UT Class and Capabilities

In order to perform link adaptation and power control, BS L3 RRC shall maintain information concerning the capabilities of the UT including the UT's ability to adjust its power and which modulation classes the UT is capable of sending and receiving.

The UT capabilities necessary to perform link adaptation and power control shall be communicated to the BS L3 RRC from the L2 MAC in three stages:

- ◆ First, the CR burst shall contain information concerning the transmit power of the burst. This information shall be communicated to BS L3 RRC via the `L2MacBsCr.ind` primitive.
- ◆ Second, at the beginning of the registration stream the UT shall communicate its class to the BS. The class shall be communicated using the RA-rreg message at the beginning of the registration stream. The class shall be communicated to BS L3 RRC with the `L2MacL3RrcBsStart.ind` which includes the `utCapability` as an argument. The capabilities for each UT class are listed in Table 63 – UT Class Mapping.
- ◆ Finally, the full UT capabilities shall be communicated to the BS in the UT capability and UT parameter messages (see clause 9.5.1.3.2, UT Capability Message and clause 9.5.1.3.3, UT Params Message) which are exchanged during registration. The BS shall use the primitive `L2MacBsUtCapability.ind` to update BS L3 RRC in the middle of the registration stream and shall use the primitive `L2MacL3RrcBsStart.ind` to provide the capabilities at the start of subsequent streams.

9.2.3.3 Modulation Class Selection

9.2.3.3.1 Initial Modulation Class Selection

The *HC-SDMA* physical layer supports different data rates by selecting among various coding and modulation schemes (see clause 4.2, Standard Modulation and Coding.) BS L3 RRC shall determine the modulation classes to start a stream.

The BS and UT shall start the downlink registration stream using modulation class zero, which all UTs are required to support. Prior to receiving the `UT_capability` message, the BS may use the downlink FACCH's recommended modulation class (RMC) field to attempt to assign an uplink modulation class other than zero. If the UT supports the recommended modulation class, it should apply the recommendation, but if not, it shall use modulation class zero. The BS will see what uplink modulation class the UT is actually using by examining the uplink FACCH's current modulation class (CMC) field.

For stream starts subsequent to the registration stream, the stream's initial modulation class shall be set via the RA/AA exchange. The modulation class selection technique is not specified but has the following inputs available from the `L2MacBsRa.ind` indication primitive:

1. The burst quality measurements SINR and DSSI for the burst carrying the RA message.
2. The UT's RA transmit power.

The selected uplink and downlink modulation classes shall be given to BS L2 MAC to insert in the AA message as part of the `L2MacBsAaInfoGet.resp` response primitive.

The first uplink and downlink bursts of a stream is RFN 1 (note RFN 0 is the RA/AA). The first frames after the RA/AA may be optionally blank meaning no bursts are transmitted. Since the FACCH field only allows modulation classes to be changed every other burst (see 7.3.6, Fast Associated Control Channel (FACCH)), the bursts shall be transmitted using the modulation classes specified in the AA message until such time as a modulation change is communicated over the FACCH. Thus, a valid non-empty downlink burst in RFN 1 shall use the modulation class set in the AA frame. If the UT sent a non-empty uplink burst in RFN 1, then the UT shall set the uplink modulation class in RFN 1. If the first non-empty uplink burst is in RFN 2 (meaning no burst was attempted to be sent in RFN 1), the RFN 2 uplink burst shall use the modulation class set in the AA message.

9.2.3.3.2 Modulation Class Adaptation

L3 RRC may change the modulation class during the stream. BS L3 RRC shall set the downlink modulation class on even numbered RFN and shall provide recommended modulation class for the uplink on odd numbered RFN. UT L3 RRC shall provide the modulation class for the uplink on odd numbered RFN.

L2 MAC and L3 RRC shall exchange an `L2MacTchTxInfoGet` indication and response on every frame. In the response, L3 RRC shall provide the modulation class to use for this frame. The BS L3 RRC response shall also contain the recommended uplink modulation class.

The L3 RRC algorithm to compute the modulation class is not specified but it should use the available inputs it receives from L2 MAC:

- ◆ `L2MacTchRxSinr.ind` provides SINR estimates for each received burst.
- ◆ `L2MacLinkAdaptData.ind` provides BS L3 RRC with the uplink modulation class and the available uplink power headroom in the UT which was reported by UT L3 RRC. UT L3 RRC receives the downlink modulation class and the uplink recommended modulation class.
- ◆ `L2MacTchRxPwrCtrlInfo.ind` provides power control commands from the peer (see sections below).

9.2.3.4 Power Control

L3 RRC coordinates power control. The power control methods are divided into two classes: *open loop* and *closed loop*. Open loop power control sets the transmit powers for the bursts containing CR, CM, RA, AA, and the initial uplink and downlink TCH bursts. Closed loop power control sets TCH burst transmit power using feedback from the peer.

9.2.3.4.1 Open-loop Power Control

Open loop power control shall be used for setting the transmit powers for the bursts containing CR, CM, RA, AA, and initial TCH bursts.

9.2.3.4.1.1 CR Burst

UT L3 RRC shall be responsible for calculating the transmit power of the configuration request burst. The UT L3 RRC receives `L2MacUtBburst.ind` every time a B burst is detected by the UT L2 MAC. The information in this primitive includes the basestation transmit power and signal strength of the received B burst. UT L2 MAC shall request a CR power with an `L2MacUtCrPwrInfoGet.ind`. Since the UT L2 MAC informs the UT L3 RRC if a CM burst has been received, the UT L3 RRC has knowledge of the number of CR attempts. UT L3 RRC should select a power for the CR burst based on the basestation transmit power, signal strength of the received B bursts, and number of failed CR attempts. UT L3 RRC should minimize network interference as an objective when deciding on the CR transmit power.

9.2.3.4.1.2 CM Burst

BS L3 RRC is responsible for calculating the power of the burst containing the CM message. BS L3 RRC should minimize network interference as an objective in calculating the CM burst power. BS L2 MAC provides in the `L2MacBsCr.ind` primitive the SINR, RSSI, and DSSI of the received CR burst, along with the UT transmit power which is included in the CR message. Based on this information, BS L3 RRC shall compute the CM burst transmit power.

The BS L3 RRC shall compute a power adjust value to send in the CM burst. This value shall be the difference between the bare minimum receive power required by the BS to reliably receive a CR burst and the actual CR burst's receive power. If the CR burst's receive power is greater than that required for reliable CR reception, the power adjustment value shall be negative. This provides an indication of how much the transmit power could have been changed while still having the CR burst successfully received.

9.2.3.4.1.3 RA Burst

UT L3 RRC shall be responsible for calculating the power of the burst containing the RA message. The UT L3 RRC should minimize network interference as an objective in calculating the RA burst power. UT L3 RRC has available to it the same information as available for calculating the CR burst power. UT L3 RRC knows not only the number of RA attempts, but also the number of RA attempts on the same conventional channel being used for this attempt. In addition, UT L3 RRC is informed of the power adjustment based on the CR burst (see above section) via the `L2MacUtCm.ind` primitive and should use that information as a basis for calculating the RA burst power. Finally, the UT L3 RRC should keep track of the number of attempts of RA and adjust the power based on number of failed attempts.

9.2.3.4.1.4 AA Burst

The BS L3 RRC shall be responsible for calculating the power of the burst containing the AA message. The BS L3 RRC should minimize network interference as an objective in calculating the AA burst power. BS L2 MAC provides BS L3 RRC with an `L2MacBsRa.ind` upon receipt of an RA burst. The indication includes the SINR, RSSI, and DSSI of the received RA burst along with the UT transmit

power which is included in the RA message. Based on this information, BS L3 RRC shall calculate the AA burst transmit power.

BS L3 RRC shall compute an AA power adjustment value to send in the AA burst. This value shall be the difference between the bare minimum receive power required by the BS to reliably receive a RA burst and the actual RA burst's receive power. If the RA burst's receive power is greater than that required for reliable RA reception, the power adjustment value shall be negative. This provides an indication of how much the transmit power could have been changed while still having the RA burst successfully received.

9.2.3.4.1.5 Initial Uplink TCH Burst Power

The UT L3 RRC is responsible for calculating the transmit power of the initial uplink TCH burst. The UT L3 RRC should minimize network interference as an objective in calculating the initial uplink TCH burst power. UT L3 RRC has available to it the same information as was available for calculating the RA burst power and additionally UT L3 RRC is informed of the power adjustment based on RA power (see above section) via the `L2MacUtAa.ind` primitive. UT L3 RRC should use that information as a basis for calculating the initial uplink TCH burst power. The choice of transmit power for the initial uplink TCH burst should account for the change in uplink modulation class between the RA and initial TCH bursts.

9.2.3.4.1.6 Initial Downlink TCH Burst Power

The BS L3 RRC is responsible for calculating the transmit power of the initial downlink TCH burst. The BS L3 RRC should minimize network interference as an objective in calculating the initial downlink TCH burst power. BS L3 RRC has available to it the same information as was available for calculating the AA burst power and should use this information as a basis for calculating the initial downlink TCH burst power. The choice of transmit power for the initial downlink TCH burst shall account for the change in downlink modulation class between the AA and initial TCH bursts.

9.2.3.4.2 Closed-loop Power Control

L3 RRC shall change the power of TCH bursts based on feedback from the peer.

9.2.3.4.2.1 Uplink Power Control

For uplink power control, BS L3 RRC gets the signal quality measurements for each received uplink TCH burst from BS L2 MAC in the `L2MacTchRxSinr.ind` primitive. Based on the signal quality measurements BS L3 RRC shall provide power adjustment commands to L2 MAC. BS L3 RRC shall provide either a one bit power adjustment command or a six bit power adjustment command. The one bit command shall instruct the UT L3 RRC to either raise the transmit power 1dB (value=1) or lower the power by 1 dB (value=0). The six bit power control command shall provide the UT L3 RRC with the amount in dB that the UT is required to change the transmit power. The range of the six bit command is plus or minus 8dB with 1/4 dB resolution.

The power adjustment commands shall be given to BS L2 MAC via the `L2MacTchTxInfoGet.resp` primitive. The L2 MAC shall place the single bit command in the downlink TCH burst's RMU header. The six bit power control command shall be communicated over the air using the single bit command in the downlink TCH burst's RMU header plus the 5 bit UM power control message

(`pwrCtrl.cmd`). The value of the six bit power control command is given by power control command = $16V/63 - 8$ dB where V is the binary expansion of c_1, \dots, c_6 , where c_1 is the 1 bit power control bit in the downlink RMU header, and c_2, \dots, c_6 are the value bits in the UM control message (c_2 is the LSB).

Figure 45 - Uplink Power Control provides an informative example of an uplink power control mechanism as pathloss, shown on the far right side, varies. The left hand side of Figure 45 - Uplink Power Control depicts the sequence of messages produced by the uplink power control and the center portion contains the time axis and graphs for the resulting transmit power of the UT and uplink SINR. The top portion of the figure shows the small scale power corrections from the BS to the UT, with either a +1 value (`RMU.pwrCtrl = 1`) or -1 value (`RMU.pwrCtrl = 0`) being issued. The UT transmit power and uplink SINR vary in response to these commands. However, when there is a sudden drop in the pathloss of 4 dB, shown at the bottom part of the figure, a `pwrCtrl.cmd` UM control message is used along with the `pwrCtrl` bit of the RMU header at the BS to make a large scale power correction needed at the UT. Refer to clause 7.3.7.2.2, Header Field Insertion for usage of the `pwrCtrl` field of the RMU header for uplink power control.

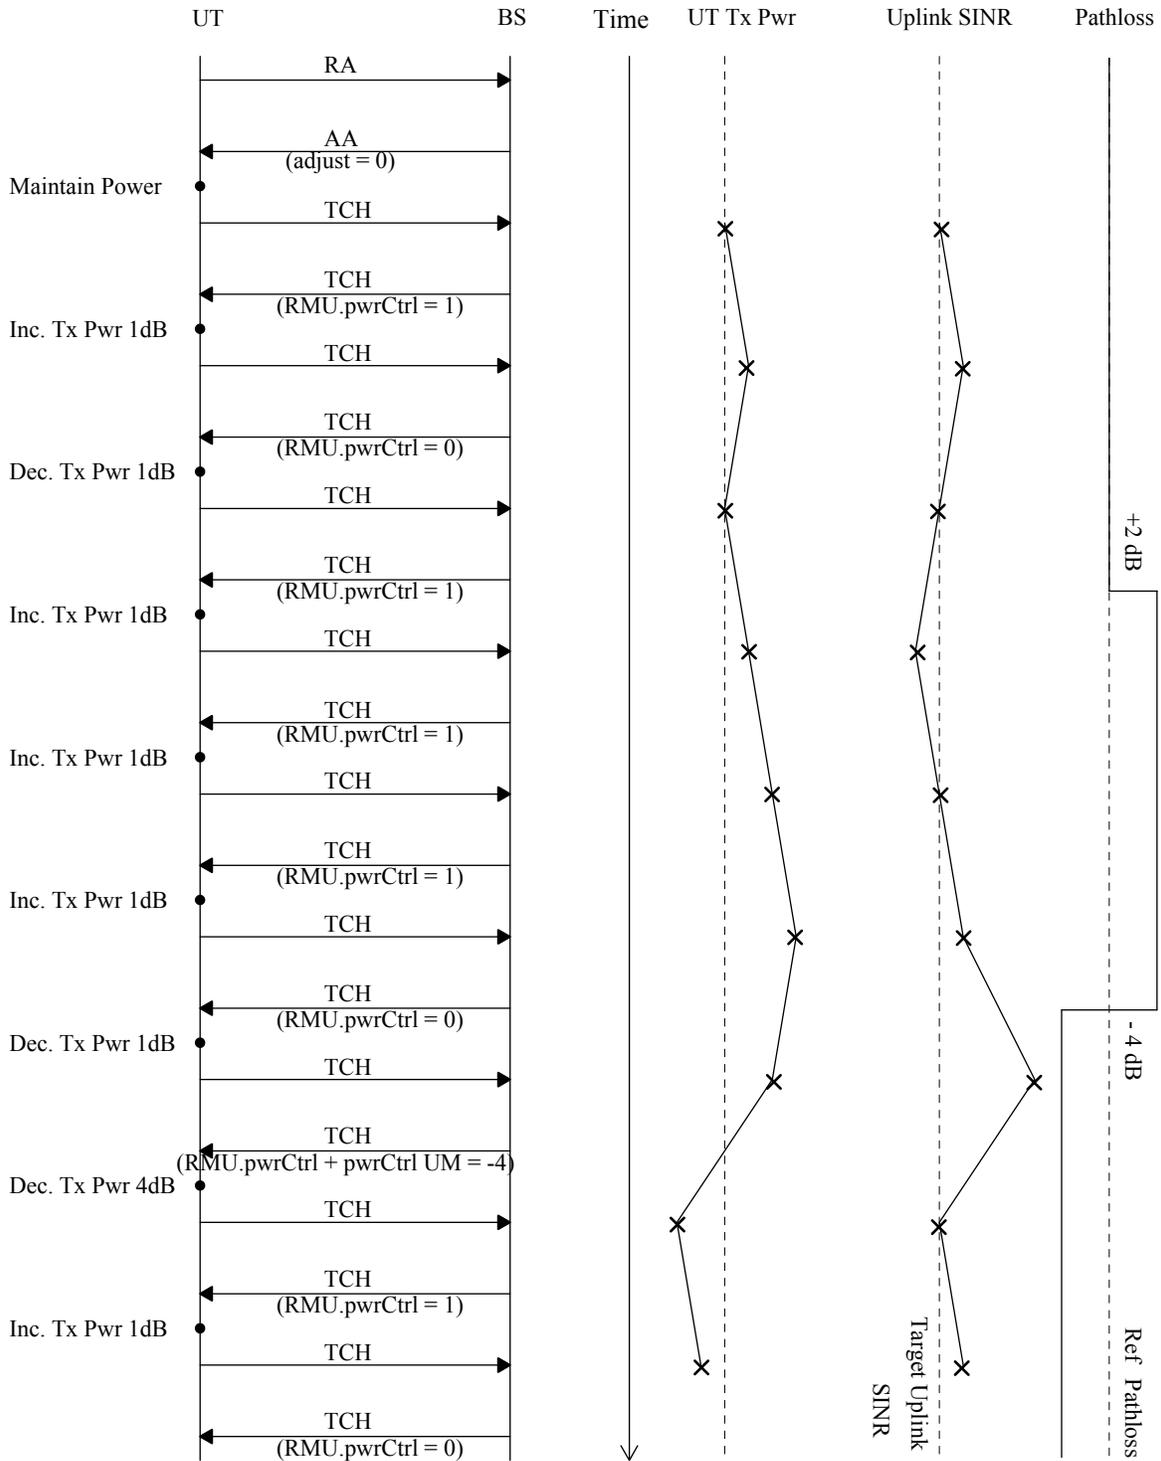


Figure 45: Uplink Power Control

The UT shall respond to the power control commands from the BS.

The CM burst contains a one bit `pwrCtrl` field and if that bit is set to zero, the one bit `pwrCtrl` field in the RMU header for the downlink TCH burst is not active. In this case, the UT shall ignore the one bit `pwrCtrl` commands in the RMU header. This provides a method for testing user terminals by

effectively disabling uplink power control. This feature is included for test purposes only and shall not be used in fielded systems.

9.2.3.4.2.2 Downlink Power Control

For downlink power control, UT L3 RRC shall receive signal quality measurement indications from UT L2 MAC for each received downlink TCH burst via the `L2MacTchRxSinr.ind` primitive. Based on the signal quality measurements UT L3 RRC shall provide SINR error reports to UT L2 MAC. The SINR error reports may be in the form of single bit or six bit values. The one bit error report shall indicate that the SINR is lower than the target or the SINR is higher than the target. The target downlink SINR for each modulation class should be as listed in Table 91 - Downlink SINR Targets. The six-bit SINR error report shall indicate to BS L3 RRC how much the SINR was above or below the target. The range of the six bit error report is plus or minus 8dB with 1/4 dB resolution.

Table 91: Downlink SINR Targets

modulation class	downlink target SINR (in dB)
0	-0.5
1	1.3
2	2.8
3	5.7
4	7.9
5	10.1
6	12.2
7	13.5
8	15.4

For a burst arriving on downlink frame **n** the SINR error report shall be sent on uplink frame **n+1**. Note that the UT's SINR error report to the BS states the difference between the target SINR for the downlink modulation class and the SINR measured by the UT. If the UT's FACCH decode delay is half a frame, the UT shall use the SINR target for the current frame's downlink modulation class as the reference for comparison against its measured SINR. If the UT's FACCH decode delay is one and a half frames, the UT shall use the SINR target for the downlink modulation class which was in use two frames before as the reference SINR, even if the downlink modulation class has since changed. Since the BS is aware of the UT's FACCH decode delay, the BS shall correct for the change in SINR reference to the best of its ability. (See clause 9.5.1.3.3, UT Params Message for a description of how the UT informs the BS of its FACCH decode delays.)

The SINR error reports shall be given to UT L2 MAC via the `L2MacTchTxInfoGet.resp`. The UT L2 MAC shall place the single bit report in the uplink TCH burst's RMU header in the `pwrCtrl` field. The six bit SINR error report shall be communicated using the single bit command in the uplink TCH burst's RMU header plus the 5 bit UM signal quality message (`sinrErr.rpt`). The value of the six bit SINR error report is given by $SINR\ error\ report = 16V/63 - 8\ dB$ where V is the binary expansion of c_1, \dots, c_6 , where c_1 is the 1 bit `pwrCtrl` bit in the uplink RMU header, and c_2, \dots, c_6 are the value bits in the UM signal quality message (c_2 is the LSB).

A positive value for the SINR report shall indicate that the SINR was below the target.

Figure 46: Downlink Power Control provides an informative example of the downlink power control mechanism as pathloss, shown on the far right side, varies. The left hand side of Figure 46: Downlink Power Control depicts the sequence of messages produced by the downlink power control and the center portion contains the time axis and graphs for the resulting transmit power of the BS and downlink SINR. The top portion of the figure depicts the small scale power control with UT reporting SINR error through the `pwrCtrl` bit of the RMU header that is either a +1 value (`RMU.pwrCtrl = 1`) or -1 value (`RMU.pwrCtrl = 0`). The BS transmit power and downlink SINR reported back from the UT vary in response to these commands. However, when there is a sudden drop in the pathloss of 4 dB, shown at the bottom part of the figure, a `sinrErr.rpt` UM control message is used along with the `pwrCtrl` bit of the RMU header at the UT to report a large scale downlink SINR error. Refer to clause 7.3.7.2.2, Header Field Insertion for usage of the `pwrCtrl` field of the RMU header for downlink power control.

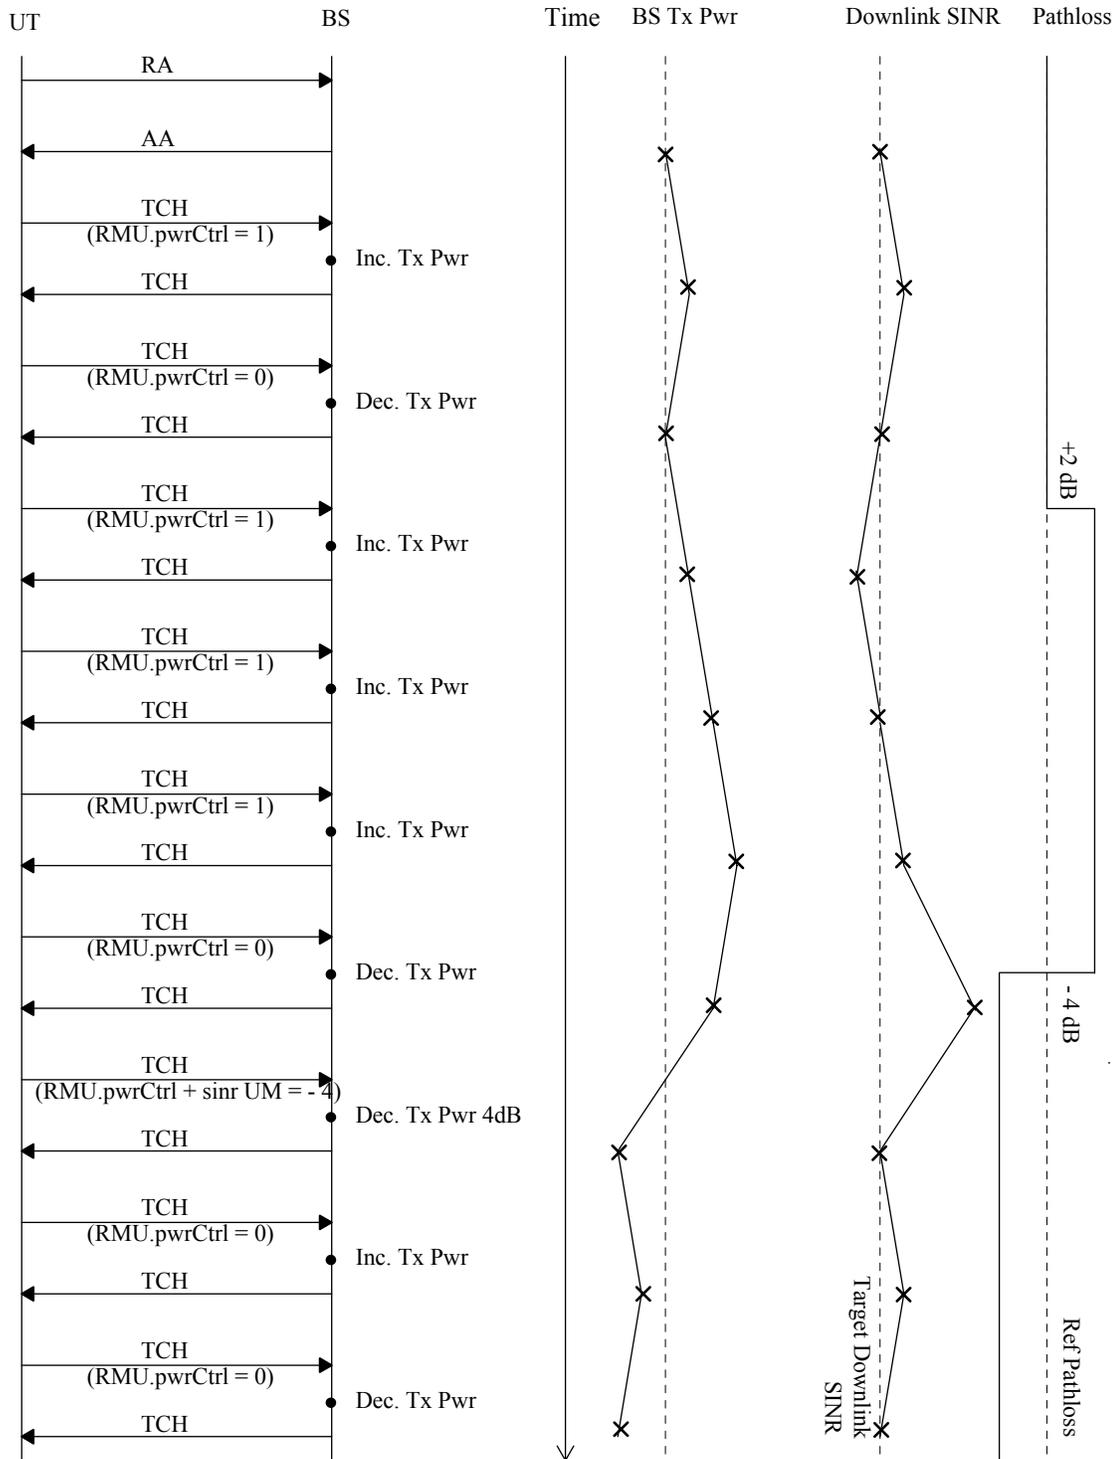


Figure 46: Downlink Power Control

9.2.4 Mobility Management and Control

The L3 MMC component shall provide mobility management and control. L3 MMC is only relevant to the UT. L3 MMC's management shall consist of selecting which base station to use and periodically reevaluating this selection to account for UT mobility. L3 MMC's control shall consist of the coupled

control problem of accounting for the distance between the UT and the BS it is communicating with, for tracking changes in this UT/BS distance, and for tracking drift in the UT's clock.

9.2.4.1 L3 MMC Interfaces

L3 MMC interacts with L1, L2 MAC, and L3 RM. Furthermore, L3 MMC and L3 RRC shall share information about the UT/BS pathloss. Refer to clause 10.3.2.17, L3 RM to L3 MMC Interface Primitives regarding the service access point provided by L3 MMC for the use of L3 RM, and clause 10.3.2.16, L3 MMC to L3 RM Interface Primitives regarding the service access point provided by L3 RM for the use of L3 MMC. Refer to clause 10.3.2.9, L3 MMC to L2 MAC Interface Primitives regarding the service access point provided by L2 MAC for the use of L3 MMC, and clause 10.3.2.8, L2 MAC to L3 MMC Interface Primitives regarding the service access point provided by L3 MMC for the use of L2 MAC. Refer to clause 10.3.2.3, L3 MMC to L1 Interface Primitives regarding the service access point provided by L1 for the use of L3 MMC. Refer to clause 10.3.2.15, L3 MMC to L3 RRC Interface Primitives and clause 10.3.2.14, L3 RRC to L3 MMC Interface Primitives for a description of how L3 MMC and L3 RRC share pathloss information with each other.

9.2.4.2 L3 MMC Management

9.2.4.2.1 L3 MMC State Machine

As L3 MMC performs downlink and uplink timing alignment, it shall transition through four states: unaligned, rough aligned (BS scan), downlink aligned, and uplink aligned. The state transitions for L3 MMC shall conform to Figure 47 - L3 MMC Timing Alignment State Transition Diagram.

Initially, L3 MMC has no notion about the network timing and shall thus be in the unaligned state. L3 MMC shall search for the BCH F burst to get an approximate idea of the network timing. However, because of the uncertainty in distance and the large granularity in this initial timing estimation, L3 MMC shall also search over a large timing window which should be 100 microseconds or greater to find the simulcast BCH T bursts from individual base stations. The BCH T bursts shall be used to maintain rough alignment in the synchronization sub-phase. L3 MMC shall keep track of the timing of each base station's BCH T burst it hears. Since the base stations are all synchronized, this also tells L3 MMC the relative distances of all the base stations it hears. For the evaluation sub-phase of the rough aligned state, L3 MMC shall scan for BCH B bursts from nearby base stations and shall keep track of the received signal qualities and payload fields each time it hears a BCH B burst. L3 MMC shall use this information to rank the BSs.

In the downlink aligned state the UT is able to listen with precise timing to any base station. However, because the absolute distance to the base station is not known, no timing advance is performed, and the UT's uplink is received late by the BS by $2 * (\text{UT} \rightarrow \text{BS distance}) / (\text{speed of light})$. The factor of two arises because the propagation delay is incurred both on the downlink and uplink legs and is not corrected in either direction (by the UT's receive delay or transmit advance) until the UT knows its distance from the base station.

L3 MMC shall choose a preferred base station and shall then send a CR burst to that base station in order to find out its distance and start doing timing advance. The CM response from the base station shall tell L3 MMC what the timing alignment of the uplink CR burst was. L3 MMC shall then compute the absolute distance to that base station. Since the relative distances to all base stations are known, knowing the absolute distance to one BS should be used to derive the absolute distances to all other base stations. The UT is now uplink aligned because it can do transmit timing advance for any base station in order to be heard at precisely the right time.

When L3 MMC reaches the downlink aligned state, it shall maintain that state by listening to downlink bursts (TCH or BCH) and correcting the timing error that accumulates due to clock drift and the UT's mobility. If the UT is unable to hear downlink data for a long enough time that it cannot predict when the BS's bursts will be heard within the allowable tolerance, the UT shall fall back to the rough aligned state. The time that L3 MMC may go without hearing downlink bursts yet remain in the downlink aligned state shall be determined by L3 MMC's estimation of the clock drift and UT mobility. See clause 9.2.4.3.1, Downlink Timing Measurements for information.

When L3 MMC reaches the uplink aligned state, it shall perform the downlink realignment described above. In addition, the UT shall maintain the uplink alignment by getting feedback from the BS regarding the uplink timing in the form of CR/CM exchanges, RA/AA exchanges, or TCH in-band messaging (the `timing.rpt` UM control message). If the UT is unable to get feedback from the BS for a long enough time that it cannot be sure its uplink bursts will be heard within the BS's timing tolerance, it shall fall back to the downlink aligned state. The time that L3 MMC may go without getting uplink timing feedback yet remain in the uplink aligned state shall be determined by L3 MMC's estimation of clock drift and UT mobility. See clause 9.2.4.3.2, Uplink Timing Measurements for information.

The UT may transition to the rough aligned state if there is an uplink or downlink realignment failure for other reasons, for example, repeated failures of RA/AA attempts on a lightly loaded BS may cause the UT to decide that the downlink alignment is off.

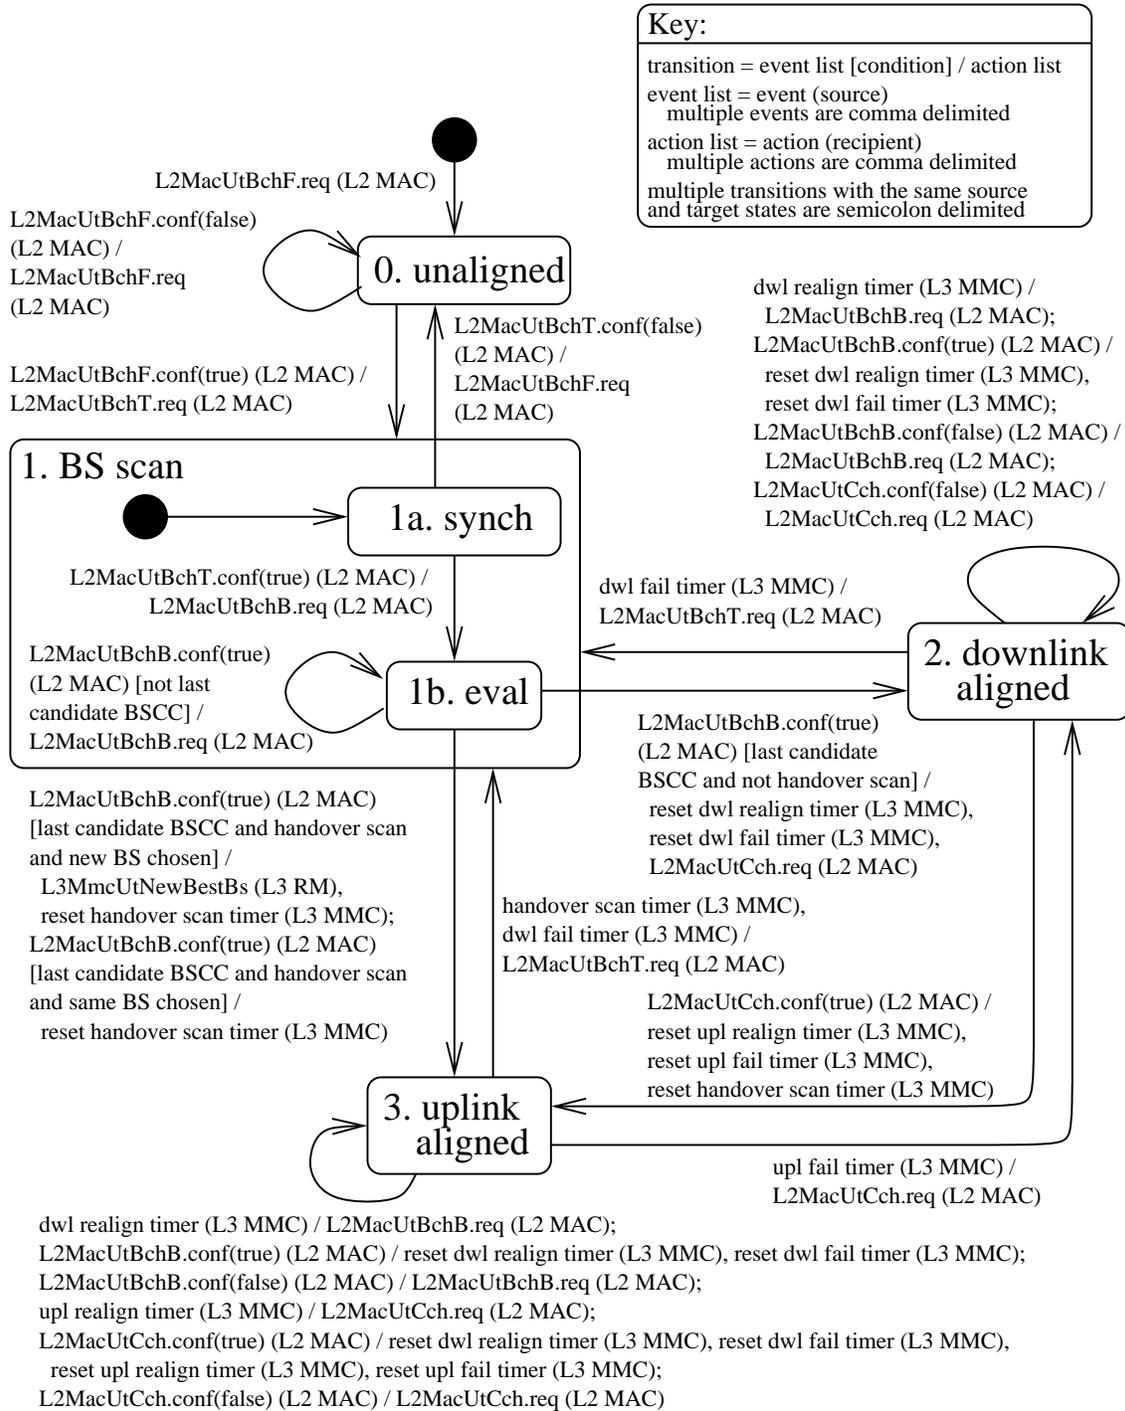


Figure 47: L3 MMC Timing Alignment State Transition Diagram

9.2.4.2.2 BS Selection

At initiation, L3 MMC shall listen to BCH from all BSs during the BS scan (rough aligned) state and it should compile information about each BS it hears. The following information should be recorded:

- ◆ BCH receive signal strength.
- ◆ Path loss to the BS (based on the difference between the BCH transmit power reported in the BCH burst and the BCH receive signal strength).

- ◆ Relative distance of the BS (based on the BCH receive timing alignment).
- ◆ BS load factor (reported in the BCH burst).

The BSs should be ranked by cost and the BS with the lowest cost should be selected as the BS to access. The “cost” of each BS should be computed by preferring BSs with high BCH receive RSSI, small UT/BS path loss, small UT/BS distance, and small BS load. (The impact of each of these components to the overall cost function should not be necessarily linear. For example, there is a maximum allowed distance; BSs that exceed this distance should receive an infinite cost.)

The sequence of events associated with initial timing acquisition to a BS shall conform to Figure 48 - L3 MMC Acquisition. L3 MMC shall sequentially request UT L2 MAC to schedule receptions of BCH F, T, and B bursts using the `L2MacUtBchF.req`, `L2MacUtBchT.req`, and `L2MacUtBchB.req` requests, respectively. The F burst provides coarse timing and frequency adjustments. The T burst provides the UT with an idea of link quality and receive timing to each base station. With these inputs, the UT shall then produce a candidate list of base stations on which to listen to B bursts. The B bursts should be used to provide more accurate receive timing to each candidate BS as well as to provide additional information (pathloss and BS load) to allow the UT to find the best BS for further communication. L3 MMC shall then issue a separate `L2MacUtBchB.req` for each base station on the candidate list. When L3 MMC has evaluated the B bursts and has decided on a best BS, it shall request a transmit of a CR burst to that BS via an `L2MacUtCch.req`. After UT L2 MAC receives the CM burst, it shall send `L2MacUtTimingErr.ind` and `L2MacUtFreqErr.ind` indications to L3 MMC informing it of timing and frequency errors so that the UT can achieve both uplink and downlink timing alignment. At this point L3 MMC shall inform UT L3 RM via an `L3MmcUtNewBestBs.ind` that there is a new best BS. If UT L3 RM succeeds in establishing a connection through this BS, it shall send L3 MMC the `L3MmcUtTrackBs.req` primitive.

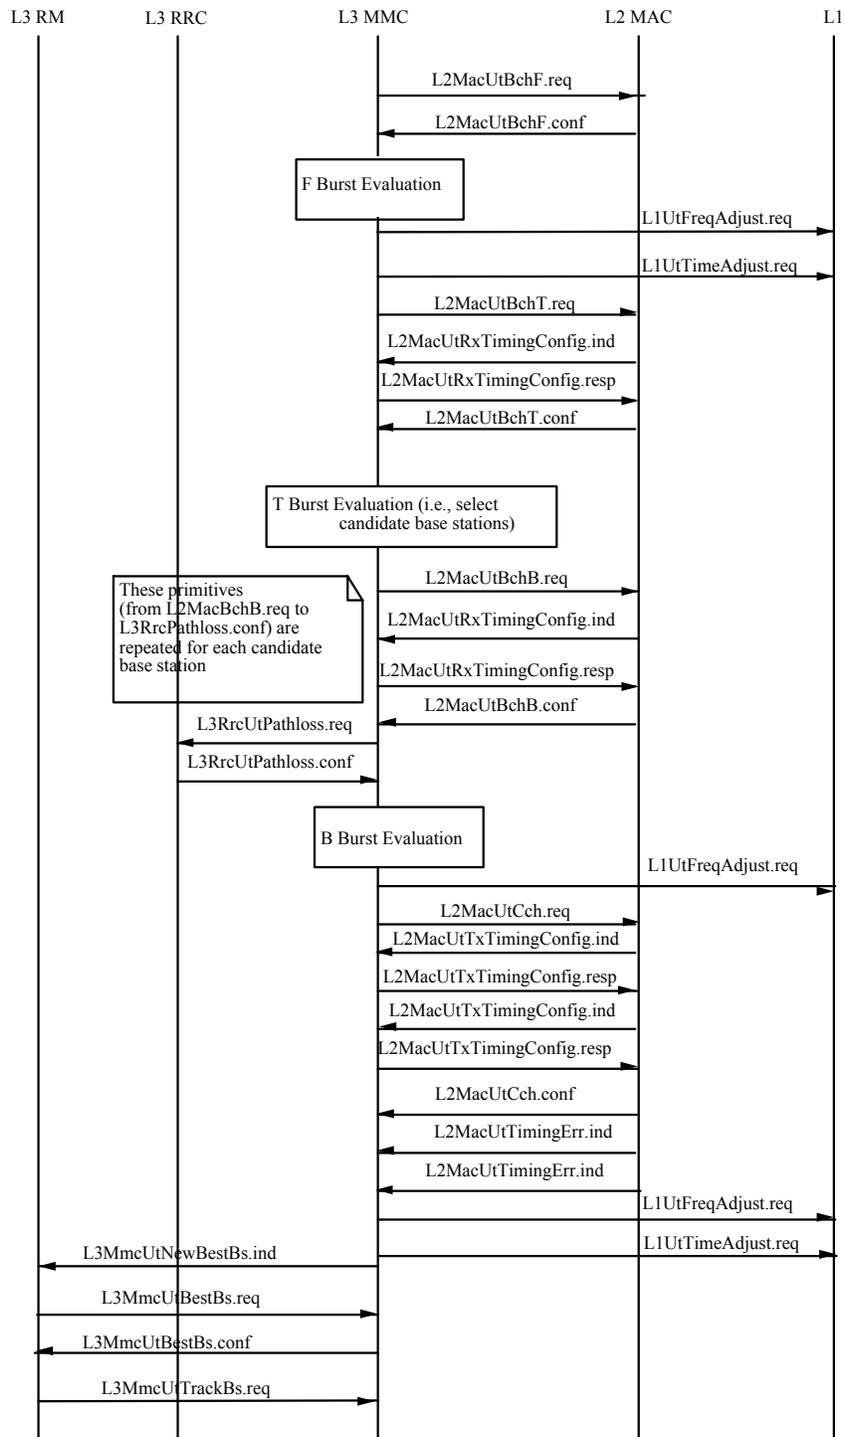


Figure 48: L3 MMC Acquisition

9.2.4.2.3 BS Monitoring

L3 MMC shall periodically listen to BCH from other BSs besides the one it has chosen to access. Note that this BS scan to determine whether it is appropriate to handover to some other BS is similar to the initial BS scan when L3 MMC was in the rough aligned state. However, in the handover scan L3 MMC should search over a shorter timing window of roughly 50 microseconds (since L3 MMC now knows the absolute distance to the current BS). L3 MMC should have a two-tiered approach to reevaluating

BSs based on their current BCH. L3 MMC should maintain a short list of BSs that rated well the last time all BSs were ranked but were not rated as the best BS. L3 MMC should listen to BCH from these few BSs frequently, stealing the BCH time slot from TCH for a frame, if necessary. On a less frequent basis, L3 MMC should listen for BCH from all BSs, again interrupting TCH to do so, if necessary.

The short list of BSs to listen to frequently should be updated based on the rankings of this universal evaluation. The two-tiered approach allows for the following:

- ◆ Frequent monitoring of BSs that are likely candidates for handover without unduly conflicting with TCH.
- ◆ Noticing BSs that are added to the system after the UT is in operation.
- ◆ Considering BSs for handover as they become viable.

The same method should be used to rank BSs as during the initial BS selection, with the addition that the BS that the UT has currently selected as the best BS receives a preferential adjustment to its cost. This hysteresis prevents the UT from continually handing over between two BSs that rate nearly the same from an RF perspective.

9.2.4.3 L3 MMC Control

L3 MMC Control shall be used to track clock drift and mobility. In order to correct for mobility, L3 MMC shall specify how much timing advance to perform each time a burst is sent or received.

HC-SDMA basestations shall support UTs with a one-way signal propagation delay less than 42.5 microseconds (between the UT and the BS), approximately equivalent to a distance of 12.75 km (in line of sight propagation). HC-SDMA basestations may receive uplink bursts upto 15 microseconds after the start of the uplink burst time for the purpose of serving UTs with a signal propagation delay larger than 42.5 microseconds, approximately equivalent to a distance of 12.75km. HC-SDMA basestations cannot support UTs with a one-way signal propagation delay larger than 50 microseconds (between the UT and the BS), approximately equivalent to a distance of 15km (in line of sight propagation).

The BS's transmission to any UT shall occur at the beginning of the downlink burst time regardless of the UT's distance. Timing advance shall be performed by the UT to correct both the uplink and downlink timing for all bursts except the CR burst. The UT shall not perform timing advance when it transmits the CR burst. (This causes the CR burst to be delayed, but this is permissible due to the extra guard time built into the CR burst.)

L3 MMC gets information separately about the Rx (downlink) timing error and the Tx (uplink) timing error. The Rx timing error shall be evaluated every time the UT receives a burst from the BS. The Tx timing error shall be determined by the feedback from the BS in the CR/CM and RA/AA exchanges as well as in `timing_rpt` UM control messages during TCH. Rx and Tx timing errors are due to UT's clock drift and UT's mobility. In order to correct for both Rx and Tx timing errors, L3 MMC shall determine which component of the errors is due to clock drift and which is due to mobility.

9.2.4.3.1 Downlink Timing Measurements

The UT shall be programmed with its own downlink timing error tolerance. L3 MMC shall use this downlink timing error tolerance to determine how often it must resynchronize its downlink timing.

If the UT has streams open on a regular basis to exchange data with the BS, L3 MMC shall use the downlink TCH bursts to obtain timing and frequency error measurements. However, if the resynchronization interval elapses without the UT hearing a downlink burst, the UT shall schedule

BCH reception to correct the downlink timing. Rather than exclusively using independent BCH receptions from the currently selected BS, L3 MMC should also use the subset of BCH receptions that come from BSs currently used for BS monitoring.

9.2.4.3.2 Uplink Timing Measurements

Based on its knowledge that the BS's uplink timing error tolerance is +/-4 microseconds (see clause 12.9.1.2, Uplink Performance Specifications), L3 MMC shall determine how often it must resynchronize its uplink timing.

The UT should determine the interval at which it needs to obtain measurement feedback regarding the uplink timing. This should be done on the basis of an estimate of the largest conceivable error in the UT's division of the downlink timing error into clock drift and distance change components and then on the basis of how quickly the uplink timing drifts outside of the BS's uplink timing tolerance, even when the downlink timing is precise. This estimation algorithm should account for any error that accumulates in the downlink timing and the granularity of uplink timing error reporting in downlink messages (1 microsecond for the CM, 0.5 microseconds for the AA and `timingErr.rpt` message).

If this interval elapses without a CR/CM, RA/AA exchange, or `timing.rpt` UM message that gives the UT feedback about its uplink timing, the UT shall initiate such a transfer. (See the `timing.rpt` entry in Table 84 - UM Message Definitions regarding the formatting of the `timing.rpt` UM control message.) If the uplink timing error on a TCH stream become too close to the 4 microsecond BS timing tolerance, the BS shall send a `timing.rpt` UM control message telling the UT the (uplink) timing error seen on the TCH stream. If it is time for the UT to obtain an uplink timing measurement with a CR/CM or RA/AA exchange but it cannot, due to having TCH in progress, the UT shall initiate a CR/CM or RA/AA exchange as soon as the conflicting TCH ends, unless a `timing.rpt` is received on the TCH. If a `timing.rpt` is received, then the UT shall make the necessary timing and distance corrections and shall recompute the time by which it requires the next uplink timing measurement.

The UT shall not schedule CR/CM or RA/AA exchanges beyond that required to maintain timing within the base station tolerance except as may be required for registration renewal, handover scanning, and initiating data transfer.

9.3 Protocol Events

9.3.1 Introduction

This section describes the various events involving L3 protocol entities and illustrates them with message sequence charts (MSCs). Refer to Message Sequence Chart (MSC) notation in clause 1.3 for a brief description of the syntax of the MSC graphical language.

9.3.2 BCH Acquisition and Selection

This process of acquiring BCH and selecting a BS shall conform to Figure 49 - BCH Acquisition and BS Selection. At power-on, L3 MMC shall scan for BCH from all BSs and shall select a suitable one. After selection, L3 MMC shall send the `L3MmcUtNewBestBs.ind` indication to UT L3 RM. Before UT L3 RM receives the initial `L3MmcUtNewBestBs.ind` indication, UT L3 RM shall respond immediately to UT L3 CM's `L3RmUtStart.req` requests with `L3RmUtStart.conf(failure)` as shown in Figure 50 - Premature Session Start Rejection. See clause 9.2.4.2.1 L3 MMC State Machine and clause 9.2.4.2.2 BS Selection for information on how L3 MMC selects a BS.

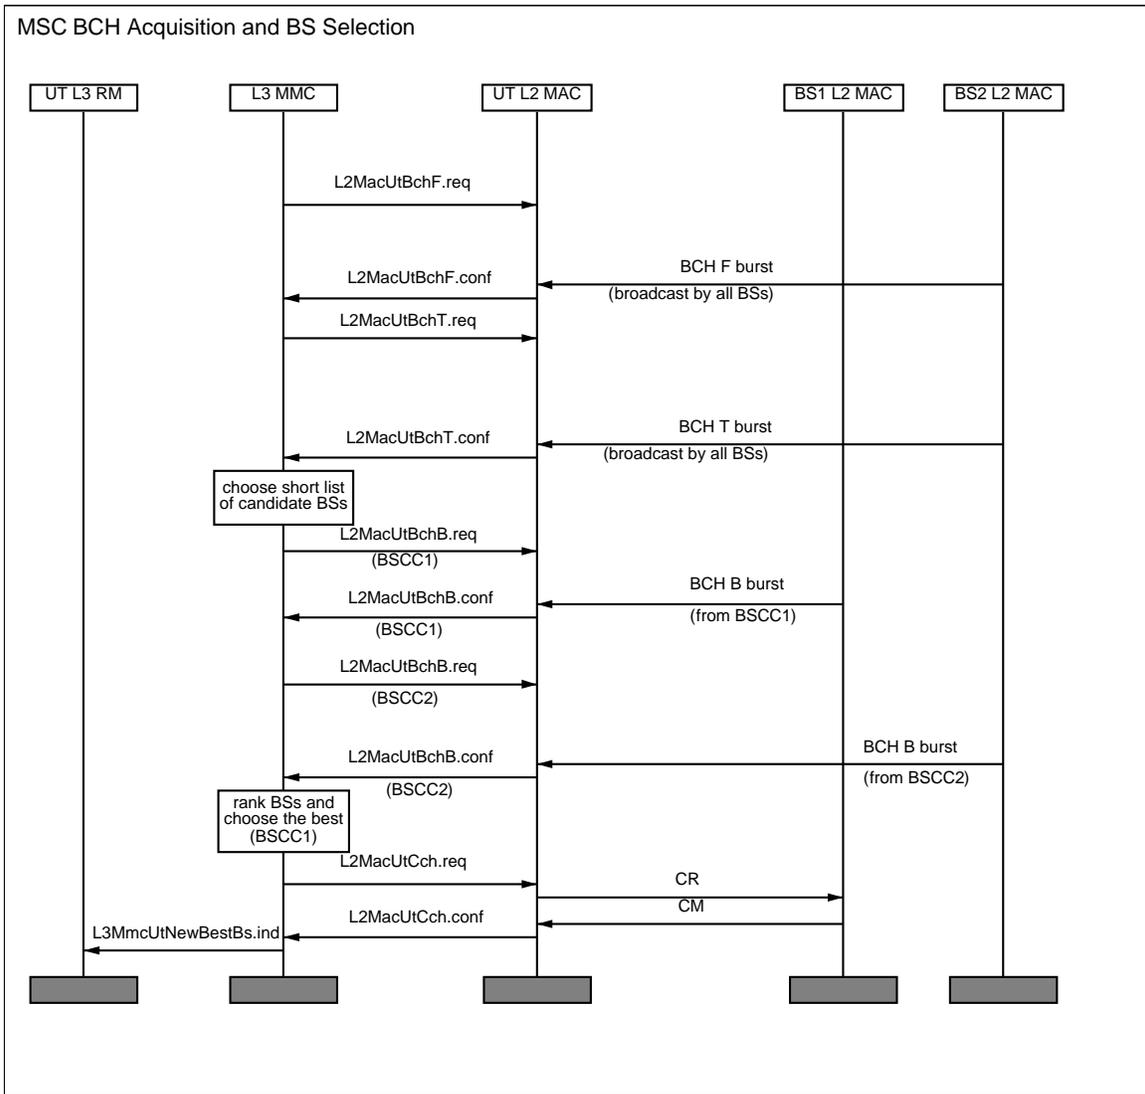


Figure 49: BCH Acquisition and BS Selection

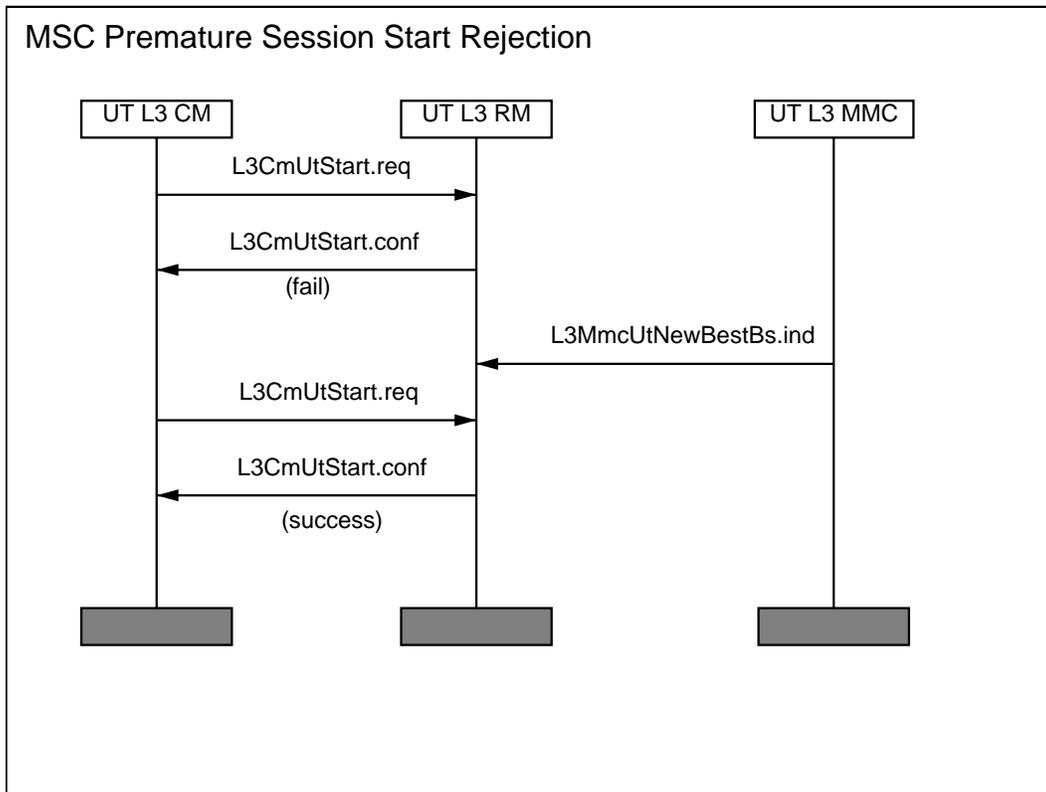


Figure 50: Premature Session Start Rejection

9.3.3 Session Start

Sessions shall start from the UT side with a session start request (`L3CmUtStart.req`) from UT L4 to UT L3 CM, as described in clause 9.2.1.3.1.1 UT L3 CM Initial State should minimize the latency of the first data by registering proactively as soon as the session is started or may minimize battery usage and network resource usage by registering reactively when the first data from UT L4 arrives.

9.3.3.1 Proactive Registration

In proactive registration, the UT shall register immediately after the request for a new session. The sequence of actions and messages involved in a session start with proactive registration shall conform to Figure 51 - Session Start (Proactive Registration).

9.3.3.2 Reactive Registration

In reactive registration, the UT shall defer registration until data arrives from UT L4. The sequence of actions and messages involved in a session start with reactive registration shall conform to Figure 52 - Session Start (Reactive Registration).

9.3.3.3 Registration Progression

When the BS hears the RA-rreg, BS L2 MAC shall inform BS L3 RM of the new registration request with an `L2MacL3RmBsStart.ind`, noting that the prospective stream is a registration stream. If BS L3 RM allows the stream to open, it shall respond by sending the `L2MacL3RmBsStart.resp` primitive back

to BS L2 MAC. With the registration stream now open, the L3 RM peers shall perform the registration message exchange handshake. Upon completion, UT L3 RM shall inform UT L3 CM with the `L3RmUtStart.conf` primitive that the requested registration was successfully created, and BS L3 RM shall inform BS L3 CM with the `L3RmBsStart.ind` primitive that a new registration exists. BS L3 CM can accept the new registration with the `L3RmBsStart.resp` primitive. UT L3 CM shall then send the `net_sess_start` message to the BS L3 CM peer. Upon receipt, BS L3 CM shall send the `L3CmBsStart.ind` primitive to BS L4, and shall wait for the `L3CmBsStart.resp`. Assuming a positive response, the session is fully established, as is shown in Figure 51 - Session Start (Proactive Registration) and Figure 52 - Session Start (Reactive Registration).

In contrast, as shown in Figure 80 - Session Rejection a negative response from BS L4 shall result in the new session being removed.

Concurrent to the establishment of the TCH stream to send the `net_sess_start` message, the registration stream shall terminate with the `L2MacL3RmStop.ind` notification from L2 MAC to L3 RM on both the BS and UT.

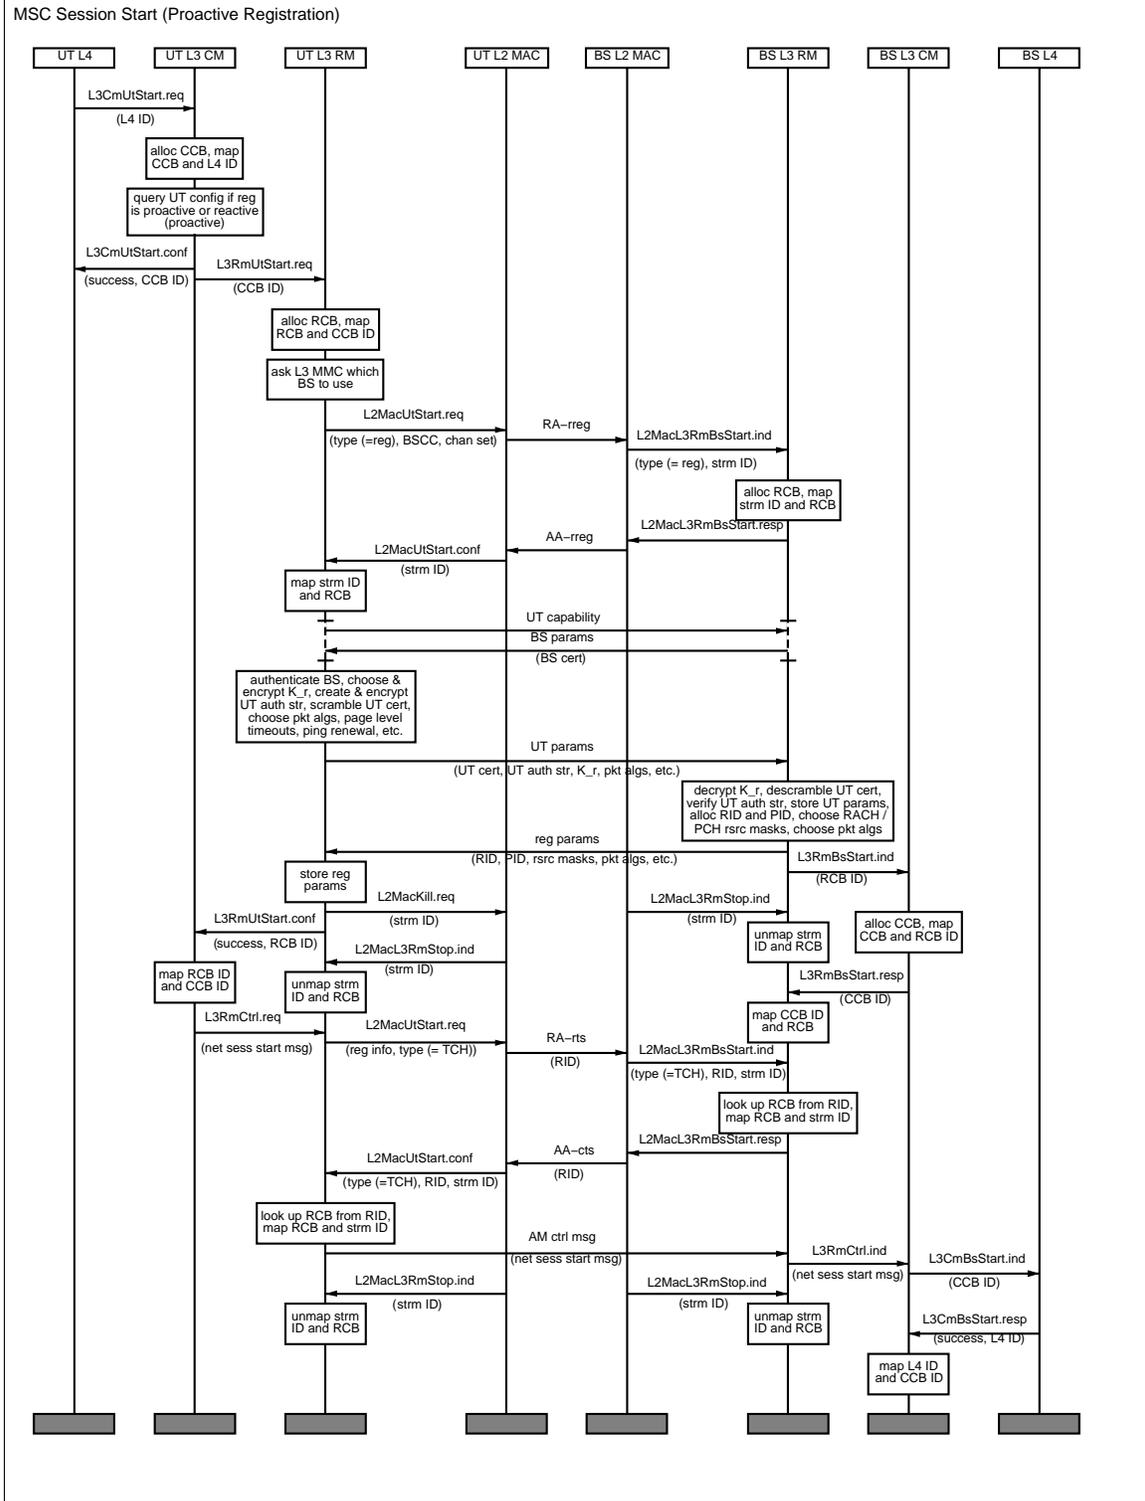


Figure 51: Session Start (Proactive Registration)

MSC Session Start (Reactive Registration)

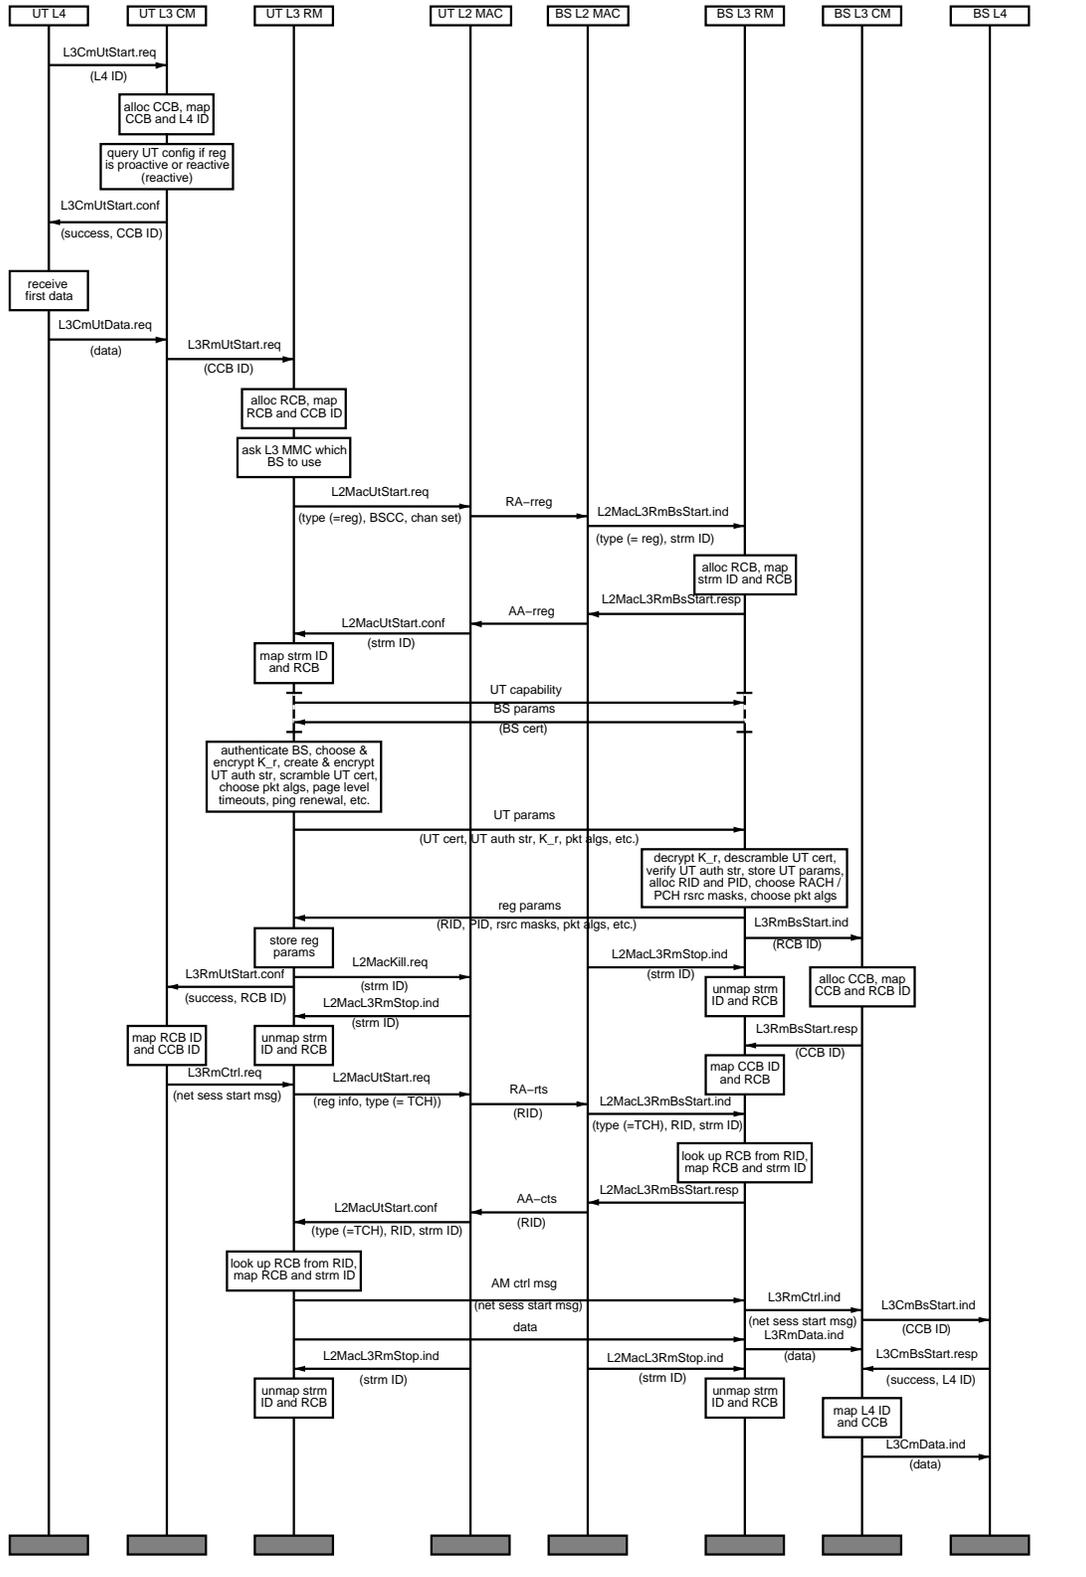


Figure 52: Session Start (Reactive Registration)

7.3.5.1.2.2, AA subtype definitions). After waiting $\max(1, \text{raDelay})$ frames, the UT shall send an RA-rts (on the same conventional channel as the AA-pending), and the BS shall send an AA-cts granting the UT a stream on the assigned channel.

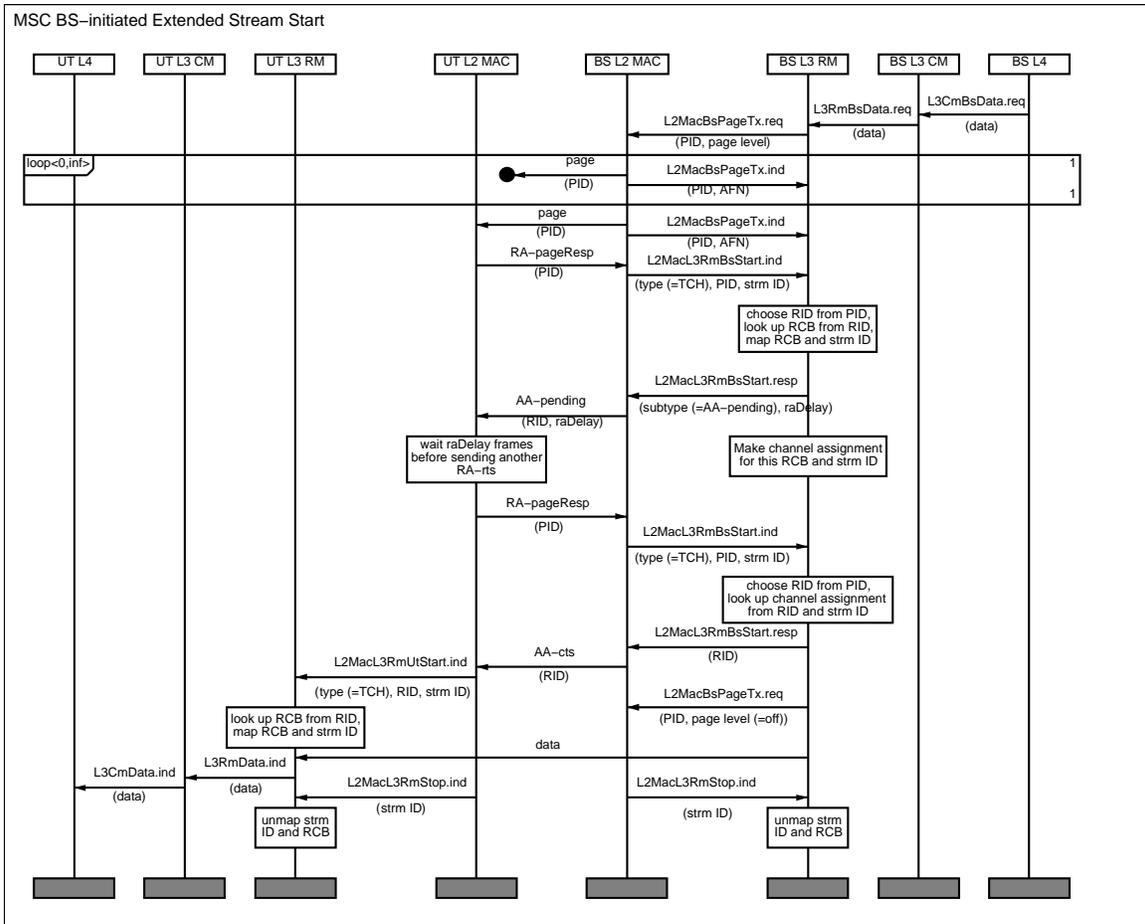


Figure 54: BS-initiated Extended Stream Start

9.3.4.1.2 Paging States and Transitions

The frequency of pages from BS L2 MAC for a particular PID shall be determined by L3 RM. There shall be discrete levels of paging activity for PIDs:

- ◆ Attempt a page every frame.
- ◆ Attempt a page every eighth frame.
- ◆ Attempt a page every sixty-fourth frame.
- ◆ Attempt a page every five hundred and twelfth frame.
- ◆ Give up attempting pages.

The paging activity level shall be maintained regardless of whether the PID currently has paging enabled. When a registration closes its last active stream, the paging activity level shall be set to every frame. After a period at this high activity level, the activity level shall fall back to every eighth frame. The activity level shall later transition to every sixty-fourth frame, then every 512th frame, and finally to no paging activity.

BS L3 RM and UT L3 RM shall agree during registration how long the registration's PID will be kept at each level before falling back to the next lower level. (See clause 9.5.1.3.3, UT Params Message for information.) A duration of zero shall be interpreted as an infinite duration. When the paging activity level falls all the way to "give up", BS L2 MAC shall not send pages for this PID any more. However, BS L2 MAC may still open a stream to send the downlink data when the UT contacts the BS. See clause 9.3.5.2, Ping Promotion to TCH for information.

Both BS L3 RM and UT L3 RM shall inform their corresponding L2 MACs when to transition the paging level through the `L2MacBsPageTx.req` and `L2MacUtPageRx.req` requests, respectively. In this way, UT L2 MAC knows when BS L2 MAC will possibly send pages and BS L2 MAC knows when UT L2 MAC will be listening for pages. BS L2 MAC shall only send pages when the BS L3 MAC sub-layer within BS L3 RM has enabled paging for the PID in question, whereas UT L2 MAC shall listen for pages according to the page hopping function when no streams are in progress.

9.3.4.1.3 Paging Selection

On a given frame, when BS L2 MAC has the physical resources available to send a page, BS L2 MAC shall find which PIDs to consider for sending their paging messages on this Paging Channel (PCH). The selection of which PID gets to use a PCH shall consist of the following steps:

1. Find all PIDs that are scheduled for the conventional channel in question according to the page hop function defined in clause 7.5, Page Hopping Sequences.
2. Find which of these PIDs have paging enabled.
3. Find which of these PIDs will be listening for pages on the current frame; that is, find the PIDs for which the following condition is true:

$$\text{afn \% (paging activity level) == PID \% (paging activity level)}$$
4. Based on fairness and quality of service, choose one of the qualified PIDs to page.
See clause 7.3.3.2, Base Station Page Procedures for full PCH details.

BS L3 RM shall assign the same PID for all registrations from the same UT. When BS L3 RM assigns the same PID for multiple RIDs and when the BS receives the RA-page-response from the UT, the BS shall determine which registration within the UT to grant the stream to, and thus which RID to put into the AA-cts. When BS L3 RM has two or more registrations that want to send pages, and they share a PID, BS L3 RM shall choose which registration gets the stream opened in response to paging. After this initial stream is open, the BS may open other streams for the other registrations by sending in-band page messages on the first registration's stream. See clause 9.3.4.1.5, In-band Page for more details.

9.3.4.1.4 Stream Aggregation (Downlink Backlog)

When the data backlog in BS L3 RM is large, BS L3 RM may choose to open additional streams to increase the registration's throughput and decrease the data backlog.

When BS L3 RM chooses to open more streams, it shall send `aggregate.cmd` UM control messages to UT L3 RM.

Examples of the stream aggregation procedure are shown in Figure 55 - Stream Aggregation (Downlink Queue, Stream Start) and Figure 56 - Stream Aggregation (Downlink Queue, New Packet Arrival). In both of these cases, the `aggregate.cmd` UM control message is depicted as a directed command that specifies a conventional channel for the UT's new stream.

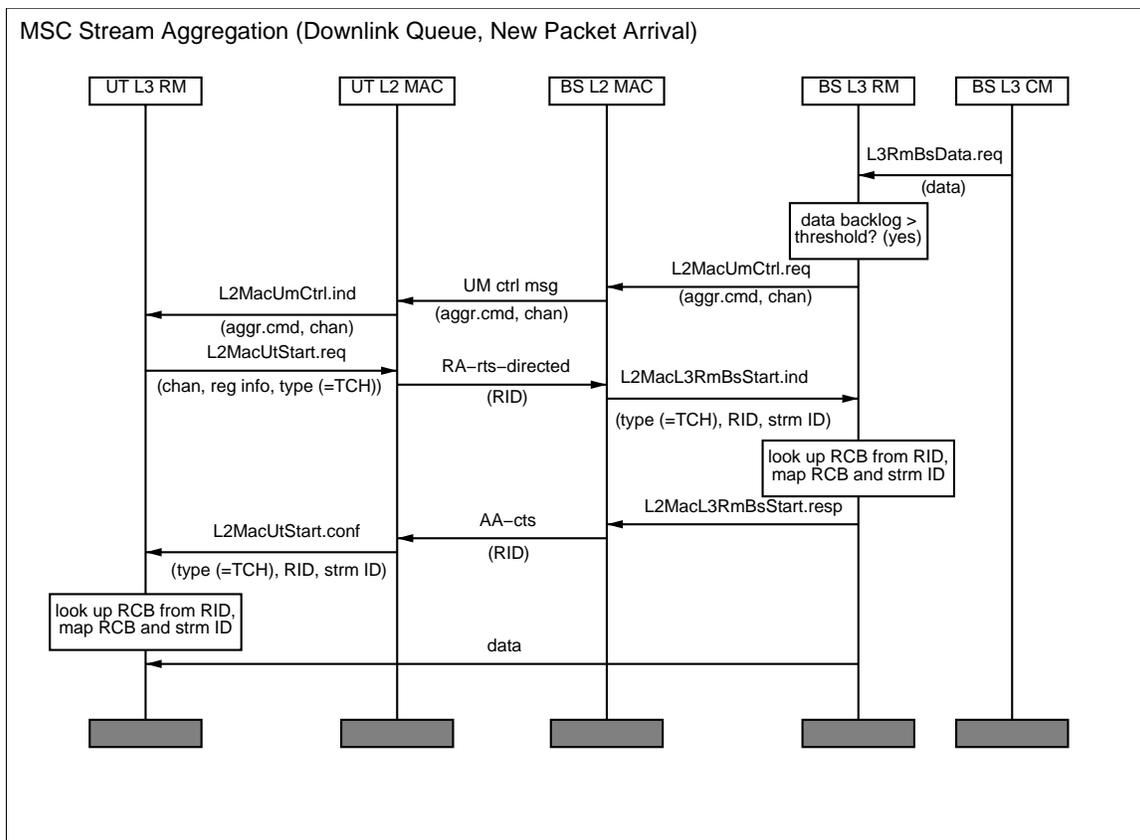


Figure 56: Stream Aggregation (Downlink Queue, New Packet Arrival)

9.3.4.1.5 In-band Page

When there are multiple registrations to the BS from the same UT, the BS shall use in-band page message on the streams of one registration to notify other registrations to open streams to receive downlink data.

The in band page message may be directed or undirected. Figure 57 - In-band Page illustrates the case where the in band page message is depicted as not carrying a specification of what channel to use for the new stream.

The BS shall construct the in band page AM control message containing the RID of the idle registration, and shall send this message on the active registration. Upon receiving this message, the active registration shall notify the idle registration, which shall then start trying to open a stream.

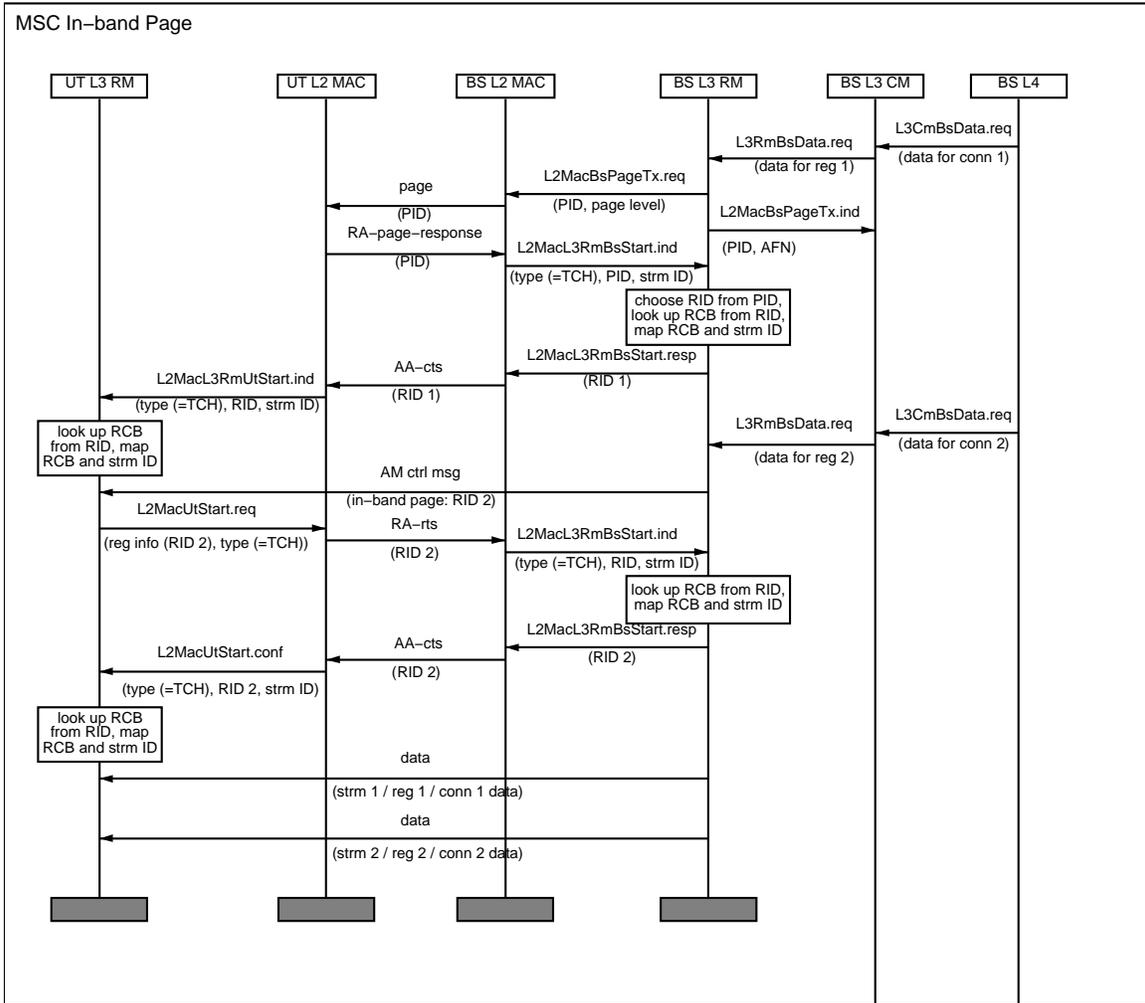


Figure 57: In-band Page

9.3.4.2 UT Side

The steps taken when UT L3 RM opens a stream are similar to the steps taken when BS L3 RM opens a stream, except that no page is sent and the RA type is an RA-*rts* rather than an RA-*page-response*. Rather than enable paging, UT L3 RM shall tell UT L2 MAC to enable stream requests or the specified registration. This messaging handshake for the UT opening a stream shall conform to Figure 58 - UT-Initiated Stream Start.

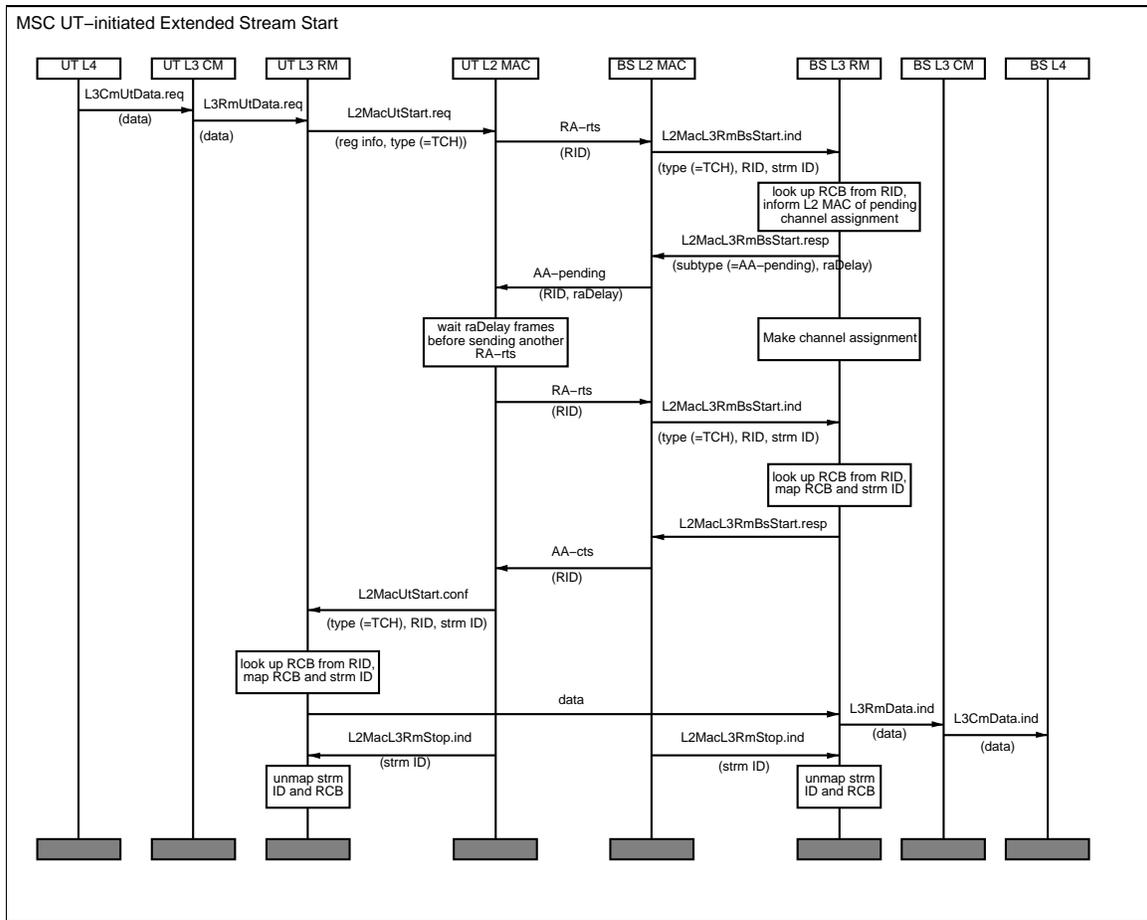


Figure 59: UT-initiated Extended Stream Start

9.3.4.2.2 Stream Retry Strategy

On the UT, each time a stream that is requested from UT L2 MAC by UT L3 RM fails to start, UT L2 MAC shall give UT L3 RM an `L2MacUtStart.conf(failure)` confirmation. When UT L3 RM receives this message, it shall determine whether it should retry. The delay before the retry shall have a capped exponential growth, according to the following formula:

$$\text{frames of delay} = \min(200 * 2^{(\text{num consecutive failures} - 1)}, 6400)$$

This delay growth is plotted in Figure 60 - UT L3 RM Stream Retry Backoff. When the number of consecutive failures exceeds a threshold, UT L3 RM shall give up and sends an `L3RmStop.ind` indication to UT L3 CM. There should be different failure thresholds for attempting registration streams rather than for attempting TCH streams. The TCH stream no-start threshold should be higher than the registration stream threshold (because UT L3 RM must work at least as much to keep a registration alive as it does to start a registration).

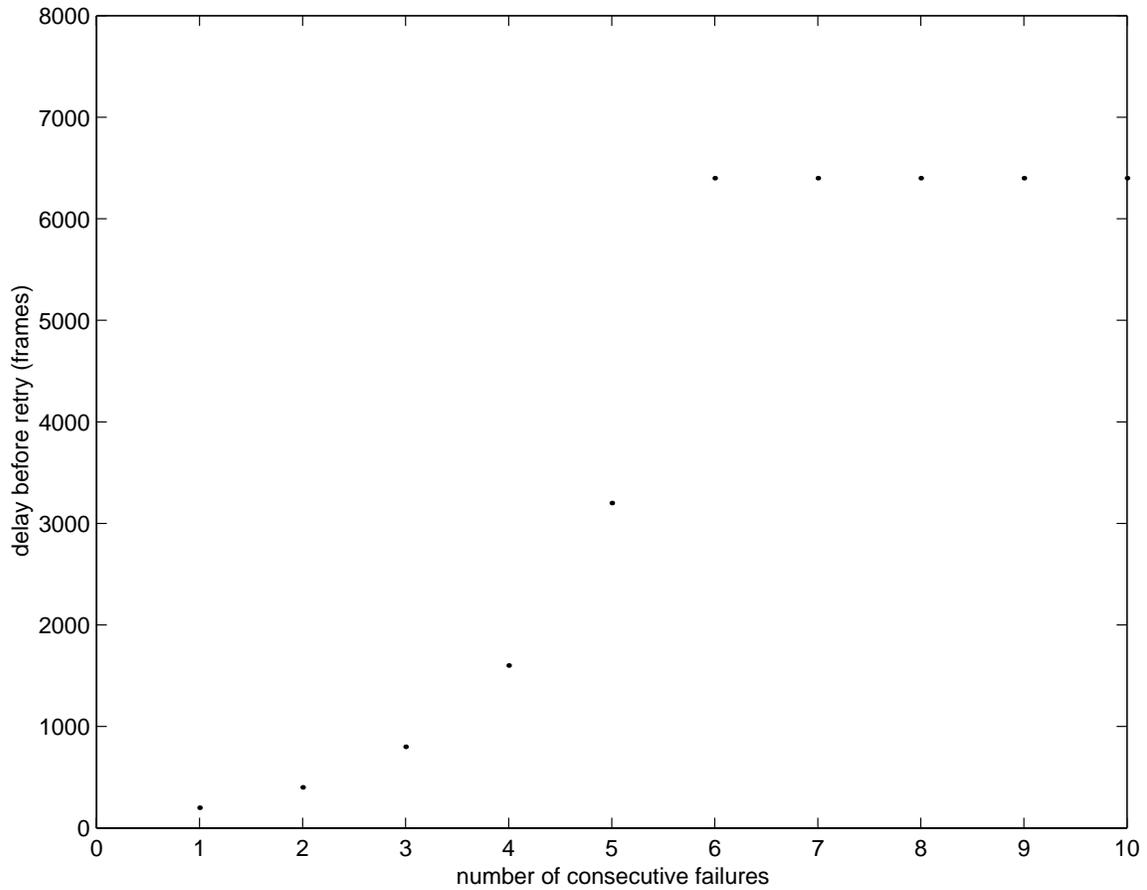


Figure 60: UT L3 RM Stream Retry Backoff

The plot in Figure 61 - UT L3 RM Stream Retry Limits shows the cumulative time as consecutive stream failures increase up to the retry limit, with the registration stream retry limit = 20 and TCH stream retry limit = 100.

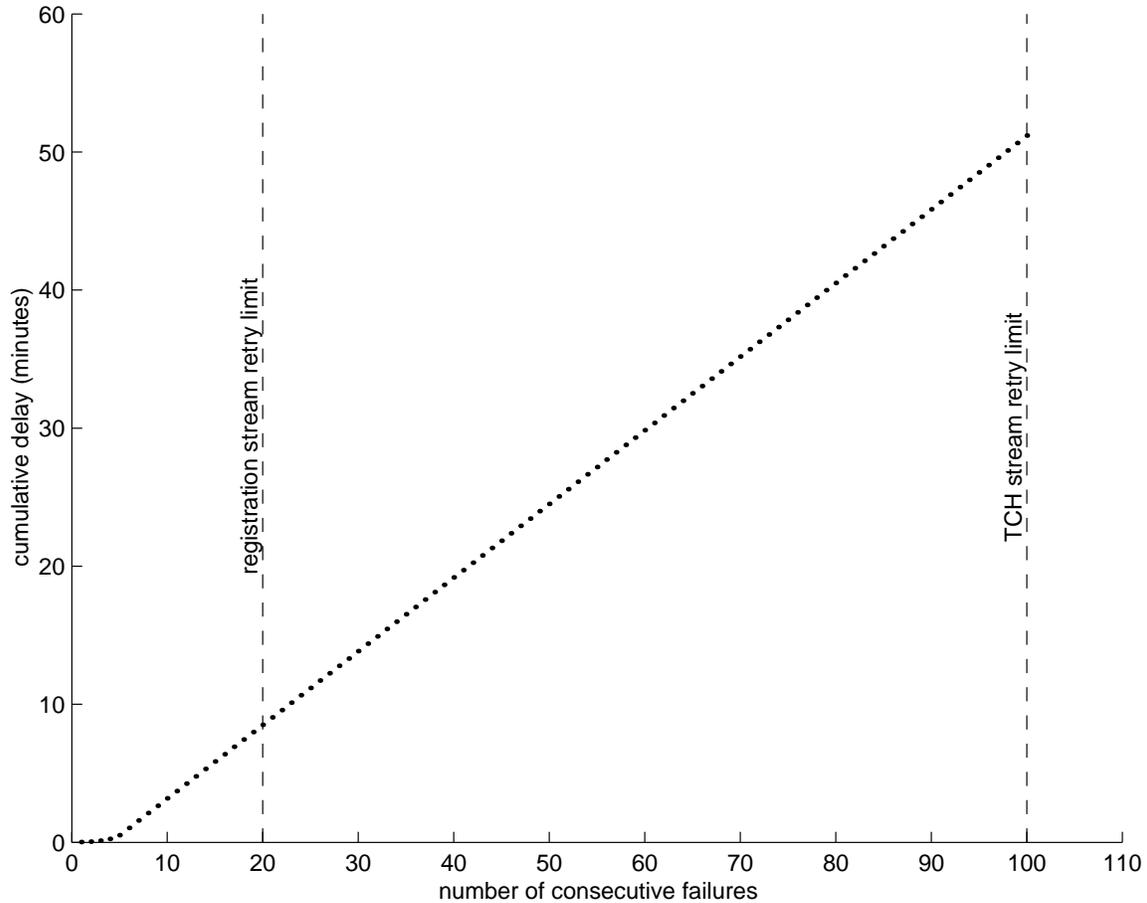


Figure 61: UT L3 RM Stream Retry Limits

9.3.4.2.3 RA Latency Messages

The BS can use knowledge of how long it is taking for the UTs to open streams to gauge how heavily loaded the RACH is. If the BS can infer the RACH load based on the stream start latencies, then it may adjust the RACH load to provide a suitable tradeoff between access latency and loading efficiency.

The BS has two methods for adjusting the RACH load. The BS may use the `hold.cmd` UM control message to tell the UTs not to attempt to start streams by sending bursts with RA-rtts too frequently. See clause 7.4, Communication Elements for more information. In the second method, the BS may respond to the UT's RA-rtts message with the AA-reject message. The AA-reject message shall cause the UT to not use the RACH until a specified time has elapsed, unless the BS pages the UT in the meantime. See clause 7.3.5.1.2, AA field definitions for more information.

The BS may use the RA latency request AM control message to ask for a stream start latency measurement from the UT. When the UT receives this message, it shall respond with an RA latency report message, as shown in Figure 62 - RA Latency Messages.

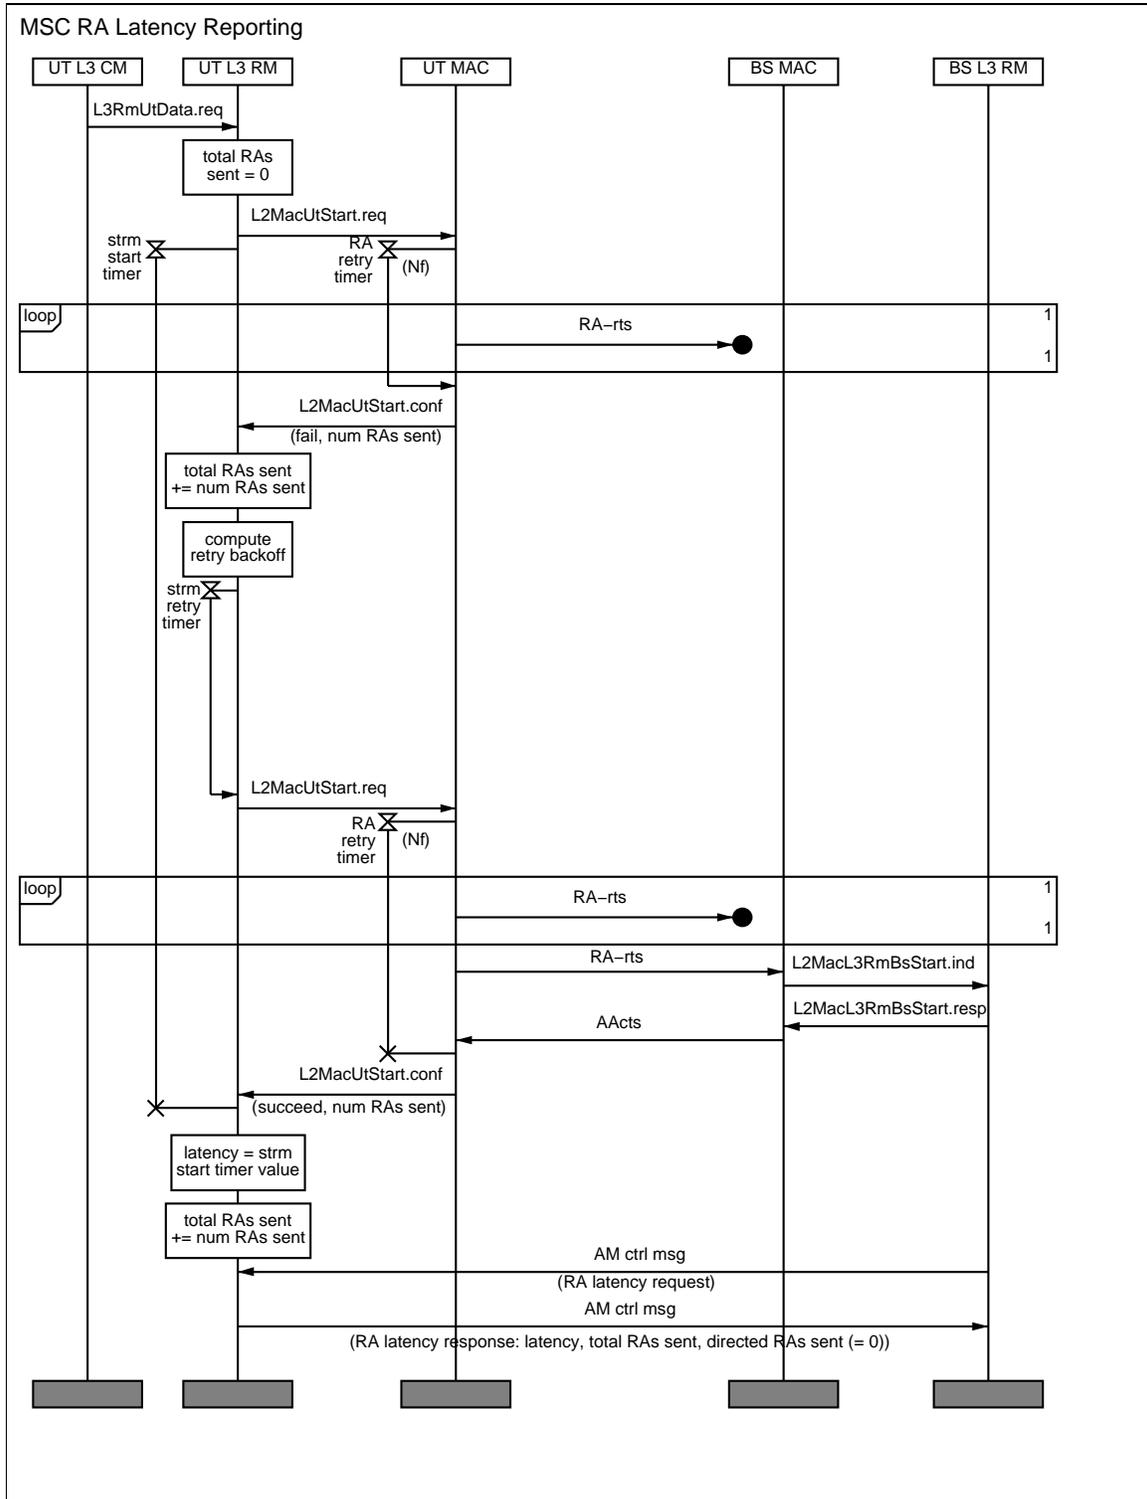


Figure 62: RA Latency Messages

9.3.4.2.4 Stream Aggregation (Uplink Backlog)

When UT L3 RM receives new data, it should send a queue . rpt UM control message to BS L3 RM to tell the BS approximately how large the uplink data backlog is. (UT L2 MAC may decide when to deliver the queue . rpt UM control messages. When there has been no queue . rpt for a long time, or

when the `queue.rpt`'s data changes significantly, the `queue.rpt` UM control message's priority should be high, otherwise it should be low.) See clause 7.3.7.2.3, UM Message Insertion for information.

When BS L3 RM gets the `queue.rpt` UM message from BS L2 MAC, it shall check whether the reported queue depth is large enough to allow opening additional streams to increase the registration's throughput and decrease the uplink queue depth.

When BS L3 RM opens additional streams, it shall send an `aggregate.cmd` UM control message to UT L3 RM. The messaging involved in stream aggregation due to an uplink data backlog is illustrated in Figure 63 - Stream Aggregation (Uplink Queue).

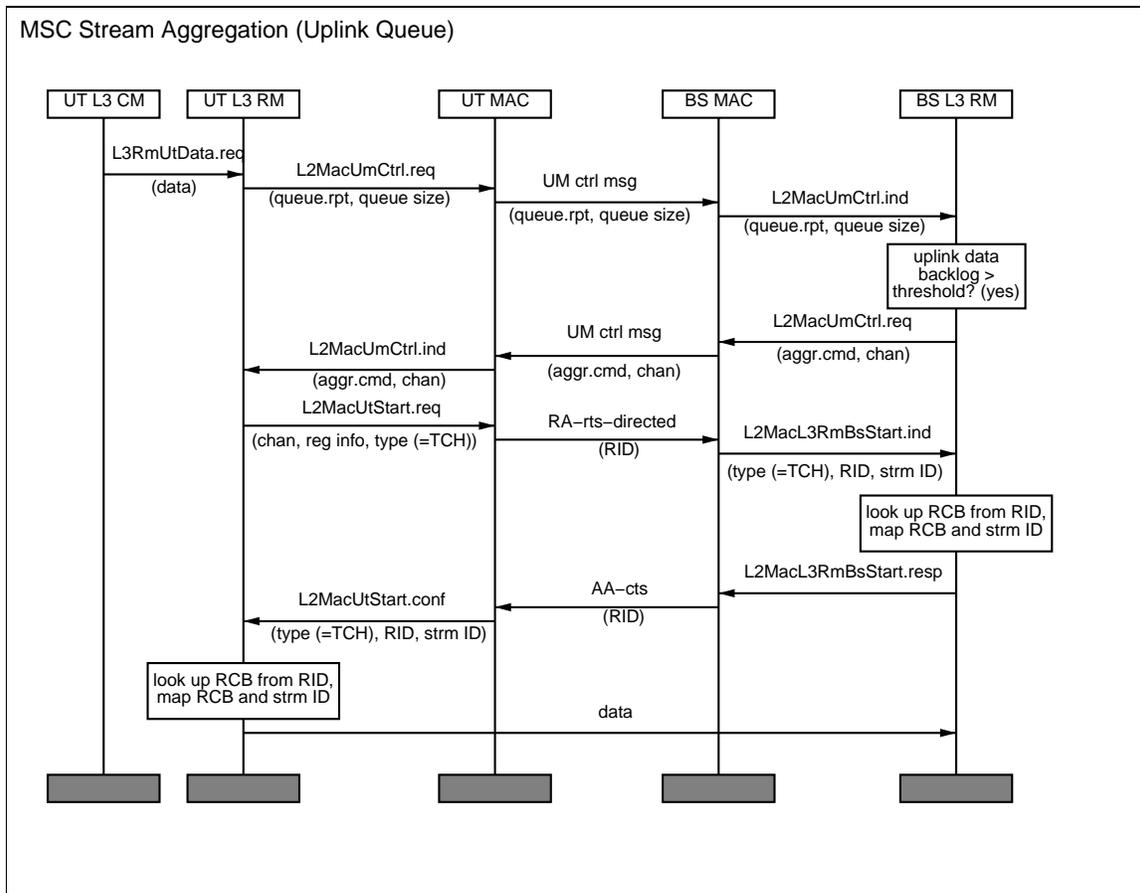


Figure 63: Stream Aggregation (Uplink Queue)

9.3.4.2.5 Resource Map Message

The BS shall initially tell the UT which conventional channels it may use to send RA-rtts during the registration handshake's `reg params` message. (See clause 9.5.1.3.4, Reg Params Message for a description of the `reg params` message.) To later change which conventional channels the UT uses to send RA-rtts messages, the BS shall send the `rsrc map` message to inform the UT of the change. (See clause 9.5.1.3.8, Rsrc Map Message for a description of the `rsrc map` message.)

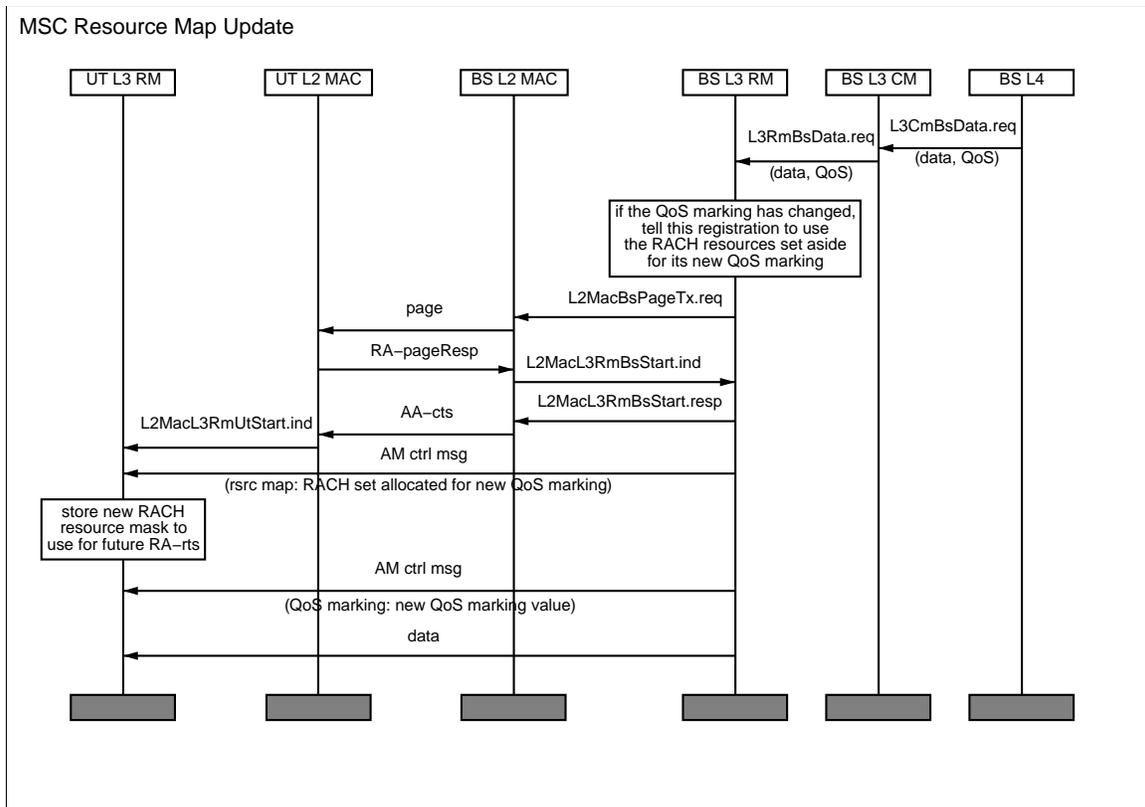


Figure 64: Resource Map Message

Figure 64 - Resource Map Message provides an a case when the BS updates the list of conventional channels the UT may use for sending RA-rtts messages. In this example, the BS is told by the network to change the session’s quality of service (QoS). As part of the new QoS, BS L3 RM allows the registration to use a different set of RACH resources (for example, a set of channels that is only lightly loaded), and uses the *rsrc map* message to inform the UT of this change. Note that the BS concurrently delivers the QoS L3 AM control message to the UT due to the registration’s change in QoS. (See clause 9.5.1.3.12, QoS Message for a description of the QoS message.)

9.3.4.2.6 Intra-UT Stream Conflicts

The BS L3 MAC sublayer of BS L3 RM shall arbitrate when to break existing streams to make room for sending pages on the downlink and for hearing RA-page-resp and RA-rtts on the uplink, so that other users can have their appropriate share of the total stream resources. However, when there are multiple connections per UT and one registration (reg 1) for one connection may be blocked from opening a stream by an existing stream for the registration (reg 2) for the other connection, the UT may use the *co reg queue report* AM message to inform the BS that reg 1 needs a stream. When the BS receives this message, BS L3 MAC may choose whether to break the stream owned by reg 2 and allow reg 1 to try to open a stream on the newly freed timeslot. Furthermore, the BS L3 MAC sublayer of BS L3 RM may explicitly direct reg 1 to a free channel by sending an *in-band page* AM message on behalf of reg 1 as the final transfer on reg 2’s stream.

Figure 65 - Intra-UT Stream Conflict Resolution illustrates this process. In the figure, it is assumed that reg 2 already has a stream open that is used to transfer the *co reg queue report* and *in-band page* AM messages.

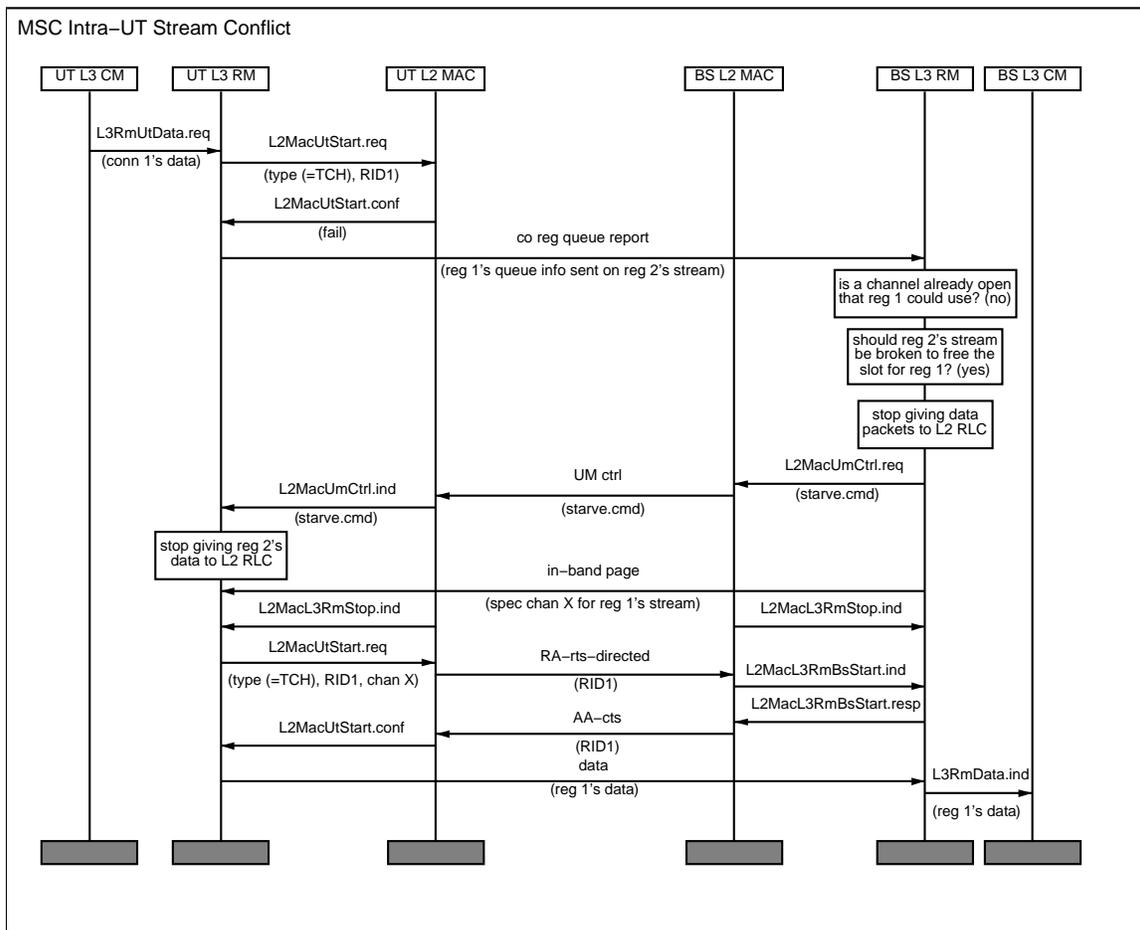


Figure 65: Intra-UT Stream Conflict Resolution

9.3.4.3 L3 initiated Stream Closure

To close a stream, L3 RM shall use either graceful stream closure or abrupt stream closure.

9.3.4.3.1 Graceful Stream Closure

The graceful shutdown process shall conform to Figure 66 - Graceful Stream Closure. The stream termination is shown originating on the UT side; exactly the reverse messaging sequence shall be followed when the stream termination originates on the BS side.

Graceful stream closure shall be accomplished by having the local L3 RM instance stop feeding data to the stream in question and having it send a `starve.cmd` UM control message on the stream telling the peer L3 RM instance to also stop feeding data to the stream on which the `starve.cmd` message was delivered. (See Table 84 - UM Message Definitions for a description of the `starve.cmd` UM control message.) Since neither L3 RM instance is sending new data to the L2 RLC endpoints of the stream, the L2 RLC instances eventually finish sending data. Shortly afterward, L2 MAC shuts down the stream due to lack of data transfer.

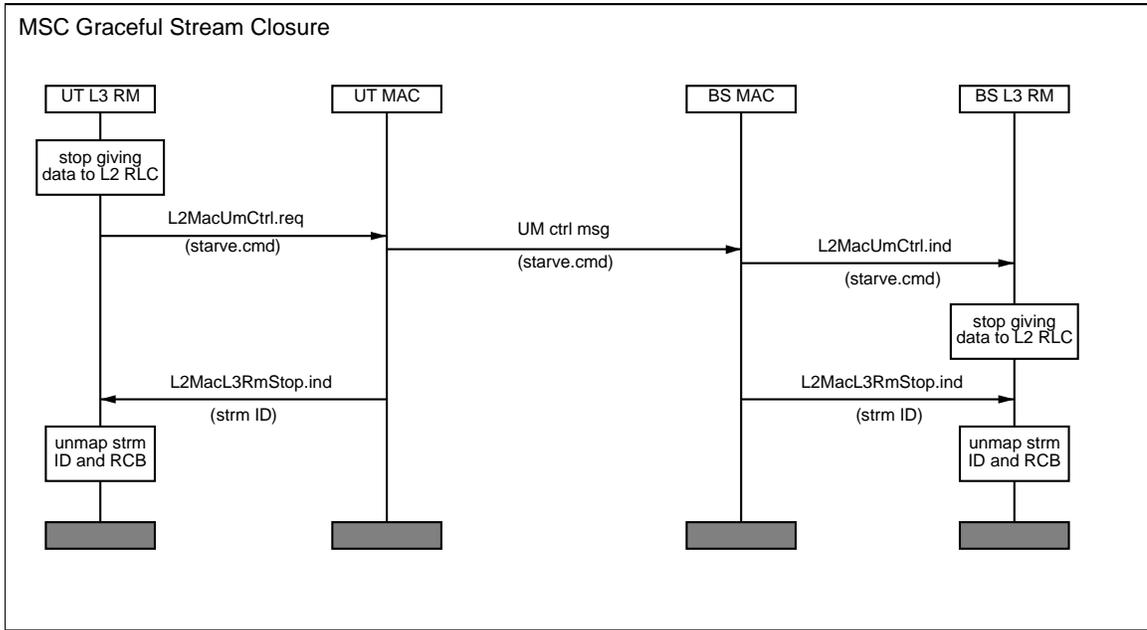


Figure 66: Graceful Stream Closure

9.3.4.3.2 Abrupt Stream Closure

Abrupt stream closure shall conform to Figure 67 - Abrupt Stream Closure. The stream termination is shown originating on the UT side; exactly the reverse messaging sequence shall be followed when the stream termination originates on the BS side. Note that stream closure does not imply that the registration owning the stream is terminated.

To close a stream abruptly, L3 RM shall give an `L2MacKill.req` request to L2 MAC. The local L2 MAC instance shall shut down the stream within a frame, although it shall send a `close.cmd` UM control message (see Table 84 - UM Message Definitions) to inform the remote L2 MAC instance that the local L2 MAC is closing the stream. If the remote L2 MAC instance receives the `close.cmd` UM control message, it shall close the stream immediately. If the message is lost, the remote L2 MAC instance shall wait until it can conclusively determine from consistent burst detection failures that the local L2 MAC instance is no longer using the stream.

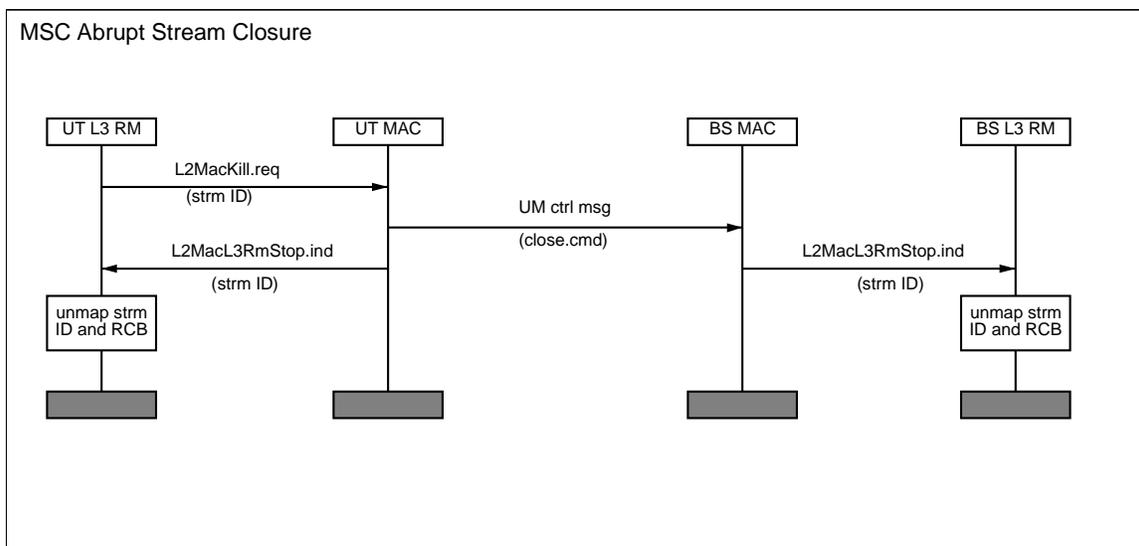


Figure 67: Abrupt Stream Closure

9.3.4.3.3 Stream Duration Limit

BS L3 MAC should limit stream durations in order to assure fairness, or ensure security (see clause 11.4.2.6, Burst Counter Test). By limiting the duration of streams, users eventually release their streams and then contend with other users who are also trying to open streams.

If BS L3 MAC limits stream durations, stream durations should be based on such factors as the user’s quality or class of service, the current load on the BS, and the nature of the data being transferred. If BS L3 MAC limits the duration of a stream, and the stream stays open for the allowed duration, BS L3 MAC should first attempt a graceful stream closure. If the stream remains open, BS L3 MAC should perform an abrupt stream closure.

9.3.5 Registration Maintenance

BS L3 RM and UT L3 RM shall have synchronized timers that tell both L3 RM entities when a registration set up between them expires. The timer duration shall be chosen by the base station and shall be communicated to the UT in the `reg params` message. This timer shall run when the registration has no streams active, shall be stopped when the registration opens a stream, and shall be reset and restarted when the last stream ends. Thus, registrations that are continually transferring data shall not expire, while those that are idle shall eventually expire. The actions taken when a registration expires are described in clause 9.3.5.3, Deregistration.

If there are multiple registrations within a UT, when any registration opens a stream the timers of all registrations (except for registrations that had the `ping renew` bit cleared in the `UT params` registration message) shall be reset. The timer for each registration which had the `ping renew` bit cleared shall be restarted when the last stream belonging to that registration closes. The timers for registrations which had the `ping renew` bit set shall be restarted when the UT as a whole closes its last stream.

9.3.5.1 Registration Renewal (Ping)

When a registration is not transferring data, it may avoid expiring by sending pings (RA-ping/AA-ping exchanges) that tell both BS L3 RM and UT L3 RM to reset the registration expiration timer.

The process of renewing the registration’s timer shall conform to Figure 68 - Registration Renewal/Ping.

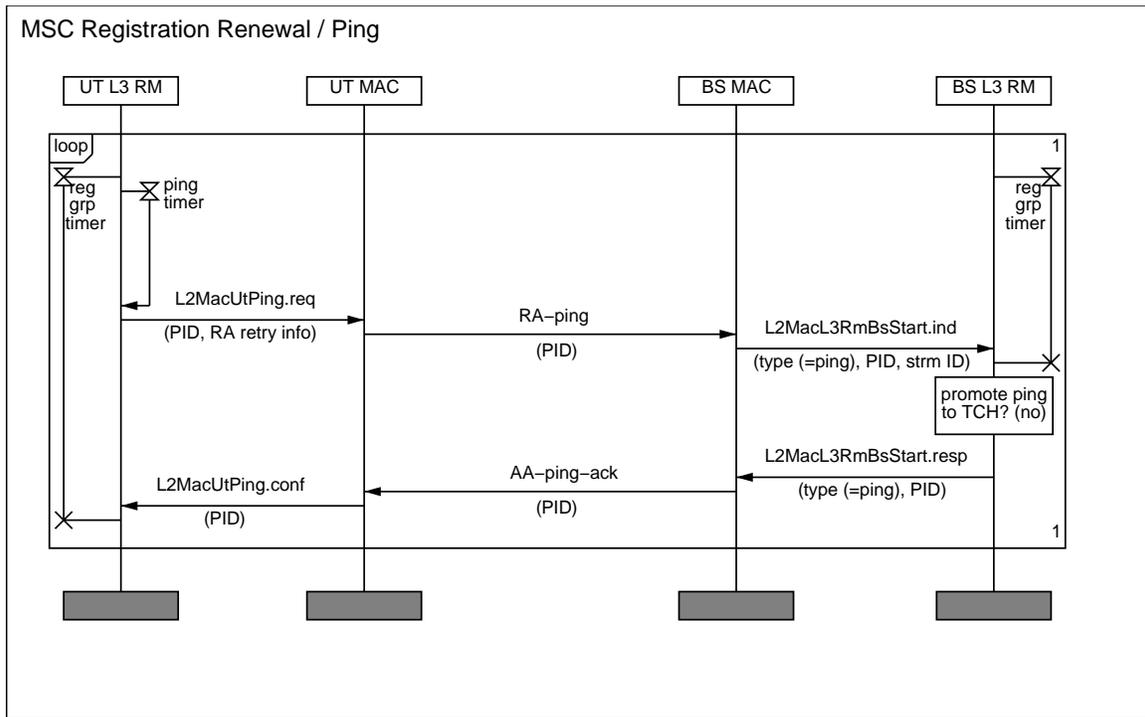


Figure 68: Registration Renewal/Ping

When UT L3 RM receives an L2MacUtPing.conf (success) confirmation from UT L2 MAC or BS L3 RM receives an L2MacBsPing.ind indication from BS L2 MAC, it shall reset the registration’s timer. There is no need to wait for an L2MacL3RmStop.ind, because a ping stream consists only of an RA/AA exchange. All registrations on the UT shall be renewed by the ping, except for registrations that have the ping renew bit cleared in the UT params registration message.

9.3.5.2 Ping Promotion to TCH

The messaging involved in promoting a ping to a TCH stream shall conform to Figure 69 - Ping to TCH Promotion.

When the BS receives a RA-ping, it shall decide whether to use the RA/AA exchange to open a stream for one of the registrations associated with the PID in the RA-ping. If the BS tries to page a PID but gets an RA-ping from the UT corresponding to this PID before it gets the RA-page-response, it should treat the RA-ping as if it had been an RA-page-response and respond with an AA-cts. In other words, the BS L3 MAC sub-layer within BS L3 RM should promote the requested ping exchange to a TCH stream. All registrations on the UT shall be renewed by the ping promotion, except for registrations that have the ping renew bit cleared in the UT params registration message.

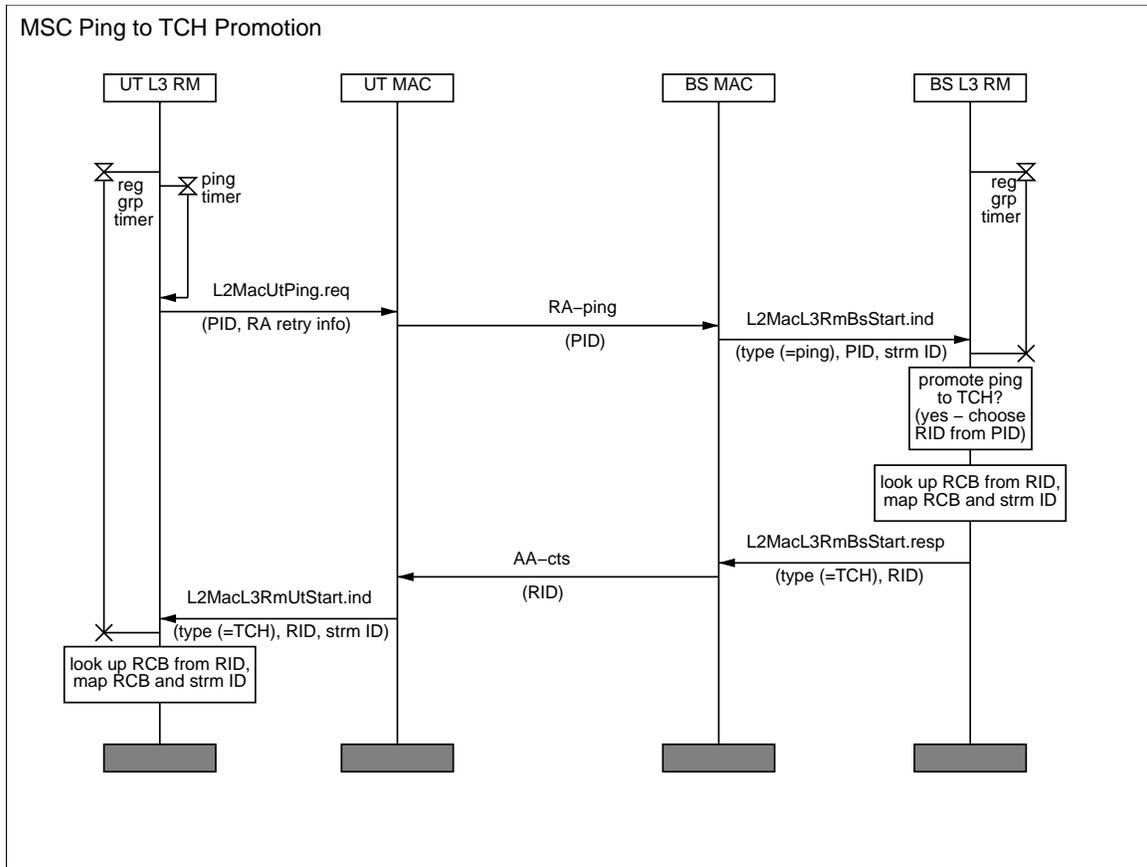


Figure 69: Ping to TCH Promotion

9.3.5.3 Deregistration

Typically, both the BS and UT are equally aware of the need to terminate a registration. As discussed in clause 9.2.2.5.1.4, UT L3 RM Registered State; clause 9.2.2.5.2.4, BS L3 RM Registered State; and clause 9.3.5, Registration Maintenance; a registration is terminated in one of three cases: in response to an L3RmStop.req request from L3 CM; an expiration of the registration timer; or autonomously due to consistent RF failure. In the first case, where either BS L3 CM or UT L3 CM issues an L3RmStop.req, the L3 CM and L3 RM instances which initiate the closure use the net_sess_stop and dereg messages to inform their remote L3 CM and L3 RM peers that the connection and registration are being terminated. In the second case, the registration timer expiration, the BS's and UT's timers expire simultaneously. In the third case, the consistent RF failure, both the BS and UT will conclude that the RF medium is unusable. These cases, where the BS and UT are both aware of the need for deregistration are covered in more detail in clause 9.3.5.3.1, Bilateral Deregistration.

However, there may be some instances when only the BS or only the UT is aware of the need for terminating a registration. For example, if the BS thought the UT should handover, it would do this by suggesting to the UT that it deregister. Another example occurs if the UT becomes aware that it will not transfer any data for an extended period of time, for example, during an extended sleep period. In this case, the UT could tell the BS to deregister it. In the latter case, the UT would reregister by means of a handover when it next needed to transfer data. These cases are covered in more detail in clause 9.3.5.3.2, Unilateral Deregistration.

9.3.5.3.1 Bilateral Deregistration

The following steps shall be taken at the point of deregistration to deallocate the registration's control block. L3 RM shall check whether there is outgoing data left to send. For BS L3 RM, this is straightforward. For UT L3 RM, checking whether there is outgoing data left to send is complicated by the possibility of being in the middle of a handover.

When there is no outgoing data left for the registration to send, L3 RM shall immediately deallocate the registration's control block. When there is outgoing data left, L3 RM shall do the following:

- ◆ Shall send `starve.cmd` UM control messages to the L3 peer on any open streams.
- ◆ Shall set the registration's L3 RM state to closing.
- ◆ Shall start a timer to limit how long the registration can remain in the closing state.

When a stream starts, L3 RM shall send a `starve.cmd` UM message on the stream to the L3 peer. When a stream ends, L3 RM shall determine whether outgoing data is left to send and whether this is the last stream. If there are no more streams and no more data, L3 RM shall deallocate the RCB. Finally, if the timer expires, L3 RM shall kill any open streams and refuse to accept any new streams. After the timer expires, L3 RM shall wait until the last stream ends and shall deallocate the registration even if there still is outgoing data left.

9.3.5.3.2 Unilateral Deregistration

If the BS or UT unilaterally decides that the registration needs to be terminated, it shall inform the peer through the `dereg` AM control message. The `dereg` message shall take the form of a command if the mandatory bit is set, or shall take the form of a suggestion if the mandatory bit is cleared. If the message is sent as a command, the sender shall immediately proceed with the registration deallocation steps presented in clause 9.3.5.3.1, Bilateral Deregistration. The peer shall also start registration deallocation as soon as it receives the message. An example of the `dereg` message being used as a command is provided in Figure 70 - Deregistration Command. In this example, the UT chooses to deregister itself to save power during an extended idle period.

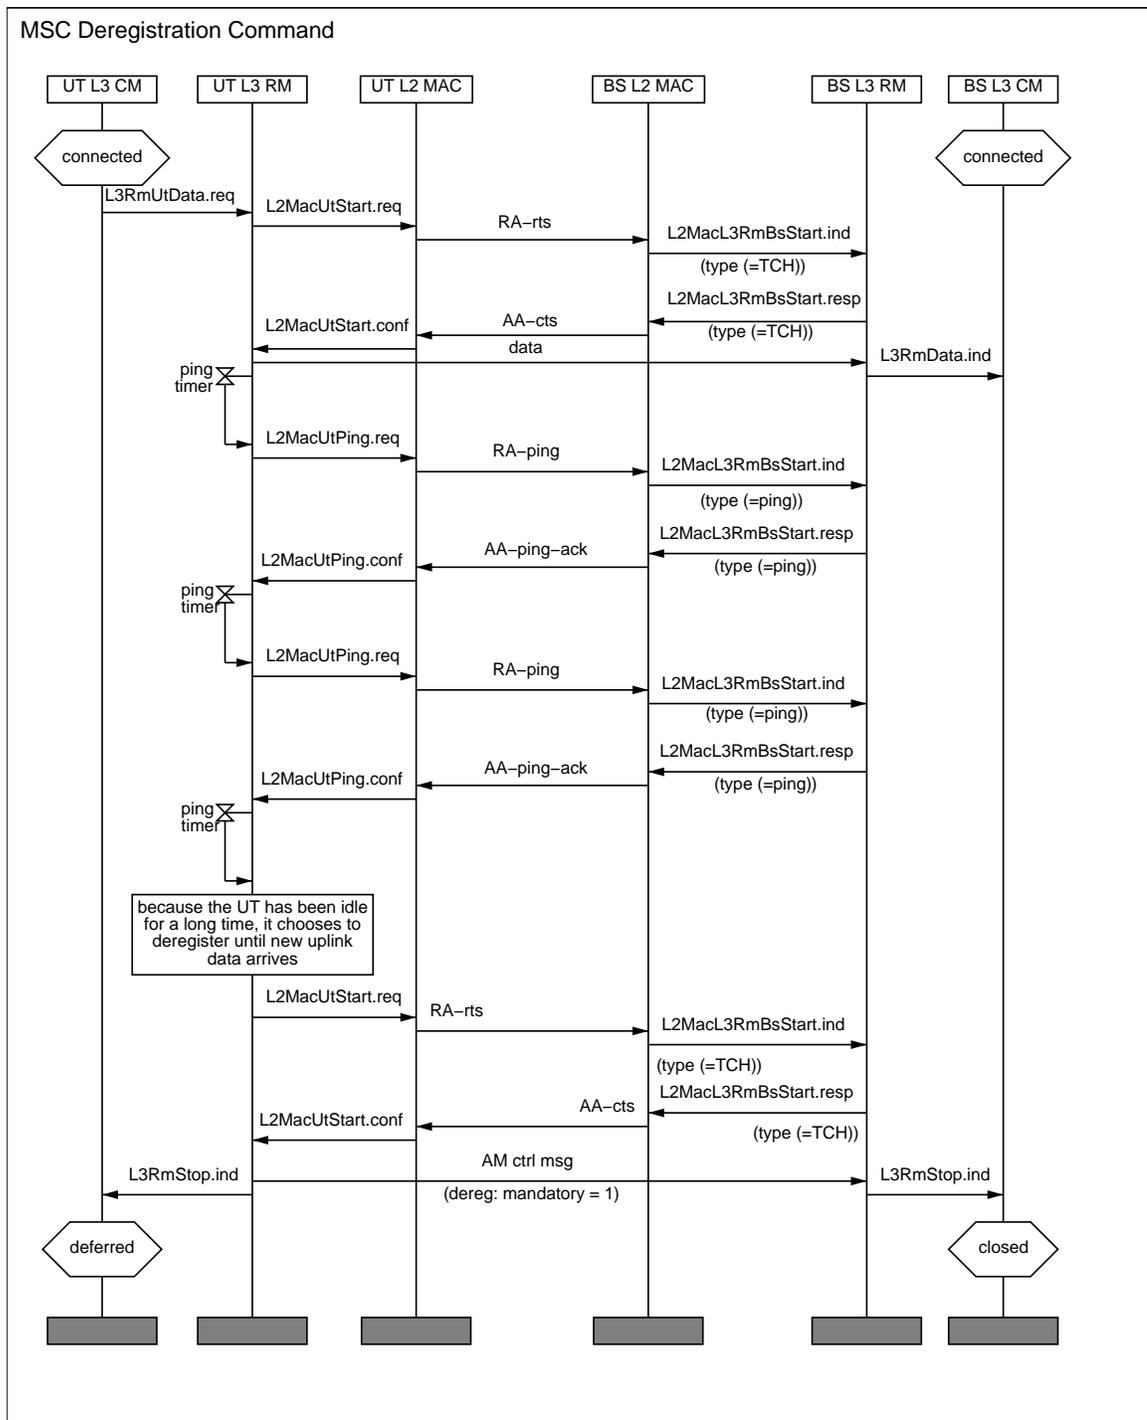


Figure 70: Deregistration Command

If the message is sent as a suggestion, then the sender may not follow through with the deregistration until some acknowledgement is received from the peer.

Figure 71 - Deregistration Suggestion demonstrates using the `dereg` message as a suggestion. In this example, the BS decides that it has to remove at least one old registration in order to make room for more new registrations. The BS suggests to the UT that owns the old registration that it deregister. The UT is configured for reactive registration -- i.e., it is permissible for the connection to have no

registration until the time when data is next transferred. In this case, the UT acknowledges the BS's deregistration suggestion by echoing the dereg message with the mandatory bit set.

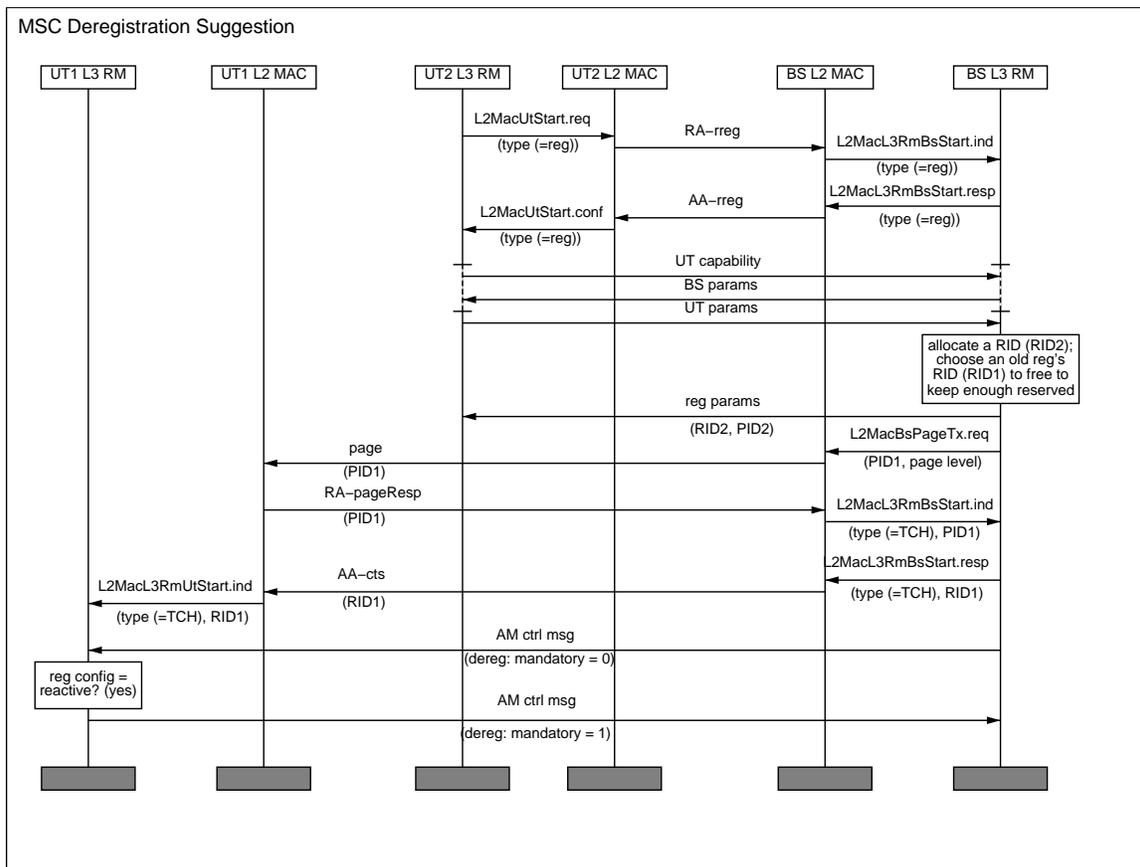


Figure 71: Deregistration Suggestion

9.3.5.4 RID Reuse Detection

Generally, the BS and UT agree whether the RID for a registration from the UT is still valid. If the registration expires, the synchronized registration timers on the BS and UT ensure that both are aware of the registration's expiration simultaneously. If the BS decides to close the registration before it expires, the BS will send a dereg message to the UT to inform it that the registration is being closed. However, it is possible that at the point the BS decides to prematurely terminate the registration, the UT is not accessible, for example due to temporarily being out of RF contact. The BS shall attempt to deliver the dereg message to the UT, but shall eventually give up and proceed with the deregistration if the UT never responds to the BS's pages. The BS may reassign the RID from the prematurely terminated registration to a new registration on a new UT. It is possible that the original UT will thereafter come back into RF coverage, and attempt to use its registration again, being unaware that the BS has terminated the registration. In this situation, the old UT will be able to open streams, but the BS will think they are from the new UT's registration rather than the old UT's registration. This situation is known as "RID Reuse".

If the BS determines that RID reuse is happening, it may direct the valid registration to perform a new registration to obtain a new RID, and thus avoid the RID conflict, or the BS may choose to take no

action, and instead rely on the UT who is no longer the valid owner of the RID to detect that its RID is invalid and do a handover to obtain a valid RID.

On the UT side, if several streams opened by the UT all get promptly closed due to L2 RLC errors or packet header checksum errors, the UT should infer that its registration has been terminated by the BS, and shall send an L3RmStop.ind primitive to UT L3 CM. UT L3 CM shall in turn attempt a handover to open a new registration, which will generate a new i-SEC secret and be given a new RID. Then the UT's session can once again transfer its data through the BS using the new RID. This process of determining that the old RID is no longer valid and performing a handover to obtain a new RID is illustrated in Figure 72 - UT RID Reuse Detection.

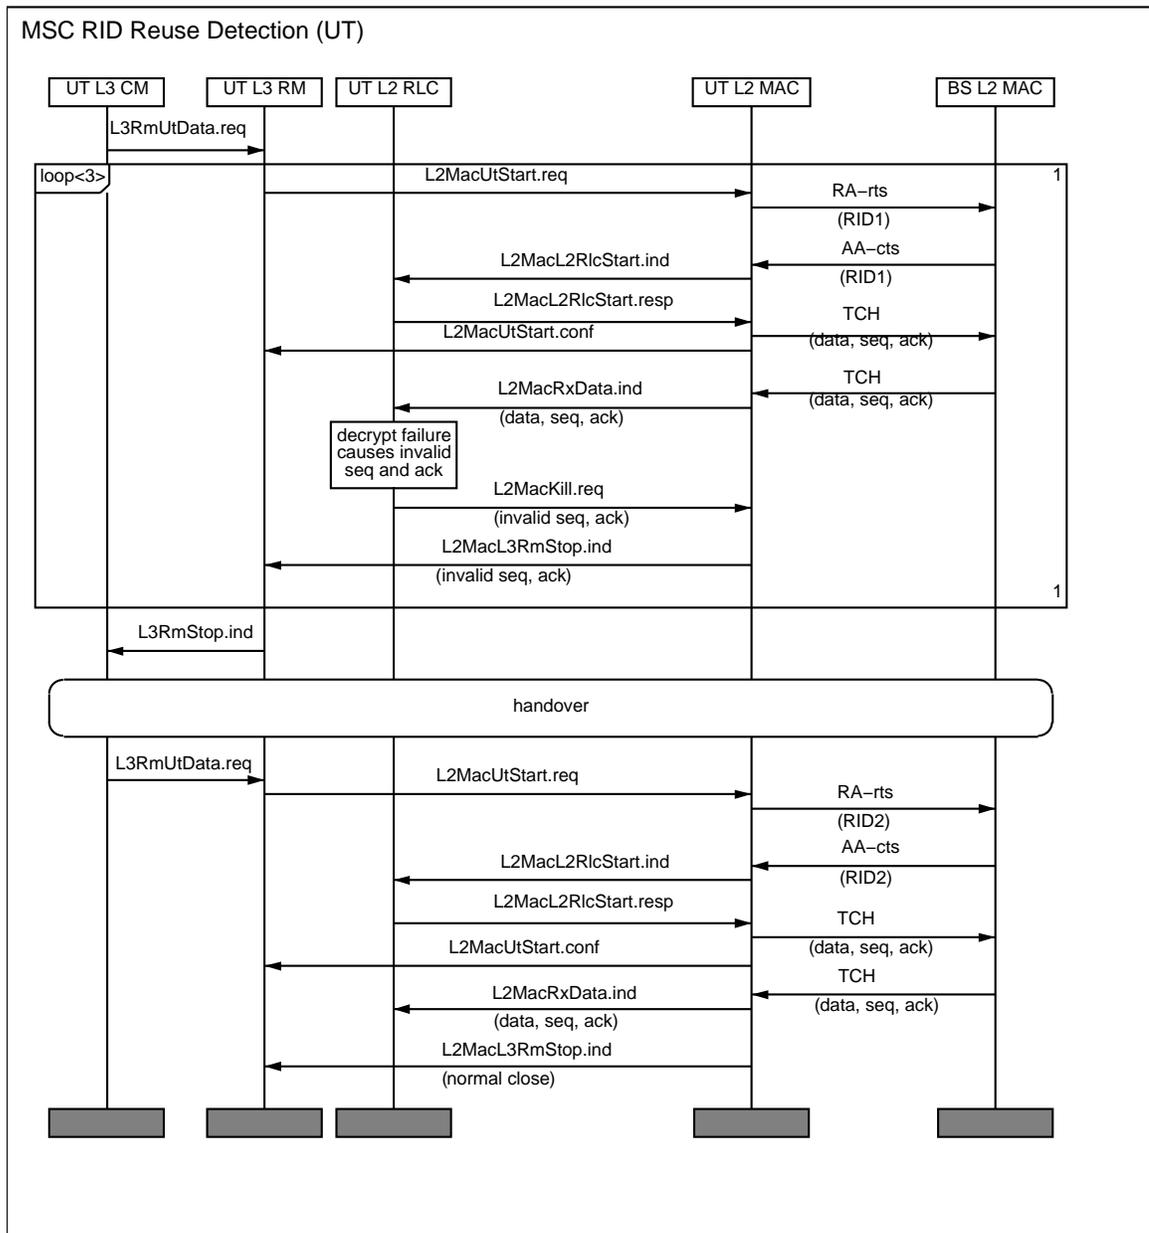


Figure 72: UT RID Reuse Detection

9.3.6 Handover

Handover shall be triggered by L3 MMC giving an `L3MmcUtNewBestBs.ind` indication to UT L3 RM. When UT L3 RM receives an `L3MmcUtNewBestBs.ind` indication from L3 MMC, it shall use the `L3MmcUtBestBs.req/conf` to find the BS recommended by L3 MMC. It then shall give an `L3RmUtHandover.ind` indication to UT L3 CM for the connections associated with each registration that are not to the best BS. UT L3 CM in turn shall give an `L3CmUtHandover.ind` indication to UT L4 for each connection for which UT L3 RM requested a handover, provided that the connection is not already in the handover process. When UT L4 allows a handover for a particular connection, it shall give an `L3CmUtHandover.resp(accept)` response to UT L3 CM.

As at the start of a session, when UT L3 CM receives an `L3CmUtHandover.resp(success)` response from UT L4, it shall choose whether to register with a new BS immediately or only when strictly necessary. UT L3 CM shall check whether the UT is configured to either minimize latency by registering proactively, or minimize battery and network resource usage by registering reactively. Both cases are described in this section.

9.3.6.1 Proactive Registration

The messaging steps for a handover with proactive registration shall conform to Figure 73 - Handover (Proactive Registration).

When the UT is configured to register proactively, UT L3 CM shall give an `L3RmUtStart.req` request to UT L3 RM as soon as it receives the `L3CmUtHandover.resp(accept)` response from UT L4. At the same time, it shall stop delivering data to UT L3 RM for that connection, though it continues to accept the connection's incoming data from UT L3 RM for the new registration.

UT L3 RM shall use the `L3MmcUtBestBs.req/conf` to ask L3 MMC which BS to register to and shall then perform the registration. When the registration is complete, UT L3 RM shall give the `L3RmUtStart.conf` confirmation to UT L3 CM. UT L3 CM shall delete the old registration by giving UT L3 RM an `L3RmStop.req` request and shall resume sending the connection's outgoing data to UT L3 RM.

9.3.6.2 Reactive Registration

The messaging steps for a handover with reactive registration shall conform to Figure 74 - Handover (Reactive Registration).

When the UT is configured to register reactively, UT L3 CM shall not give an `L3RmUtStart.req` request to UT L3 RM until it next receives outgoing data from UT L4 for the connection. At that point, the registration shall proceed as described above, with UT L3 RM asking L3 MMC which BS to register to, and UT L3 CM holding the outgoing data until the registration completes.

9.3.6.3 Registration Teardown During Handover

When the UT hands over to a new BS, the Packet Services Switch (PSS) finds that the connection through the new BS is a handover of the connection previously routed through the old BS, rather than a new connection. The PSS instructs the old BS to delete its connection. This sequence of events is shown in Figure 75 - Inter-BS Handover.

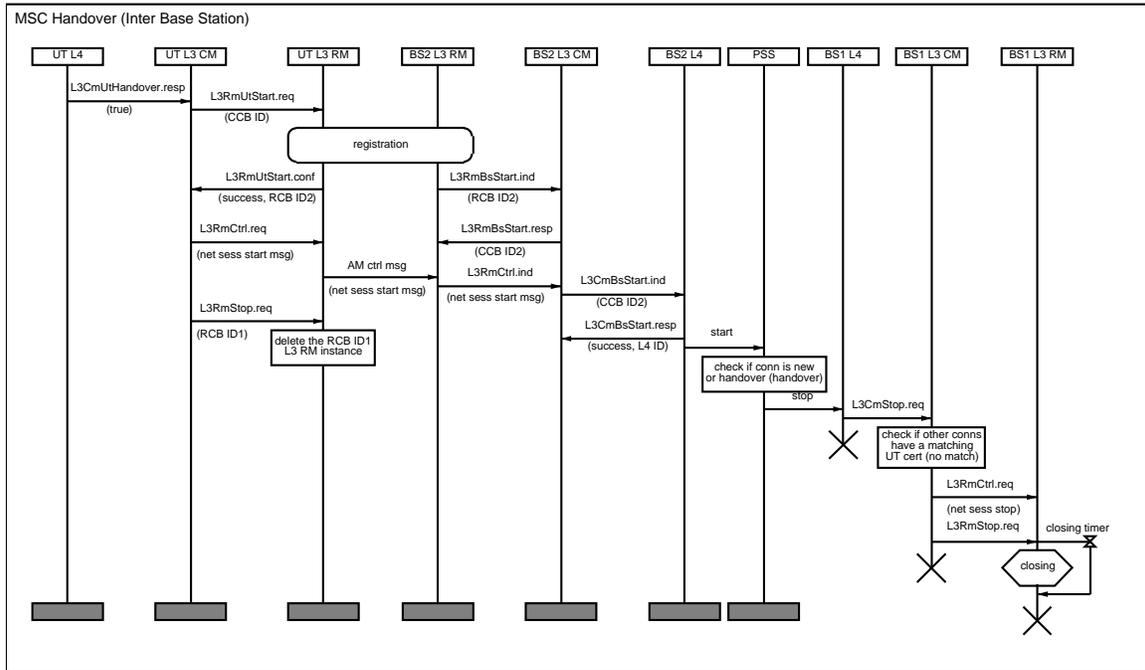


Figure 75: Inter-BS Handover

In rare cases, the UT may “hand over” within a BS, rather than from an old BS to a new BS. There are a variety of reasons this intra-BS handover could occur, such as the following:

- ◆ The UT finds that its i-SEC decryption is consistently failing, indicating either a pre-empted RID or a loss of AFN synchronization.
- ◆ The UT has opened the maximum number of streams allowed by i-HAP/i-SEC per registration (see clause 11.4.2.2, Stream Counter Test).
- ◆ The BS attempts to move the UT to a new resource partition by sending a Dereg L3 AM message and using the CM carrier/timeslot masks to direct the UT to the new resource set when the UT hands over.
- ◆ The BS has multiple independent broadcast and configuration channels (BCH/CCH), so the UT gives the BS multiple distinct scores for each BCH/CCH, and may hand over from one BCH/CCH/TCH resource set to another.
- ◆ The BS may appear to have multiple BCH/CCH as the UT goes in and out of coverage of a frequency-shifting repeater. This may cause the UT to hand over as it goes in or out of repeater coverage.

There are slight differences between the inter-BS and intra-BS handover cases. When the old connection is told to terminate, it shall check if there are any other connections with a matching UT certificate. If so, it shall infer that the UT is doing an intra-BS handover, and that it is not necessary to send a net sess

stop message to the UT before terminating the old connection and its associated registration. This sequence of events shall conform to Figure 76 - Intra-BS Handover.

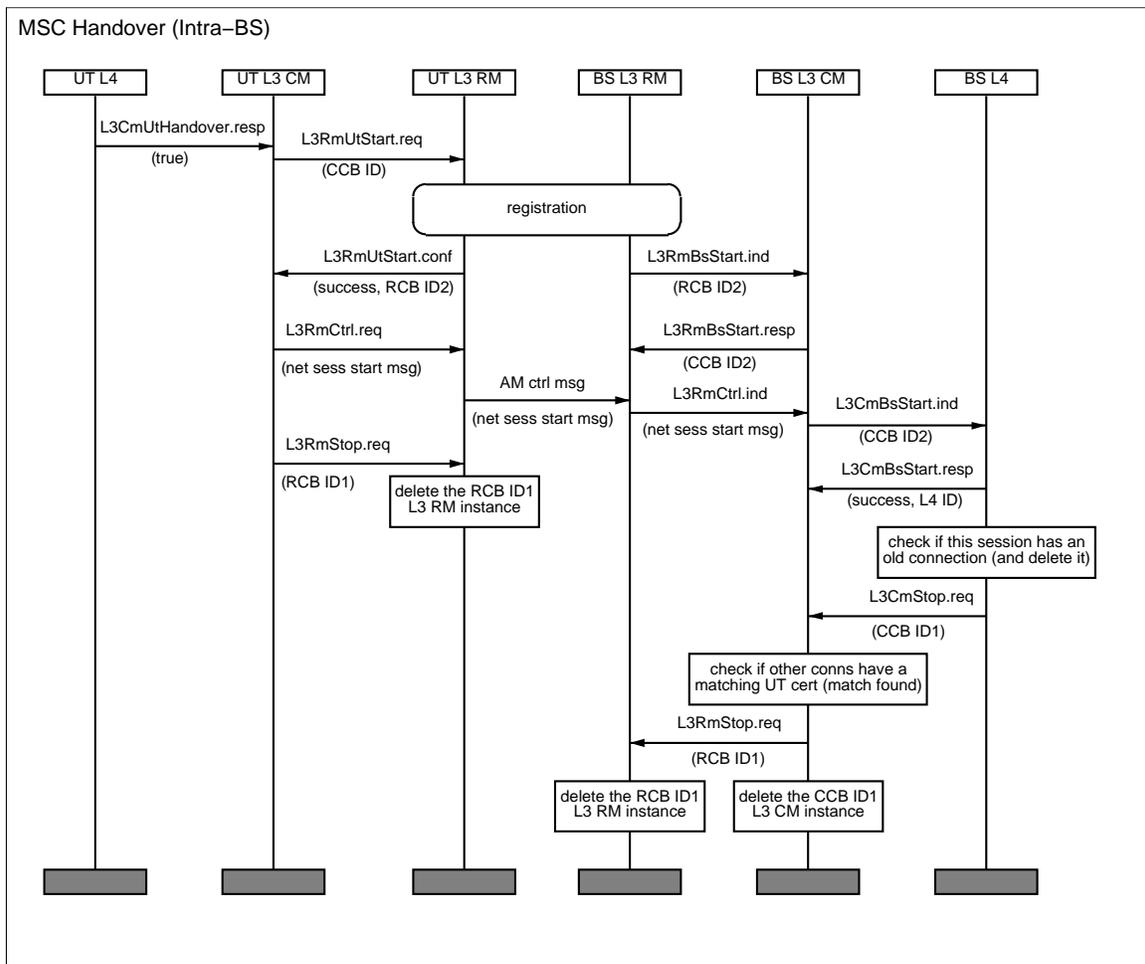


Figure 76: Intra-BS Handover

9.3.7 Session Stop

Messaging for shutdown shall conform to either Figure 77 - Session Stop (No Top-level Handshake) or Figure 78 - Session Stop (With Top-level Handshake). Figure 77 - Session Stop (No Top-level Handshake) depicts the endpoint on the UT side initiating the shutdown without sending a message to inform its peer that it is closing the connection. In this case, the remote (BS) side of the protocol stack shall be cleanly terminated.

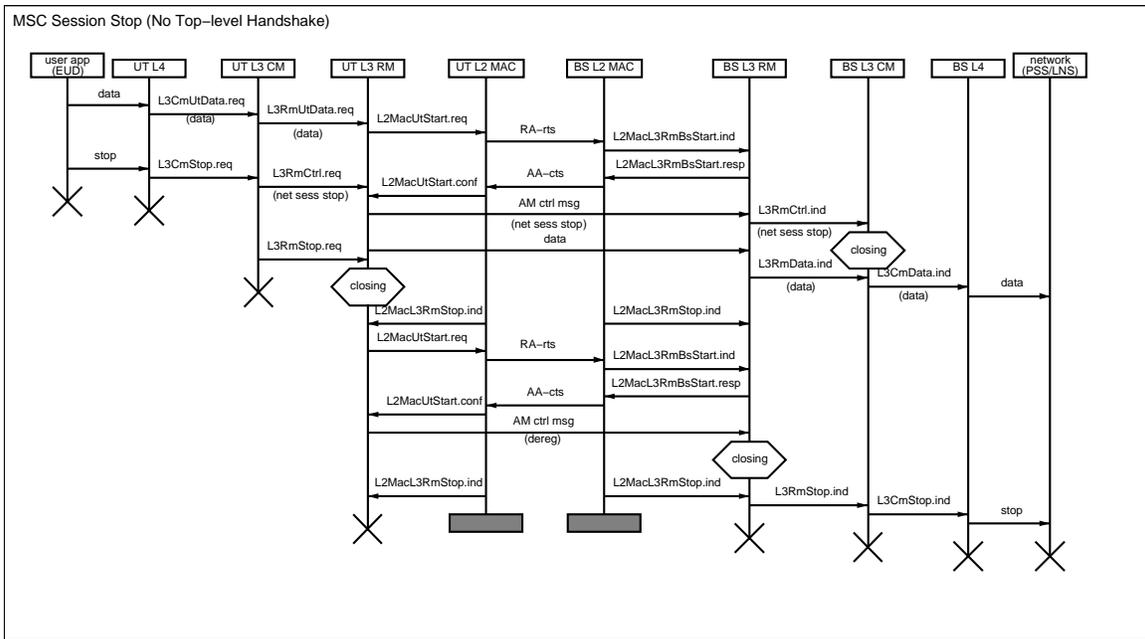


Figure 77: Session Stop (No Top-level Handshake)

Figure 78 - Session Stop (With Top-level Handshake) shows the UT side initiating the shutdown with an acknowledgement from its network peer before terminating the UT's side of the protocol stack. In this scenario, the top-level UT application starts cleaning up the UT's protocol stack top-down at the same time that UT L3 RM starts cleaning up the protocol stack bottom-up. The figure arbitrarily depicts the top-down and bottom-up cleanups meeting at the UT L3 RM level.

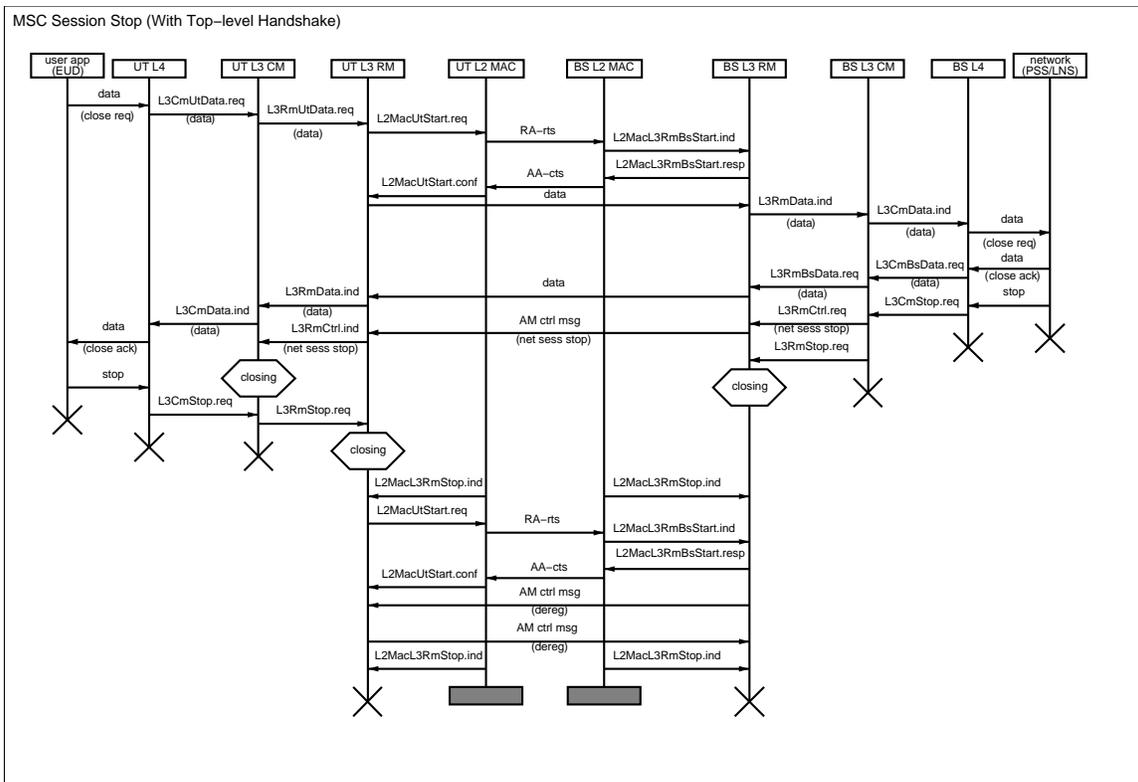


Figure 78: Session Stop (With Top-level Handshake)

In both figures, note that the choice of user-initiated shut-down versus network-initiated shut-down is arbitrary, since the shutdown messaging is symmetric.

Shutdown messaging for network initiated shutdown shall follow the symmetric duplicate of Figure 77: Session Stop (No Top-level Handshake) or Figure 78: Session Stop (With Top-level Handshake).

9.3.8 Session Abort

The event sequence for Registration Failure shall conform to Figure 79 - Registration Failure for the case of proactive registration. The case of reactive registration shall conform similarly.

As discussed in clause 9.2.2.5.1.4, UT L3 RM Registered State and clause 9.2.2.5.2.4, BS L3 RM Registered State there are various reasons for a registration being terminated by L3 RM autonomously rather than in response to an L3RmStop.req request from L3 CM.

The two most common are:

1. The RID becomes invalid because it expires or is pre-empted for a different registration.
2. There is a loss of RF connectivity between the BS and UT.

It is likely that the RID expires before L3 RM gives up on the TCH RF connectivity. However, lack of RF connectivity when UT L3 RM is trying to open a registration stream can occur (See clause 9.3.4.2.2, Stream Retry Strategy for information).

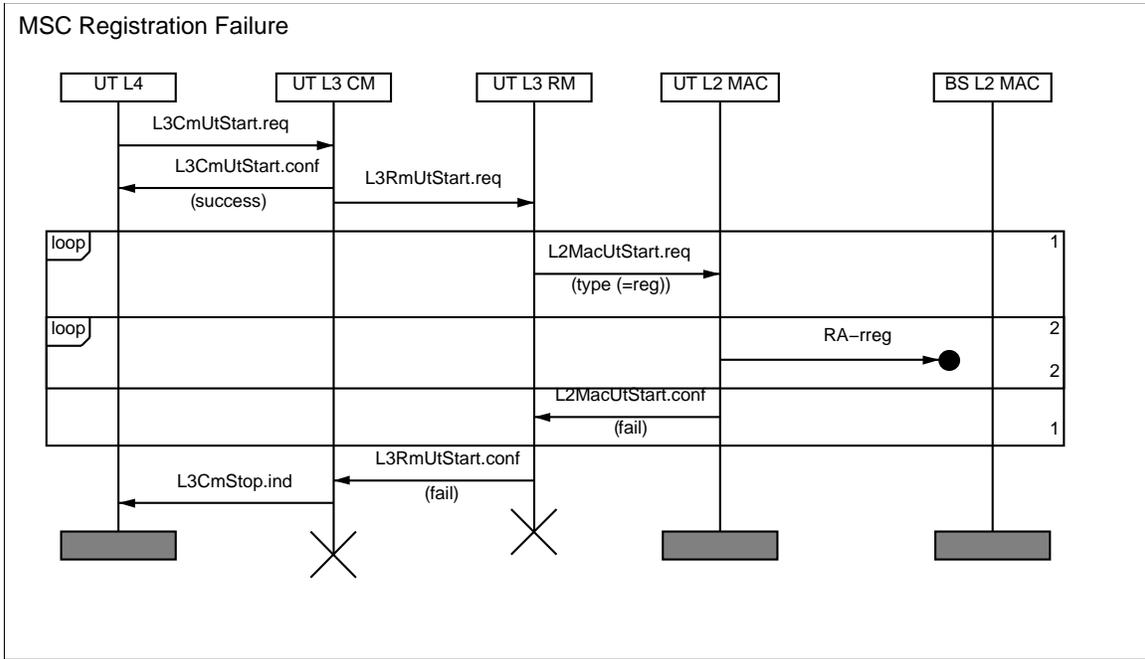


Figure 79: Registration Failure

See clause 9.3.5.3, Deregistration for a description of the steps involved in deallocating a registration's CCB.

Session Rejection occurs when a new session is rejected by BS L4 or the network (PSS/LNS/AAA). The event sequence for Session Rejection shall conform to Figure 80 - Session Rejection.

Table 92: L3 Packet Tag Bytes

L3 packet tag bytes100	byte values	added to
DATA_TAG	0xd0	data packets beginning with DATA_TAG (0xd0)or MSG_TAG (0xa7)
MSG_TAG	0xa7	all control packets
L4_MSG_TAG	0x7e	all L4 control packets
L3_MSG_TAG	0xff	L3 control packets beginning with L4_MSG_TAG (0x7e)or L3_MSG_TAG (0xff)

9.4.2 Outgoing Data Flow

The steps L3 takes to process outgoing data and control packets shall conform to Figure 81 - Outgoing Data Flow.

Outgoing (from the network to the air interface), L3 shall accept L4 data and L4 control packets and shall generate L3 control packets of its own, and shall then send them to L2 RLC.

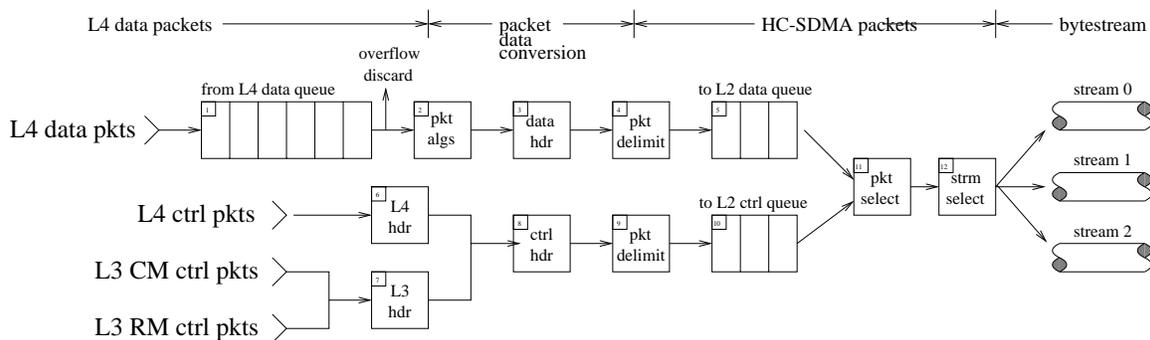


Figure 81: Outgoing Data Flow

The purpose of each of the numbered blocks in the illustration is described below:

- From L4 Data Queue:** The L4 data packets shall be queued until there are streams available to deliver them across the air interface to the peer L3. If too many packets are queued, the packets shall be discarded from the front of the queue. See clause 9.4.7, Packet Discarding for information.
- Packet Algorithms Block:** During registration, BS L3 RM and UT L3 RM agree on a set of optional packet data conversion layer algorithms to apply to the registration’s data packets. The negotiated set of algorithms shall be applied to the data packets at this point. See clause 9.4.6, Packet Data Conversion for information.
- Data Header Block:** The first byte of each data packet shall be examined to see if it is one of two special byte values used to distinguish L4 data packets from L3 and L4 control packets: MSG_TAG (0xa7) and DATA_TAG (0xd0). If the L4 data packet begins with either of these special bytes, a one byte DATA_TAG header shall be prepended to the L4 data packet to indicate that it is an L4 data packet. Thus, if the first byte is random, an L4 data packet is expanded by 1 byte exactly $2/256 = 0.8\%$ of the time.

4. **Packet Delimiting Block:** An indication of the packet boundary shall be added to the packet so that when the peer receives a data byte stream from L2 RLC, it segments the incoming data byte stream into packets. See clause 9.4.5, Packet Delimiting and Reconstruction for details.
5. **To L2 Data Queue:** By this point in the data flow, L4 packets have been fully converted into *HC-SDMA* data packets and are ready to be sent over the air interface. The *HC-SDMA* data packets shall be queued and pulled from the head of the queue when streams have room to accept them. If a stream terminates without delivering all its *HC-SDMA* data packets, the packets shall be pushed back onto the head of the queue.
6. **L4 Header Block:** When an L4 control message packet is received from L4, a one-byte `L4_MSG_TAG` header shall be prepended to it to indicate that the packet contains an L4 control message rather than an L3 control message.
7. **L3 Header Block:** If the first byte of an L3 control message packet is `L3_MSG_TAG` (`0xff`) or `L4_MSG_TAG` (`0x7e`), a one-byte `L3_MSG_TAG` header shall be prepended to the packet to indicate that it contains an L3 control message rather than an L4 control message.
8. **Control Header Block:** Each control message packet shall have a one-byte `MSG_TAG` header prepended to indicate that it is a control message packet rather than a data packet.
9. **Packet Delimiting Block:** The packet delimiting applied to control packets shall be the same as that applied to data packets.
10. **To L2 Control Queue:** At this point, the control packets have been fully converted into *HC-SDMA* control packets and are ready to be sent over the air interface. The *HC-SDMA* control packets shall be queued and shall be pulled from the head of the queue when streams have room to accept them. If a stream terminates without delivering all its *HC-SDMA* control packets, they shall be pushed back onto the head of this queue.
11. **Packet Selection Block:** If there are *HC-SDMA* control packets in the control queue, they shall be sent. If no control packets are present, *HC-SDMA* data packets from the data queue shall be sent.
 Note that these are AM control packets. In contrast, UM control messages are routed differently -- they are sent directly to L2 MAC via the `L2MacUmCtrl.req` primitive. See clause 9.4.1, Introduction for further explanation.
12. **Stream Selection Block:** Packets shall be routed onto streams based on the packet's size, the stream's queue depth, and the stream's throughput. See clause 9.4.8, Packet Routing for information.

9.4.3 Incoming Data Flow

The steps L3 takes to process incoming data shall be as illustrated in Figure 82 - Incoming Data Flow. In the incoming direction (from the air interface to the network), L3 shall accept byte streams from L2 RLC, shall form *HC-SDMA* packets from the byte streams by determining packet boundaries, shall determine whether the packet is a data packet, an L3 control packet, or an L4 control packet, and shall route the packet accordingly.

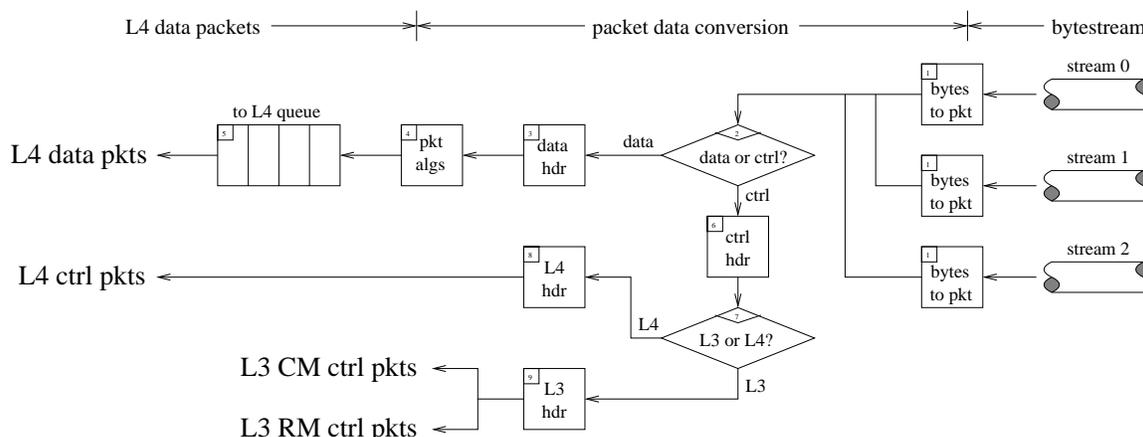


Figure 82: Incoming Data Flow

The purpose of each of the numbered blocks in the diagram is described below:

1. **Bytes to Packet Block:** L2 shall deliver a stream of bytes to L3. The first step of incoming processing shall be to break the byte stream into discrete packets based on the packet delimitation and reconstruction scheme. See clause 9.4.5, Packet Delimiting and Reconstruction.
2. **Data or Control Block:** The first byte of the packet shall be examined. If the first byte is the special byte MSG_TAG, the packet shall be routed as a control message packet, otherwise it shall be routed as a data packet.

Note that these are AM control packets. In contrast, UM control messages are received from L2 MAC through the L2MacUmCtrl.ind primitive. See clause 9.4.1, Introduction for further explanation.

3. **Data Header Block:** If the data packet begins with the special byte DATA_TAG, this header byte shall be removed from the packet.
4. **Packet Algorithms Block:** The receive part of the packet data conversion layer algorithms set agreed upon during registration shall be performed here. The receive part of these algorithms shall be the inverse of the transmit part, so the data packets coming out of the algorithms on the incoming side of the air interface are identical to the data packets going into the algorithms on the outgoing side. See clause 9.4.6, Packet Data Conversion for details.
5. **To L4 Data Queue:** Data that has gone through the receive side of the packet algorithms shall be queued until L4 is ready to receive it. If L3 CM is in the connected state, it is known that L4 is ready to receive data. However, L3 RM may deliver data across the air interface while BS L3 CM is still in the connecting state, in which case the data shall be queued until BS L4 confirms the connection.
6. **Control Header Block:** The MSG_TAG header byte shall be removed from the packet.
7. **L3 or L4 Block:** If the first byte of the control message is the special byte, L4_MSG_TAG, the packet shall be routed as an L4 control message; otherwise it shall be routed as an L3 control message.
8. **L4 Header Block:** The L4_MSG_TAG header byte shall be removed from the packet.
9. **L3 Header Block:** If the L3 control message packet begins with the special byte L3_MSG_TAG, this header byte shall be removed from the packet.

9.4.4 Control vs. Data

As shown in the preceding data flow diagrams, the optional PDCL algorithms shall be applied to data packets but not control packets. Every control message (L3 or L4) fits into a single HC-SDMA packet.

9.4.5 Packet Delimiting and Reconstruction

The reliable transfer provided by the L2 RLC (data link) layer operates on bytes (without a notion of how sequences of bytes are grouped into packets). Therefore, the sending L3 shall provide a mechanism for the receiving L3 to recognize where packet boundaries are in the incoming byte stream it receives from L2. The packet delimiting and the distinction of data packets from control packets as described in this section shall constitute the mandatory subset of L3’s packet data conversion layer (PDCL).

There are three *HC-SDMA* packet header formats. The header format shall be selected based on the packet length. The header formats shall conform to Table 93 - Header Formats. The initial bits (MSBs) of the first byte shall indicate which header format is used, the final bits shall indicate the length of the packet, and the middle bits shall provide a checksum of the length (by negating the length bits or some portion thereof). The table shows the actual bits used to indicate header format and uses “1” and “n” to mark the bit fields that specify the length and length negation (checksum), respectively.

Table 93: Header Formats

	Format 1	Format 2	Format 3
Packet Length	1-16 bytes	17-144 bytes	1-2048 bytes
Header Format	hdr[0]=0nnnllll	hdr[0]=10nnnnnnn, hdr[1]=nlllllll	hdr[0]=11nnnnnnn, hdr[1]=nnnnnlll, hdr[2]=llllllll
Length Bitfield (lenBits)	pkt len - 1	pkt len - 17	pkt len - 1
Checksum Bitfield (chkBits)	~(lenBits >> 1)	~lenBits	~lenBits
Checksum Comparison	~chkBits == lenBits >> 1	~chkBits == lenBits	~chkBits == lenBits

In order to support transfer of L4 packets larger than 2048 bytes, the packet fragmentation and recombination algorithm must be turned on so that large L4 data packets are fragmented into multiple *HC-SDMA* data packets. *HC-SDMA* packets shall not be longer than 2048 bytes. See clause 9.4.6.2, Packet Fragmentation and Recombination for information.

If the checksum comparison fails, L3 RM shall kill the stream in question, as it is no longer able to distinguish the packet boundaries.

9.4.6 Packet Data Conversion

There are a number of optional algorithms that BS L3 RM and UT L3 RM may use as part of the transformation of L4 data packets into *HC-SDMA* data packets. These algorithms constitute the optional portion of L3’s packet data conversion layer (PDCL).

During the registration process, UT L3 RM shall request the set of optional algorithms that it wishes to use. BS L3 RM shall confirm the use of a (possibly full) subset of the requested algorithms. This

negotiation shall be performed by using an `alg_choices` bit mask field in the `UT_params` and `reg_params` registration messages to specify which algorithms to use. If a bit in the mask is set, the corresponding PDCL algorithm shall be enabled. The bits in this bit mask shall be assigned to the following algorithms:

Table 94: Packet Data Conversion Algorithms

Bit	Algorithm
0	Packet fragmentation/recombination
1	Packet checksumming
2	Packet sequencing/reordering
3..255	Reserved

These algorithms are described in the following sections.

9.4.6.1 Packet Checksumming

Although L1 (the physical layer) uses a CRC to verify that each burst is received with no errors, L3 has the option of adding a checksum to each data packet. This option is communicated to the UT and BS via the bit indications as shown in Table 97 - L3 AM Messages. Primarily this is useful in conjunction with packet fragmentation/reassembly to detect packets that are corrupted by missing fragments. In addition, the checksumming provides a second level of security against corruption of data as it is transferred on the physical layer, and a guard against any corruption of data either before the data is given to L1 on the sending side or after the data is obtained from L1 on the receiving side.

The checksum is a 16-bit CRC that shall be computed by the following steps:

1. The CRC shall be initialized to 0xffff.
2. For each data byte, a 256 entry CRC table shall be indexed by the lower 8 bits of the current CRC XOR'ed with the data byte.
3. The lower eight bits of the value from the CRC table shall be XOR'd with the upper 8 bits of the CRC value.
4. The XOR result shall be used as the lower 8 bits of the new CRC value, and the upper 8 bits of the value from the CRC table shall be used as the upper 8 bits of the new CRC value. The CRC table shall be obtained from RFC 1662, *PPP in HDLC-like Framing*.

The CRC shall be appended to the data packet with the most-significant byte first.

9.4.6.2 Packet Fragmentation and Recombination

In order to minimize latency of large L4 data packets, it may be desirable to fragment them and send the pieces in parallel over multiple streams. If packet fragmentation and recombination is used, packet sequencing and reordering and packet checksumming shall be used. Furthermore, on the sending side the checksumming shall be done first, followed by the fragmentation, followed by the sequence numbering. On the receiving side, the reordering shall be done first, followed by recombination, followed by checksum verification.

Packet fragmentation and recombination shall piggyback its header information with the packet sequencing and reordering header. Specifically, bit 7 (the MSB) of the byte shall be appended by packet sequencing and reordering shall be used by packet fragmentation and recombination to indicate whether the packet fragment concludes a fragmented L4 data packet. If the bit is set, the fragment shall not be the final one. If the bit is cleared, the fragment shall conclude an L4 data packet.

9.4.6.3 Packet Sequencing and Reordering

For this optional PDCL algorithm, the sender shall append a one-byte packet sequence number. The packet sequence number shall be 0 for the first packet transferred by the registration and shall increment once for each packet. The packet sequence number shall wrap back to 0 after the sequence number 127. Thus, only seven bits are required for the packet sequence number. These seven bits shall be bits [0..6] of the appended byte; bit 7 shall be set to the value specified by the packet fragmentation and recombination algorithm if packet fragmentation is being used, or shall be cleared otherwise. The receiver shall strip this extra byte off the packet before passing it on, but if packet fragmentation and recombination is used, the packet recombination algorithm shall be told whether bit 7 of the extra byte was set or cleared.

The specific method to be used by the receiver to reorder packets should conform to the following. If the packet sequence number is not too much greater than expected (accounting for sequence number wrap-around), the packet sequencing algorithm should hold on to the packet until the intervening packets arrive. If the packet sequence number is smaller than expected (again, accounting for wrap-around), the packet should be discarded. This might occur when the packet sequencing algorithm waiting for the packet gives up and releases subsequent packets anyway, or a when duplicate packet is received. If a packet is discarded, an `L3CmPktDrop.ind` indication for the incoming direction shall be given to L4.

9.4.7 Packet Discarding

In order to prevent consumption of an unbounded amount of memory, L3 RM shall limit the number of data packets and total number of data bytes that it holds for a registration. If the outgoing data queue becomes too large, L3 RM shall discard data packets. Control packets shall not be discarded. If L3 RM discards outgoing data packets, it should use the `L3RmPktDrop.ind(outgoing)` primitive to inform the local L3 CM of the packet loss. L3 CM should in turn use the `L3CmPktDrop.ind(outgoing)` primitive to notify the local L4 of the packet loss.

Similarly, at the start of a session, when BS L3 CM is queuing up outgoing L4 data and/or L4 control packets while waiting for a confirm from L4, BS L3 CM shall discard packets if its queue becomes excessively large. Only data packets may be discarded. If BS L3 CM discards packets, it should use the `L3CmPktDrop(indcoming)` primitive to inform BS L4.

9.4.8 Packet Routing

On the outgoing (sending) side, L3 RM shall choose how to distribute *HC-SDMA* packets between all the streams that are open.

9.4.9 Data Recovery from L2 RLC

When a stream ends, any outgoing packets which were not completely sent by the local L2 RLC and acknowledged by the peer L2 RLC shall be reclaimed by L3 RM and resent on a new stream. See clause 9.2.2.2 L3 RM Interface with L2 RLC for a description of this cleanup handshake that recovers unsent/unacknowledged packets.

L2 RLC shall inform L3 RM how much data has not been sent to the L2 RLC peer or has been sent but not acknowledged by the L2 RLC peer

A partially transmitted and acknowledged HC-SDMA packet shall be reconstructed so it can be re-sent on a different stream. L3 RM shall add back the portion of the packet that was sent and acknowledged to the unacknowledged part returned by L2 RLC.

Note that because of this, it is possible for the L3 RM peer to receive the packet in full but to have the L2 RLC peer's acknowledgement lost and the stream itself shut down because of RF problems. In this case, the L3 RM sender recovers the unacknowledged data from L2 and re-sends it on a new stream. This causes the receiving L3 RM to get two copies of the *HC-SDMA* packet. A way to discover and correct this is to enable the packet sequencing and reordering algorithm. See clause 9.4.6.3, Packet Sequencing and Reordering for information.

9.4.10 Stream Aggregation

When the L3 RM data backlog gets too large, L3 RM may request new streams to increase the registration's throughput and relieve the backlog. BS L3 RM should make this decision based directly on its own outgoing data backlog. UT L3 RM shall inform BS L3 RM about the state of its outgoing data backlog by using the `queue.rpt` UM control message. See clause 9.3.4.1.4, Stream Aggregation (Downlink Backlog) and clause 9.3.4.2.4, Stream Aggregation (Uplink Backlog) for descriptions of this messaging.

The `queue.rpt` UM control message shall have a five-bit data field. The values 0 and 31 are special values that shall report that the queue has no data or that more than the maximum amount that can be closely specified, respectively. The intermediate thirty values shall give an indication of how large the queue of uplink packets on the UT is: `queue.rpt` message value = minimum of 30 and floor $((\text{queue depth} - 1) / 512) + 1$ as shown in Table 95 - Queue Depth.

Table 95: Queue Depth

Field Value	Queue Depth is At Least	Queue Depth is Less Than
0	queue depth = 0	n/a
1	1 byte	513 bytes
2	513 bytes	1025 bytes
3	1025 bytes	1537 bytes
4	1537 bytes	2049 bytes
5	2049 bytes	2561 bytes
6	2561 bytes	3073 bytes
7	3073 bytes	3585 bytes
8	3585 bytes	4097 bytes
9	4097 bytes	4609 bytes
10	4609 bytes	5121 bytes
11	5121 bytes	5633 bytes
12	5633 bytes	6145 bytes
13	6145 bytes	6657 bytes
14	6657 bytes	7169 bytes
15	7169 bytes	7681 bytes
16	7681 bytes	8193 bytes
17	8193 bytes	8705 bytes
18	8705 bytes	9217 bytes
19	9217 bytes	9729 bytes
20	9729 bytes	10241 bytes
21	10241 bytes	10753 bytes
22	10753 bytes	11265 bytes
23	11265 bytes	11777 bytes
24	11777 bytes	12289 bytes
25	12289 bytes	12801 bytes
26	12801 bytes	13313 bytes
27	13313 bytes	13825 bytes
28	13825 bytes	14337 bytes
29	14337 bytes	14849 bytes
30	14849 bytes	15361 bytes
31	15361 bytes	n/a

9.5 Message Definitions

9.5.1 Acknowledged Mode (AM) Messages

Acknowledged Mode (AM) messages shall be transported via the L2 RLC service.

9.5.1.1 Message Structure

The structure of an AM message shall conform to Figure 83 - AM Message Structure. AM messages shall contain tag length value-groups (TLVGs) with data fields.

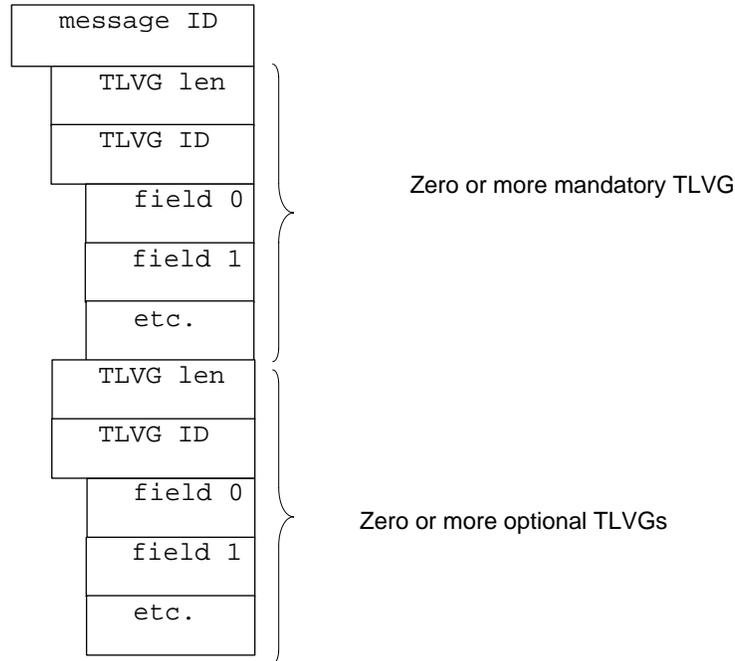


Figure 83: AM Message Structure

9.5.1.1.1 Message Top-level Structure

Each AM message shall consist of a message header and message data. The message header shall be a one-byte message ID and the message data shall consist of (possibly zero) mandatory TLVGs and (possibly zero) optional TLVGs (in no specified order). The message ID shall implicitly indicate which TLVGs are mandatory for the message.

9.5.1.1.2 Message TLVG Structure

Each TLVG shall consist of a header and data. The header shall be a length indication followed by the one-byte ID. The length indication shall be an extensible field that uses the seven LSBs of each byte to encode the length of the TLVG in bytes, and that uses the MSB of each byte to indicate if more bytes follow. For example, TLVG lengths of 1-127 bytes shall be encoded in a single byte with the MSB cleared; TLVG lengths of 128-16383 shall be encoded in two bytes with the MSB of the first byte set and the MSB of the second byte cleared.

The TLVG length shall include the TLVG ID byte and all TLVG data field bytes but shall not include the TLVG length field itself. Because the length field is at a fixed (zero) offset within the TLVG, L3 may skip over a corrupted TLVG provided that the corruption is not in the length field itself. Similarly, L3 may skip over any unrecognized TLVGs.

If the receiving L3 parses to the end of the TLVG but finds that the amount of data parsed does not complete all the expected fields in the TLVG, the receiving L3 shall fill in the initial fields of the TLVG with the data sent by the peer transceiver, and shall then use default values for the remaining fields.

If the receiving L3 parses to the end of the TLVG but finds that the amount of data parsed does not complete all the expected fields in the TLVG, the receiving L3 shall fill in the initial fields of the TLVG with the data sent by the peer transceiver, and shall then use default values for the remaining fields.

In contrast, if the receiving L3 parses all the fields in the TLVG that it expects but find that it has not reached the end of the TLVG in the packet, the receiving L3 shall ignore the remaining data in the current TLVG, and jump to the start of the subsequent TLVG to resume parsing.

If the packet size of a message that L3 parses is inconsistent with the sum of the specified lengths of all TLVGs in the packet, the packet has been malformed or corrupted, so the entire message packet shall be discarded.

9.5.1.1.3 Message Field Structure

Most fields in an L3 AM message TLVG are a fixed number of bits long (N), with options for regarding these bits as a scalar or as an array of M scalars of N/M bits each. For fields interpreted as arrays of scalars, the first element of the array shall occupy the most significant bits of the N bits. Each scalar is either a signed or unsigned number.

The L3 AM message library also allows variable-length fields in which a sub-field preceding the data specifies the length. Each variable length field shall use a fixed number of bits to represent the message length and a fixed number of bits as the length quantum. For example, a one variable length field may have an 8-bit length specification with units of 1 bit, allowing the data to be 0, 1, 2, ..., 254, or 255 bits long. Another variable length field may have a 5-bit length specification with units of 4 bits, allowing the data itself to be 0, 4, 8, ..., 120, or 124 bits long. All variable-length fields shall be followed by padding up to the next byte boundary. This allows the receiver of the message to know in advance at which bit-offset each information element starts, simplifying parsing the L3 AM messages.

All padding bits in the message shall be set to zero.

9.5.1.2 Message Formatting

L3 message shall follow the following formatting rules:

- ◆ L3 AM messages shall use network byte order, that is, the most significant byte first.
- ◆ The MSB of data for a field shall go into the MSB of the reserved portion of the first byte reserved for the field.
- ◆ Fields in a TLVG shall be packed together with no padding (other than the inherent padding following variable-length fields, or padding explicitly added as a distinct field).
- ◆ TLVGs shall start on byte boundaries.
- ◆ Any padding required to make a TLVG use an integer number of bytes shall come after all data fields of the TLVG.

Table 96 - Example Message Coding illustrates how a two field TLVG is encoded. The first, data field A, is 6 bits long (a5..a0, where a5 is the MSB and a0 is the LSB). The second, data field B, is 16 bits long (b15..b0, where b15 is the MSB and b0 is the LSB).

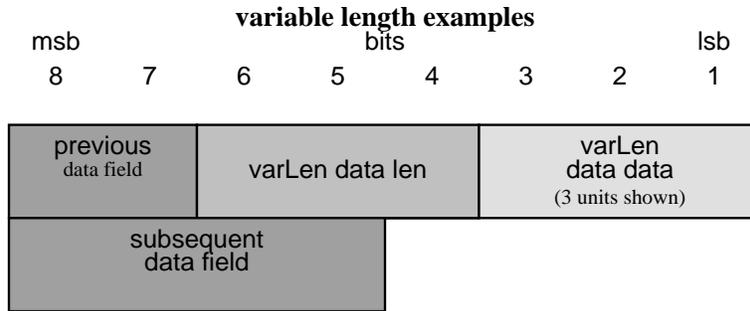


Figure 85: Variable-length Field Padding Case 2

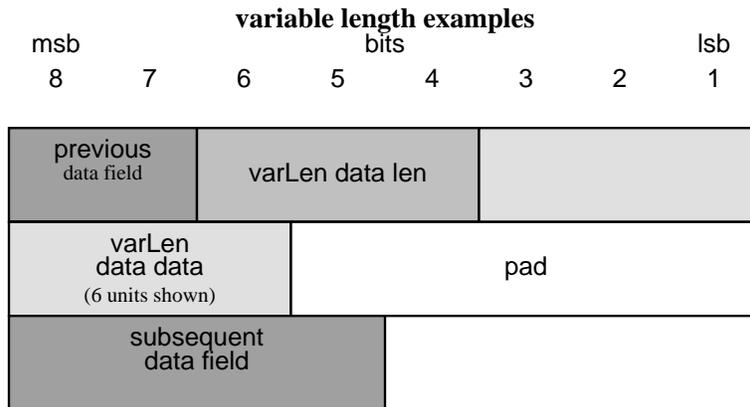


Figure 86: Variable-length Field Padding Case 3

9.5.1.3 Individual Messages

The directionality of L3 AM messages between L3 peer components shall be as indicated in Table 97 - L3 AM Messages. See clause 9.3, Protocol Events for information about how messages are sequenced to form protocol handshakes.

Table 97: L3 AM Messages

Message	Endpoints
BS params	UT L3 RM <--- BS L3 RM
UT capability	UT L3 RM ---> BS L3 RM
UT params	UT L3 RM ---> BS L3 RM
Reg params	UT L3 RM <--- BS L3 RM
network session start	UT L3 CM ---> BS L3 CM
network session stop	UT L3 CM <--> BS L3 CM
in band page	UT L3 RM <--- BS L3 RM
rsrc map	UT L3 RM <--- BS L3 RM
RA latency request	UT L3 RM <--- BS L3 RM
RA latency report	UT L3 RM ---> BS L3 RM
Dereg	UT L3 RM <--> BS L3 RM

Message	Endpoints
QoS	UT L3 RM <--- BS L3 RM
AFN verify	UT L3 RM <--- BS L3 RM
co reg queue report	UT L3 RM ---> BS L3 RM
RACH retry params	UT L3 RM <--- BS L3 RM

The individual messages are described in this section. Each message has a textual description of its components, and a graphical depiction of its layout. In the figure depicting the message, the octets are numbered to the extent to which the layout is statically determined. Following a variable-length field, the octet number is dynamically determined based on the actual length of the variable-length field. Hence, in the figures the octet numbering is shown as a variable N1 after the first variable-length field in the message, N2 after the second, and so on. The figures choose an arbitrary number of elements to depict for the variable length fields.

9.5.1.3.1 BS Params Message

Message ID: 1

TLVG: BS params std

TLVG ID: 1

Information Elements:

- ◆ IE: current HC-SDMA protocol version, 8 bits

description:

The current protocol version field shall indicate which revision of the HC-SDMA protocol the sending party is currently using. This shall be used by the UT in the UT capability message and by the BS in the BS params message to inform each other of which protocol version is being used during the registration stream. In contrast, the final protocol version used after the registration stream shall get assigned by the BS in the reg params message. The translation from the values of this field and the corresponding protocol versions shall conform to Table 98 - Protocol Version Field Settings .

- ◆ IE: noise floor, 4 bits

description:

The noise floor shall indicate how far above -116.5 dBm the unit's noise floor is, in 0.5 dB steps. i.e., the unit's noise floor is -116.5 dBm + 0.5 * the value of this field.

- ◆ IE: TCH decode

array of 3 1-bit elements

description:

The TCH decode time shall specify whether the unit can complete decoding the received burst's payload before it has to build the subsequent transmit burst (regardless of mod class), or if it has to wait until after preparing the transmit burst to finish decoding the received payload. This capability is specified for the unit assuming a full-rate stream -- i.e., a stream with TCH bursts granted every AFN. The TCH decode times shall be stored individually for the 3 timeslots, since the timing requirements are different for the different timeslots. The first array element shall correspond to the first timeslot. A value of 0 shall mean the unit is capable of decoding and processing the received TCH before building the next transmit burst. Note that the BS and UT start the registration stream assuming that the peer cannot decode TCH prior to the subsequent transmission. During the course of the registration stream, the BS and UT find out each other's actual capabilities, but for uniformity they must continue to use the initial assumptions until the end of the registration stream.

◆ IE: FACCH decode

array of 3 1-bit elements

description:

The FACCH decode time shall specify whether the unit can complete decoding the received burst's FACCH before it has to build the subsequent transmit burst, or if it has to wait until after preparing the transmit burst to finish decoding the received burst's FACCH. The FACCH decode times shall be stored individually for the 3 timeslots, since the timing requirements are different for the different timeslots. The first array element shall correspond to the first timeslot. A value of 0 shall mean the unit is capable of decoding and processing the received FACCH before building the next transmit burst. Note that the BS and UT start the registration stream assuming that the peer cannot decode FACCH prior to the subsequent transmission. During the course of the registration stream, the BS and UT find out each other's actual capabilities, but for uniformity they must continue to use the initial assumptions until the end of the registration stream.

◆ IE: AFN high bits, 38 bits

description:

The UT knows the 10 LSBs of the AFN from the CM burst. This field shall supply the 38 MSBs of the AFN of the frame of the RA-rreg/AA-rreg exchange that started the registration stream. This allows the UT to have complete AFN synchronization with the BS.

◆ IE: BS cert

variable length: 8 bits to specify length, in units of 8 bits

no subsequent padding

description:

The complete base-station certificate. The certificate shall be comprised of a message and the signature that proves the message's veracity. Refer to clause 11.2.2, Base Station Transmits Its Digital Certificate for a description of the contents of the BS certificate.

Table 98: Protocol Version Field Settings

Value	Meaning
0	unknown / non-compliant
1	Unused
2	Reserved
3	Reserved
4	Reserved
5	1.0
6-127	reserved for future protocol revisions
128-255	reserved for vendor-specific protocol modifications

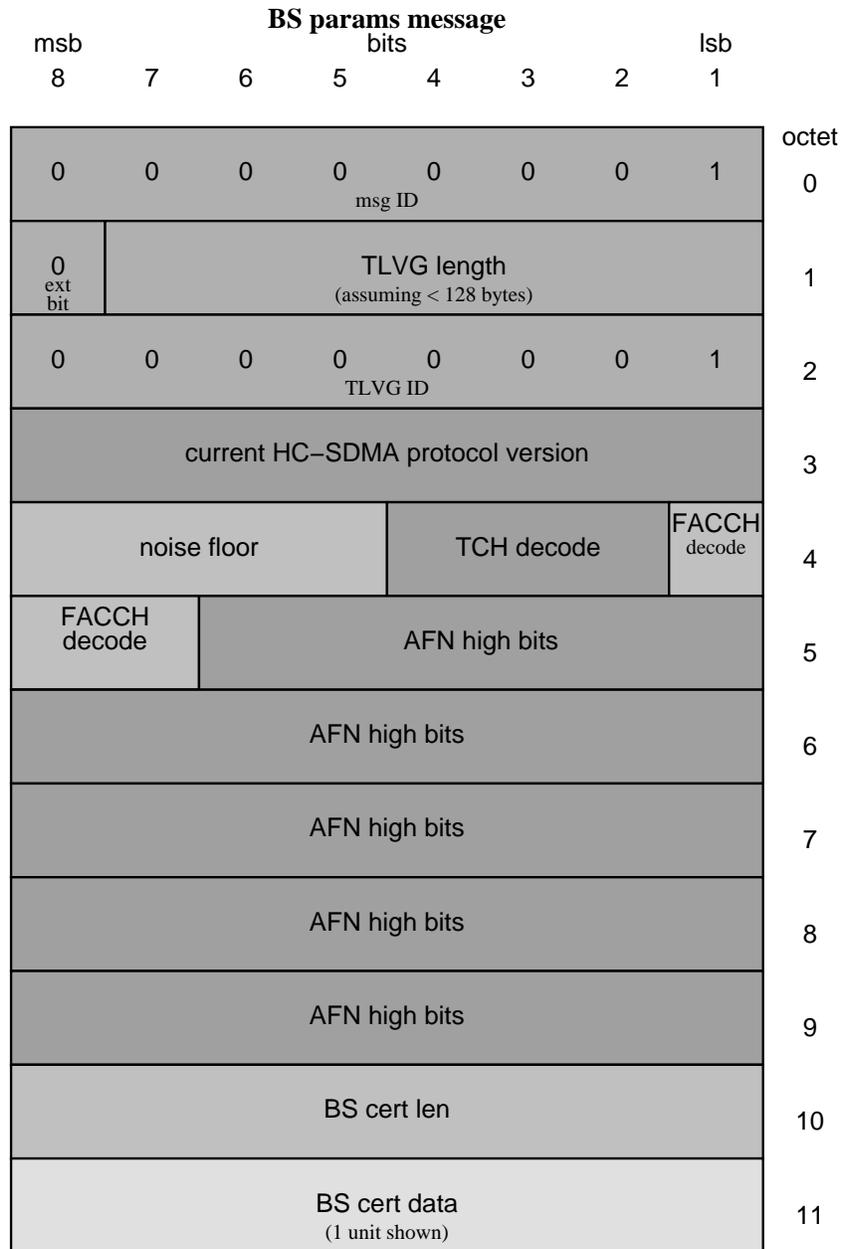


Figure 87: BS Params Message

9.5.1.3.2 UT Capability Message

Message ID: 2

TLVG: UT capability std

TLVG ID: 2

Information Elements:

- ◆ IE: current HC-SDMA protocol version, 8 bits

description:

The current protocol version field shall indicate which revision of the HC-SDMA protocol the sending party is currently using. This is used by the UT in the UT capability message and by the BS in the BS params message to inform each other of which protocol version is being used during the

registration stream. In contrast, the final protocol version used after the registration stream shall be assigned by the BS in the reg params message. The translation from the values of this field and the corresponding protocol versions shall conform to Table 98 - Protocol Version Field Settings.

- ◆ IE: upl mod classes, 16 bits

description:

The upl mod classes field shall be a bit map of the mod classes that the UT is able to transmit on the uplink. The LSB of the bit map shall correspond to mod class 0.

- ◆ IE: dwl mod classes, 16 bits

description:

The dwl mod classes field shall be a bit map of the mod classes that the UT is able to receive on the downlink. The LSB of the bit map shall correspond to mod class 0.

- ◆ IE: tx pwr diff mod01, 4 bits

description:

The tx pwr diff mod01 field shall state how much above the maximum transmit power for the modulation scheme with the highest peak-to-average ratio the maximum transmit power for modulation classes 0 or 1 is. (Modulation classes 0 and 1 have the same peak-to-average ratio.) This field shall be unsigned with half dB steps. For example, if the transmitter transmits a maximum power of 24.5 dBm with mod class 0 and 20.0 dBm with the maximum peak-to-average modulation scheme it supports, then the value of this field shall be $\text{floor}((24.5 - 20.0) * 2) = 9$.

- ◆ IE: tx pwr diff mod2345, 4 bits

description:

The tx pwr diff mod2345 field shall state how much above the maximum transmit power for the modulation scheme with the highest peak-to-average ratio the maximum transmit power for modulation classes 2-5 is. (Modulation classes 2-5 all have roughly the same peak-to-average ratio.) This field shall be unsigned with half dB steps. For example, if the transmitter transmits a maximum power of 21.5 dBm with mod class 4 and 20.0 dBm with the maximum peak-to-average modulation scheme it supports, then the value of this field shall be $\text{floor}((21.5 - 20.0) * 2) = 3$.

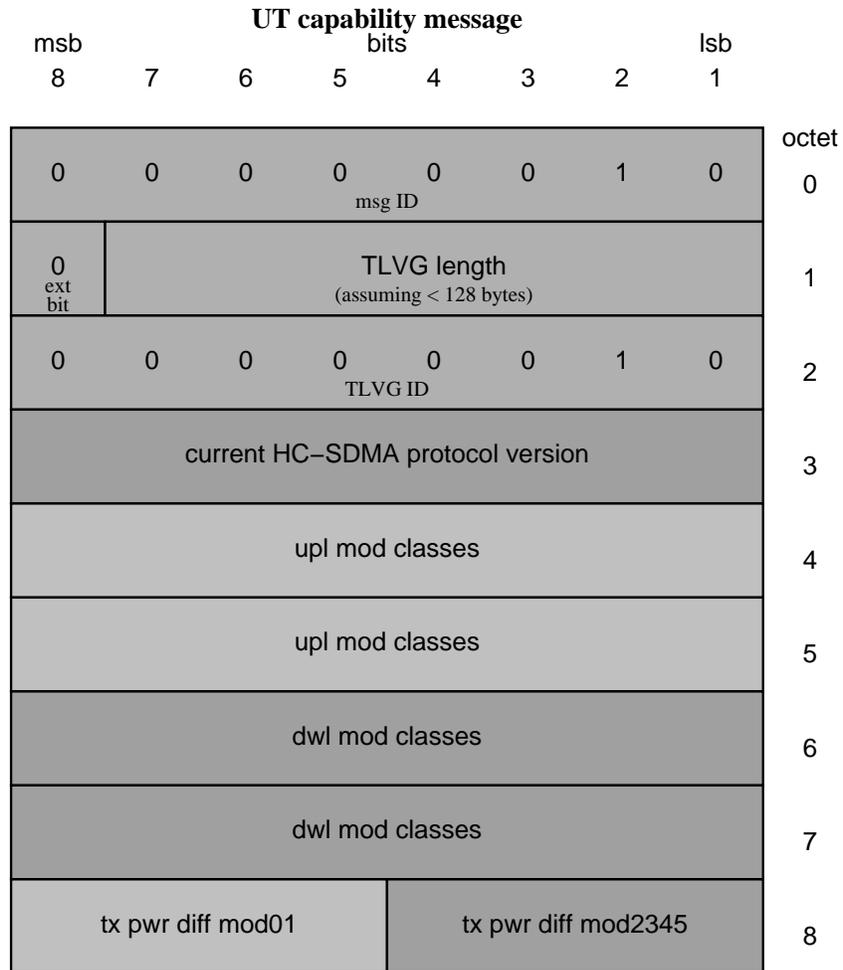


Figure 88: UT Capability Message

9.5.1.3.3 UT Params Message

Message ID: 3

TLVG: UT params std

TLVG ID: 3

Information Elements:

- ◆ IE: min HC-SDMA protocol version, 8 bits

description:

The minimum protocol version field shall indicate the earliest revision in a contiguous range of HC-SDMA protocol revisions the sending party is able to support. The translation from the values of this field and the corresponding protocol versions shall conform to Table 98 - Protocol Version Field Settings.

- ◆ IE: max HC-SDMA protocol version, 8 bits

description:

The maximum protocol version field indicates the latest revision in a contiguous range of HC-SDMA protocol revisions the sending party is able to support. The translation from the values of this field and the corresponding protocol versions shall conform to Table 98 - Protocol Version Field Settings.

- ◆ IE: pkt protocol, 16 bits

description:

This field shall be used to pass from UT L3 RM to BS L3 RM what type of data (protocol) a new session will carry. Generally, the HC-SDMA system doesn't need to know the protocol of the packets that it is tunneling. However, there are possibly PDCL algorithms such as VJ header compression that only apply for certain protocol encapsulations.

- ◆ IE: page timeout 1, 16 bits

description:

This field shall give the number of frames the UT will take to transition from listening for pages on every frame to listening every 8th frame. A value of 0 shall mean that the UT never transitions. The UT shall choose the same value for this parameter for all concurrent registrations with the same BS.

- ◆ IE: page timeout 2, 16 bits

description:

This field shall give the number of frames the UT takes to transition from listening for pages on every 8th frame to listening every 64th frame. A value of 0 shall mean that the UT never transitions. The UT shall choose the same value for this parameter for all concurrent registrations with the same BS.

- ◆ IE: page timeout 3, 16 bits

description:

This field shall give the number of frames the UT takes to transition from listening for pages on every 64th frame to listening every 512th frame. A value of 0 shall mean that the UT never transitions. The UT shall choose the same value for this parameter for all concurrent registrations with the same BS.

- ◆ IE: page timeout 4, 16 bits

description:

This field shall give the number of frames the UT takes to transition from listening for pages on every 512th frame to no longer listening. A value of 0 shall mean that the UT never transitions. The UT shall choose the same value for this parameter for all concurrent registrations with the same BS.

- ◆ IE: lowest freq, 5 bits

description:

The lowest freq field shall indicate to the BS what the lowest frequency is that the UT is capable of using. This field shall be formatted as a twos-complement signed number ranging from -16 to +15, where zero shall represent the BS's base carrier. Thus, the UT shall specify its lowest carrier anywhere within the 10 MHz (16 carrier) maximum carrier range handled by the current BS, or within the 10 MHz range immediately below the current BS's 10 MHz maximum carrier range. (It is possibly useful for the current BS to know if the UT is capable of operating in the 10 MHz immediately below the BS's frequency allocation, so that the BS may suggest to the UT whether it should hand over to a co-located BS operating in the 10 MHz below the current BS.) If the UT is capable of using a frequency more than 10 MHz below the BS's base frequency, it shall set the lowest freq field to -16.

- ◆ IE: highest freq, 5 bits

description:

The highest freq field shall indicate to the BS what the highest frequency is that the UT is capable of using. This field shall be formatted as an unsigned number ranging from 0 to 31, where zero represents the BS's base carrier. Thus, the UT may specify its highest carrier anywhere within the 10 MHz (16 carrier) maximum carrier range handled by the current BS, or within the 10 MHz range immediately above the current BS's 10 MHz maximum carrier range. (It is useful for the current BS to know if the UT is capable of operating in the 10 MHz immediately above the BS's frequency allocation, so that the BS may suggest to the UT whether it should hand over to a co-located BS

operating in the 10 MHz above the current BS.) If the UT is capable of using a frequency more than 20 MHz above the BS's base frequency, it shall set the highest freq field to 31.

- ◆ IE: total streams, 6 bits

description:

This field shall indicate to the BS how many streams the UT can support simultaneously. This field shall be limited by the instantaneous frequency span of the UT's radio, the number of separate frequency up/down-conversion channels the radio supplies, and the processing cycles available to the UT. The total streams handled by the UT shall be less than or equal to sum (streams per timeslot). Note that the total streams field is a measure of how many streams the UT processes with the most aggressive mod class within the set of mod classes supported by the UT; that is, the mod class with the largest processing requirements. As such, this provides a conservative limit for the BS of how many streams the UT is able to process concurrently.

- ◆ IE: streams per timeslot

array of 3 4-bit elements

description:

The streams per timeslot array shall indicate for each timeslot, how many streams the UT can process simultaneously on that timeslot, in the absence of streams on other timeslots. The timeslot number shall be used as the array index -- i.e., array element 0 states the maximum number of streams the UT can support on timeslot 0. This reflects both the number of separate frequency up/down-conversion channels supplied by the UT's radio (which is constant across timeslots), and the processing cycles available to the UT to meet its hard real-time processing deadlines (which are different for the different timeslots). Note that this is a measure of how many streams the UT is able to process with the most aggressive mod class within the set of mod classes supported by the UT; that is, the mod class with the largest processing requirements. As such, this provides a conservative limit for the BS of how many streams the UT is able to process concurrently on the timeslot.

- ◆ IE: inst freq span, 4 bits

description:

The inst freq span field shall indicate the maximum possible carrier separation the UT can support for any streams on the same timeslot. Note that this information is still relevant even for UTs that are not able to process multiple streams per timeslot due to other limitations, such as only having a single frequency up/down-conversion channel. For example, a UT that supports concurrent streams within a band of 4 adjacent carriers sets this value to 3. A UT that is able to support concurrent streams anywhere within the 16 carriers of the BS's 10 MHz range sets this value to 15. A UT that is able to only support 1 stream due to having a single frequency up/down-conversion channel, but which is able to set the up/down-conversion frequency to address any carrier in the BS's 10 MHz range sets this value to 0.

- ◆ IE: interslot tuning span, 4 bits

description:

The interslot tuning span shall indicate by how many carriers the UT can move its instantaneous frequency span band during the guard time between timeslots. For example, a UT which is able to concurrently support two streams within a four carrier block might for timeslot zero have this 4-carrier span tuned to carriers 0-3, and for timeslot 1 be able to re-tune the 4-carrier span as high as carriers 4-7. In this case, the UT's interslot tuning span is 4.

- ◆ IE: noise floor, 4 bits

description:

The noise floor shall indicate how far above -116.5 dBm the unit's noise floor is, in 0.5 dB steps -- i.e., the unit's noise floor shall be $-116.5 \text{ dBm} + 0.5 * \text{the uint value of this field}$.

- ◆ IE: 1-bit pad, 1bits

description:

1bit of padding (to improve byte-alignment of the subsequent fields) - this field shall be set to 0.

- ◆ IE: UT max tx pwr, 7bits

description:

The UT max tx pwr field shall indicate the UT's maximum possible transmit power, in half dB steps above 0 dBm - i.e., the UT's maximum transmit power shall be $0 \text{ dBm} + 0.5 * \text{the uint value of this field}$.

- ◆ IE: ping renew, 1 bits

description:

The ping renew field shall indicate whether the registration gets refreshed when the BS receives a ping from the UT owning the registration. Some registrations on the UT may choose to be "always on", and use pings to keep the registration valid even if the registration has no data to transfer, while other registrations may choose to let the registration expire if there is not data to transfer (and thus free up the RID and other BS resources associated with a registration).

- ◆ IE: CCH decode, 1 bits

description:

The CCH decode time shall specify whether the UT can finish decoding a CM burst before it has to build a burst for transmit on the same uplink timeslot on the next frame. A value of 0 shall mean the UT is capable of decoding the CM burst before preparing a burst to transmit in the next frame's same uplink timeslot.

- ◆ IE: RACH decode

array of 3 1-bit elements

description:

The RACH decode time shall specify whether the UT completes decoding an AA burst before it has to build a burst for transmit for the corresponding uplink timeslot of the next frame. The RACH decode times shall be stored individually for the 3 timeslots, since the timing requirements shall be different for the different timeslots. The first array element shall correspond to the first timeslot. A value of 0 shall mean the unit is capable of decoding and processing the received AA burst before building the transmit burst for the corresponding timeslot of the uplink frame.

- ◆ IE: TCH decode

array of 3 1-bit elements

description:

The TCH decode time shall specify whether the unit can complete decoding the received burst's payload before it has to build the subsequent transmit burst (regardless of mod class), or if it has to wait until after preparing the transmit burst to finish decoding the received payload. This capability shall be specified for the unit assuming a full-rate stream -- i.e., a stream with TCH bursts granted every AFN. The TCH decode times shall be stored individually for the 3 timeslots, since the timing requirements are different for the different timeslots. The first array element shall correspond to the first timeslot. A value of 0 shall mean the unit is capable of decoding and processing the received TCH before building the next transmit burst. Note that the BS and UT shall start the registration stream assuming that the peer cannot decode TCH prior to the subsequent transmission. During the course of the registration stream, the BS and UT shall find out each other's actual capabilities, but for uniformity they continue to use the initial assumptions until the end of the registration stream.

- ◆ IE: FACCH decode

array of 3 1-bit elements

description:

The FACCH decode time shall specify whether the unit completes decoding the received burst's

FACCH before it has to build the subsequent transmit burst, or if it has to wait until after preparing the transmit burst to finish decoding the received burst's FACCH. The FACCH decode times shall be stored individually for the 3 timeslots, since the timing requirements are different for the different timeslots. The first array element shall correspond to the first timeslot. A value of 0 means the unit is capable of decoding and processing the received FACCH before building the next transmit burst. Note that the BS and UT shall start the registration stream assuming that the peer cannot decode FACCH prior to the subsequent transmission. During the course of the registration stream, the BS and UT shall find out each other's actual capabilities, but for uniformity they shall continue to use the initial assumptions until the end of the registration stream.

- ◆ IE: PCH decode

array of 3 1-bit elements

description:

The PCH decode time shall specify whether the UT finishes reception of PCH in time to transmit a RA-page-response on the corresponding uplink timeslot of the following frame, or if the UT has to wait an extra frame before sending the RA. The PCH decode times shall be stored individually for the 3 timeslots, since the timing requirements are different for the different timeslots. The first array element shall correspond to the first timeslot. A value of 0 shall mean the unit is capable of decoding the PCH in time to transmit the RA-page-response in the uplink of the next frame.

- ◆ IE: iSEC alg, 2bits

description:

The i-SEC alg field shall be used by the UT to specify its choice of i-SEC bulk encryption algorithms from those that the BS indicates in the i-HAP BS certificate that it supports. Refer to clause 11.2.2, Base Station Transmits Its Digital Certificate.

- ◆ IE: iSEC key size, 5 bits

description:

The i-SEC key size field shall be used by the UT to specify how many bytes of the (decrypted) i-HAP secret will be used as the i-SEC encryption secret. If this field is non-zero, the key size shall be the value of this field plus four, in bytes. If the field is zero, this shall indicate that for testing purposes i-SEC is disabled for the registration in question. This value shall be less than or equal to the "i-SEC Bulk Encryption Key Size" value from the BS certificate for the bulk encryption algorithm selected by the UT. Refer to clause 11.2.2, Base Station Transmits Its Digital Certificate and clause 11.4, *HC-SDMA Secure Communications Protocol, i-SEC*.

- ◆ IE: 2-bit pad, 2 bits

description:

2 bits of padding (to improve byte-alignment of the subsequent fields) - this field shall be set to 0.

- ◆ IE: AFN CA ID req

variable length: 1 bits to specify length, in units of 8 bits

no subsequent padding

description:

The AFN CA ID, if present, shall tell the BS which CA to forward the (decrypted) UT authentication string to, so that the CA can sign the authentication string to prove to the UT that the AFN contained therein is correct. If the AFN CA ID is not present (length = 0), the UT doesn't require an AFN verification from a trusted CA. If the AFN CA ID is present, it shall be scrambled based on the i-HAP shared secret as documented in clause 11.3.4, UT Parameter Scrambling.

- ◆ IE: CA ID, 8 bits

description:

The CA ID shall tell the BS which Certificate Authority signed the UT's certificate. This field shall be scrambled based on the i-HAP shared secret as documented in clause 11.3.4, UT Parameter

Scrambling.

◆ IE: UT cert

variable length: 8 bits to specify length, in units of 8 bits

no subsequent padding

description:

The UT certificate signature. The UT certificate can be recovered from this field when the BS knows which CA signed the UT's certificate. This field shall be scrambled based on the i-HAP shared secret as documented in clause 11.3.4, UT Parameter Scrambling. Refer to clause 11.3.1 UT Certificates for the types and respective contents of the UT certificate.

◆ IE: UT auth str

variable length: 8 bits to specify length, in units of 8 bits

no subsequent padding

description:

The UT-encrypted authentication string shall prove to the BS that the UT possesses the private key associated with the UT certificate that the UT purports to own. The BS shall verify the authentication string by decrypting it with the UT's public key from the UT certificate, and checking that the information in the authentication string is correct. Refer to clause 11.3.2 UT Authenticator Message and String Generation for a full description of the data in the authentication string and its format, and the methods for encrypting/decrypting the authentication string.

◆ IE: auth msg AFN, 48 bits

description:

The authentication message AFN shall be the cleartext version of the AFN used to generate the UT authentication message, which is then hashed and encrypted to create the UT authentication string. The BS shall decrypt the UT's authentication string, and then build its own version based on the UT ID and session reference from the UT certificate, and the AFN reported here. By comparing the decrypted authentication string from the UT against the one it has generated, the BS shall verify that the UT has the private key associated with the UT's certificate. The authentication message AFN shall be approximately the AFN at which the registration stream started, but need not be precisely accurate. Refer to clause 11.3.2 UT Authenticator Message and String Generation regarding the authentication message/string. The auth msg AFN field shall be scrambled based on the i-HAP shared secret as documented in clause 11.3.4, UT Parameter Scrambling.

◆ IE: iHAP secret

variable length: 10 bits to specify length, in units of 1 bits

followed by padding to byte boundary

description:

The i-HAP secret shall be the encrypted and CRC protected form of the secret generated by the UT and passed to the BS for use in i-SEC encryption and i-TAP descrambling.

◆ IE: alg choices

variable length: 8 bits to specify length, in units of 1 bits

followed by padding to byte boundary

description:

The alg choices field shall identify a set of PDCL optional algorithms. The LSB of this field shall correspond to algorithm 0. Refer to Table 94 - Packet Data Conversion Algorithms for a listing of the PDCL optional algorithms.

◆ IE: opt PID

variable length: 1 bits to specify length, in units of 15 bits

followed by padding to byte boundary

description:

The optional paging ID (PID) shall be used to either specify that the UT currently has no PID assigned by this BS (specifically, by the registration domain in question within the BS), or if a PID has already been assigned to another concurrent registration on the UT, what the value is.

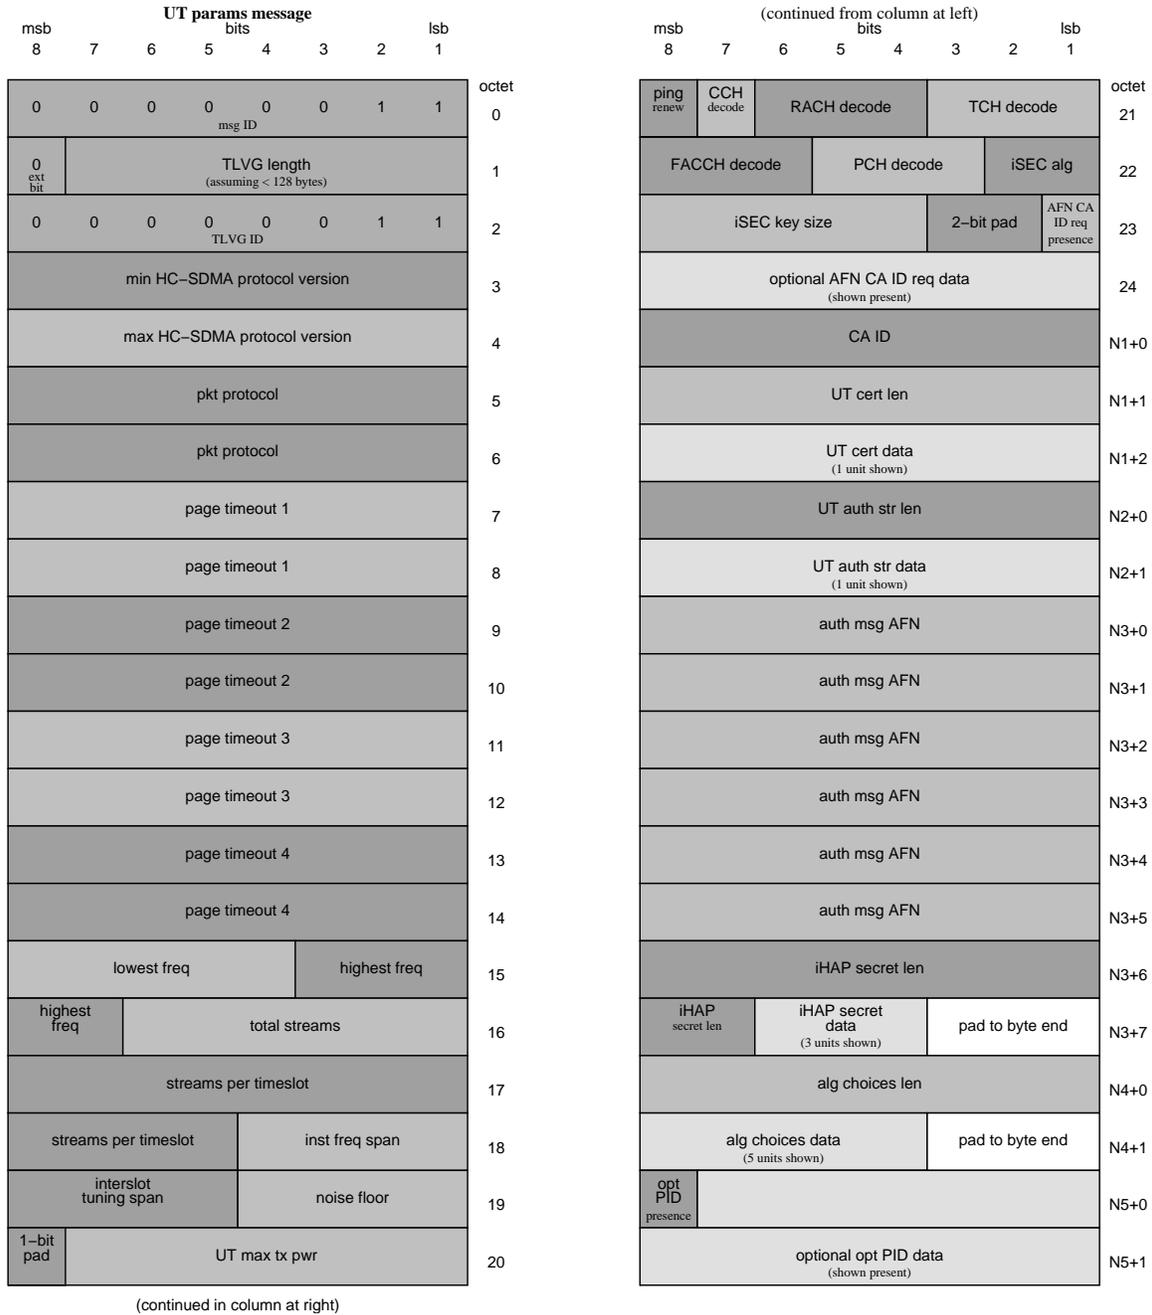


Figure 89: UT Params Message

9.5.1.3.4 Reg Params Message

Message ID: 4

TLVG: reg params std

TLVG ID: 4

Information Elements:

- ◆ IE: assigned HC-SDMA protocol version, 8 bits

description:

The assigned protocol version field shall indicate which revision of the HC-SDMA protocol the sending party is requiring the receiving party to use. This shall be used by the BS in the reg params message to tell the UT what protocol version to use throughout the registration, apart from the registration stream itself. The translation from the values of this field and the corresponding protocol versions shall conform to Table 98 - Protocol Version Field Settings.

- ◆ IE: RACH rsrc mask, 48 bits

description:

The RACH resource mask shall specify which conventional channels of the BS the UT is allowed to use for RA-rts. The LSB of this bit map shall correspond to conventional channel ID 0. See Glossary clause A.25 for a definition of the conventional channel ID.

- ◆ IE: RID expire frames, 24 bits

description:

This field shall indicate after how many idle frames the RID will become invalid, causing the registration to terminate.

- ◆ IE: RID, 15 bits

description:

The registration ID (RID) shall be used in various of the RA/AA messages that open streams to identify which registration the stream is for.

- ◆ IE: PCH conv mask

variable length: 1 bits to specify length, in units of 48 bits

no subsequent padding

description:

The PCH conventional mask shall specify which conventional channels page bursts may be transmitted on. If the length of this field is zero, then the PCH conventional mask shall be identical to the RACH rsrc mask. Otherwise, the data shall be a bit map of conventional channel IDs to be used for PCH (see Glossary clause A.25). The LSB shall correspond to conventional channel ID 0. This mask shall determine the subslot structure for each carrier as defined in clause 7.3.3.2, Base Station Page Procedures. The PCH conventional mask shall be applied to all concurrent registrations between the BS and UT. Therefore, the PCH conventional mask in the reg params message shall be the same for all concurrent registrations between the BS and a single UT. The PCH conventional mask in the rsrc map message may be different than the PCH conventional mask in the reg params message, and it shall update the PCH conventional mask for all concurrent registrations on the UT. See clause 7.3.3, Paging Channel (PCH) for a complete description of how the PCH conventional mask is used.

- ◆ IE: PID, 15 bits

description:

The paging ID (PID) shall tell the UT on what frames it will possibly be paged, how to identify whether a paging message is intended for this UT, and how to identify itself to the BS when it responds to a page. The BS shall choose the same value for this parameter for all concurrent registrations with the same UT.

- ◆ IE: PCH logical mask

variable length: 1 bits to specify length, in units of 48 bits

no subsequent padding

description:

The PCH logical mask shall specify which paging resources (see clause 7.3.3.2, Base Station Page Procedures) are valid for this UT. The mapping from the PCH logical mask to paging resources on each frequency shall conform to Table 99- Mapping of Frequency and Page Resource to PCH logical mask (0 is LSB). The valid paging resources shall be used to define the page hopping sequence for the UT (see clause 7.5, Page Hopping Sequences). If the length of this field is zero, then the PCH logical mask shall be identical to the PCH conventional mask. The PCH logical mask shall be applied to all concurrent registrations between the BS and UT. Therefore, the PCH logical mask in the reg params message shall be the same for all concurrent registrations between the BS and a single UT. The PCH logical mask in the rsrc map message may be different than the PCH logical mask in the reg params message, and it shall update the PCH logical mask for all concurrent registrations on the UT. See clause 7.3.3, Paging Channel (PCH) for a complete description of how the PCH logical mask is used.

- ◆ IE: N page resp impl, 4 bits

description:

The N page resp impl field, Nprirp, shall indicate how many RA/AA attempts the UT shall make in the absence of successful AA reception. These RA/AA attempts shall occur on the resource implied by the location of the received page burst before (possibly) going into the full RA-retry resource scan. A value of 0 shall indicate that the UT should immediately go into the full scan. The BS shall choose the same value for this parameter for all concurrent registrations with the same UT.

- ◆ IE: N aggr resp impl, 4 bits

description:

The N aggr resp impl field, Nrddrp, shall indicate how many RA/AA attempts the UT shall make in the absence of successful AA reception. These RA/AA attempts shall occur on the resource stated in the aggregation message, if there is a resource assignment. In contrast to paging, the number of attempts that are made on the directed resource shall be the value of this field plus one, since an undirected aggregation attempt is explicitly specified, rather than implied by setting the number of directed tries to zero. Note that the UT's response to an in-band page message containing a resource specification shall also use the Nrddrp value to decide how many times the UT should attempt to open a stream on the indicated resource.

- ◆ IE: Na, 3 bits

description:

Shall be the number of resources per RA-retry sub-scan (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nb, 3 bits

description:

Shall be the number of frames to wait between RA tries within a sub-scan (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Ne, 2 bits

description:

Shall be the maximum frames of backoff between RA tries (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nc, 3 bits

description:

Shall be the first parameter for determining the number of frames to wait between RA-retry sub-scans (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nd, 3 bits

description:

Shall be the second parameter for determining the number of frames to wait between RA-retry sub-

scans (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nf, 2 bits

description:

Shall be how many frames to pursue RA retry strategy before conceding failure (see Table 100 - L2 MAC RA retry policy parameter settings) .

- ◆ IE: RA page rsrc, 1 bits

description:

The RA page rsrc field shall indicate whether the UT shall go into its full RA-retry resource scan if it fails in its RA/AA attempts on the conventional channel it was directed to through the location of the page burst. If the bit is set, then the UT shall attempt the full RA-retry resource scan. The BS shall choose the same value for this parameter for all concurrent registrations with the same UT.

- ◆ IE: hold page, 1 bits

description:

The hold page field shall indicate whether the UT may be paged during the time it is held out from making RA-rts access attempts. If this bit is set, the BS may page the UT even while the UT is being held out of the system, and the UT should listen and respond according to the RA-page-response strategy. Else, the UT shall not listen for pages during the time it is being held out of the system. Note that this parameter shall only apply to hold period initiated by a hold.cmd UM control message. The alternate method of hold period initiation, the AA-reject message, contains an explicit specification of whether the UT may be paged during that particular hold period. The BS shall choose the same value for this parameter for all concurrent registrations with the same UT.

- ◆ IE: reject

variable length: 1 bits to specify length, in units of 5 bits

followed by padding to byte boundary

description:

This field shall indicate whether the BS chose to reject the registration, and if so, why. If the length specification for this field is zero, then the registration is not rejected. If the length specification's value is 1, then the data bits shall indicate the reason for the rejection. The rejection causes shall conform to Table 101 - Connection Rejection Values.

- ◆ IE: alg choices

variable length: 8 bits to specify length, in units of 1 bits

followed by padding to byte boundary

description:

The alg choices field shall identify a set of PDCL optional algorithms. The LSB of this field shall correspond to algorithm 0. The PDCL optional algorithms shall conform to Table 94 - Packet Data Conversion Algorithms.

Table 99: Mapping of Frequency and Page Resource to PCH logical mask (0 is LSB)

Timeslot	Frequency				
	f0	f1	f2	...	f_{N_f-1}
0	0	3	6	...	$3N_f-3$
1	1	4	7	...	$3N_f-2$
2	2	5	8	...	$3N_f-1$

Table 100: L2 MAC RA retry policy parameter settings

Na		Nb		Ne	
000	1	000	1	00	9
001	3	001	2	01	15
010	4	010	3	10	66
011	6	011	4	11	120
100	8	100	5		
101	18	101	8		
110	36	110	16		
111	All	111	32		
Nc		Nd		Nf	
000	1	000	1	00	50
001	2	001	2	01	100
010	3	010	3	10	200
011	4	011	4	11	400
100	5	100	5		
101	8	101	8		
110	16	110	16		
111	32	111	32		

Table 101: Connection Rejection Values

Value	Rejection/Cause
0	unknown/ unspecified
1	the UT's protocol version is too old
2	the UT's protocol version is too new
3	the UT's class/capabilities are not supported by the BS
4	the BS's current traffic load is too high
5	the BS's i-HAP decryption of the shared secret failed
6	the i-TAP AFN authentication CA requested by the UT is unknown to the BS
7	the UT's i-TAP certificate is signed by a CA unknown to the BS
8	the BS's i-TAP recovery of the UT certificate message from the signature failed
9	the UT's i-TAP certificate version is not supported by the BS
10	the UT's i-TAP certificate type is not supported by the BS
11	the UT's i-TAP certificate has expired
12	the UT's i-TAP certificate has been revoked
13	the BS's i-TAP decryption of the UT's authentication string failed
14	the BS's i-TAP verification of the UT's authentication message failed
15	the BS has no i-HAP certificate matching the parameters requested by the UT (CA, signature generation method, public key alg, shared secret exchange method, bulk encryption alg)
16	the bulk encryption algorithm requested by the UT is invalid (it was not advertised in the BS certificate)
17	the bulk encryption key length requested by the UT is invalid (either larger than the maximum supported by the BS, or smaller than the minimum)
18-31	Reserved

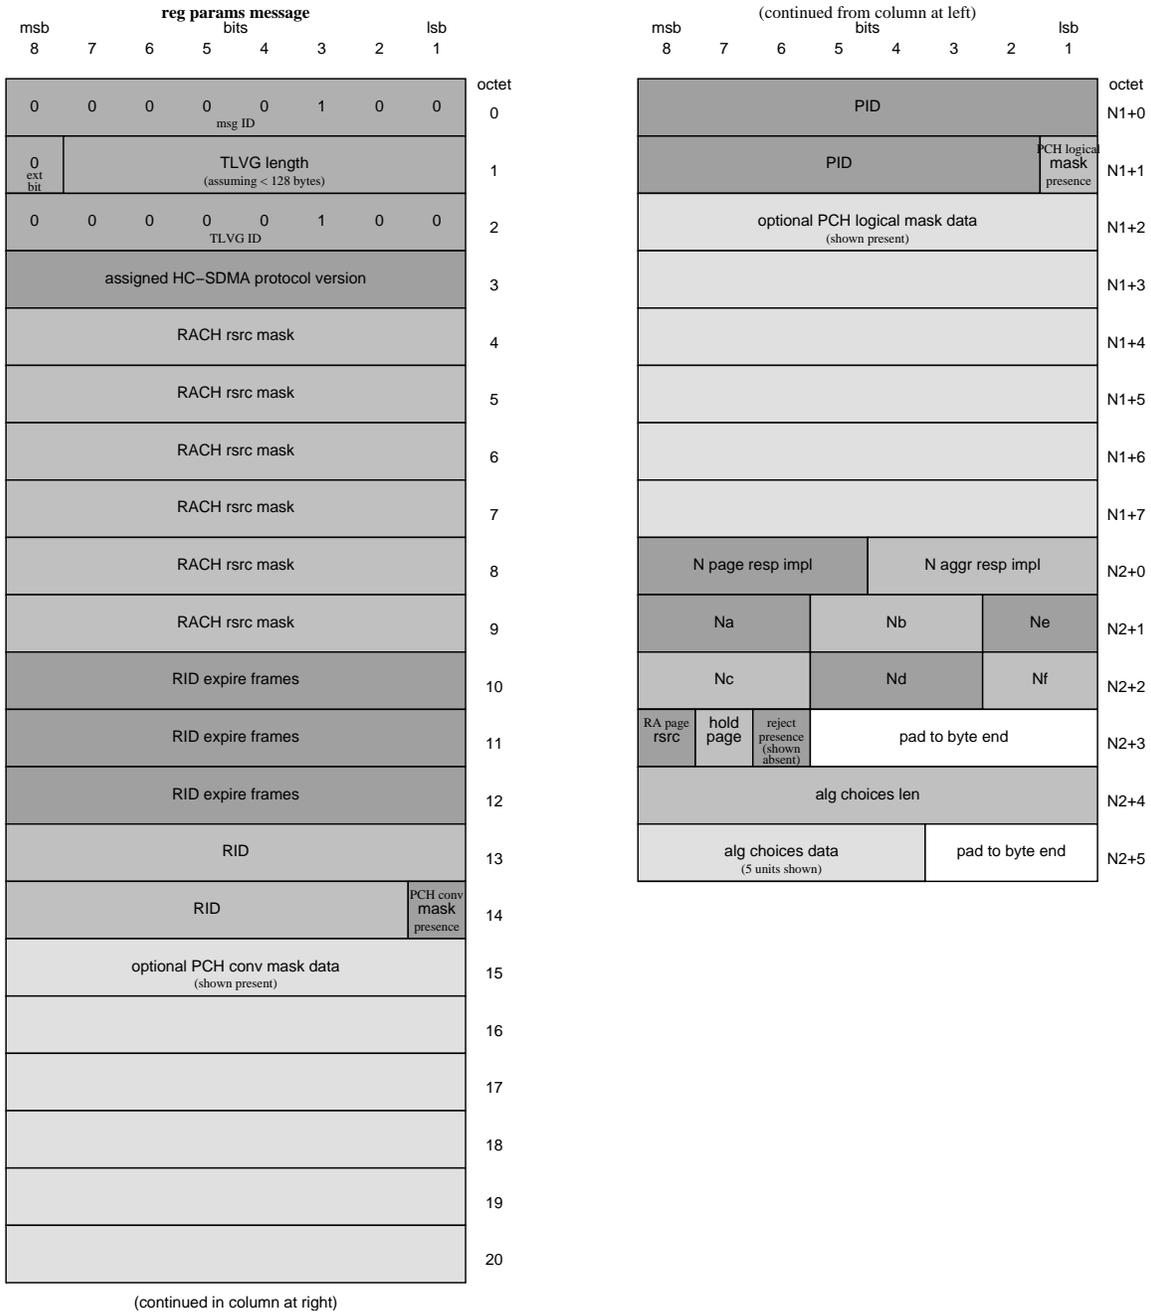


Figure 90: Reg Params Message

9.5.1.3.5 Net Sess Start Message

Message ID: 5
(no required TLVGs)

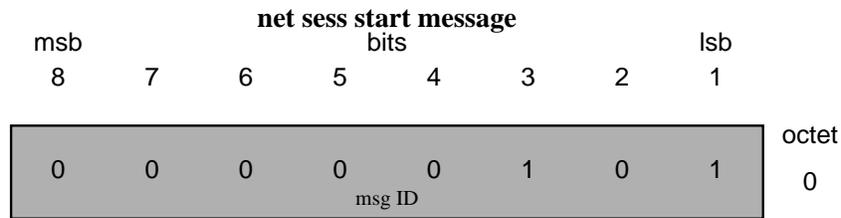


Figure 91: Net Sess Start Message

9.5.1.3.6 Net Sess Stop Message

Message ID: 6
(no required TLVGs)

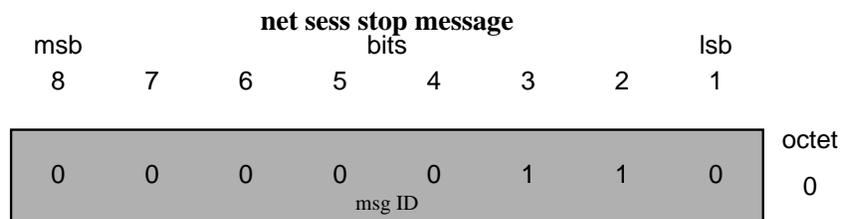


Figure 92: Net Sess Stop Message

9.5.1.3.7 In Band Page Message

Message ID: 7
TLVG: in band page std
TLVG ID: 7
Information Elements:

- ◆ IE: rsrc hint

variable length: 2 bits to specify length, in units of 6 bits
followed by padding to byte boundary
description:

The resource hint list shall provide zero or more indications to the UT of specific channels that it should attempt to open streams on, and zero or one specifications of how many undirected streams it should try to open with a random channel scanning strategy. Each set of 6 bits in the resource hint list field shall be a separate resource hint. If a resource hint has a value in the range 0-47, it shall specify that the UT should try to open a stream on the conventional channel with the corresponding number. If a resource hint has a value in the range of 48-63, it shall specify that the UT should try to open (value-47) undirected streams. If the UT receives a subsequent message with a rsrc hint list, any directed resource hints shall be applied in addition to previous directed resource hints, and if there is an undirected specification, it shall be applied instead of any previous undirected specification. This resource hint shall be interpreted the same way as the aggregate.cmd UM control message's resource hint. See clause 9.5.2.2, aggregate.cmd.

- ◆ IE: opt RID

variable length: 1 bits to specify length, in units of 15 bits
followed by padding to byte boundary

description:

The optional registration ID (RID) shall be used either to specify a RID (for registrations that want to open a TCH stream), or to show that the registration receiving the message has no valid RID (for new registrations that want to open a registration stream). In the latter case, since no Nrddrp value has been assigned by the BS yet, the UT shall make up to a constant number of directed RA-rreg attempts. See clause 7.3.5.2.2, Directed RACH retry policy for further information.

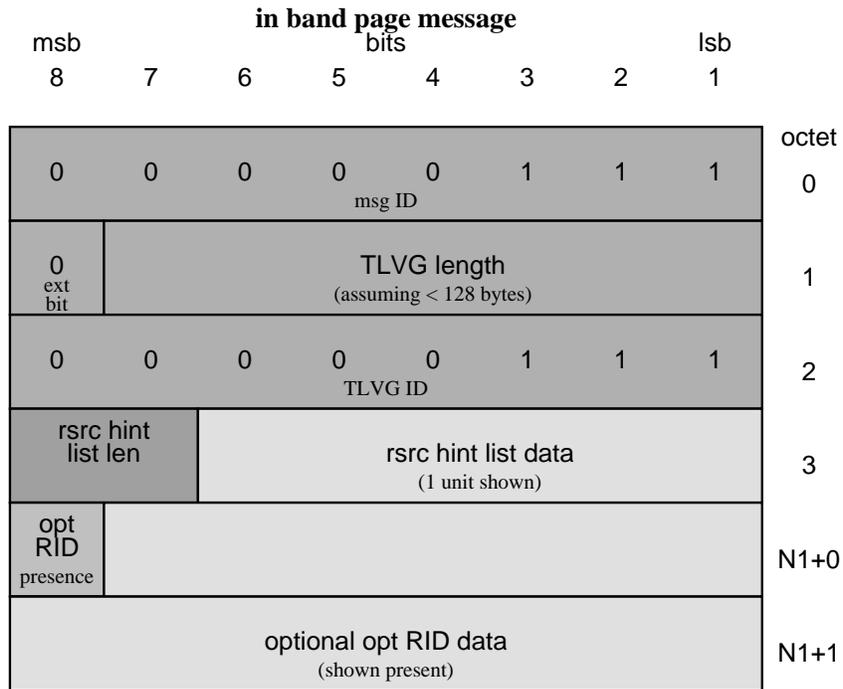


Figure 93: In Band Page Message

9.5.1.3.8 Rsrc Map Message

Message ID: 8

TLVG: rsrc map std

TLVG ID: 8

Information Elements:

- ◆ IE: RACH rsrc mask, 48 bits

description:

The RACH resource mask shall specify which conventional channels of the BS the UT is allowed to use for RA-rts. The LSB of this bit map shall correspond to conventional channel ID 0. See Glossary clause A.25 for a definition of the conventional channel ID.

- ◆ IE: PCH conv mask

variable length: 1 bits to specify length, in units of 48 bits followed by padding to byte boundary

description:

The PCH conventional mask shall specify which conventional channels page bursts may be transmitted on. If the length of this field is zero, then the PCH conventional mask shall be identical to the RACH rsrc mask. Otherwise, the data shall be a bit map of conventional channel IDs to be

used for PCH (see Glossary clause A.25). The LSB is conventional channel ID 0. This mask shall determine the subslot structure for each carrier as defined in clause 7.3.3.2, Base Station Page Procedures. The PCH conventional mask shall be applied to all concurrent registrations between the BS and UT. Therefore, the PCH conventional mask in the reg params message shall be the same for all concurrent registrations between the BS and a single UT. The PCH conventional mask in the rsrc map message may be different than the PCH conventional mask in the reg params message, and it shall update the PCH conventional mask for all concurrent registrations on the UT. See clause 7.3.3, Paging Channel (PCH) for a complete description of how the PCH conventional mask is used.

◆ IE: PCH logical mask

variable length: 1 bits to specify length, in units of 48 bits

followed by padding to byte boundary

description:

The PCH logical mask shall specify which paging resources (see clause 7.3.3.2, Base Station Page Procedures) are valid for this UT. The mapping from the PCH logical mask to paging resources on each frequency shall conform to Table 99- Mapping of Frequency and Page Resource to PCH logical mask (0 is LSB). The valid paging resources shall form the page hopping sequence for the UT (see clause 7.5, Page Hopping Sequences). If the length of this field is zero, then the PCH logical mask shall be identical to the PCH conventional mask. The PCH logical mask shall be applied to all concurrent registrations between the BS and UT. Therefore, the PCH logical mask in the reg params message shall be the same for all concurrent registrations between the BS and a single UT. The PCH logical mask in the rsrc map message may be different than the PCH logical mask in the reg params message, and it shall update the PCH logical mask for all concurrent registrations on the UT. See clause 7.3.3, Paging Channel (PCH) for a complete description of how the PCH logical mask is used.

ATIS-0700004.2007 (S2017)

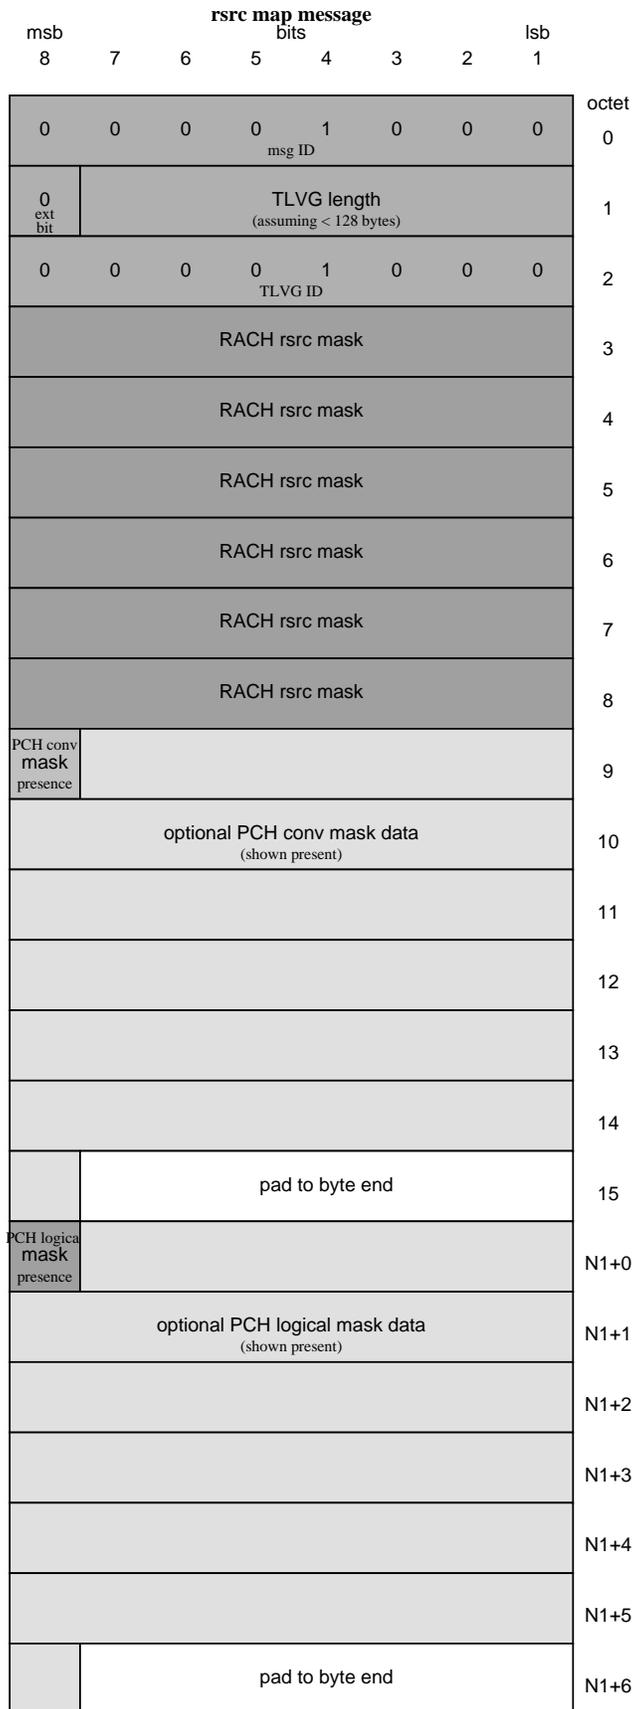


Figure 94: Rsrc Map Message

9.5.1.3.9 RA Latency Request Message

Message ID: 9
(no required TLVGs)

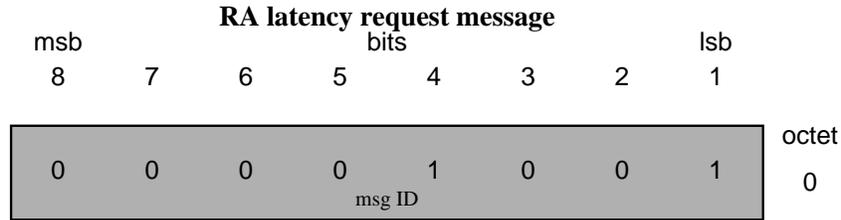


Figure 95: RA Latency Request Message

9.5.1.3.10 RA Latency Report Message

Message ID: 10
TLVG: RA latency report std
TLVG ID: 10
Information Elements:

- ◆ IE: RA latency, 10 bits

description:

The RA latency field shall indicate the number of frames elapsed from the first RA attempt to the reception of the AA-cts. Note that this shall include the backoff delays when UT L3 RM waits after receiving an L2MacUtStart.conf(fail) primitive before sending the next L2MacUtStart.req. The interpretation of this field shall conform to Table 102 - RA Latency Encoding.

- ◆ IE: total RAs, 10 bits

description:

The total RAs field shall indicate the number of RAs sent by the UT in the course of opening a stream. This shall include both directed and undirected RAs. In the event that multiple L2 MAC RA scans were required to open the stream, this field shall sum the number of RAs across all the scans. The interpretation of this field shall conform to Table 103 - RA Attempts Encoding.

- ◆ IE: direct RAs, 4 bits

description:

The direct RAs field shall state how many directed RAs were sent on the channel specified by the BS while opening a stream in response to a page or aggregation message. This field shall be formatted as an unsigned number. There is a possibility that more than 15 RAs can be sent on a specified channel in the process of opening a stream -- e.g., when the UT hears multiple pages but gets no AA-cts in response to its RA-page-response. In this overflow case, the UT shall set the direct RAs field to its maximum value, 15.

Table 102: RA Latency Encoding

encoded value	latency (frames)
0000xxxxx	0:1:31 ^a
00001xxxxx	32:2:95
0001xxxxxxx	96:4:351
001xxxxxxx	352:8:1375
01xxxxxxx	1376:32:9567
1xxxxxxx	9568:256:140383
111111111	>= 140384

Table 103: RA Attempts Encoding

encoded value	total RA attempts
0000xxxxx	0:1:31 ^a
00001xxxxx	32:2:95
0001xxxxxxx	96:4:351
001xxxxxxx	352:4:863
01xxxxxxx	864:4:1887
1xxxxxxx	1888:8:5975
111111111	>= 5976

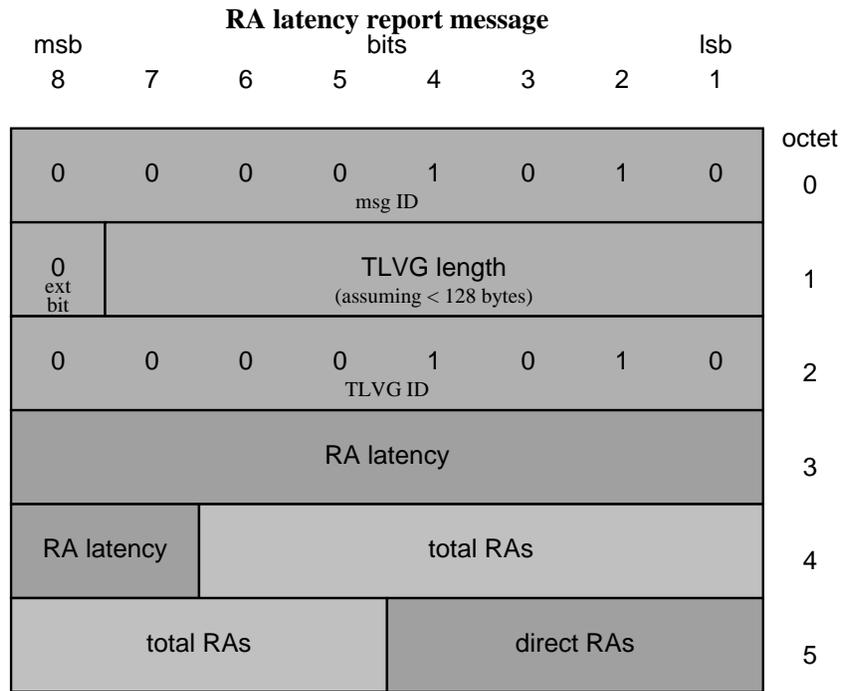


Figure 96: RA Latency Report Message

9.5.1.3.11 Dereg Message

Message ID: 11

TLVG: dereg std

TLVG ID: 11

Information Elements:

- ◆ IE: mandatory, 1 bits

description:

If this field has the value 1, the specified action shall be required; else, it shall be optional.

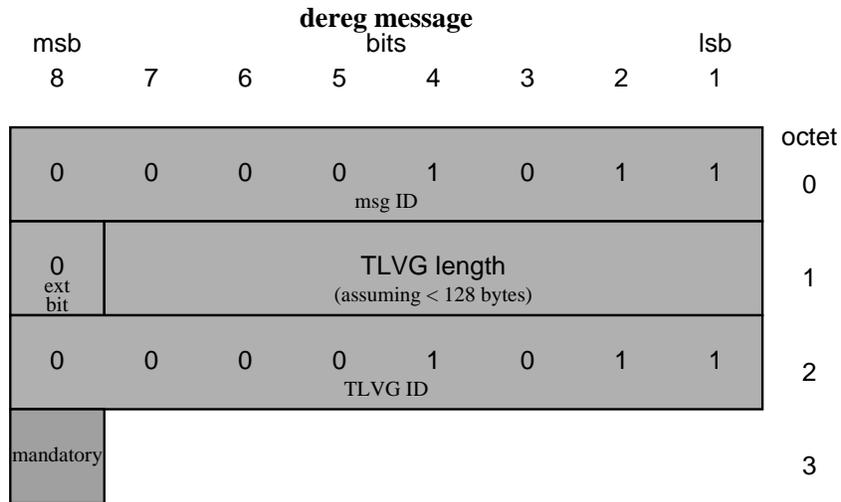


Figure 97: Dereg Message

9.5.1.3.12 QoS Message

Message ID: 12

TLVG: QoS std

TLVG ID: 12

Information Elements:

- ◆ IE: QoS marking, 8 bits

description:

This QoS marking shall indicate the quality of service that the carrier registration should receive -- similar to a Differentiated Services Code Point (DSCP). Rather than directly specifying service quality, this marking shall be used as an index into a configurable database of service parameter values.

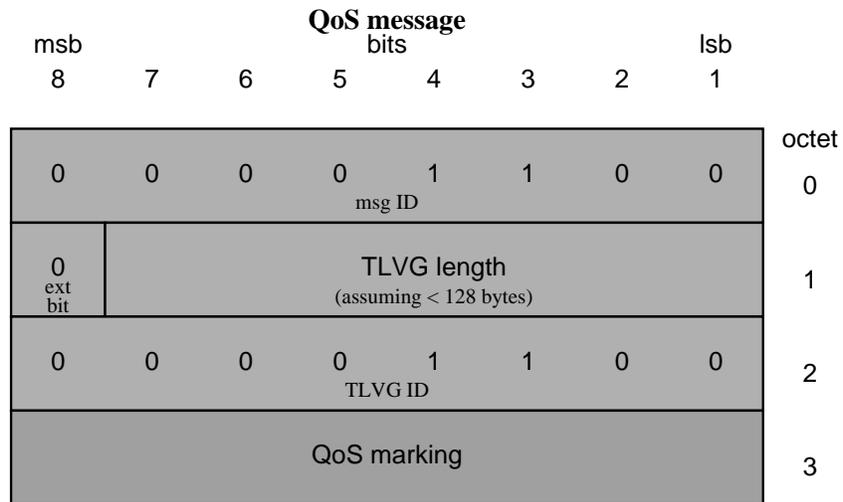


Figure 98: QoS Message

9.5.1.3.13 AFN Verify Message

Message ID: 13

TLVG: AFN verify std

TLVG ID: 13

Information Elements:

- ◆ IE: AFN CA ID, 8 bits

description:

The AFN CA ID shall tell the UT which Certificate Authority signed and returned the UT's authentication string.

- ◆ IE: CA auth str

variable length: 8 bits to specify length, in units of 8 bits

no subsequent padding

description:

The CA-encrypted authentication string shall prove to the UT that the AFN reported in the authentication string is correct, within some small margin.

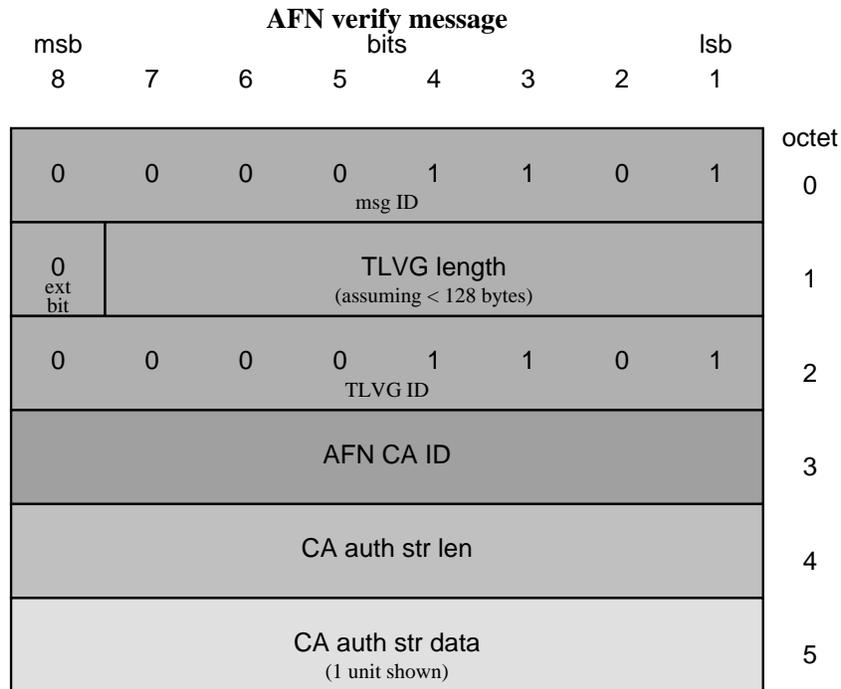


Figure 99: AFN Verify Message

9.5.1.3.14 Co Reg Queue Report Message

Message ID: 14

(no required TLVGs)

optional TLVG (typically contained in this message): co reg queue report opt

TLVG ID: 14

Information Elements:

- ◆ IE: RID, 15 bits

description:

The registration ID (RID) shall be used to identify which registration generated this report.

- ◆ IE: queue report, 5 bits

description:

The queue report field shall tell the BS approximately how many bytes of uplink data the registration in question has queued on the UT. The BS can use this information to choose to break streams to make room for a stream for the registration in question. The format of this field shall be the same as the format of the queue.rpt UM control message's data field, as is described in clause 9.4.10, Stream Aggregation.

If an existing registration wants to send the co reg queue report AM control message, it shall use the co reg queue report opt TLVG to identify itself and to state how much of an uplink data backlog it has. If a new registration is beginning, it has no way to identify itself, but it may still send the co reg queue report AM control message without the co reg queue report opt TLVG to inform the BS that a new registration on the UT would like to open a registration stream.

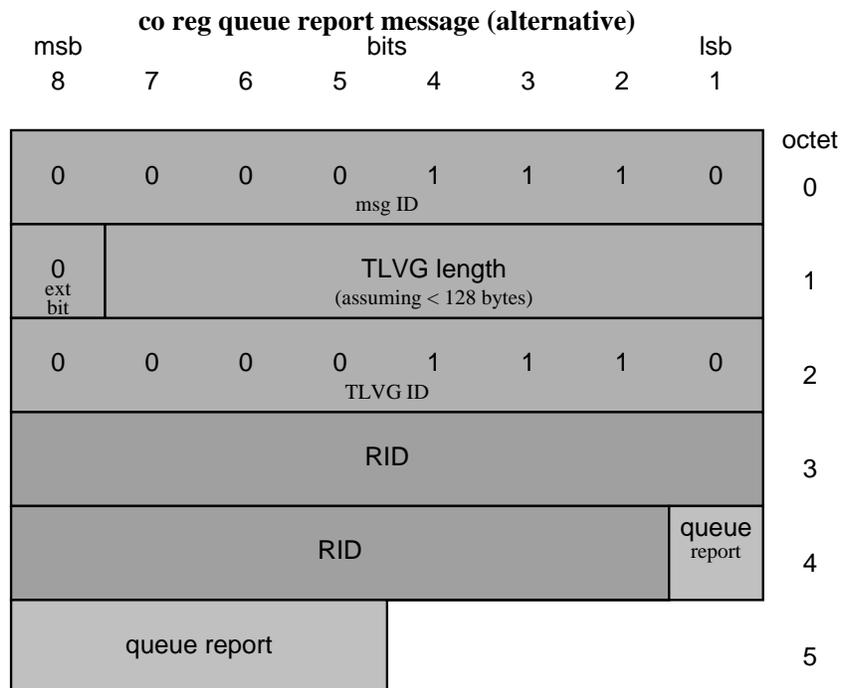


Figure 100: Co Reg Queue Report Message (Existing Registration)

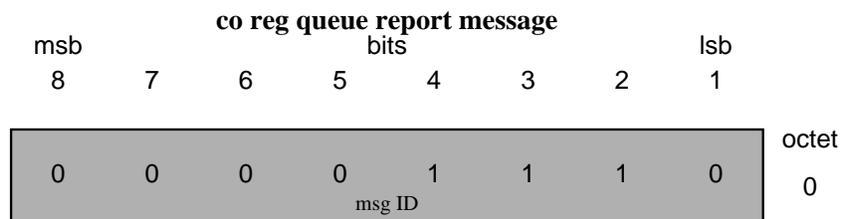


Figure 101: Co Reg Queue Report Message (New Registration)

9.5.1.3.15 RACH Retry Params Message

Message ID: 15

TLVG: RACH retry params std

TLVG ID: 15

Information Elements:

- ◆ IE: Na, 3 bits

description:

Shall be the number of resources per RA-retry sub-scan (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nb, 3 bits

description:

Shall be the number of frames to wait between RA tries within a sub-scan (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Ne, 2 bits

description:

Shall be the maximum frames of backoff between RA tries (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nc, 3 bits

description:

Shall be the first parameter for determining the number of frames to wait between RA-retry sub-scans (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nd, 3 bits

description:

Shall be the second parameter for determining the number of frames to wait between RA-retry sub-scans (see Table 100 - L2 MAC RA retry policy parameter settings).

- ◆ IE: Nf, 2 bits

description:

Shall be how many frames to pursue RA retry strategy before conceding failure (see Table 100 - L2 MAC RA retry policy parameter settings).

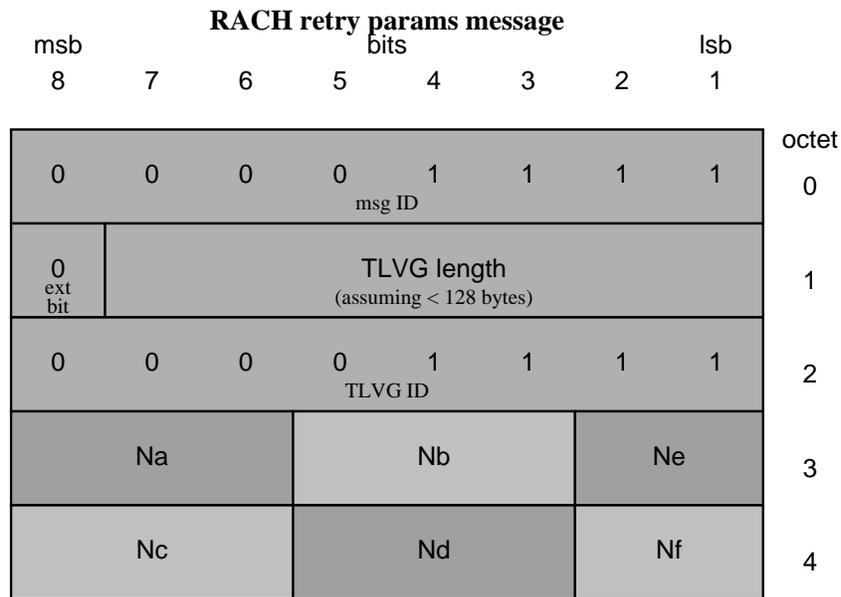


Figure 102: RACH Retry Params Message

9.5.2 Unacknowledged Mode (UM) Messages

Unacknowledged mode messages are used for time sensitive control messages or messages that do not otherwise require reliable delivery. Unacknowledged mode messages shall bypass L2 RLC and directly utilize the L2 MAC.

The directionality of L3 UM messages between L3 peer components shall be as indicated in Table 104 - L3 UM Messages. Note that there are other UM messages within the *HC-SDMA* protocol, but only the ones listed below involve L3.

Table 104: L3 UM Messages

Message	Endpoints
queue.rpt	UT L3 RM --->□BS L3 RM
aggregate.cmd	UT L3 RM <---□BSL3 RM
starve.cmd	UT L3 RM <--> BS L3 RM
hold.cmd	UT L3 RM <--- BS L3 RM
sinrErr.rpt	UT L3 RRC --->□BS L3 RRC
pwrCtrl.cmd	UT L3 RRC <---□BS L3 RRC

Each L3 UM message is described below. See Table 84 - UM Message Definitions for a listing of UM message tags and formatting and clause 7.3.7.2.3, UM Message Insertion for a description of UM message prioritization.

9.5.2.1 queue.rpt

The `queue.rpt` message shall tell BS L3 RM how much data the UT has backlogged. The payload of this message shall be five bits long. See clause 9.4.10, Stream Aggregation for a description of how this message’s data is interpreted.

9.5.2.2 aggregate.cmd

The `aggregate.cmd` message shall be tell UT L3 RM to open a new stream. The payload of this message shall be 6 bits long. The payload shall contain a resource hint to indicate to the UT which conventional channel to use. Values 0 through 47 shall be used to indicate a particular conventional channel -- these are directed aggregation messages. Values 48 through 63 shall indicate that the UT should use the undirected RACH retry policy to open (*value* - 47) streams -- these are undirected aggregation messages. Directed aggregation messages shall be treated independently, in parallel. That is, the BS may send the UT multiple directed aggregation commands even within the same burst, and the UT should open one stream for each aggregation command, on the conventional channel specified in the command. In contrast, only one undirected aggregation command shall be in effect at a time. Each new undirected aggregation command shall cancel the previous one. For example, if the BS were to send the UT an `aggregate.cmd` message with a value of 49, the UT shall attempt to open two streams with the undirected RACH retry policy. If, after opening one stream, the UT received another `aggregate.cmd` message with a value of 50, the UT shall attempt to open three more streams, resulting in four new streams total, rather than the two new streams originally specified.

9.5.2.3 starve.cmd

The `starve.cmd` message shall tell the recipient to not send any more packets on the stream, but to continue to receive packets. The recipient may complete transmission of all packets already queued for the stream in L2 RLC. There shall be one bit of data associated with this message - a Boolean indication of whether the cessation of data on the stream is intended to be unidirectional or bidirectional. The conclusion of data transfer shall be unidirectional if the bit is cleared, and bidirectional if the bit is set. In the latter case, the recipient L3 RM should use the `L2RlcConfig.req` primitive to make L2 RLC’s idle shutdown limit small, so that the stream closes quickly once all the data already enqueued in the stream has been transferred.

9.5.2.4 `hold.cmd`

The `hold.cmd` message shall tell the UT how long it must wait before requesting the next stream. When its last stream closes, UT L3 RM shall check if it received any `hold.cmd` messages during the period when it continuously had streams open, and if so waits for the specified time before requesting any new streams for the registration. However, if the UT receives a page during this holding period, it shall respond immediately. Furthermore, if this page response results in a stream being opened, the UT shall cancel the hold period. When more than one `hold.cmd` UM control messages are sent to the UT, UT L3 RM shall remember only the hold time from the most recently received `hold.cmd` message. The hold time shall be applied only once. After the UT again closes its last stream, it shall not impose the hold time again unless it is told to again by the BS during the most recent set of concurrent streams. See the `hold.cmd` entry in Table 84 - UM Message Definitions regarding how the UT determines the hold time from the contents of the `hold.cmd` message.

9.5.2.5 `sinrErr.rpt`

The `sinrErr.rpt` message shall tell BS L3 RRC what the UT's downlink SINR is. The payload of this message shall be five bits long. See clause 9.2.3.4.2, Closed-loop Power Control for a description of how the payload is interpreted.

9.5.2.6 `pwrCtrl.cmd`

The `pwrCtrl.cmd` message shall tell UT L3 RRC what adjustment to make to the UT's transmit (uplink) power for that stream. The payload of this message shall be five bits long. See clause 9.2.3.4.2, Closed-loop Power Control for a description of how the payload is interpreted.

10 PROTOCOL LAYER PRIMITIVES (INFORMATIVE)

10.1 Overview

This chapter lists the primitives used by the HC-SDMA protocol layers to interact with each other. A general description is provided for each primitive, along with a list of the data arguments for the primitive, and a description of each data argument. The lists of primitives are preceded by a summary of the data types that are used as arguments for any of the primitives. Since these primitives concern messages between protocol layers within the HC-SDMA UT entity and within the HC-SDMA BS entity, text in this chapter is informative.

Note that in the definition of these primitives, it is frequently assumed that the primitive can be directed to a specific recipient. Otherwise, each primitive would require an additional "intended recipient(s)" argument. For example, a BS will generally have many registrations active simultaneously; it is assumed that there is an underlying method to direct a `L2R1cData.ind` primitive to the registration which owns the stream producing the incoming data.

10.2 Data Types

There are three categories of data types used as arguments for the protocol layer primitives:

1. Fundamental types which are general in nature and widely recognized.

2. Enumeration types which list symbolic names representing a set of values an object can have.
3. Synonym types which are equivalent to a fundamental or enumeration type, but which have been given a more demonstrative name.

Table 105: Fundamental Data Types

Type	Description
Bit	a binary digit
Bit_String	a sequence of bits
Boolean	logical true/false
Duration	a time span
Integer	a whole number
Natural	a non-negative whole number
Octet_String	a sequence of octets/bytes
Process ID	a unique identification of a process / active class (a process is an independent execution instance)
PktQueue	an ordered list of packets
Real	a real number

Table 106: Enumerated Data Types

Enumeration type	Description	Enumeration values
Direction	enumeration of data flow directions The directions are defined in two ways. First, with respect to the air interface. "Incoming" means from the air interface, "outgoing" means to the air interface. Second, with respect to the system as a whole. "Uplink" is in the direction from the end user device (EUD) to the network, "downlink" is in the direction from the network to the EUD. On the UT, "incoming" is "downlink" and "outgoing" is "uplink". On the BS, "incoming" is "uplink" and "outgoing" is "downlink".	incoming, outgoing, uplink, downlink, bidirectional
ibBurstTypeSort	burst types in HC-SDMA	F, T, B, CR, CM, RA, AA, TCH, PCH
ibSubRateSort	with what fraction of the frames does the specified sub-channel use its channel	ibSubRateInvalid, ibSubRateFull, ibSubRateHalf, ibSubRateThird, ibSubRateQuarter, ibSubRateFifth
itapUtCertTypeSort	distinguishes between the different kinds of UT certificates	itapUtCertTypeIdentity, itapUtCertTypeServiceAuth, itapUtCertTypeNetAccess, itapUtCertTypeAny
L2MacL3RrcOneOrSixBitSort	indication from L3 RRC of whether the dedicated 1-bit RMU header field is sufficient, or if it would prefer to send a UM ctrl message with a 5-bit payload too	OneBitType, SixBitType
L2MacPageLevel	how frequently to send/listen for pages <ul style="list-style-type: none"> • pageLv1 -- every frame • pageLv2 -- every 8th frame • pageLv3 -- every 64th frame • pageLv4 -- every 512th frame • pagePassive -- don't send or listen for pages, but if the BS gets a RA-ping, respond with AA-cts • pageOff -- no paging 	pageLv1, pageLv2, pageLv3, pageLv4, pagePassive, pageOff
L2MacPrioritySort	indication of whether BCH/CCH should supercede PCH/RACH/TCH	AboveTchRachPch, BelowTchAboveRachPch, BelowTchRachAbovePch, BelowTchRachPch
L2MacStopCauseSort	indication of why a stream should be / was stopped	L2MacNormalClose, L2MacInitialFrameErrs, L2MacExcessFrameErrs, L2MacPeerClose, L3RmStop, L3RmPktHdrErr, L2RlcBadSeq, L2RlcBadAck, L2RlcTxStuck, L2RlcRxStuck

Enumeration type	Description	Enumeration values
L2MacStreamType	types of streams supported by L2 MAC	StreamTypeNone, StreamTypeReg, StreamTypeTchUpl, StreamTypeTchDwl, StreamTypeShort, StreamTypePing
L2MacUmCtrlType	what types of control messages may be sent on the UM control channel. The UM control messages are classified as either reports or commands.	queue.rpt, aggregate.cmd, hold.cmd, starve.cmd, sinrErr.rpt, pwrCtrl.cmd, done.rpt, close.cmd, timing.rpt
L2MacUtRachScanModeSort	specification of whether L2 MAC is sending RACH messages on a specific channel that the BS directed it to use, or randomly scanning through the channels	directed, unDirected
L2MacUtReqResultSort	indication of whether a request for an RA/AA exchange (either to open a stream or do a ping) succeeded (the UT sent a RA and received an AA) or failed; and if it failed, why (RA sent but no AA received, or unable to send RA). Note that the stream/ping can still fail even if the RA/AA exchange succeeded if the AA was a negative-acknowledgement or an indication of an invalid ID.	noGrant, rachFail, reqSuccess
L2RlcChan	differentiation between the reliable (AM) versus datagram (UM) data services offered by L2 RLC	AM, UM
L2RlcStrmState	enumeration of the externally-relevant states within L2 RLC's stream state machine (L2 MAC needs some visibility into the stream state so it can gracefully close idle streams)	noClose, rlcIdle, closeIfReady
PathlossBurstType	specification of what types of measurements should be used when reporting a pathloss estimate	BBurst, MostAccurate

Table 107: Synonym Data Types

Synonym type	Equivalent type	Description
AaDlModClassSort	Integer	4 bit field in AA burst stating the mod class of the first downlink TCH burst
AaFrameDecSort	Integer	4 bit field in AA burst. Defines fractional-rate channels
AaPwrAdjustSort	Integer	4 bit field in AA burst that indicates the power adjustment the UT must apply to the initial uplink burst of the subsequent stream
AaRaDelaySort	Integer	4 bit field that indicates the amount of time (frames) after which the UT should send another RA for the extended RA mechanism
AaSubTypeSort	Integer	4 bit field in AA burst
AaTchIndexSort	Integer	3 bit field in AA burst that indicates the TCH index for the subsequent stream
AaUlModClassSort	Integer	4 bit field in AA burst stating the mod class of the first uplink TCH burst
AfnDiffSort	Integer	difference between two absolute frame numbers (AFNs). The AFN is 48 bits, so AfnDiffSort's values range from $-(2^{48}-1):2^{48}-1$
AfnSort	Integer	data sort to hold the absolute frame number. The protocol limits the AFN to 48 bits, so AfnDiffSort's values range from $0:2^{48}-1$
BitString	Bit_String	
BsLoadSort	Integer	2 bit field in BCH B-burst indicating how heavily the BS is currently being utilized
BsTxPwrSort	Integer	4 bit field in BCH B-burst used for open-loop power control
ByteString	Octet_String	
CarrierMaskSort	Integer	data sort specifying which carriers within a 10 MHz band are enabled/disabled. The LSB (bit 0) specifies whether carrier 0 is in use, bit 1 specifies whether carrier 1 is in use, etc. If the bit is set, the corresponding carrier is enabled.
CmNetworkIdSort	Integer	8 bit field identifies which operator owns the BS
CmPwrAdjustSort	Integer	4 bit field in the CM burst which indicates what power adjustment the UT should apply to future configuration request and random access bursts
CmUlPwrCtrlActiveSort	Boolean	indicates whether uplink closed-loop power control is active
CrAttemptNumSort	Integer	the total number of attempts for a CR
CrRandomIdSort	Integer	8 bit field in CR burst; a random identifier used to distinguish multiple simultaneous CR bursts
CrTxPwrSort	Integer	5 bit field in the CR burst used for open-loop power control
DecodeFailSort	Boolean	did the decoded data have a correct CRC
DetectFailSort	Boolean	was the burst present, or did the detection fail
DlModClassSort	Integer	which mod classes may be used on the downlink
DssiDbmSort	Real	data sort for specifying DSSI in dBm units
DuplexTimeSlotSort	Integer	uplink/downlink timeslot pairs
FacchDataSort	Integer	4 bit data used to be sent in FACCH

Synonym type	Equivalent type	Description
FacchErrSort	Boolean	is it probable that the FACCH decode was correct, or is it virtually certain that the FACCH decode was in error. Since FACCH has no CRC, the only indication of whether the FACCH was decoded correctly is from checking the correlation (normalized by DSSI) of the received FACCH symbols with the symbols generated by encoding the presumed FACCH data
ibBsc	Integer	HC-SDMA Base Station Color Code
ibChanNum	Integer	an identification of a particular RF channel within an HC-SDMA system. Although a BS may cover more than 10 MHz, a UT will at most cover 10 MHz, so the number of channels in an HC-SDMA channel set is $(10 \text{ MHz} / 625 \text{ KHz}) * 3$ (timeslots) = 48
ibFreqAdjHzSort	Real	how much to change the transceiver unit's clock frequency (applies to UT only). The adjustment is with respect to the HC-SDMA baud frequency (500 KHz). A positive adjustment means to increase the transceiver's clock frequency. For example, if the RF (radio frequency) is nominally 1900 MHz but needs to be 300 Hz higher, then the corresponding <code>ibFreqAdjustHzSort</code> is $300 / 1900e6 * 500,000 \text{ Hz} = 0.079 \text{ Hz}$.
ibFreqErrHzSort	Real	the difference between the expected frequency at which a burst will be received and the actual frequency measured. The error is with respect to the HC-SDMA baud frequency (500 KHz). A positive frequency error means that the burst was received at a higher frequency than expected. (Or, in the case of a UT, that the frequency downconversion to baseband was less than expected.) For example, if a burst is expected to be received at 1900 MHz and it is actually 300 Hz higher (1900.0003 MHz), the corresponding <code>ibFreqErrorSort</code> value would be $(1900.0003 - 1900) / 1900 * 500,000 \text{ Hz} = 0.079 \text{ Hz}$.
ibFreqKhzSort	Real	generic sort for frequency values in KHz
ibPageId	Integer	HC-SDMA paging ID The protocol limits the paging ID to 15 bits, so <code>ibPageIds</code> range from $0:2^{15}-1$
ibPktProtocol	Integer	encapsulation protocol layers of packets being tunneled across the HC-SDMA air interface
ibQos	Integer	HC-SDMA quality of service marking
ibRegId	Integer	HC-SDMA registration ID The protocol limits the registration ID to 15 bits, so <code>ibRegIds</code> range from $0:2^{15}-1$
ibRegPageId	Integer	a generic ID that could be either a registration ID or a paging ID, depending on the context
ibSubChan	Integer	different subchannel rates and temporal occupancy patterns. The sub-chan classification has 4 bits in the AA messages. The resulting 16 values are used to indicate whether the channel's rate is decimated by a factor of 1, 2, 3, 4, or 5, and for the chans which are not full-rate, which sub-channel is intended.
ihapCaIdSort	Integer	identification of a certificate authority
itapAuthMsgCleartxtSort	Bit_String	i-TAP cleartext authentication message (concatenation of authentication message fields)
itapAuthStrEncrSort	Bit_String	i-TAP encrypted authentication string (ECC-163 encryption of cleartext authentication string)

Synonym type	Equivalent type	Description
itapUtCertDataSort	Bit_String	i-TAP UT certificate contents/message
L1ScrambleIdSort	Integer	info used to scramble a burst All bursts except BCH F are scrambled. All the scrambled bursts except PCH use the BSCC (and other info, for some burst types) to create the scrambling pattern. PCH uses the paging ID and some LSBs of the AFN.
L2MacGrantSort	Natural	counter of how many times the request to receive/transmit a burst was granted
L2MacL3RrcTchPwrCtrlInfo1BitSort	Boolean	single-bit report that the downlink SINR is above/below the target, or command to increase/decrease the uplink transmit power by 1 dB
L2MacL3RrcTchPwrCtrlInfo6BitSort	Integer	6-bit report of the downlink SINR, or command to adjust the uplink transmit power
L2MacUmCtrlData	Integer	sort for data component of UM control messages Some UM control messages have no parameters, but most do. This sort provides storage for parameters for those UM control message types which do have parameters.
L2MacUmCtrlMaxTx	Integer	how many times a UM ctrl message may be transmitted
L2MacUmCtrlValidFrms	Integer	how many frames a UM ctrl message will be valid
L2MacValidFrmsSort	Natural	validity time of outstanding burst requests in frames
L2RlcIdleLimit	Integer	the number of frames that a stream may remain idle (no data being transferred in either direction) before it gets shut down as part of the graceful stream closure
LinkAdaptDataSort	Integer	4 bit data used by link adaptation for reporting AP(Available Power) , RMC (Recommended Mod Class), or CMC (Current Mod Class)
MaxCrAttemptsSort	Integer	max number of attempts in a CR scan
ModClassSort	Integer	data sorts for specifying uplink and downlink HC-SDMA mod classes
PathlossDbSort	Real	data sort to store pathloss (in dB units)
Pkt	Octet_String	synonym for octet string, used for generic packets throughout the HC-SDMA protocol stack
PowerDbmSort	Real	data sort for specifying powers in dBm units
PowerDbSort	Real	data sort for specifying powers in dB units
RaTxPwrSort	Integer	5 bit field in RA burst used for open-loop power control
RfnSort	Integer	data sort to hold the relative frame number The protocol limits the AFN to 48 bits, so RfnSort's values range from 0:2 ⁴⁸ -1
RmuErrSort	Boolean	sort specifying whether the frame failed the CRC check
RssiDbmSort	Real	data sort for specifying RSSI in dBm units
SinrDbSort	Real	data sort for specifying SINR in dB units
SubSlotSort	Integer	for PCH, which half of the timeslot does the paging burst occupy
TchAckSort	Integer	12 bits in the RMU header that indicate the sequence number of the last AM octet received by the peer

Synonym type	Equivalent type	Description
TchSeqNumSort	Integer	13 bits in the RMU header that indicate the sequence number of the first AM octet in the payload. If the AM payload size is 0 octets, this value is ignored.
TimeslotMaskSort	Integer	data sort specifying which timeslots are enabled or disabled. The LSB (bit 0) specifies whether timeslot 0 is in use, bit 1 -> timeslot 1, and bit 2 -> timeslot 2. If the bit is set, the corresponding timeslot is enabled.
TimingAdjSort	Duration	a timing adjustment used either to correct for clock error or for detailed specification of when to transmit or receive a burst. A positive adjustment means to jump ahead in time with respect to the timing reference; a negative adjustment means to jump back in time.
TimingErrSort	Duration	the difference between the expected timing of an event (primarily reception of a burst) and the actual timing measured. A positive timing error means that the event occurred later than expected.
UlModClassSort	Integer	which mod classes may be used on the uplink

10.3 Primitive Lists and Descriptions

10.3.1 Interface List

The interfaces between the layers in the HC-SDMA protocol stack are listed below. Figure 103 - Protocol Layer Interactions depicts these interfaces:

- ◆ 10.3.2.1, L1 to L2 MAC Interface Primitives
- ◆ 10.3.2.2, L2 MAC to L1 Interface Primitives
- ◆ 10.3.2.3, L3 MMC to L1 Interface Primitives
- ◆ 10.3.2.4, L2 MAC to L2 RLC Interface Primitives
- ◆ 10.3.2.5, L2 RLC to L2 MAC Interface Primitives
- ◆ 10.3.2.5, L2 RLC to L2 MAC Interface Primitives
- ◆ 10.3.2.6, L2 MAC to L3 RRC Interface Primitives
- ◆ 10.3.2.7, L3 RRC to L2 MAC Interface Primitives
- ◆ 10.3.2.8, L2 MAC to L3 MMC Interface Primitives
- ◆ 10.3.2.9, L3 MMC to L2 MAC Interface Primitives
- ◆ 10.3.2.10, L2 MAC to L3 RM Interface Primitives
- ◆ 10.3.2.11, L3 RM to L2 MAC Interface Primitives
- ◆ 10.3.2.12, L2 RLC to L3 RM Interface Primitives
- ◆ 10.3.2.13, L3 RM to L2 RLC Interface Primitives
- ◆ 10.3.2.14, L3 RRC to L3 MMC Interface Primitives
- ◆ 10.3.2.15, L3 MMC to L3 RRC Interface Primitives
- ◆ 10.3.2.16, L3 MMC to L3 RM Interface Primitives
- ◆ 10.3.2.17, L3 RM to L3 MMC Interface Primitives
- ◆ 10.3.2.18, L3 RM to L3 CM Interface Primitives
- ◆ 10.3.2.19, L3 CM to L3 RM Interface Primitives
- ◆ 10.3.2.20, L3 CM to L4 Interface Primitives
- ◆ 10.3.2.21, L4 to L3 CM Interface Primitives

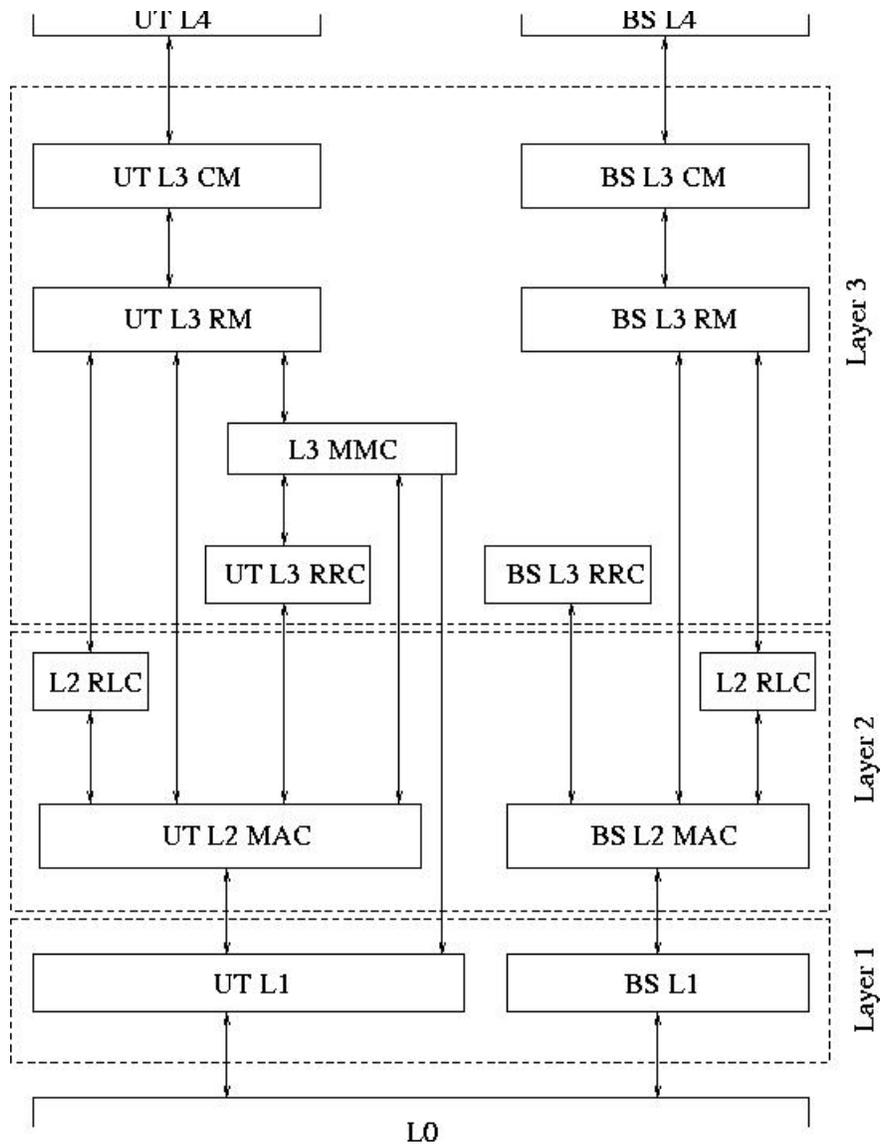


Figure 103: Protocol Layer Interactions

10.3.2 Individual Interfaces

10.3.2.1 L1 to L2 MAC Interface Primitives

Primitive list:

- ◆ L1TxSched.ind
- ◆ L1RxSched.ind
- ◆ L1RxDetect.ind
- ◆ L1RxDecode.ind
- ◆ L1Frm.ind

Indication primitive descriptions

<i>L1TxSched.ind</i>	
Applies to:	BS, UT
Associated primitives:	L1TxSched.resp
<p>Description: indication from L1 to L2 MAC to prompt L2 MAC to provide bursts to transmit, should L2 MAC want to do so. L2 MAC must respond with a L1TxSched.resp. If L2 MAC doesn't want to transmit, then the L1TxSched.resp will have an empty list of bursts. Note that the L1TxSched.ind is for all transmit timeslots. That is, L2 MAC should schedule the bursts to be transmitted on all 3 timeslots when it receives the L1TxSched.ind. Note that L1 will only send L1TxSched.ind to UT L2 MAC for the uplink timeslots, and will only send L1TxSched.ind to BS L2 MAC for the downlink timeslots.</p>	

<i>L1RxSched.ind</i>	
Applies to:	BS, UT
Associated primitives:	L1RxSched.resp
<p>Description: indication from L1 to L2 MAC to prompt L2 MAC to provide specifications of what bursts to listen for, should L2 MAC want to receive bursts. L2 MAC must respond with a L1RxSched.resp. If L2 MAC doesn't want to listen for bursts, then the L1RxSched.resp will have an empty list of bursts. Note that the L1RxSched.ind is for all receive timeslots. That is, L2 MAC should schedule the bursts to be listened for on all 3 timeslots when it receives the L1RxSched.ind. Note that L1 will only send L1RxSched.ind to BS L2 MAC for the uplink timeslots, and will only send L1RxSched.ind to UT L2 MAC for the downlink timeslots.</p>	

<i>L1RxDetect.ind</i>	
Applies to:	BS, UT
Associated primitives:	L1RxSched.resp, L1RxDecode.ind
<p>Description:</p> <p>indication from L1 to L2 MAC stating whether the burst L2 MAC requested to listen for was actually heard. Every burst that L2 MAC requests to listen for will cause a corresponding L1RxDetect.ind.</p>	
<p>Parameters:</p> <ol style="list-style-type: none"> 1. L1RxBurstSort -- the burst that L2 MAC requested to listen for (struct) 2. L1RxBurstSort::L1BurstIdSort -- (burstId) (struct) 3. L1RxBurstSort::L1BurstIdSort::ibBurstTypeSort -- (burstType) 4. L1RxBurstSort::L1BurstIdSort::L1ScrambleIdSort -- (scrambleId) applies to all but BCH F 5. L1RxBurstSort::L1BurstIdSort::AaTchIndexSort -- (tchIndex) applies to TCH bursts only 6. L1RxBurstSort::L1BurstIdSort::ModClassSort -- (modClass) applies to TCH bursts only 7. L1RxBurstSort::L1BurstIdSort::ibBsc -- (bscc) affects selection of training sequence 8. L1RxBurstSort::L1BurstIdSort::Integer -- (burstIndex) This field is used to distinguish multiple CRs, CMs and RAs and AAs in the same conventional channel 9. L1RxBurstSort::Boolean -- (modClassFromFacch) For certain TCH bursts (uplink bursts in odd RFN and downlink bursts in even RFN), the Rx modclass has to be obtained by decoding the FACCH 10. L1RxBurstSort::L1RxBurstRfInfoSort -- (rf) (struct) 11. L1RxBurstSort::L1RxBurstRfInfoSort::ibFreqKhzSort -- (rxFreq) what frequency is the unit listening at 12. L1RxBurstSort::L1RxBurstRfInfoSort::DuplexTimeSlotSort -- (timeSlot) the timeslot on which the burst reception should be attempted 13. L1RxBurstSort::L1RxBurstRfInfoSort::SubSlotSort -- (subSlot) the sub-slot within the timeslot (applies only to PCH bursts) 14. L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort -- (rxWin) when might the burst actually be heard, w.r.t. the nominal timeslot timing (struct) 15. L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort -- (min) what is the earliest burst that can be heard, w.r.t. the nominal timeslot timing 16. L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort -- (max) what is the latest burst that can be heard, w.r.t. the nominal timeslot timing 17. L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort -- (expect) when is the burst expected, w.r.t. the nominal timeslot timing 18. DetectFailSort -- whether the detect failed 19. L1RxBurstRfResultSort -- RF characteristics of the burst (struct) 20. L1RxBurstRfResultSort::SinrDbSort -- (sinr) 21. L1RxBurstRfResultSort::Boolean -- (sinrValid) whether the SINR is valid. For example, when this struct is sent during a L1RxDetect.ind, L1 might choose to not provide a valid SINR until after decode to give a better estimate 22. L1RxBurstRfResultSort::RssiDbmSort -- (rssi) 23. L1RxBurstRfResultSort::Boolean -- (rssiValid) whether the RSSI is valid 24. L1RxBurstRfResultSort::DssiDbmSort -- (dssi) 25. L1RxBurstRfResultSort::Boolean -- (dssiValid) whether the DSSI is valid 26. L1RxBurstRfResultSort::ibFreqErrHzSort -- (freqErr) positive error means the burst's frequency was higher than expected 27. L1RxBurstRfResultSort::Boolean -- (freqErrValid) whether freqErr is valid 28. L1RxBurstRfResultSort::TimingErrSort -- (timeErr) positive error means the burst was at a later time than expected 29. L1RxBurstRfResultSort::Boolean -- (timeErrValid) whether timeErr is valid 	

<i>L1RxDecode.ind</i>	
Applies to:	BS, UT
Associated primitives:	L1RxSched.resp, L1RxDetect.ind
<p>Description:</p> <p>indication that L1 has finished decoding the data contained in a burst that L2 MAC asked to listen for. Any bursts which were successfully detected will result in a L2RxDecode.ind for each distinct payload they contain, as described in the following list:</p> <ol style="list-style-type: none"> 1. no payload - thus no L2RxDecode.ind: BCH F, BCH T, PCH 2. single payload - thus one L2RxDecode.ind per successful detect: BCH B, CR 3. dual payload (FACCH) - thus two L2RxDecode.inds per successful detect: CM, RA, AA, TCH 	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>L1RxBurstSort</i> -- the burst that L2 MAC requested to listen for (struct) • <i>L1RxBurstSort::L1BurstIdSort</i> -- (burstId) (struct) • <i>L1RxBurstSort::L1BurstIdSort::ibBurstTypeSort</i> -- (burstType) • <i>L1RxBurstSort::L1BurstIdSort::L1ScrambleIdSort</i> -- (scrambleId) applies to all but BCH F • <i>L1RxBurstSort::L1BurstIdSort::AaTchIndexSort</i> -- (tchIndex) applies to TCH bursts only • <i>L1RxBurstSort::L1BurstIdSort::ModClassSort</i> -- (modClass) applies to TCH bursts only • <i>L1RxBurstSort::L1BurstIdSort::ibBsc</i> -- (bscc) affects selection of training sequence • <i>L1RxBurstSort::Boolean</i> -- (modClassFromFacch) For certain TCH bursts (uplink bursts in odd RFN and downlink bursts in even RFN), the Rx modclass has to be obtained by decoding the FACCH • <i>L1RxBurstSort::L1RxBurstRfInfoSort</i> -- (rf) (struct) • <i>L1RxBurstSort::L1RxBurstRfInfoSort::ibFreqKhzSort</i> -- (rxFreq) what frequency is the unit listening at • <i>L1RxBurstSort::L1RxBurstRfInfoSort::DuplexTimeSlotSort</i> -- (timeSlot) the timeslot on which the burst reception should be attempted • <i>L1RxBurstSort::L1RxBurstRfInfoSort::SubSlotSort</i> -- (subSlot) the sub-slot within the timeslot (applies only to PCH bursts) • <i>L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort</i> -- (rxWin) when might the burst actually be heard, w.r.t. the nominal timeslot timing (struct) • <i>L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort</i> -- (min) what is the earliest burst that can be heard, w.r.t. the nominal timeslot timing • <i>L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort</i> -- (max) what is the latest burst that can be heard, w.r.t. the nominal timeslot timing • <i>L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort</i> -- (expect) when is the burst expected, w.r.t. the nominal timeslot timing • <i>DecodeFailSort</i> -- whether the data decode failed: true -> decode fail, false -> decode okay • <i>Bit</i> -- value 0 means the time for decode was less than half a frame time, value 1 means more than half a frame time • <i>Boolean</i> -- indication of whether this decode result is for the primary or secondary (FACCH) payload (true -> primary, false -> secondary/FACCH) • <i>BitString</i> -- the decoded data • <i>L1RxBurstRfResultSort</i> -- RF characteristics of the burst (struct) • <i>L1RxBurstRfResultSort::SinrDbSort</i> -- (sinr) • <i>L1RxBurstRfResultSort::Boolean</i> -- (sinrValid) whether the SINR is valid. For example, when this struct is sent during a L1RxDetect.ind, L1 might choose to not provide a valid SINR until after decode to give a better estimate • <i>L1RxBurstRfResultSort::RssiDbmSort</i> -- (rssi) • <i>L1RxBurstRfResultSort::Boolean</i> -- (rssiValid) whether the RSSI is valid • <i>L1RxBurstRfResultSort::DssiDbmSort</i> -- (dssi) • <i>L1RxBurstRfResultSort::Boolean</i> -- (dssiValid) whether the DSSI is valid • <i>L1RxBurstRfResultSort::ibFreqErrHzSort</i> -- (freqErr) positive error means the burst's frequency was higher than expected • <i>L1RxBurstRfResultSort::Boolean</i> -- (freqErrValid) whether freqErr is valid • <i>L1RxBurstRfResultSort::TimingErrSort</i> -- (timeErr) positive error means the burst was at a later time than expected • <i>L1RxBurstRfResultSort::Boolean</i> -- (timeErrValid) whether timeErr is valid 	

<i>L1Frm.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L1 to L2 MAC that a new HC-SDMA frame boundary has arrived	

10.3.2.2 L2 MAC to L1 Interface Primitives

Primitive list:

- ◆ L1RxSched.resp
- ◆ L1TxSched.resp

Response primitive descriptions

<i>L1RxSched.resp</i>	
Applies to:	BS, UT
Associated primitives:	L1RxSched.ind
Description: response from L2 MAC to L1 indicating what bursts L2 MAC wants to listen for. This primitivePrimitive must be sent in response to the L1RxSched.ind.	
Parameters:	
<ul style="list-style-type: none"> • L1RxBurstListSort -- the list of bursts that L2 MAC wants to listen for, or an empty list if L2 MAC doesn't want to listen for any bursts set element type: L1RxBurstSort (struct) • L1RxBurstListSort::L1RxBurstSort::L1BurstIdSort -- (burstId) (struct) • L1RxBurstListSort::L1RxBurstSort::L1BurstIdSort::ibBurstTypeSort -- (burstType) • L1RxBurstListSort::L1RxBurstSort::L1BurstIdSort::L1ScrambleIdSort -- (scrambleId) applies to all but BCH F • L1RxBurstListSort::L1RxBurstSort::L1BurstIdSort::AaTchIndexSort -- (tchIndex) applies to TCH bursts only • L1RxBurstListSort::L1RxBurstSort::L1BurstIdSort::ModClassSort -- (modClass) applies to TCH bursts only • L1RxBurstListSort::L1RxBurstSort::L1BurstIdSort::ibBsc -- (bscc) affects selection of training sequence • L1RxBurstListSort::L1RxBurstSort::Boolean -- (modClassFromFacch) For certain TCH bursts (uplink bursts in odd RFN and downlink bursts in even RFN), the Rx modclass has to be obtained by decoding the FACCH • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort -- (rf) (struct) • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::ibFreqKhzSort -- (rxFreq) what frequency is the unit listening at • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::DuplexTimeSlotSort -- (timeSlot) the timeslot on which the burst reception should be attempted • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::SubSlotSort -- (subSlot) the sub-slot within the timeslot (applies only to PCH bursts) • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort -- (rxWin) when might the burst actually be heard, w.r.t. the nominal timeslot timing (struct) • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort -- (min) what is the earliest burst that can be heard, w.r.t. the nominal timeslot timing • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort -- (max) what is the latest burst that can be heard, w.r.t. the nominal timeslot timing • L1RxBurstListSort::L1RxBurstSort::L1RxBurstRfInfoSort::RxWinSort::TimingAdjSort -- (expect) when is the burst expected, w.r.t. the nominal timeslot timing 	

<i>L1TxSched.resp</i>	
Applies to:	BS, UT
Associated primitives:	L1TxSched.ind
<p>Description: response from L2 MAC to L1 indicating what bursts L2 MAC wants to transmit. This primitivePrimitive must be sent in response to the L1TxSched.ind.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • L1TxBurstListSort -- the list of bursts that L2 MAC wants to transmit, or an empty list if L2 MAC doesn't want to transmit any bursts set element type: L1TxBurstSort (struct) • L1TxBurstListSort::L1TxBurstSort::L1BurstIdSort -- (burstId) (struct) • L1TxBurstListSort::L1TxBurstSort::L1BurstIdSort::ibBurstTypeSort -- (burstType) • L1TxBurstListSort::L1TxBurstSort::L1BurstIdSort::L1ScrambleIdSort -- (scrambleId) applies to all but BCH F • L1TxBurstListSort::L1TxBurstSort::L1BurstIdSort::AaTchIndexSort -- (tchIndex) applies to TCH bursts only • L1TxBurstListSort::L1TxBurstSort::L1BurstIdSort::ModClassSort -- (modClass) applies to TCH bursts only • L1TxBurstListSort::L1TxBurstSort::L1BurstIdSort::ibBsc -- (bscc) affects selection of training sequence • L1TxBurstListSort::L1TxBurstSort::L1TxBurstRfInfoSort -- (rf) (struct) • L1TxBurstListSort::L1TxBurstSort::L1TxBurstRfInfoSort::ibFreqKhzSort -- (txFreq) the frequency at which the burst should be transmitted • L1TxBurstListSort::L1TxBurstSort::L1TxBurstRfInfoSort::DuplexTimeSlotSort -- (timeSlot) the timeslot on which the burst should be transmitted • L1TxBurstListSort::L1TxBurstSort::L1TxBurstRfInfoSort::SubSlotSort -- (subSlot) the sub-slot within the timeslot (applies only to PCH bursts) • L1TxBurstListSort::L1TxBurstSort::L1TxBurstRfInfoSort::TimingAdjSort -- (txTime) the time at which the burst should be transmitted (w.r.t. the nominal timing for the timeslot) • L1TxBurstListSort::L1TxBurstSort::L1TxBurstRfInfoSort::PowerDbmSort -- (txPwrDbm) the power at which the burst should be transmitted • L1TxBurstListSort::L1TxBurstSort::BitString -- (data) applies to all bursts except BCH_F and BCH_T • L1TxBurstListSort::L1TxBurstSort::FacchDataSort -- (facchData) applies to CM, RA, AA, and TCH bursts (FACCH) 	

10.3.2.3 L3 MMC to L1 Interface Primitives

Primitive list:

- ◆ L1UtTimeAdjust.req
- ◆ L1UtFreqAdjust.req

Request primitive descriptions

<i>L1UtTimeAdjust.req</i>	
Applies to:	UT only
Associated primitives:	none
<p>Description: request from L3 MMC to UT L1 to adjust the UT's clock value. The adjustment from L3 MMC should drive the frame timing of UT L1 to match the common frame timing of all BSs.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • TimingAdjSort -- how much the timing needs to be adjusted - a positive adjustment means to increase (delay) the UT's time reference, in essence jumping backwards in time. 	

<i>L1UtFreqAdjust.req</i>	
Applies to:	UT only
Associated primitives:	none
Description: request from L3 MMC to UT L1 to adjust the UT's clock frequency. The adjustment from L3 MMC should drive the UT to have its frequency synchronized with the BSs'.	
Parameters: <ul style="list-style-type: none"> • <i>ibFreqAdjHzSort</i> -- how much the frequency needs to be adjusted (at RF) - a positive adjustment means the UT should increase its clock rate. 	

10.3.2.4 L2 MAC to L2 RLC Interface Primitives

Primitive list:

- ◆ *L2MacL2RlcUtStart.ind*
- ◆ *L2MacRxData.ind*
- ◆ *L2MacTxDataInfoGet.ind*

Indication primitive descriptions

<i>L2MacL2RlcUtStart.ind</i>	
Applies to:	UT only
Associated primitives:	<i>L2MacL2RlcUtStart.resp</i>
Description: indication from UT L2 MAC to L2 RLC to start a connection to transfer data across the air.	
Parameters: <ul style="list-style-type: none"> • <i>process ID</i> -- ID of the L2 MAC process associated with the stream • <i>process ID</i> -- ID of the L3 RM process associated with the stream 	

<i>L2MacRxData.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L2 MAC to L2 RLC that incoming data bytes are available	
Parameters: <ul style="list-style-type: none"> • <i>RfnSort</i> -- the RFN which this burst is for • <i>TchSeqNumSort</i> -- sequence number of the first AMoctet in the payload • <i>TchAckSort</i> -- the sequence number of the first unacknowledged AM octet by the peer • <i>ByteString</i> -- a series of bytes received from the peer L2 MAC • <i>L2RlcChan</i> -- indication of whether the data was received on the AM (acknowledged mode) or UM (unacknowledged mode) channel 	

<i>L2MacTxDataInfoGet.ind</i>	
Applies to:	BS, UT
Associated primitives:	L2MacTxDataInfoGet.resp
Description: indication from L2 MAC to L2 RLC that a burst is to be sent and how many AM or UM data bytes can be accommodated in this burst	
Parameters: <ul style="list-style-type: none"> • RfnSort -- the RFN which this burst is for • Integer -- the number of bytes of data that the L2 MAC can accommodate in this burst (n_max) • Integer -- the number of bytes of data that the L2 MAC wants for medium priority UM control messages (n_reserve) • L2RlcChan -- indication of whether this data is for the AM (acknowledged mode) or UM (unacknowledged mode) channel 	

10.3.2.5 L2 RLC to L2 MAC Interface Primitives

Primitive list:

- ◆ L2MacL2RlcUtStart.resp
- ◆ L2MacTxDataInfoGet.resp

Response primitive descriptions

<i>L2MacL2RlcUtStart.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacL2RlcUtStart.ind
Description: response from UT L2 RLC to UT L2 MAC that L2 RLC was able to create a new instance to service the new stream.	
Parameters: <ul style="list-style-type: none"> • process ID -- ID of the L2 RLC process associated with the stream 	

<i>L2MacTxDataInfoGet.resp</i>	
Applies to:	BS, UT
Associated primitives:	L2MacTxDataInfoGet.ind
Description: response from L2 RLC to L2 MAC with the data that the current burst is to send for outgoing data	
Parameters: <ul style="list-style-type: none"> • TchSeqNumSort -- sequence number of the first AM octet in the payload • TchAckSort -- the sequence number of the first AM octet not received contiguously • ByteString -- a series of bytes to be sent to the peer • L2RlcChan -- indication of whether this is AM (acknowledged mode) or UM (unacknowledged mode) data • L2RlcStrmState -- the current state of stream from the local RLC perspective 	

10.3.2.6 L2 MAC to L3 RRC Interface Primitives

Primitive list:

- ◆ L2MacL3RrcUtStart.ind
- ◆ L2MacL3RrcBsStart.ind
- ◆ L2MacL3RrcStop.ind

- ◆ L2MacLinkAdaptData.ind
- ◆ L2MacTchRxPwrCtrlInfo.ind
- ◆ L2MacUtRaPwrInfoGet.ind
- ◆ L2MacBsRa.ind
- ◆ L2MacUtOverrideModClass.ind
- ◆ L2MacUtCm.ind
- ◆ L2MacTchRxSinr.ind
- ◆ L2MacBsAaInfoGet.ind
- ◆ L2MacBsOverrideDlModClass.ind
- ◆ L2MacBsBburstInfoGet.ind
- ◆ L2MacBsUtCapability.ind
- ◆ L2MacUtCrPwrInfoGet.ind
- ◆ L2MacTchTxInfoGet.ind
- ◆ L2MacUtBburst.ind
- ◆ L2MacUtAa.ind
- ◆ L2MacBsOverrideUlModClass.ind
- ◆ L2MacBsCr.ind

Indication primitive descriptions

<i>L2MacL3RrcUtStart.ind</i>	
Applies to:	UT only
Associated primitives:	<i>L2MacL3RrcUtStart.resp</i>
Description: indication from UT L2 MAC to UT L3 RRC that a stream has started or is about to start	
Parameters:	
<ul style="list-style-type: none"> • <i>IbL3RrcUtCapabilitySort</i> -- the UT's pwr ctrl / link adapt capabilities. Note that if the stream is not full-rate, L2 MAC modifies the UT's decode delay capabilities to reflect the apparent 0.5 frame decode delay. (struct) • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (minTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::Bit</i> -- (cch) cch decode delay information • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (rachMap) rach decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (pchMap) pch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (ulModClasses) which mod classes can the UT tx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (dlModClasses) which mod classes can the UT rx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (maxTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod01) how much higher is the mod class 0 and 1 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod2345) how much higher is the mod class 2-5 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • <i>IbBsCapabilitySort</i> -- the capabilities of the BS at the remote side of the stream (struct) • <i>IbBsCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>ibBsc</i> -- color code of the associated BS 	

<i>L2MacL3RrcBsStart.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacL3RrcBsStart.resp
Description: indication from BS L2 MAC to BS L3 RRC that a stream is starting	
Parameters:	
<ul style="list-style-type: none"> • <i>IbBsCapabilitySort</i> -- the BS's capabilities (struct) • <i>IbBsCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort</i> -- the pwr ctrl / link adapt capabilities of the UT on the remote side of this stream. Note that if the stream is not full-rate, L2 MAC modifies the UT's decode delay capabilities to reflect the apparent 0.5 frame decode delay. (struct) • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (minTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::Bit</i> -- (cch) cch decode delay information • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (rachMap) rach decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (pchMap) pch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> - (ulModClasses) which mod classes can the UT tx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (dlModClasses) which mod classes can the UT rx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (maxTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod01) how much higher is the mod class 0 and 1 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod2345) how much higher is the mod class 2-5 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • <i>ibSubRateSort</i> -- indicates the ratio of frames to RFNs (1:1 for full rate, 2:1 for half, etc.) and thus allows L3 RRC to make an informed estimate of how much the RF channel characteristics may change between RFNs 	

<i>L2MacL3RrcStop.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: Indication from L2 MAC to L3 RRC that the stream has completed. Upon receipt of this signal, the L3 RRC instance will terminate itself.	

<i>L2MacLinkAdaptData.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L2 MAC to L3 RRC that an attempt to decode a FACCH payload has finished. Although the FACCH contents will alternate each frame between CMC and RMC/AP, both CMC and RMC/AP information is sent to L3 RRC on every frame.	
Parameters: <ul style="list-style-type: none"> • FacchErrSort -- whether the FACCH decode seems to have been successful • LinkAdaptDataSort -- (CMC) 4 bit FACCH data used by link adaptation and power control • LinkAdaptDataSort -- (RMC/AP) 4 bit FACCH data used by link adaptation and power control • RfnSort -- the RFN on which the burst was received 	

<i>L2MacTchRxPwrCtrlInfo.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L2 MAC to L3 RRC that an attempt to decode a TCH burst has finished. L2 MAC provides L3 RRC with the link adaptation and power control data from the burst, if the burst was decoded successfully.	
Parameters: <ul style="list-style-type: none"> • RmuErrSort -- whether a frame error occurred • L2MacL3RrcOneOrSixBitSort -- whether L2 MAC received 1-bit or 6-bit data • L2MacL3RrcTchPwrCtrlInfo1BitSort -- 1 bit data (valid only when OneBitType is indicated) • L2MacL3RrcTchPwrCtrlInfo6BitSort -- 6 bit data (valid only when SixBitType is indicated) • RfnSort -- the RFN on which the burst was received 	

<i>L2MacUtRaPwrInfoGet.ind</i>	
Applies to:	UT only
Associated primitives:	L2MacUtRaPwrInfoGet.resp
Description: indication from UT L2 MAC to UT L3 RRC that it needs to know at what power it should send an RA	
Parameters: <ul style="list-style-type: none"> • <i>ibBsc</i> -- color code of the BS the RA is directed to • <i>L2MacUtRachScanInfoSort</i> -- information about UT's RACH scan (struct) • <i>L2MacUtRachScanInfoSort::L2MacUtRachScanModeSort</i> -- (<i>rachMode</i>) directed or unDirected • <i>L2MacUtRachScanInfoSort::Integer</i> -- (<i>dirRaAttemptNum</i>) How many directed RA attempts so far; includes the current attempt -- valid only for directed • <i>L2MacUtRachScanInfoSort::Integer</i> -- (<i>unDirRaTransmitNum</i>) How many undirected RA transmits so far; includes the current transmit -- valid only for unDirected • <i>L2MacUtRachScanInfoSort::Integer</i> -- (<i>unDirRaAttemptNum</i>) How many undirected RA attempts so far; includes the current attempt -- valid only for unDirected • <i>L2MacUtRachScanInfoSort::Integer</i> -- (<i>raTransmitsPerSubScan</i>) Number of RA transmits per sub-scan -- valid only for unDirected • <i>L2MacUtRachScanInfoSort::Integer</i> -- (<i>numRachRsrcs</i>) Number of channels (resources) on which a unDir RA could be sent; useful in case sweep number is necessary -- valid only for unDirected • <i>L2MacUtRachScanInfoSort::Integer</i> -- (<i>maxRaTransmits</i>) the maximum number of RA transmits allowed for the whole RACH scan - note that for unDirected scan this is an estimate, which will be continually more accurate as L2 MAC nears the end of scan 	

<i>L2MacBsRa.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacBsAaInfoGet.ind, L2MacBsAaInfoGet.resp
Description: indication from BS L2 MAC to BS L3 RRC that a RA has been received. The transmit power from the RA burst is passed to BS L3 RRC, along with the DSSI and SINR measured during reception of the RA.	
Parameters: <ul style="list-style-type: none"> • <i>DssiDbmSort</i> -- the desired signal strength indicator measured for the RA burst • <i>SinrDbSort</i> -- the SINR of the RA burst • <i>RaTxPwrSort</i> -- RA data field: what power did the UT transmit the RA burst at 	

<i>L2MacUtOverrideModClass.ind</i>	
Applies to:	UT only
Associated primitives:	L2MacUtOverrideModClass.resp
Description: indication from UT L2 MAC to UT L3 RRC that L2 MAC wants to override the uplink mod class suggested by L3 RRC	
Parameters: <ul style="list-style-type: none"> • <i>ModClassSort</i> -- uplink mod class suggested by L2 MAC • <i>RfnSort</i> -- the RFN on which this suggested uplink mod class change would be applied 	

<i>L2MacUtCm.ind</i>	
Applies to:	UT only
Associated primitives:	none
Description: indication from UT L2 MAC to UT L3 RRC that a CM was received. L2MacUtCm.ind information is used when L2MacUtRaPwrInfoGet.resp is first issued.	
Parameters: <ul style="list-style-type: none"> • ibBsc -- color code of the BS that sent the CM • CmPwrAdjustSort -- CM data field: recommended power adjustment (based on the expected and actual receive powers of the CR burst) • CmUIPwrCtrlActiveSort -- CM data field: is uplink closed-loop power control active (this tells the UT whether it should respect/ignore the power control bit in the downlink TCH bursts' RMU headers) 	

<i>L2MacTchRxSinr.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L2 MAC to L3 RRC that a TCH burst was received (but not necessarily decoded)	
Parameters: <ul style="list-style-type: none"> • DetectFailSort -- whether the burst was successfully detected • SinrDbSort -- the SINR of the received burst • RfnSort -- the RFN on which the burst was received 	

<i>L2MacBsAaInfoGet.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacBsRa.ind, L2MacBsAaInfoGet.resp
Description: indication from BS L2 MAC to BS L3 RRC asking for L3 RRC to choose the power for transmitting the AA burst, and the values for the relevant data fields of the AA.	

<i>L2MacBsOverrideDlModClass.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacBsOverrideDlModClass.resp
Description: indication from BS L2 MAC to BS L3 RRC that L2 MAC wants to override the downlink mod class suggested by L3 RRC	
Parameters: <ul style="list-style-type: none"> • DlModClassSort -- downlink mod class suggested by L2 MAC • RfnSort -- the RFN on which this suggested downlink mod class change would be applied 	

<i>L2MacBsBburstInfoGet.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacBsBburstInfoGet.resp
Description: indication from BS L2 MAC to BS L3 RRC that L2 MAC needs to know at what power to transmit the BCH B burst	

ATIS-0700004.2007

<i>L2MacBsUtCapability.ind</i>	
Applies to:	BS only
Associated primitives:	none
<p>Description:</p> <p>indication from BS L2 MAC to BS L3 RRC informing L3 RRC of the actual UT capabilities. This signal is only used during the registration stream, where the actual capabilities of the UT are discovered during the stream. For regular TCH streams, BS L2 MAC tells BS L3 RRC of the UT's capabilities through the L2MacL3RrcBsStart.ind.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>IbL3RrcUtCapabilitySort</i> -- capabilities of the UT related to link adaptation and power control (struct) • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (minTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) <ul style="list-style-type: none"> • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::Bit</i> -- (cch) cch decode delay information • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (rachMap) rach decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (pchMap) pch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (ulModClasses) which mod classes can the UT tx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (dlModClasses) which mod classes can the UT rx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (maxTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod01) how much higher is the mod class 0 and 1 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod2345) how much higher is the mod class 2-5 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? 	

<i>L2MacUtCrPwrInfoGet.ind</i>	
Applies to:	UT only
Associated primitives:	L2MacUtCrPwrInfoGet.resp
<p>Description:</p> <p>indication from UT L2 MAC to UT L3 RRC asking at what power it should transmit a CR burst</p>	

Parameters:	
<ul style="list-style-type: none"> • <i>ibBsc</i> -- color code of the BS the UT is sending the CR to • <i>CrAttemptNumSort</i> -- how many times the UT has already tried sending the CR • <i>MaxCrAttemptsSort</i> -- how many times the UT will try sending the CR before giving up 	

<i>L2MacTchTxInfoGet.ind</i>	
Applies to:	BS, UT
Associated primitives:	<i>L2MacTchTxInfoGet.resp</i>
Description: indication from L2 MAC to L3 RRC requesting the information needed to transmit a TCH burst	
Parameters:	
<ul style="list-style-type: none"> • <i>RfnSort</i> -- the RFN on which the burst will be sent 	

<i>L2MacUtBburst.ind</i>	
Applies to:	UT only
Associated primitives:	none
Description: indication from UT L2 MAC to UT L3 RRC that a BCH B burst was received.	
Parameters:	
<ul style="list-style-type: none"> • <i>ibBsc</i> -- color code of the BS which sent the B burst • <i>BsTxPwrSort</i> -- B burst data field: transmitted power from the BS • <i>DssiDbmSort</i> -- the desired signal strength indicator measured during the B burst rx 	

<i>L2MacUtAa.ind</i>	
Applies to:	UT only
Associated primitives:	none
Description: indication from UT L2 MAC to UT L3 RRC that an AA has been received. The power adjustment and modclass assignments from the AA burst are passed to UT L3 RRC, along with the SINR measured during reception of the AA. Information from <i>L2MacUtAa.ind</i> is used when <i>L2MacTchTxInfoGet.resp</i> is first sent.	
Parameters:	
<ul style="list-style-type: none"> • <i>AaSubTypeSort</i> -- AA data field: the subtype of the AA • <i>AaPwrAdjustSort</i> -- AA data field: recommended uplink power adjustment (note that this does not include the power adjustment due to use of a different mod class for TCH than for RA) • <i>AaUIModClassSort</i> -- AA data field: the uplink mod class recommended by the BS for the first TCH frames • <i>AaDlModClassSort</i> -- AA data field: the downlink mod class the BS intends to use for the first TCH frames • <i>AaFrameDecSort</i> -- AA data field: the subrate information of the TCH stream -- valid only for <i>AaCts</i>, <i>AaCtsShort</i> and <i>AaRreg</i> • <i>SinrDbSort</i> -- the SINR of the AA burst 	

<i>L2MacBsOverrideUlModClass.ind</i>	
Applies to:	BS only
Associated primitives:	<i>L2MacBsOverrideUlModClass.resp</i>

Description: indication from BS L2 MAC to BS L3 RRC that L2 MAC wants to override the uplink mod class suggested by L3 RRC. This assumes that BS L2 MAC has some information regarding how much uplink data the UT has queued. Currently there is no mechanism for BS L2 MAC to receive such information, but a method may be added in the future. e.g., the queue depth reports received by BS L3 RM might be shared with BS L2 MAC.	
Parameters: <ul style="list-style-type: none"> • UIModClassSort -- uplink mod class suggested by L2 MAC • RfnSort -- the RFN on which this suggested uplink mod class change would be applied 	

<i>L2MacBsCr.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacBsCm.req
Description: indication from BS L2 MAC to BS L3 RRC providing the CR data fields and signal quality measurements	
Parameters: <ul style="list-style-type: none"> • DssiDbmSort -- desired signal strength indicator for the CR rx • SinrDbSort -- SINR of the CR rx • CrTxPwrSort -- CR data field: 5-bit quantized representation of the UT's CR tx power • CrRandomIdSort -- CR data field: 8-bit random value 	

10.3.2.7 L3 RRC to L2 MAC Interface Primitives

Primitive list:

- ◆ L2MacBsCm.req
- ◆ L2MacL3RrcBsStart.resp
- ◆ L2MacL3RrcUtStart.resp
- ◆ L2MacUtOverrideModClass.resp
- ◆ L2MacUtRaPwrInfoGet.resp
- ◆ L2MacBsBurstInfoGet.resp
- ◆ L2MacUtCrPwrInfoGet.resp
- ◆ L2MacTchTxInfoGet.resp
- ◆ L2MacBsAaInfoGet.resp
- ◆ L2MacBsOverrideUIModClass.resp
- ◆ L2MacBsOverrideDlModClass.resp

Request primitive descriptions

<i>L2MacBsCm.req</i>	
Applies to:	BS only
Associated primitives:	L2MacBsCr.ind
Description: request from BS L3 RRC to BS L2 MAC providing information to use to transmit the CM burst	
Parameters: <ul style="list-style-type: none"> • CrRandomIdSort -- 8-bit random value which was in the corresponding CR • CmUIPwrCtrlActiveSort -- CM data field: is uplink power control active (this tells the UT whether it should respect/ignore the power control bit in the downlink TCH bursts' RMU headers) • CmPwrAdjustSort -- CM data field: how much should the UT adjust its transmit power based on the expected and actual receive powers of the CR burst • PowerDbmSort -- power to be used for transmitting the CM burst 	

<i>L2MacL3RrcBsStart.resp</i>	
Applies to:	BS only
Associated primitives:	L2MacL3RrcBsStart.ind
Description: response from BS L3 RRC to BS L2 MAC that a L3 RRC instance has been created for the stream	
Parameters: <ul style="list-style-type: none"> process ID -- process ID of the L3 RRC stream instance 	

<i>L2MacL3RrcUtStart.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacL3RrcUtStart.ind
Description: response from UT L3 RRC to UT L2 MAC that a L3 RRC instance has been created for the stream	
Parameters: <ul style="list-style-type: none"> process ID -- process ID of the L3 RRC stream instance 	

<i>L2MacUtOverrideModClass.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacUtOverrideModClass.ind
Description: response from UT L3 RRC to UT L2 MAC stating whether L3 RRC accepted L2 MAC's uplink mod class suggestion, and if so, what revised transmit power to use and FACCH data to report. This assumes that L2MacUtOverrideModClass.ind came after L2MacTchTxInfoGet.ind. Further, this primitive is only valid for the odd RFNs, when the UT may modify its mod class. Therefore, there is no need for L3 RRC to also recalculate the additional available power, since the UT won't report the available power until the next even RFN.	
Parameters: <ul style="list-style-type: none"> ModClassSort -- what uplink mod class to actually use PowerDbmSort -- what power to transmit at with the new mod class 	

<i>L2MacUtRaPwrInfoGet.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacUtRaPwrInfoGet.ind
Description: response from UT L3 RRC to UT L2 MAC that the specified power should be used to transmit the RA burst	
Parameters: <ul style="list-style-type: none"> RaTxPwrSort -- RA data field: 5-bit quantized representation of the UT's RA transmit power PowerDbmSort -- absolute power in dBm at which to transmit the RA burst 	

<i>L2MacBsBburstInfoGet.resp</i>	
Applies to:	BS only
Associated primitives:	L2MacBsBburstInfoGet.ind
Description: response from BS L3 RRC to BS L2 MAC with the power at which to transmit the BCH B burst	
Parameters: <ul style="list-style-type: none"> PowerDbmSort -- power to be used to transmit B burst BsTxPwrSort -- B burst data field: 4-bit quantized representation of transmit power 	

<i>L2MacUtCrPwrInfoGet.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacUtCrPwrInfoGet.ind
Description: response from UT L3 RRC to UT L2 MAC stating the power which should be used to transmit the CR burst	
Parameters: <ul style="list-style-type: none"> • CrTxPwrSort -- CR data field: 5 bit quantized representation of the UT's CR transmit power • PowerDbmSort -- absolute power in dBm at which to transmit the CR 	

<i>L2MacTchTxInfoGet.resp</i>	
Applies to:	BS, UT
Associated primitives:	L2MacTchTxInfoGet.ind
Description: response from L3 RRC to L2 MAC stating the power at which to transmit the TCH burst and the link adaptation and power control data that should be communicated in the burst.	
Parameters: <ul style="list-style-type: none"> • ModClassSort -- uplink (UT) or downlink (BS) mod class to be used (CMC could be obtained from this) • LinkAdaptDataSort -- AP (UT) or RMC (BS) - on each even RFN, the AP given in this parameter is sent in the uplink FACCH payload; on each odd RFN, the RMC given in this parameter is sent in the downlink FACCH payload • PowerDbmSort -- the power to be used to transmit the TCH burst • L2MacUmCtrlType -- which UM ctrl type, sinr.rpt (UT) or pwrCtrl.cmd (BS) • L2MacL3RrcOneOrSixBitSort -- whether L3 RRC prefers to send 1-bit or 6-bit data • L2MacL3RrcTchPwrCtrlInfo1BitSort -- 1 bit data (includes the fallback case of 6-bit being intended but not possible) • L2MacL3RrcTchPwrCtrlInfo6BitSort -- 6 bit data (valid only when SixBitType is indicated) 	

<i>L2MacBsAaInfoGet.resp</i>	
Applies to:	BS only
Associated primitives:	L2MacBsAaInfoGet.ind
Description: response from BS L3 RRC to BS L2 MAC stating at what power to transmit the AA, and what values to put in the link adaptation / power control AA data fields	
Parameters: <ul style="list-style-type: none"> • PowerDbmSort -- pwr to be used for transmitting the AA burst • AaUIModClassSort -- AA data field: the uplink mod class recommended by the BS for the first TCH frames • AaDIModClassSort -- AA data field: the downlink mod class the BS intends to use for the first TCH frames • AaPwrAdjustSort -- AA data field: recommended uplink power adjustment (note that this does not include the power adjustment due to use of a different mod class for TCH than for RA) 	

<i>L2MacBsOverrideUlModClass.resp</i>	
Applies to:	BS only
Associated primitives:	L2MacBsOverrideUlModClass.ind

Description: response from BS L3 RRC to BS L2 MAC stating whether L3 RRC accepted L2 MAC's uplink mod class suggestion, and if so, what revised power control data to send the UT.	
Parameters:	
<ul style="list-style-type: none"> • UIModClassSort -- the final RMC information to be sent via FACCH • L2MacL3RrcOneOrSixBitSort -- whether a 1-bit or 6-bit power control command should be sent • L2MacL3RrcTchPwrCtrlInfo1BitSort -- 1 bit power control command (includes the fallback case of 6-bit being intended but not possible) • L2MacL3RrcTchPwrCtrlInfo6BitSort -- 6 bit power control command (valid only when SixBitType is indicated) 	

<i>L2MacBsOverrideDlModClass.resp</i>	
Applies to:	BS only
Associated primitives:	L2MacBsOverrideDlModClass.ind
Description: response from BS L3 RRC to BS L2 MAC stating whether L3 RRC accepted L2 MAC's downlink mod class suggestion, and if so, what revised transmit power to use and FACCH data to report.	
Parameters:	
<ul style="list-style-type: none"> • DModClassSort -- the final CMC information to be sent via FACCH • PowerDbmSort -- power to be used for transmitting the TCH burst 	

10.3.2.8 L2 MAC to L3 MMC Interface Primitives

Primitive list:

- ◆ L2MacUtTxTimingConfig.ind
- ◆ L2MacUtTimingErr.ind
- ◆ L2MacUtRxTimingConfig.ind
- ◆ L2MacUtFreqErr.ind
- ◆ L2MacUtBchF.conf
- ◆ L2MacUtBchT.conf
- ◆ L2MacUtBchB.conf
- ◆ L2MacUtCch.conf

Indication primitive descriptions

<i>L2MacUtTxTimingConfig.ind</i>	
Applies to:	UT only
Associated primitives:	L2MacUtTxTimingConfig.resp
Description: indication from UT L2 MAC to L3 MMC that MAC needs to know how much to advance/delay the transmit time of a specified burst so that the burst will arrive at the correct time at the specified BS. The timing advance/delay accounts for the UT->BS propagation delay and any intentional timing offset, e.g., the random delay used for CR bursts' time diversity.	
Parameters:	
<ul style="list-style-type: none"> • ibBurstTypeSort -- what kind of burst is going to be transmitted. (The CR burst has no timing advance, but does have a random delay. Other bursts have timing advance.) • ibBscc -- color code of the BS the burst is being sent to • ibFreqKhzSort -- carrier frequency the burst will travel on (in case frequency-dependent multipath causes different propagation delays on different carriers) 	

<i>L2MacUtTimingErr.ind</i>	
Applies to:	UT only
Associated primitives:	none
<p>Description: indication from UT L2 MAC to L3 MMC of the downlink timing error in the received burst, and also the uplink timing error for bursts containing an in-band uplink timing error message from the BS</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>ibBscc</i> -- color code of the BS that sent the burst • <i>ibBurstTypeSort</i> -- the burst type received • <i>ibFreqKhzSort</i> -- the frequency of the carrier at which burst was detected • <i>TimingErrSort</i> -- (downlink) positive means that burst arrived too late relative to UT clock • <i>TimingErrSort</i> -- (uplink) positive means that burst arrived too late relative to BS clock • <i>Boolean</i> -- denotes whether the uplink timing in this primitive is valid. For TCH, the uplink timing may not be present. For any burst, there may have been a decode error • <i>SinrDbSort</i> -- SINR of the received burst • <i>DecodeFailSort</i> -- whether the burst decode failed 	

<i>L2MacUtRxTimingConfig.ind</i>	
Applies to:	UT only
Associated primitives:	<i>L2MacUtRxTimingConfig.resp</i>
<p>Description: indication from UT L2 MAC to L3 MMC that MAC needs to know how much to delay the reception time of a specified burst so that the UT will listen at the right time to hear the specified burst from the specified BS. The timing delay accounts for the BS->UT propagation delay.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>ibBurstTypeSort</i> -- what kind of burst is being received. (The same timing delay is used for all burst types, but the extent of the rx window depends on the burst type. For example, the UT will listen much longer to hear BCH T bursts than other bursts, since it wants to hear all BSs' T bursts, which will have significantly different propagation delays.) • <i>ibBscc</i> -- color code of the BS sending the burst • <i>ibFreqKhzSort</i> -- carrier frequency the burst will travel on (in case frequency-dependent multipath causes different propagation delays on different carriers) 	

<i>L2MacUtFreqErr.ind</i>	
Applies to:	UT only
Associated primitives:	none
<p>Description: indication from UT L2 MAC to L3 MMC of the frequency error measured in the received burst</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>ibBscc</i> -- color code of the BS that sent the burst • <i>ibBurstTypeSort</i> -- the burst type received. • <i>ibFreqErrHzSort</i> -- frequency err - positive means the frequency was higher than expected • <i>SinrDbSort</i> -- SINR of the received burst • <i>DecodeFailSort</i> -- whether the burst decode failed • <i>Boolean</i> -- is the <i>DecodeFailSort</i> parameter valid (the <i>FreqErr</i> indication can be sent either immediately after the detect phase of rx, when the decode result is not known, or after the decode phase) 	

Confirmation primitive descriptions

<i>L2MacUtBchF.conf</i>	
Applies to:	UT only
Associated primitives:	L2MacUtBchF.req
Description: confirmation from UT L2 MAC to L3 MMC with the results of the requested attempt to listen for a BCH F burst	
Parameters: <ul style="list-style-type: none"> • L2MacGrantSort -- indicates whether L2 MAC scheduled a rx attempt for the BCH F burst • <i>ibRfDataSort</i> -- RF signal metrics (struct) • <i>ibRfDataSort::ibBsc</i> -- (bscc) base-station color code • <i>ibRfDataSort::DssiDbmSort</i> -- (dssi) desired signal strength • <i>ibRfDataSort::RssiDbmSort</i> -- (rssi) total received signal strength • <i>ibRfDataSort::SinrDbSort</i> -- (sinr) signal to noise + interference ratio • <i>ibRfDataSort::TimingErrSort</i> -- (timeOffset) timing error (positive means the burst was later than expected) • <i>ibRfDataSort::ibFreqErrHzSort</i> -- (freqOffset) frequency error (positive means the freq was higher than expected) • <i>ibRfDataSort::DetectFailSort</i> -- (detectFail) indicates whether burst is detected • <i>ibFreqKhzSort</i> -- the frequency on which the UT listened to hear the BCH F burst 	

<i>L2MacUtBchT.conf</i>	
Applies to:	UT only
Associated primitives:	L2MacUtBchT.req
Description: confirmation from UT L2 MAC to L3 MMC with the results of the requested attempt to listen for BCH T bursts	
Parameters: <ul style="list-style-type: none"> • L2MacGrantSort -- indicates whether L2 MAC scheduled a receive of this burst type. • L2MacTBurstListSort -- array of info on all T bursts detected array index type: <i>ibBsc</i> array element type: <i>ibRfDataSort</i> (struct) • <i>L2MacTBurstListSort::ibRfDataSort::ibBsc</i> -- (bscc) base-station color code • <i>L2MacTBurstListSort::ibRfDataSort::DssiDbmSort</i> -- (dssi) desired signal strength • <i>L2MacTBurstListSort::ibRfDataSort::RssiDbmSort</i> -- (rssi) total received signal strength • <i>L2MacTBurstListSort::ibRfDataSort::SinrDbSort</i> -- (sinr) signal to noise + interference ratio • <i>L2MacTBurstListSort::ibRfDataSort::TimingErrSort</i> -- (timeOffset) timing error (positive means the burst was later than expected) • <i>L2MacTBurstListSort::ibRfDataSort::ibFreqErrHzSort</i> -- (freqOffset) frequency error (positive means the freq was higher than expected) • <i>L2MacTBurstListSort::ibRfDataSort::DetectFailSort</i> -- (detectFail) indicates whether burst is detected • <i>ibFreqKhzSort</i> -- the frequency on which the UT listened to hear the BCH T burst 	

<i>L2MacUtBchB.conf</i>	
Applies to:	UT only
Associated primitives:	L2MacUtBchB.req
<p>Description: confirmation from UT L2 MAC to L3 MMC with the results of the requested attempt to listen for a BCH B burst</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • L2MacGrantSort -- indicates whether L2 MAC scheduled a receive of this burst type. • ibRfDataSort -- RF signal metrics (struct) • <i>ibRfDataSort::ibBsc</i> -- (bscc) base-station color code • <i>ibRfDataSort::DssiDbmSort</i> -- (dssi) desired signal strength • <i>ibRfDataSort::RssiDbmSort</i> -- (rssi) total received signal strength • <i>ibRfDataSort::SinrDbSort</i> -- (sinr) signal to noise + interference ratio • <i>ibRfDataSort::TimingErrSort</i> -- (timeOffset) timing error (positive means the burst was later than expected) • <i>ibRfDataSort::ibFreqErrHzSort</i> -- (freqOffset) frequency error (positive means the freq was higher than expected) • <i>ibRfDataSort::DetectFailSort</i> -- (detectFail) indicates whether burst is detected • BsLoadSort -- BS load information field • BsTxPwrSort -- BS transmitted power information field • ibFreqKhzSort -- the frequency on which the UT listened to hear the BCH B burst 	

<i>L2MacUtCch.conf</i>	
Applies to:	UT only
Associated primitives:	L2MacUtCch.req
<p>Description: confirmation from UT L2 MAC to L3 MMC with the results of the requested CR/CM exchange</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • L2MacGrantSort -- indicates whether L2 MAC scheduled a CR/CM exchange • DecodeFailSort -- was a CM burst received successfully (true -> burst was not detected, or was detected but not decoded; false -> burst was detected and decoded with no errors) • ibBsc -- color code of the BS that sent the CM burst • ibFreqKhzSort -- the frequency on which CM burst was received • ibChanSet -- set of conventional channels UT may use for registration streams. set element type: ibChanNum • CmNetworkIdSort -- network operator id 	

10.3.2.9 L3 MMC to L2 MAC Interface Primitives

Primitive list:

- ◆ L2MacUtBchF.req
- ◆ L2MacUtBchT.req
- ◆ L2MacUtBchB.req
- ◆ L2MacUtCch.req
- ◆ L2MacUtRxTimingConfig.resp
- ◆ L2MacUtTxTimingConfig.resp

Request primitive descriptions

<i>L2MacUtBchF.req</i>	
Applies to:	UT only
Associated primitives:	L2MacUtBchF.conf
Description: request from L3 MMC to UT L2 MAC to listen for a BCH F burst	
Parameters: <ul style="list-style-type: none"> • L2MacValidFrmsSort -- how long this request is valid (if the F burst rx has not successfully completed by the end of this time, UT L2 MAC will generate a L2MacUtBchF.conf signal to inform L3 MMC that either the F burst rx attempt was not granted, or that it was granted but was not successful) • L2MacPrioritySort -- whether this request should take precedence over or defer to competing requests to receive TCH, RACH, or PCH • ibFreqKhzSort -- the frequency at which to listen for the BCH F burst 	

<i>L2MacUtBchT.req</i>	
Applies to:	UT only
Associated primitives:	L2MacUtBchT.conf
Description: request from L3 MMC to UT L2 MAC to listen for a BCH T burst. Note that this signal implicitly asks UT L2 MAC to listen for the T burst from every BSCC.	
Parameters: <ul style="list-style-type: none"> • L2MacValidFrmsSort -- how long this request is valid (if the T burst rx has not successfully completed by the end of this time, UT L2 MAC will generate a L2MacUtBchT.conf signal to inform L3 MMC that either the T burst rx attempt was not granted, or that it was granted but was not successful) • L2MacPrioritySort -- whether this request should take precedence over or defer to competing requests to receive TCH, RACH, or PCH • ibFreqKhzSort -- the frequency at which to listen for the BCH T burst 	

<i>L2MacUtBchB.req</i>	
Applies to:	UT only
Associated primitives:	L2MacUtBchB.conf
Description: request from L3 MMC to UT L2 MAC to listen for a BCH B burst	
Parameters: <ul style="list-style-type: none"> • L2MacValidFrmsSort -- how long this request is valid (if the B burst rx has not successfully completed by the end of this time, UT L2 MAC will generate a L2MacUtBchB.conf signal to inform L3 MMC that either the B burst rx attempt was not granted, or that it was granted but was not successful) • L2MacPrioritySort -- whether this request should take precedence over or defer to competing requests to receive TCH, RACH, or PCH • ibBscC -- which BSCC's BCH B burst to listen for • ibFreqKhzSort -- the frequency at which to listen for the BCH B burst 	

<i>L2MacUtCch.req</i>	
Applies to:	UT only
Associated primitives:	L2MacUtCch.conf
Description: request from L3 MMC to UT L2 MAC to transmit a CR burst to the specified BS and listen for the answering CM burst	
Parameters: <ul style="list-style-type: none"> • L2MacValidFrmsSort -- how long this request is valid (if a CM has not been received by the end of this time, UT L2 MAC will send a L2MacUtCch.conf signal to inform L3 MMC that either the CR/CM exchange was not granted, or that it was granted but was not successful) • L2MacPrioritySort -- whether this request should take precedence over or defer to competing requests to receive TCH, RACH, or PCH • ibBscc -- Type -- which BS to do the CR/CM exchange with • ibFreqKhzSort -- the frequency on which to do the CR/CM exchange 	

Response primitive descriptions

<i>L2MacUtRxTimingConfig.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacUtRxTimingConfig.ind
Description: response from L3 MMC to UT L2 MAC telling MAC how much to delay burst reception	
Parameters: <ul style="list-style-type: none"> • ibBurstTypeSort -- what burst type the rx timing configuration is for • ibBscc -- color code of the BS sending the burst • ibFreqKhzSort -- carrier frequency the burst will travel on • RxWinSort -- start, end, and reference (i.e., expected burst rx time) of receive timing window (struct) • RxWinSort::TimingAdjSort -- (min) what is the earliest burst that can be heard, w.r.t. the nominal timeslot timing • RxWinSort::TimingAdjSort -- (max) what is the latest burst that can be heard, w.r.t. the nominal timeslot timing • RxWinSort::TimingAdjSort -- (expect) when is the burst expected, w.r.t. the nominal timeslot timing 	

<i>L2MacUtTxTimingConfig.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacUtTxTimingConfig.ind
Description: response from L3 MMC to UT L2 MAC telling MAC how much to advance / delay burst transmission	
Parameters: <ul style="list-style-type: none"> • ibBurstTypeSort -- what burst type the tx timing adjustment is meant for • ibBscc -- color code of the BS the tx timing adjustment is meant for • ibFreqKhzSort -- carrier frequency the burst will travel on • TimingAdjSort -- tx timing adjustment - positive means to increase (i.e., delay) the timing 	

10.3.2.10 L2 MAC to L3 RM Interface Primitives

Primitive list:

- ◆ L2MacL3RmUtStart.ind
- ◆ L2MacL3RmBsStart.ind
- ◆ L2MacL3RmStop.ind
- ◆ L2MacUmCtrl.ind

- ◆ L2MacBsPageTx.ind
- ◆ L2MacUtPageRx.ind
- ◆ L2MacUtStart.conf
- ◆ L2MacUtPing.conf

Indication primitive descriptions

<i>L2MacL3RmUtStart.ind</i>	
Applies to:	UT only
Associated primitives:	L2MacL3RmUtStart.resp (Note that the L2RlcStart.req and L2RlcStart.conf are also triggered by the L2MacL3RmUtStart.ind primitive.)
Description: indication from UT L2 MAC to UT L3 RM that a new stream has opened (due to being paged)	
Parameters:	
<ul style="list-style-type: none"> • process ID -- ID of the L2 MAC process associated with this stream • L2MacUtRegInfoSort -- which registration the stream is for (struct) • L2MacUtRegInfoSort::ibRegId -- (regId) which registration ID within the registration domain • L2MacUtRegInfoSort::L2MacUtPageDomainSort -- (pageDomain) which paging domain (this identifies the reg domain) (struct) • L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibPageId -- (pageId) which paging ID within the registration domain • L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibBsc -- (bsc) which BS • L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibChanSet -- (convPchChans) which chans the BS has enabled for paging - this determines the mapping between paging subslots and implied resources set element type: ibChanNum • L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibChanSet -- (logicPchChans) which chans the BS will use to page this UT - this determines the set of channels this paging ID will hop over set element type: ibChanNum • L2MacUtRegInfoSort::ibChanSet -- (rachChans) the chan set usable for RACH set element type: ibChanNum • L2MacUtRegInfoSort::L2MacUtRaPolicySort -- (raPolicy) when to resend RA bursts (struct) • L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer -- (Na) number of channels per sub-scan • L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer -- (Nb) intrasubscan delay • L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer -- (Nc) intersubscan delay, scaling term • L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer -- (Nd) intersubscan delay, exponential base term • L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer -- (Ne) intersubscan delay limit • L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer -- (Nf) total time limit • L2MacStreamType -- what kind of a stream was opened (e.g., TCH or TCH-short) • AfnSort -- when did the stream start (the RA/AA frame) • ibChanNum -- what channel is the stream on • ibSubChan -- is the stream full-rate or partial-rate • Integer -- how many RA bursts were sent while trying to open the stream • L2MacRejectSort -- (this parameter is relevant only for L2MacStreamType = None, i.e., an AA-reject, AA-invalid-id, or AA-pending) why was the stream rejected - was the page ID invalid, or did the BS simply choose to not assign a stream? If the latter, how long does this reg need to wait before issuing the next stream request, and might the BS page the UT in the meantime? (struct) • L2MacRejectSort::Boolean -- (noAa) Whether or not the BS sends an AA • L2MacRejectSort::AaSubTypeSort -- (aaSubType) AaInvalidId, AaReject or AaPending • L2MacRejectSort::Integer -- (holdMsg) encoded hold frames -- only valid for aaSubType = AaReject • L2MacRejectSort::Boolean -- (noPage) if true, the BS is certain to not page the UT during the specified hold period -- only valid for aaSubType = AaReject • L2MacRejectSort::AaRaDelaySort -- (raDelay) indicates the amount of time (frames) after which the UT should send another RA for the extended RA mechanism -- only valid for aaSubType = AaPending 	

<i>L2MacL3RmBsStart.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacL3RmBsStart.resp (Note that the L2RlcStart.req and L2RlcStart.conf are also triggered by the L2MacL3RmBsStart.ind primitive.)
Description: indication from BS L2 MAC to BS L3 RM that a new stream has opened	
Parameters: <ul style="list-style-type: none"> • process ID -- ID of the L2 MAC process associated with this stream • L2MacStreamType -- what kind of a stream did the UT request (e.g., reg, TCH, TCH-short) • ibRegPageId -- reg ID or paging ID, based on stream type: <ul style="list-style-type: none"> ping: which paging ID sent the ping page response: which paging ID sent the RA-page-response else: which reg ID sent the RA-rts (or RA-directed, or RA-rtsShort...) • AfnSort -- when did the stream start (the RA/AA frame) • ibChanNum -- what conventional channel did the RA happen on 	

<i>L2MacL3RmStop.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L2 MAC to L3 RM that a stream has closed	
Parameters: <ul style="list-style-type: none"> • L2MacStopCauseSort -- why the stream was closed 	

<i>L2MacUmCtrl.ind</i>	
Applies to:	BS, UT
Associated primitives:	L2MacUmCtrl.req (on peer)
Description: indication from L2 MAC to a higher protocol layer that a UM control message was received	
Parameters: <ul style="list-style-type: none"> • L2MacUmCtrlType -- which UM control message • L2MacUmCtrlData -- the message's data (may not be relevant, depending on message type) 	

<i>L2MacBsPageTx.ind</i>	
Applies to:	BS only
Associated primitives:	L2MacBsPageTx.req
Description: indication from BS L2 MAC to BS L3 RM that a paging burst was sent to the specified user (paging ID) at the specified time. Note that the channel on which the paging burst was sent can be determined by providing the paging ID and frame on which the paging burst was sent as inputs to the page-hopping function.	
Parameters: <ul style="list-style-type: none"> • ibPageId -- which user was paged • AfnSort -- when was the page sent 	

<i>L2MacUtPageRx.ind</i>	
Applies to:	UT only
Associated primitives:	L2MacUtPageRx.req, L2MacUtPing.req (The Ping.req is triggered by the PageRx.ind.)
<p>Description: indication from UT L2 MAC to UT L3 RM that a page was received, but a stream could not be opened on the implied channel</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>L2MacUtPageDomainSort</i> -- which paging domain received the page (struct) • <i>L2MacUtPageDomainSort::ibPageId</i> -- (pageId) which paging ID within the registration domain • <i>L2MacUtPageDomainSort::ibBsc</i> -- (bsc) which BS • <i>L2MacUtPageDomainSort::ibChanSet</i> -- (convPchChans) which chans the BS has enabled for paging - this determines the mapping between paging subslots and implied resources set element type: <i>ibChanNum</i> • <i>L2MacUtPageDomainSort::ibChanSet</i> -- (logicPchChans) which chans the BS will use to page this UT - this determines the set of channels this paging ID will hop over set element type: <i>ibChanNum</i> • <i>Integer</i> -- how many RA bursts were sent while responding to the page on the implied resource. This should typically be equal to the number of directed attempts provided in the PageRx.req, (which will be <i>Nprirp</i> from the reg params msg), but in cases where a 2nd page is received while the UT is attempting to respond to an initial page, the total number of RAs sent will be more than <i>Nprirp</i>. 	

Confirmation primitive descriptions

<i>L2MacUtStart.conf</i>	
Applies to:	UT only
Associated primitives:	L2MacUtStart.req
Description: confirmation from UT L2 MAC to UT L3 RM of a requested new stream	
Parameters:	
<ul style="list-style-type: none"> • L2MacUtReqResultSort -- was L2 MAC able to open a stream in response to L3 RM's request, and if not, why not • process ID -- ID of the L2 MAC process associated with this stream • process ID -- ID of the L2 RLC process associated with this stream • L2MacStreamType -- what kind of a stream was opened (e.g., reg, TCH, TCH-short, none). If the stream request failed, then this should be set to the StreamType specified in the associated L2MacUtStart.req. • AfnSort -- when did the stream start (the RA/AA frame) • ibChanNum -- what channel is the stream on. If the stream request failed, then this should be set to the channel specified in the associated L2MacUtStart.req, if a channel was specified. If no channel was specified, then it should be set to ibChanNumInvalid. • ibSubChan -- is the stream full-rate or partial-rate • Integer -- how many RA bursts were sent while trying to open the stream • L2MacRejectSort -- (this parameter is relevant only for L2MacStreamType = None, i.e., an AA-reject, AA-invalid-id, or AA-pending) why was the stream rejected - was the page ID invalid, or did the BS simply choose to not assign a stream? If the latter, how long does this reg need to wait before issuing the next stream request, and might the BS page the UT in the meantime? (struct) • L2MacRejectSort::Boolean -- (noAa) Whether or not the BS sends an AA • L2MacRejectSort::AaSubTypeSort -- (aaSubType) AaInvalidId, AaReject or AaPending • L2MacRejectSort::Integer -- (holdMsg) encoded hold frames -- only valid for aaSubType = AaReject • L2MacRejectSort::Boolean -- (noPage) if true, the BS is certain to not page the UT during the specified hold period -- only valid for aaSubType = AaReject • L2MacRejectSort::AaRaDelaySort -- (raDelay) indicates the amount of time (frames) after which the UT should send another RA for the extended RA mechanism -- only valid for aaSubType = AaPending 	

<i>L2MacUtPing.conf</i>	
Applies to:	UT only
Associated primitives:	L2MacUtPing.req
Description: confirmation from UT L2 MAC to UT L3 RM of a requested ping	
Parameters: <ul style="list-style-type: none"> • L2MacUtReqResultSort -- was the requested ping successful • L2MacUtPageDomainSort -- which registration group (sharing a paging ID within a paging domain) was the ping for (struct) • L2MacUtPageDomainSort::ibPageId -- (pageId) which paging ID within the registration domain • L2MacUtPageDomainSort::ibBsc -- (bscc) which BS • L2MacUtPageDomainSort::ibChanSet -- (convPchChans) which chans the BS has enabled for paging - this determines the mapping between paging subslots and implied resources set element type: ibChanNum • L2MacUtPageDomainSort::ibChanSet -- (logicPchChans) which chans the BS will use to page this UT - this determines the set of channels this paging ID will hop over set element type: ibChanNum • Integer -- how many RAs were sent before an AA was heard • Boolean -- was the page ID used in the ping invalid (true -> invalid, false -> valid) 	

10.3.2.11 L3 RM to L2 MAC Interface Primitives

Primitive list:

- ◆ L2MacUtStart.req
- ◆ L2MacUmCtrl.req
- ◆ L2MacUtAfnAdj.req
- ◆ L2MacUtPing.req
- ◆ L2MacBsRachMasks.req
- ◆ L2MacBsPageTx.req
- ◆ L2MacBsUtCapability.req
- ◆ L2MacBsPageChans.req
- ◆ L2MacUtPageRx.req
- ◆ L2MacKill.req
- ◆ L2MacIsec.req
- ◆ L2MacL3RmUtStart.resp
- ◆ L2MacL3RmBsStart.resp

Request primitive descriptions

<i>L2MacUtStart.req</i>	
Applies to:	UT only
Associated primitives:	<i>L2MacUtStart.conf</i> (Note that the <i>L2RlcUtStart.ind</i> and <i>L2RlcUtStart.resp</i> are also triggered by the <i>L2MacUtStart.req</i> primitive.)
Description: request from UT L3 RM to UT L2 MAC for a new stream	
Parameters:	
<ul style="list-style-type: none"> • <i>process ID</i> -- ID of the L3 RM process requesting the stream • <i>L2MacUtRegInfoSort</i> -- which registration the stream is for (struct) • <i>L2MacUtRegInfoSort::ibRegId</i> -- (<i>regId</i>) which registration ID within the registration domain • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort</i> -- (<i>pageDomain</i>) which paging domain (this identifies the reg domain) (struct) • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibPageId</i> -- (<i>pageId</i>) which paging ID within the registration domain • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibBscc</i> -- (<i>bscc</i>) which BS • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibChanSet</i> -- (<i>convPchChans</i>) which chans the BS has enabled for paging - this determines the mapping between paging subslots and implied resources set element type: <i>ibChanNum</i> • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibChanSet</i> -- (<i>logicPchChans</i>) which chans the BS will use to page this UT - this determines the set of channels this paging ID will hop over set element type: <i>ibChanNum</i> • <i>L2MacUtRegInfoSort::ibChanSet</i> -- (<i>rachChans</i>) the chan set usable for RACH set element type: <i>ibChanNum</i> • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort</i> -- (<i>raPolicy</i>) when to resend RA bursts (struct) • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (<i>Na</i>) number of channels per subscan • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (<i>Nb</i>) intrasubscan delay • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (<i>Nc</i>) intersubscan delay, scaling term • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (<i>Nd</i>) intersubscan delay, exponential base term • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (<i>Ne</i>) intersubscan delay limit • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (<i>Nf</i>) total time limit • <i>L2MacUtChanPref</i> -- is there a specific channel to try (struct) • <i>L2MacUtChanPref::ibChanNum</i> -- (<i>prefChan</i>) which chan is preferred? (<i>ibChanNumInvalid</i> signifies no preference) • <i>L2MacUtChanPref::Integer</i> -- (<i>prefAttempts</i>) how many times to try on the pref chan? • <i>L2MacStreamType</i> -- what kind of a stream is being requested (e.g., reg, TCH, TCH-short) • <i>IbBsCapabilitySort</i> -- capabilities of the BS which the stream is requested to target (needed for open-loop power control's selection of RA-rts power) (struct) • <i>IbBsCapabilitySort::Real</i> -- (<i>noiseFloorDbm</i>) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort</i> -- (<i>decode</i>) decode delays for different channels (struct) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (<i>tchMap</i>) tch decode delay information array index type: <i>DuplexTimeSlotSort</i> array element type: <i>Bit</i> • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (<i>facchMap</i>) facch decode delay information array index type: <i>DuplexTimeSlotSort</i> array element type: <i>Bit</i> 	

<i>L2MacUmCtrl.req</i>	
Applies to:	BS, UT
Associated primitives:	L2MacUmCtrl.ind (on peer)
<p>Description: request by a higher protocol layer for L2 MAC to deliver a UM control message across the air interface to the peer. The requestor gives L2 MAC an indication of how long L2 MAC should attempt to deliver the UM ctrl datagram message by specifying how long the message is valid, and the maximum number of times it should be transmitted. Once the message has been transmitted the maximum number of times, it will be discarded, even if its validity time has not expired. Once the validity time expires, the message will be discarded, even if it has not been sent the maximum number of times (or at all).</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • L2MacUmCtrlType -- which UM control message • L2MacUmCtrlData -- the message's data (may not be relevant, depending on message type). For messages that contain data, the data provided in this parameter must already be in the appropriately encoded format; L2 MAC will send this data without reformatting it. • L2MacUmCtrlMaxTx -- the max number of times the UM ctrl message may be transmitted (the value 0 indicates no limit on the number of transmissions) • L2MacUmCtrlValidFrms -- how long (frames) the UM ctrl message is valid (the value 0 indicates the message is always valid) 	

<i>L2MacUtAfnAdj.req</i>	
Applies to:	UT only
Associated primitives:	none
<p>Description: request from UT L3 RM to L2 UT MAC to adjust the AFN by the specified amount</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • AfnDiffSort -- the adjustment to be applied to the AFN - L2 MAC must add this adjustment to the current AFN 	

<i>L2MacUtPing.req</i>	
Applies to:	UT only
Associated primitives:	L2MacUtPing.conf
Description: request from UT L3 RM to UT L2 MAC to send the BS a ping to renew the registration	
Parameters:	
<ul style="list-style-type: none"> • L2MacUtRegInfoSort -- which RACH set and RA policy should be used, and which paging domain is the ping for (struct) • <i>L2MacUtRegInfoSort::ibRegId</i> -- (regId) which registration ID within the registration domain • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort</i> -- (pageDomain) which paging domain (this identifies the reg domain) (struct) • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibPageId</i> -- (pageId) which paging ID within the registration domain • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibBsc</i> -- (bsc) which BS • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibChanSet</i> -- (convPchChans) which chans the BS has enabled for paging - this determines the mapping between paging subslots and implied resources set element type: <i>ibChanNum</i> • <i>L2MacUtRegInfoSort::L2MacUtPageDomainSort::ibChanSet</i> -- (logicPchChans) which chans the BS will use to page this UT - this determines the set of channels this paging ID will hop over set element type: <i>ibChanNum</i> • <i>L2MacUtRegInfoSort::ibChanSet</i> -- (rachChans) the chan set usable for RACH set element type: <i>ibChanNum</i> • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort</i> -- (raPolicy) when to resend RA bursts (struct) • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (Na) number of channels per subscan • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (Nb) intrasubscan delay • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (Nc) intersubscan delay, scaling term • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (Nd) intersubscan delay, exponential base term • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (Ne) intersubscan delay limit • <i>L2MacUtRegInfoSort::L2MacUtRaPolicySort::Integer</i> -- (Nf) total time limit • <i>IbBsCapabilitySort</i> -- capabilities of the BS which the ping is requested to target (needed for open-loop power control's selection of RA-ping power) (struct) • <i>IbBsCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: <i>DuplexTimeSlotSort</i> array element type: Bit • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: <i>DuplexTimeSlotSort</i> array element type: Bit 	

<i>L2MacBsRachMasks.req</i>	
Applies to:	BS only
Associated primitives:	none
<p>Description: request from BS L3 RM to BS L2 MAC to update the RACH carrier and timeslot masks being sent in the CMs</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>CarrierMaskSort</i> -- the carriers enabled for RACH (specifically RA-rreg). Note that the bit-ordering within this field is the same as the bit-ordering of the RACH carrier mask in the CM; no adjustment is needed. • <i>TimeslotMaskSort</i> -- the timeslots enabled for RACH (specifically RA-rreg). Note that the bit-ordering within this field is the same as the bit-ordering of the RACH slot mask in the CM; no adjustment is needed. 	

<i>L2MacBsPageTx.req</i>	
Applies to:	BS only
Associated primitives:	<i>L2MacBsPageChans.req</i> , <i>L2MacBsPageTx.ind</i>
<p>Description: request from L3 RM to L2 MAC to send paging bursts</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>ibPageId</i> -- whom to page • <i>L2MacPageLevel</i> -- how often to page • <i>ibChanSet</i> -- what conventional channels (implied resources) the paging ID in question is hopping over - having identified the implied resource for the page, the BS uses the mask of conventional channels being used for paging (from the <i>L2MacBsPageChans.req</i> msg) to determine which subslots to transmit the page on. set element type: <i>ibChanNum</i> • <i>Integer</i> -- paging priority - if L2 MAC could page either of two users, it must select the one with the higher priority. A lower value indicates a higher priority, that is 0 -> highest priority, 1 -> second highest, etc. • <i>Boolean</i> -- will the UT try to respond on the implied channel? (true->yes, false->no) This has implications on how many pages L2 MAC can send concurrently. 	

<i>L2MacBsUtCapability.req</i>	
Applies to:	BS only
Associated primitives:	L2MacBsUtCapability.ind
<p>Description:</p> <p>request from BS L3 RM to BS L2 MAC to pass the provided specification of the UT's capability on to the L3 RRC instance associated with this stream. This is only used during the registration stream, where BS L3 RM first finds out the UT's capabilities through messages received while the stream is in progress. For all later streams, BS L3 RM will use the L2MacL3RmBsStart.resp to inform L2 MAC of the UT's capabilities.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>IbL3RrcUtCapabilitySort</i> -- pwr ctrl / link adapt capabilities of the UT on the remote side of this stream (struct) • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (minTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::Real</i> -- (noiseFloorDbm) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort</i> -- (decode) decode delays for different channels (struct) • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::Bit</i> -- (cch) cch decode delay information • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (rachMap) rach decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (pchMap) pch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (ulModClasses) which mod classes can the UT tx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::ModClassSetSort</i> -- (dlModClasses) which mod classes can the UT rx set element type: ModClassSort • <i>IbL3RrcUtCapabilitySort::PowerDbmSort</i> -- (maxTxPwrDbm) • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod01) how much higher is the mod class 0 and 1 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • <i>IbL3RrcUtCapabilitySort::PowerDbSort</i> -- (txPwrDiffMod2345) how much higher is the mod class 2-5 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? 	

<i>L2MacBsPageChans.req</i>	
Applies to:	BS only
Associated primitives:	L2MacBsPageTx.req
Description: request from L3 RM to L2 MAC to identify which channels to use for paging	
Parameters: <ul style="list-style-type: none"> • <i>ibChanSet</i> -- what conventional channels the BS is allowed to use to send pages and receive page response RAs. This channel mask determines the mapping of implied resources to paging subslots and vice versa. set element type: <i>ibChanNum</i>	

<i>L2MacUtPageRx.req</i>	
Applies to:	UT only
Associated primitives:	L2MacUtPageRx.ind, L2MacL3RmUtStart.ind
Description: request from UT L3 RM to UT L2 MAC to listen for paging bursts	
Parameters: <ul style="list-style-type: none"> • <i>L2MacUtPageDomainSort</i> -- when/where to listen for pages, what paging ID to listen for (struct) • <i>L2MacUtPageDomainSort::ibPageId</i> -- (<i>pageId</i>) which paging ID within the registration domain • <i>L2MacUtPageDomainSort::ibBsc</i> -- (<i>bsc</i>) which BS • <i>L2MacUtPageDomainSort::ibChanSet</i> -- (<i>convPchChans</i>) which chans the BS has enabled for paging - this determines the mapping between paging subslots and implied resources set element type: <i>ibChanNum</i> • <i>L2MacUtPageDomainSort::ibChanSet</i> -- (<i>logicPchChans</i>) which chans the BS will use to page this UT - this determines the set of channels this paging ID will hop over set element type: <i>ibChanNum</i> • <i>L2MacPageLevel</i> -- how frequently to listen for pages • <i>Integer</i> -- how many times to send RA-page-response on the implied resource before giving up and sending a L2MacUtPageRx.ind to UT L3 RM • <i>IbBsCapabilitySort</i> -- capabilities of the BS sending the page (this is needed to decide at what power to send the RA-page-response after the page is heard) (struct) • <i>IbBsCapabilitySort::Real</i> -- (<i>noiseFloorDbm</i>) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort</i> -- (<i>decode</i>) decode delays for different channels (struct) • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (<i>tchMap</i>) <i>tch</i> decode delay information array index type: <i>DuplexTimeSlotSort</i> array element type: <i>Bit</i> • <i>IbBsCapabilitySort::IbBsDecodeDelaysSort::DecodeDelayBitMapSort</i> -- (<i>facchMap</i>) <i>facch</i> decode delay information array index type: <i>DuplexTimeSlotSort</i> array element type: <i>Bit</i> 	

<i>L2MacKill.req</i>	
Applies to:	BS, UT
Associated primitives:	L2MacL3RmStop.ind
<p>Description: request from L3 RM or L2 RLC to forcibly close a stream It will take a finite time for L2 MAC to complete this request to close the stream, since it will send a single best-effort UM control close.cmd message to inform the peer L2 MAC that the stream is being forcibly closed. When L2 MAC has completed the stream closure, it should follow the regular process for informing L3 RM with a L2MacL3RmStop.ind, which will in turn send a L2RlcStop.req to L2 RLC. The closure cause in the L2MacL3RmStop.ind signal will tell L3 RM whether the stream closed because it had no more data or if it was expressly killed.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> L2MacStopCauseSort -- reason this stream was killed (this is reported to L3 RM in the L2MacL3RmStop.ind) 	

<i>L2MacIsec.req</i>	
Applies to:	BS, UT
Associated primitives:	none
<p>Description: request from L3 RM to L2 MAC to start encrypting/decrypting the data using the i-SEC protocol. This primitive must always be sent from L3 RM to L2 MAC at the start of a stream. For streams that don't use encryption (registration streams), L3 RM will use this primitive to explicitly tell L2 MAC to not do encryption.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> Boolean -- whether encryption/decryption should be performed BitString -- the shared secret to use in encryption/decryption 	

Response primitive descriptions

<i>L2MacL3RmUtStart.resp</i>	
Applies to:	UT only
Associated primitives:	L2MacL3RmUtStart.ind
<p>Description: response from UT L3 RM to UT L2 MAC stating whether UT L3 RM accepted the new stream (opened by a page)</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> Boolean -- was the stream accepted process ID -- ID of the L2 RLC process associated with this stream process ID -- ID of the L3 RM process associated with this stream 	

<i>L2MacL3RmBsStart.resp</i>	
Applies to:	BS only
Associated primitives:	<i>L2MacL3RmBsStart.ind</i>
Description: response from L3 RM to L2 MAC stating whether L3 RM accepted the new stream	
Parameters:	
<ul style="list-style-type: none"> • process ID -- ID of the L2 RLC process associated with this stream • process ID -- ID of the L3 RM process associated with this stream • L2MacBsStreamAssignSort -- what kind of stream, which chan, for whom (struct) • L2MacBsStreamAssignSort::ibRegId -- (regId) who is the stream for • L2MacBsStreamAssignSort::ibChanNum -- (chan) which conventional channel should the stream use • L2MacBsStreamAssignSort::ibSubChan -- (subChan) full-rate vs. fractional-rate channel • L2MacBsStreamAssignSort::L2MacStreamType -- (strmType) what kind of stream is it, e.g., reg, TCH, TCH-short, or none • L2MacBsStreamAssignSort::L2MacRejectSort -- (this parameter is relevant only for L2MacStreamType = None, i.e., an AA-reject, AA-invalid-id, or AA-pending) why was the stream rejected - was the page ID invalid, or did the BS simply choose to not assign a stream? If the latter, how long does this reg need to wait before issuing the next stream request, and might the BS page the UT in the meantime? (struct) • L2MacBsStreamAssignSort::L2MacRejectSort::Boolean -- (noAa) Whether or not the BS sends an AA • L2MacBsStreamAssignSort::L2MacRejectSort::AaSubTypeSort -- (aaSubType) AaInvalidId, AaReject or AaPending • L2MacBsStreamAssignSort::L2MacRejectSort::Integer -- (holdMsg) encoded hold frames -- only valid for aaSubType = AaReject • L2MacBsStreamAssignSort::L2MacRejectSort::Boolean -- (noPage) if true, the BS is certain to not page the UT during the specified hold period -- only valid for aaSubType = AaReject • L2MacBsStreamAssignSort::L2MacRejectSort::AaRaDelaySort -- (raDelay) indicates the amount of time (frames) after which the UT should send another RA for the extended RA mechanism -- only valid for aaSubType = AaPending • IbL3RrcUtCapabilitySort -- pwr ctrl / link adapt capabilities of the UT opening the stream (struct) • IbL3RrcUtCapabilitySort::PowerDbmSort -- (minTxPwrDbm) • IbL3RrcUtCapabilitySort::Real -- (noiseFloorDbm) • IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort -- (decode) decode delays for different channels (struct) • IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::Bit -- (cch) cch decode delay information • IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort -- (rachMap) rach decode delay information array index type: DuplexTimeSlotSort array element type: Bit • IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort -- (tchMap) tch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort -- (facchMap) facch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • IbL3RrcUtCapabilitySort::IbUtDecodeDelaysSort::DecodeDelayBitMapSort -- (pchMap) pch decode delay information array index type: DuplexTimeSlotSort array element type: Bit • IbL3RrcUtCapabilitySort::ModClassSetSort -- (ulModClasses) which mod classes can the UT tx set element type: ModClassSort • IbL3RrcUtCapabilitySort::ModClassSetSort -- (dlModClasses) which mod classes can the UT rx set element type: ModClassSort • IbL3RrcUtCapabilitySort::PowerDbmSort -- (maxTxPwrDbm) • IbL3RrcUtCapabilitySort::PowerDbSort -- (txPwrDiffMod01) how much higher is the mod class 0 and 1 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? • IbL3RrcUtCapabilitySort::PowerDbSort -- (txPwrDiffMod2345) how much higher is the mod class 2-5 maximum transmit power than the minimum across all supported mod classes of the max transmit powers? 	

10.3.2.12 L2 RLC to L3 RM Interface Primitives

Primitive list:

- ◆ L2RlcData.ind
- ◆ L2RlcHungry.ind
- ◆ L2RlcStart.conf
- ◆ L2RlcStop.conf
- ◆ L2RlcTxStatus.conf

Indication primitive descriptions

<i>L2RlcData.ind</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcData.req (on peer)
Description: indication from L2 RLC to L3 RM that incoming data bytes are available	
Parameters: <ul style="list-style-type: none"> • ByteString -- a series of bytes received from the peer L2 RLC • L2RlcChan -- indication of whether the data was received on the AM (acknowledged mode) or UM (unacknowledged mode) channel 	

<i>L2RlcHungry.ind</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcData.req
Description: indication from L2 RLC to L3 RM that it will soon run out of outgoing data. L3 RM is free to choose how to respond to this indication, but typically L3 RM would respond with a L2RlcData.req if L3 RM has outgoing data available.	
Parameters: <ul style="list-style-type: none"> • Integer -- how many bytes of outgoing data does L2 RLC have enqueued but not transmitted (note that this does not include bytes which have been transmitted but not acknowledged) • Integer -- how many bytes of outgoing data have been acknowledged since the report in the previous Hungry.ind or TxStatus.conf primitive 	

Confirmation primitive descriptions

<i>L2RlcStart.conf</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcStart.req
Description: confirmation from L2 RLC to L3 RM that L2 RLC was able to open a communications channel to service the L3 RM registration This primitive is used by BS L3 RM for every stream start, and by UT L3 RM for streams which are started by paging.	
Parameters: <ul style="list-style-type: none"> • Boolean -- was the request for a communications channel satisfied • process ID -- ID of the L2 RLC process associated with the stream 	

<i>L2RlcStop.conf</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcStop.req
Description: confirmation from L2 RLC to L3 RM that L2 RLC has closed its communications channel which was servicing the L3 RM registration	
Parameters: <ul style="list-style-type: none"> Integer -- how many bytes of outgoing data L2 RLC has gotten acknowledged since the report in the previous Hungry.ind or TxStatus.conf primitive 	

<i>L2RlcTxStatus.conf</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcTxStatus.req, L2RlcHungry.ind
Description: confirmation of L3 RM's request for L2 RLC's tx status information This primitive provides the same information as the L2RlcHungry.ind primitive plus the third argument that describes how much space (in bytes) the TX buffer of L2 RLC has left. This signal can be requested by L3 RM, whereas the L2RlcHungry.ind is generated autonomously and asynchronously by L2 RLC.	
Parameters: <ul style="list-style-type: none"> Integer -- how many bytes of outgoing data does L2 RLC have enqueued but not transmitted (note that this does not include bytes which have been transmitted but not acknowledged) Integer -- how many bytes of outgoing data have been acknowledged since the report in the previous Hungry.ind or TxStatus.conf primitive Integer -- how much space (in bytes) the TX buffer of L2 RLC has left - a negative value indicates that L2 RLC imposes no buffer limit. 	

10.3.2.13 L3 RM to L2 RLC Interface Primitives

Primitive list:

- ◆ L2RlcStart.req
- ◆ L2RlcStop.req
- ◆ L2RlcData.req
- ◆ L2RlcTxStatus.req
- ◆ L2RlcConfig.req
- ◆ L2RlcData.resp

Request primitive descriptions

<i>L2RlcStart.req</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcStart.conf, L2RlcConfig.req
Description: request from L3 RM for a L2 RLC connection to transfer data across the air This primitive is used by BS L3 RM for all streams and by UT L3 RM for streams started by paging.	
Parameters: <ul style="list-style-type: none"> process ID -- ID of the L2 MAC process associated with the stream process ID -- ID of the L3 RM process associated with the stream 	

<i>L2RlcStop.req</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcStop.conf
<p>Description: request from L3 RM to close a L2 RLC connection Note that unlike most of the HC-SDMA protocol layer interfaces, which have just a unilateral Stop.req with no Stop.conf, the L3 RM / L2 RLC stop is a two-part Stop.req + Stop.conf handshake. L2 RLC is allowed to still send L2RlcData.ind signals to L3 RM between the times when the L2RlcStop.req is received and the L2RlcStop.conf is sent.</p>	

<i>L2RlcData.req</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcData.ind (on peer)
<p>Description: request from L3 RM for L2 RLC to deliver data bytes across the air interface</p> <p>Parameters:</p> <ul style="list-style-type: none"> • ByteString -- data segment to deliver across the air interface to the peer • Boolean -- does the data segment complete a packet or is there still an outstanding fragment? (true -> complete packet, false -> outstanding fragment) • L2RlcChan -- indication of whether the data should be sent via AM (acknowledged mode, i.e., reliable delivery) or UM (unacknowledged mode, i.e., best effort delivery) 	

<i>L2RlcTxStatus.req</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcTxStatus.conf, L2RlcHungry.ind
<p>Description: request from L3 RM for a summary of how much untransmitted data L2 RLC has, and how much data has been acknowledged by the peer This signal provides a way for L3 RM to request the same information provided by the L2RlcHungry.ind signal, which is generated asynchronously by L2 RLC, plus the space available information of the TX buffer.</p>	

<i>L2RlcConfig.req</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcStart.req, L2MacL3RmBsStart.ind, L2MacL3RmUtStart.ind (The L2RlcConfig.req primitive is sent from L3 RM to L2 RLC immediately after one of these associated primitives is received.)
<p>Description: request from L3 RM to L2 RLC providing parameters that affect L2 RLC's retransmission and stream closure policies</p> <p>Parameters:</p> <ul style="list-style-type: none"> • L2RlcIdleLimit -- how many frames should L2 RLC allow the stream to stay idle before shutting it down • Integer -- how many frames after the local L2 RLC transmits data the peer L2 RLC should transmit the associated acknowledgement. Note that this includes the TCH decode delay of the peer L2 RLC, but not of the local one. 	

Response primitive descriptions

<i>L2RlcData.resp</i>	
Applies to:	BS, UT
Associated primitives:	L2RlcData.ind
Description: response from L3 RM to L2 RLC stating whether the data received in the associated L2RlcData.ind signal completes a L3 packet, or if more incoming data is expected to complete the packet. This info is used by L2 RLC to decide when it should close the stream.	
Parameters: <ul style="list-style-type: none"> • Boolean -- did the L2RlcData.ind data end with a complete L3 packet 	

10.3.2.14 L3 RRC to L3 MMC Interface Primitives

Primitive list:

- ◆ L3RrcUtPathloss.conf

Confirmation primitive descriptions

<i>L3RrcUtPathloss.conf</i>	
Applies to:	UT only
Associated primitives:	L3RrcUtPathloss.req
Description: confirmation from UT L3 RRC to L3 MMC stating the pathloss measured between the UT and the BS in question.	
Parameters: <ul style="list-style-type: none"> • iBScc -- color code of the BS whose pathloss is reported • PathlossDbSort -- pathloss in dB 	

10.3.2.15 L3 MMC to L3 RRC Interface Primitives

Primitive list:

- ◆ L3RrcUtPathloss.req

Request primitive descriptions

<i>L3RrcUtPathloss.req</i>	
Applies to:	UT only
Associated primitives:	L3RrcUtPathloss.conf
Description: request from L3 MMC to UT L3 RRC for the measured pathloss between the UT and a specified BS. Note that the pathloss measurement could come from a BCH B burst rx, a CM rx, or an AA rx. L3 MMC can use the PathlossBurstType argument to partially specify which pathloss measurement it wants.	
Parameters: <ul style="list-style-type: none"> • iBScc -- color code of the BS whose pathloss is desired • PathlossBurstType -- should the pathloss estimate be based on the B-burst, or on the most recent/accurate estimate (the UT can make more precise pathloss estimates from the CM and AA bursts than from the B burst, though any measurement becomes increasingly suspect as it ages) 	

10.3.2.16 L3 MMC to L3 RM Interface Primitives

Primitive list:

- ◆ L3MmcUtNewBestBs.ind
- ◆ L3MmcUtBestBs.conf

Indication primitive descriptions

<i>L3MmcUtNewBestBs.ind</i>	
Applies to:	UT only
Associated primitives:	L3MmcUtBestBs.req
<p>Description: indication from L3 MMC to UT L3 RM that a new BS has been selected as the best one to use. This should trigger a handover, assuming that the UT is allowed to perform handovers. L3 MMC's BS ranking algorithm should include hysteresis so that this primitive will not be generated repeatedly as the UT crosses the boundary from one BS's cell into another's cell.</p>	

Confirmation primitive descriptions

<i>L3MmcUtBestBs.conf</i>	
Applies to:	UT only
Associated primitives:	L3MmcUtBestBs.req
<p>Description: confirmation from L3 MMC to UT L3 RM's query of which BS is best to register with</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>ibBsc</i> -- color code of the recommended basestation • <i>ibChanSet</i> -- set of channels the UT may use for registration streams set element type: <i>ibChanNum</i> • <i>Boolean</i> -- whether any BS meets the criteria for use 	

10.3.2.17 L3 RM to L3 MMC Interface Primitives

Primitive list:

- ◆ L3MmcUtRejectBs.req
- ◆ L3MmcUtBestBs.req
- ◆ L3MmcUtTrackBs.req

Request primitive descriptions

<i>L3MmcUtRejectBs.req</i>	
Applies to:	UT only
Associated primitives:	none
<p>Description: request from UT L3 RM to L3 MMC to reject a BS, due to conditions that L3 MMC is not aware of, such as i-HAP incompatibilities between the UT and the BS in question. Because the BSCC of the rejected BS will be reused throughout the network, it is important that this rejection will not be permanent. L3 MMC will autonomously re-enable the BSCC in question once L3 MMC has determined that the UT may be seeing a different BS with the same BSCC.</p>	
<p>Parameters:</p> <ul style="list-style-type: none"> • <i>ibBsc</i> -- color code of the BS to be ignored 	

<i>L3MmcUtBestBs.req</i>	
Applies to:	UT only
Associated primitives:	L3MmcUtNewBestBs.ind, L3MmcUtBestBs.conf
Description: request from UT L3 RM to L3 MMC to be told which BS is most suitable	

<i>L3MmcUtTrackBs.req</i>	
Applies to:	UT only
Associated primitives:	none
Description: request from UT L3 RM to L3 MMC to track the distance/freq/timing to the indicated BS. Only one BS is tracked at any given time	
Parameters: <ul style="list-style-type: none"> • <i>ibBsc</i> -- color code of the new BS to track 	

10.3.2.18 L3 RM to L3 CM Interface Primitives

Primitive list:

- ◆ L3RmBsStart.ind
- ◆ L3RmStop.ind
- ◆ L3RmData.ind
- ◆ L3RmCtrl.ind
- ◆ L3RmPktDrop.ind
- ◆ L3RmUtHandover.ind
- ◆ L3RmUtStart.conf
- ◆ L3RmUtCheckBs.conf
- ◆ L3RmPktRecover.conf

Indication primitive descriptions

<i>L3RmBsStart.ind</i>	
Applies to:	BS only
Associated primitives:	L3RmBsStart.resp
Description: indication from BS L3 RM to BS L3 CM that a new registration has been created, and requires an associated L3 CM connection	
Parameters: <ul style="list-style-type: none"> • <i>itapUtCertDataSort</i> -- (authenticated) UT certificate, of interest to BS L3 CM because it contains the UT ID (hardware ID) and for some certificate types, also the session ID (IMSI and session reference). BS L3 CM passes this information to BS L4 to use while establishing the connection to the PSS. 	

<i>L3RmStop.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L3 RM to L3 CM that the registration could not be maintained	

<i>L3RmData.ind</i>	
Applies to:	BS, UT
Associated primitives:	L3RmUtData.req/L3RmBsData.req (on peer)
Description: indication from L3 RM to L3 CM that incoming data packets are available	
Parameters: <ul style="list-style-type: none"> PktQueue -- a list of data packets received across the air interface from the peer L3 CM entity 	

<i>L3RmCtrl.ind</i>	
Applies to:	BS, UT
Associated primitives:	L3RmCtrl.req (on peer)
Description: indication from L3 RM to L3 CM that an incoming control packet is available	
Parameters: <ul style="list-style-type: none"> Pkt -- a control packet received across the air interface from the peer L3 CM entity 	

<i>L3RmPktDrop.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L3 RM to L3 CM that one or more packets have been dropped while in transit across the air interface. There are two reasons why the packet could have been dropped. On the sending side, a packet may have been discarded because of excessive queue length. In this case, L3 RM will notify the local L3 CM with a PktDrop(outgoing) signal. On the receiving side, the PDCL algorithms may discard packets which are corrupted (detected by the checksum algorithm) or irrecoverably out of order (detected by the packet reordering algorithm). In this case, L3 RM will notify the local L3 CM with a PktDrop(incoming) signal.	
Parameters: <ul style="list-style-type: none"> Direction -- which direction the dropped packet(s) was(were) traveling Because the direction is specified with respect to the air interface, the interpretation of the direction is different on BS versus UT: BS: incoming -> uplink, outgoing -> downlink UT: incoming -> downlink, outgoing -> uplink 	

<i>L3RmUtHandover.ind</i>	
Applies to:	UT only
Associated primitives:	L3RmStop.req (for reg to old BS), L3RmUtStart.req (for reg to new BS)
Description: indication from UT L3 RM to UT L3 CM that the connection in question should handover to a new BS to obtain better service than the current registration can provide	

Confirmation primitive descriptions

<i>L3RmUtStart.conf</i>	
Applies to:	UT only
Associated primitives:	L3RmUtStart.req
Description: confirmation from UT L3 RM to UT L3 CM stating whether UT L3 RM was able to establish the requested registration	

Parameters:
<ul style="list-style-type: none"> • Boolean -- was the registration established

<i>L3RmUtCheckBs.conf</i>	
Applies to:	UT only
Associated primitives:	L3RmUtCheckBs.req
Description: confirmation from UT L3 RM to UT L3 CM regarding whether the BS in use by the L3 RM instance is the one currently regarded as best by L3 MMC	
Parameters:	
<ul style="list-style-type: none"> • Boolean -- is the BS in use by the UT L3 RM instance the one currently selected by L3 MMC? 	

<i>L3RmPktRecover.conf</i>	
Applies to:	BS, UT
Associated primitives:	L3RmPktRecover.req
Description: confirmation from L3 RM to L3 CM returning the unsent data and L3 CM or L4 control packets	
Parameters:	
<ul style="list-style-type: none"> • PktQueue -- unsent L3 CM and L4 control packets • PktQueue -- unsent data packets (which have not gone through the PDCL transformation) 	

10.3.2.19 L3 CM to L3 RM Interface Primitives

Primitive list:

- ◆ L3RmUtStart.req
- ◆ L3RmStop.req
- ◆ L3RmUtData.req
- ◆ L3RmBsData.req
- ◆ L3RmCtrl.req
- ◆ L3RmUtCheckBs.req
- ◆ L3RmPktRecover.req
- ◆ L3RmBsStart.resp

Request primitive descriptions

<i>L3RmUtStart.req</i>	
Applies to:	UT only
Associated primitives:	L3RmUtStart.conf
Description: request from UT L3 CM to UT L3 RM for a registration to allow communication across the air interface	
Parameters:	
<ul style="list-style-type: none"> • ibPktProtocol -- protocol of packets being tunneled across the HC-SDMA air interface • itapUtCertTypeSort -- what type of UT certificate the session should use - network access (for user data transfer) or identity (for UT management) 	

<i>L3RmStop.req</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: request from L3 CM to L3 RM to terminate the registration	
Parameters: <ul style="list-style-type: none"> • Boolean -- stop gracefully or abruptly: true -> attempt to send all remaining data and then send a Dereg message to inform the peer that the registration is being terminated; false -> kill any streams and terminate immediately, discarding any remaining data 	

<i>L3RmUtData.req</i>	
Applies to:	UT only
Associated primitives:	L3RmData.ind (on BS)
Description: request from UT L3 CM to UT L3 RM to send data packets across the air interface to BS L3 CM	
Parameters: <ul style="list-style-type: none"> • PktQueue -- a list of data packets to send across the air interface to the BS L3 CM entity 	

<i>L3RmBsData.req</i>	
Applies to:	BS only
Associated primitives:	L3RmData.ind (on UT)
Description: request from BS L3 CM to BS L3 RM to send data packets across the air interface to UT L3 CM	
Parameters: <ul style="list-style-type: none"> • PktQueue -- a list of data packets to send across the air interface to the UT L3 CM entity • ibQoS -- the quality of service that should be associated with these packets. Note that this QoS marking will probably also be applied to all preceding outgoing packets that have not yet been sent. 	

<i>L3RmCtrl.req</i>	
Applies to:	BS, UT
Associated primitives:	L3RmCtrl.ind (on peer)
Description: request from L3 CM to L3 RM to send a control packet across the air interface to the peer L3 CM	
Parameters: <ul style="list-style-type: none"> • Pkt -- a control packet to send across the air interface to the peer L3 CM entity 	

<i>L3RmUtCheckBs.req</i>	
Applies to:	UT only
Associated primitives:	L3RmUtCheckBs.conf
Description: request from UT L3 CM to UT L3 RM to verify that the BS that the L3 RM instance is registered with is the BS considered best by L3 MMC	

<i>L3RmPktRecover.req</i>	
Applies to:	BS, UT
Associated primitives:	L3RmPktRecover.conf
Description: request from L3 CM to L3 RM to return any unsent L3 CM and L4 ctrl pkts and any data pkts that have not yet undergone the PDCL transformations. L3 CM can use this signal to recover data from one registration and re-send it on another registration, e.g., on the UT during handover.	

Response primitive descriptions

<i>L3RmBsStart.resp</i>	
Applies to:	BS only
Associated primitives:	L3RmBsStart.ind
Description: response from BS L3 CM to BS L3 RM stating whether BS L3 CM was able to start a connection in response to the RmStart.ind from BS L3 RM	
Parameters: <ul style="list-style-type: none"> • Boolean -- was the connection started 	

10.3.2.20 L3 CM to L4 Interface Primitives

Primitive list:

- ◆ L3CmBsStart.ind
- ◆ L3CmStop.ind
- ◆ L3CmData.ind
- ◆ L3CmCtrl.ind
- ◆ L3CmUtHandover.ind
- ◆ L3CmPktDrop.ind
- ◆ L3CmUtStart.conf

Indication primitive descriptions

<i>L3CmBsStart.ind</i>	
Applies to:	BS only
Associated primitives:	L3CmBsStart.resp
Description: indication from BS L3 CM to BS L4 to start a connection to the BS-side network	
Parameters: <ul style="list-style-type: none"> • itapUtCertDataSort -- (authenticated) UT certificate, of interest to BS L4 because it contains the UT ID (hardware ID) and for some certificate types, also the session ID (IMSI and session reference). BS L4 uses this information while establishing the connection to the PSS. 	

<i>L3CmStop.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L3 CM to L4 that the air interface connection could not be maintained	

<i>L3CmData.ind</i>	
Applies to:	BS, UT
Associated primitives:	L3CmUtData.req/L3CmBsData.req (on peer)
Description: indication from L3 CM to L4 that incoming data packets are available	
Parameters: <ul style="list-style-type: none"> • PktQueue -- a list of data packets received across the air interface from the peer L4 entity 	

<i>L3CmCtrl.ind</i>	
Applies to:	BS, UT
Associated primitives:	L3CmCtrl.req (on peer)
Description: indication from L3 CM to L4 that an incoming control packet is available	
Parameters: <ul style="list-style-type: none"> • Pkt -- a control packet received across the air interface from the peer L4 entity 	

<i>L3CmUtHandover.ind</i>	
Applies to:	UT only
Associated primitives:	L3CmUtHandover.resp
Description: indication from UT L3 CM to UT L4 that a handover is needed to maintain a reliable connection. UT L4 must answer with a L3CmUtHandover.resp to inform UT L3 CM whether a handover is permissible.	

<i>L3CmPktDrop.ind</i>	
Applies to:	BS, UT
Associated primitives:	none
Description: indication from L3 CM to L4 that a packet has been dropped while in transit across the air interface	
Parameters: <ul style="list-style-type: none"> • Direction -- which direction the dropped packet(s) was(were) traveling Because the direction is specified with respect to the air interface, the interpretation of the direction is different on BS versus UT: BS: incoming -> uplink, outgoing -> downlink UT: incoming -> downlink, outgoing -> uplink 	

Confirmation primitive descriptions

<i>L3CmUtStart.conf</i>	
Applies to:	UT only
Associated primitives:	L3CmUtStart.req
Description: confirmation from UT L3 CM to UT L4 stating whether UT L3 CM was able to establish the requested connection across the air interface	
Parameters: <ul style="list-style-type: none"> • Boolean -- was the connection established 	

10.3.2.21 L4 to L3 CM Interface Primitives

Primitive list:

- ◆ L3CmUtStart.req
- ◆ L3CmStop.req
- ◆ L3CmBsData.req
- ◆ L3CmUtData.req
- ◆ L3CmCtrl.req
- ◆ L3CmBsSrv.req
- ◆ L3CmBsStart.resp
- ◆ L3CmUtHandover.resp

Request primitive descriptions

<i>L3CmUtStart.req</i>	
Applies to:	UT only
Associated primitives:	L3CmUtStart.conf
Description: request from UT L4 to UT L3 CM to start a connection across the air interface	
Parameters: <ul style="list-style-type: none"> • <i>ibPktProtocol</i> -- protocol of packets being tunneled across the HC-SDMA air interface • <i>itapUtCertTypeSort</i> -- what type of UT certificate the session should use - network access (for user data transfer) or identity (for UT management) 	

<i>L3CmStop.req</i>	
Applies to:	BS, UT
Associated primitives:	none
Description request from L4 to L3 CM to terminate the connection across the air interface	

<i>L3CmBsData.req</i>	
Applies to:	BS only
Associated primitives:	L3CmData.ind (on UT)
Description: request from BS L4 to BS L3 CM to send data packets across the air interface to UT L4	
Parameters: <ul style="list-style-type: none"> • <i>PktQueue</i> -- a list of data packets to send across the air interface to the UT L4 entity • <i>ibQoS</i> -- the quality of service that should be associated with these packets. Note that this QoS marking will probably also be applied to all preceding outgoing packets that have not yet been sent. 	

<i>L3CmUtData.req</i>	
Applies to:	UT only
Associated primitives:	L3CmData.ind (on BS)
Description: request from UT L4 to UT L3 CM to send data packets across the air interface to BS L4	
Parameters: <ul style="list-style-type: none"> • <i>PktQueue</i> -- a list of data packets to send across the air interface to the BS L4 entity 	

<i>L3CmCtrl.req</i>	
Applies to:	BS, UT
Associated primitives:	L3CmCtrl.ind (on peer)
Description: request from L4 to L3 CM to send a control packet across the air interface to the peer L4	
Parameters: <ul style="list-style-type: none"> • Pkt -- a control packet to send across the air interface to the peer L4 entity 	

<i>L3CmBsSrv.req</i>	
Applies to:	BS only
Associated primitives:	none
Description: request from BS L4 to BS L3 CM to universally allow/deny new connections. This primitive is used by BS L4 to tell BS L3 CM whether to universally reject new registrations, or to attempt to create a new connection. If the registrations are initially accepted by the BS, and new connections are started, these connections may still be rejected on an individual basis by the network through the L3CmBsStart.resp primitive. If the BS has been servicing registrations but then receives the L3CmBsSrv.req(false) primitive from BS L4, all existing connections' registrations must be removed, though the "net sess stop" message must NOT be sent, so that the UT will have an opportunity to hand over to another BS.	
Parameters: <ul style="list-style-type: none"> • Boolean -- should the BS accept or categorically reject all new registrations? (true -> accept, false -> reject) 	

Response primitive descriptions

<i>L3CmBsStart.resp</i>	
Applies to:	BS only
Associated primitives:	L3CmBsStart.ind
Description: response from BS L4 to BS L3 CM stating whether BS L4 was able to start a connection to the BS network in response to the L3CmBsStart.ind from BS L3 CM	
Parameters: <ul style="list-style-type: none"> • Boolean -- was the connection accepted (true) or rejected (false) 	

<i>L3CmUtHandover.resp</i>	
Applies to:	UT only
Associated primitives:	L3CmUtHandover.ind
Description: response from UT L4 to UT L3 CM granting UT L3 CM permission to attempt a handover	
Parameters: <ul style="list-style-type: none"> • Boolean -- is the handover allowed 	

11 PRIVACY AND AUTHENTICATION

The goals of the HC-SDMA privacy and authentication protocols are to provide:

- ◆ *BS Authentication:* In order to ensure that User Terminals (UT) communicate only with authorized HC-SDMA Base Stations (BS).
- ◆ *UT Authentication:* In order to ensure that only authorized UTs have access to the HC-SDMA network and prevent session theft.
- ◆ *Confidentiality:* In order to maintain privacy of all user data over the HC-SDMA air interface.

Authentication and confidentiality are provided by the following protocols:

- ◆ *i-HAP* protocol defined in clause 11.2, *HC-SDMA Handshake and Authentication Protocol*, *i-HAP* is responsible for BS authentication and the establishment of a shared secret (K_r) between a UT and a BS.

The *i-HAP* protocol requires the digital signature of a BS (generated by a CA) to be verified at the user terminal. The digital signatures are formatted using: (+) ISO/IEC 9796 First Edition 1991-09-15: *Information Technology-Security Techniques-Digital Signature Scheme Giving Message Recovery*, September 15, 1991.

The signature generation/verification algorithm is the RSA algorithm as defined in Chapter 11.3 of: (+) A.J. Menezes, P. van Oorschot, and S. Vanstone, *Handbook of Applied Cryptography*, CRC Press Series on Discrete Mathematics and Its Applications, 1996.

Furthermore, *i-HAP* requires elliptic curve encryption primitives on the Koblitz curve K-163 for shared secret transmission from the user terminal to the base station. Elliptic curve encryption and decryption routines are described in: (+) Darrel Hankerson, Julio Lopez Hernandez, Alfred Menezes, "Software Implementation of Elliptic Curve Cryptography Over Binary Fields," *Lecture Notes in Computer Science*, 2000.

- ◆ *i-TAP* protocol defined in clause 11.3, *HC-SDMA Terminal Authentication Protocol*, *i-TAP* permits only authorized user terminals to access an HC-SDMA network.

The *i-TAP* protocol requires digital signatures of the UT generated by the CA to be verified at the BS. The algorithms for this are exactly the same as those in *i-HAP*.

In addition, *i-TAP* protocol requires elliptic curve digital signature generation at the user terminal and verification at the base station in addition to the digital signature verification tasks already described for *i-HAP*. The following reference defines elliptic curve digital signature generation on Koblitz curve K-163: (+) FIPS PUB 186-2: Federal Standards Publication: *Digital Signature Standard (DSS)*, January 27, 2000.

- ◆ *i-SEC* protocol defined in clause 11.4, *HC-SDMA Secure Communications Protocol*, *i-SEC* is responsible for symmetric encryption of data and control messages based on the shared secret K_r .

The symmetric key encryption engine utilized is the RC4 algorithm as described in: (+) Bruce Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, Second Edition, pages 397-398.

11.1 Overview

The flow of operations in the *HC-SDMA* privacy and authentication protocol shall conform to Figure 104 - General Description of Privacy and Authentication Protocols.

Upon power-up, the UT shall search⁵ for the BCH channel -- refer to clause 7.3.2. After the completion of the search, the UT shall establish a link with an appropriate BS. The *HC-SDMA* Handshake and Authentication Protocol (*i-HAP*) shall start after the BS selection is completed by the UT. The initial step of the *i-HAP* protocol shall be the transmission of the BS digital certificate to the UT. The BS certificate shall provide information about the public key, security parameters and the identity of the

⁵ BCH search is not a part of the *HC-SDMA* privacy and authentication protocol. During this search, the UT gathers minimal information about the *HC-SDMA* network parameters.

BS. The signature of the certificate shall be verified⁶ by the UT using the public key of the CA subsequently.

The UT shall generate random bits to be used as the shared secret K_r . The UT shall encrypt the shared secret using the BS public key and transmit the encrypted shared secret to the BS. The BS shall decrypt the message from the UT and determine the shared secret K_r . The length of the transmitted shared secret shall depend on the field size⁷ used for BS public key encryption.

The shared secret K_r shall be used by the *HC-SDMA* Secure Communication Protocol (*i-SEC*) in order to encrypt burst payloads that contain user data and air interface control messages and by *HC-SDMA* Terminal Authentication Protocol (*i-TAP*) for authenticator message (A_m) generation and scrambling UT specific parameters prior to transmission.

The UT Params message shall transmit the encrypted shared secret to the BS. In addition, the UT Params message shall include the UT digital certificate that provides information about the public key, security parameters and identity of the UT. This certificate is signed by the CA. Authenticity of the UT certificate shall be verified by the BS using the public key of the CA. The UT shall scramble⁸ its certificate together with UT specific parameters before transmitting them to the BS. Cryptographic authentication of the UT at the BS ensures that the identity presented by UT can be trusted by other entities in an *HC-SDMA* network without additional authentication steps.

The UT shall also transmit the UT authenticator string (A_s) to the BS in the same message as the UT certificate. The UT authenticator message (A_m) shall include a time stamp and parameters related to the UT identity. The hash value⁹ of the UT authenticator message shall be computed and encrypted with the UT private key¹⁰ to produce the authenticator string A_s . The authenticator string A_s shall be used to prevent UT certificate replay attacks.

In the case of a handover, the BCH search process is activated to determine the best candidate base station, after which *i-HAP* and *i-TAP* protocols shall be carried out with the new BS.

⁶ *HC-SDMA* privacy and authentication protocol does not require a specific method for certificate generation/distribution, but it does require that the public key of the CA is known at the UT for verification. *The HC-SDMA* protocol does not require a specific method for loading the CA public key in the UT, this can be accomplished during the production of the UT or before the sale of the UT.

⁷ Currently, BS public key encryption shall employ the finite field GF(2163) as defined in the field (7) of BS certificate and the shared secret K_r is transferred using two elements from GF(2163) as defined in field (8) of BS certificate, resulting in 326 bits for shared secret establishment (refer to clause 11.2.2, Base Station Transmits Its Digital Certificate).

⁸ Scrambling shall be performed using a 39 bit LFSR seeded with the last 39 bits of the shared secret K_r . Scrambling is intended to make UT tracking difficult for unauthorized parties.

⁹ Using the SHA-1 hash function. Refer to [SCHNEIER].

¹⁰ The private key shall be associated with the UT public key enclosed in the UT digital certificate. UT public key encryption shall employ the finite field GF(2¹⁶³) and the key shall be represented by its 164 bit minimal representation.

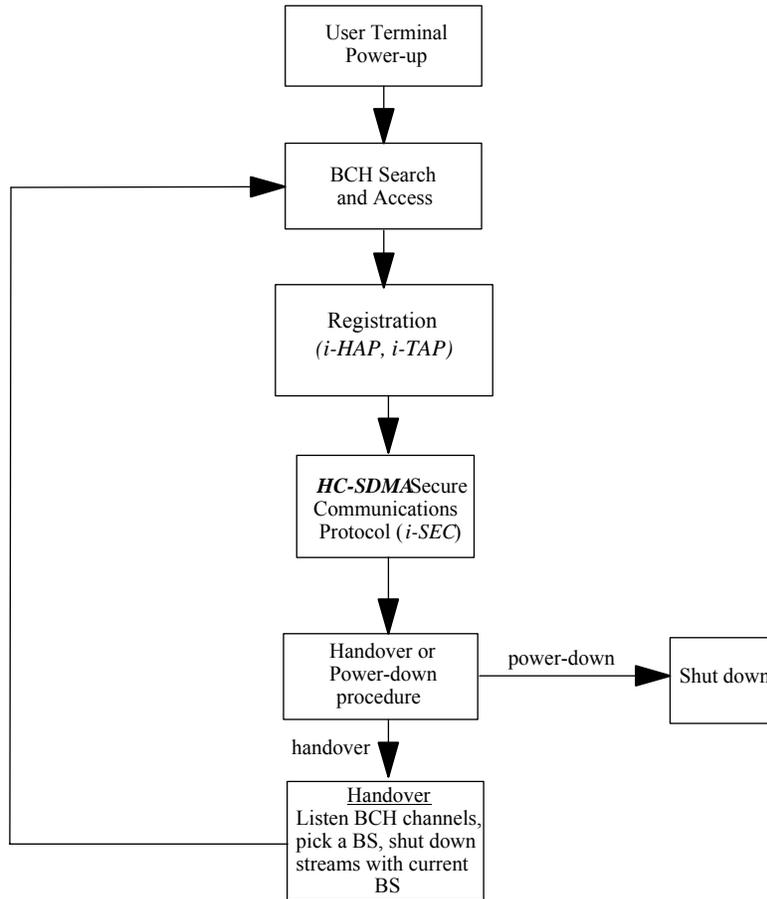


Figure 104: General Description of Privacy and Authentication Protocols

11.2 HC-SDMA Handshake and Authentication Protocol, i-HAP

The services within *i-HAP* shall be performed following the initial access to the selected base station. The *i-HAP* protocol shall conform to Figure 105 - Baseline *i-HAP* Event Diagram. The steps in the *i-HAP* protocol shall be as follows: a UT registers with a BS (Figure 106 - BCH Search and Registration), requests a digital certificate from the BS, verifies BS certificate authenticity and determines the BS public key (Figure 107 - HC-SDMA BS certificate message fields. Mandatory field values are identified in square brackets. Not to scale.), and finally establishes a shared secret K_r for symmetric key encryption of user data and control messages and user terminal authentication (Figure 112 - Encryption in HC-SDMA Secure Communication protocol, *i-SEC*).

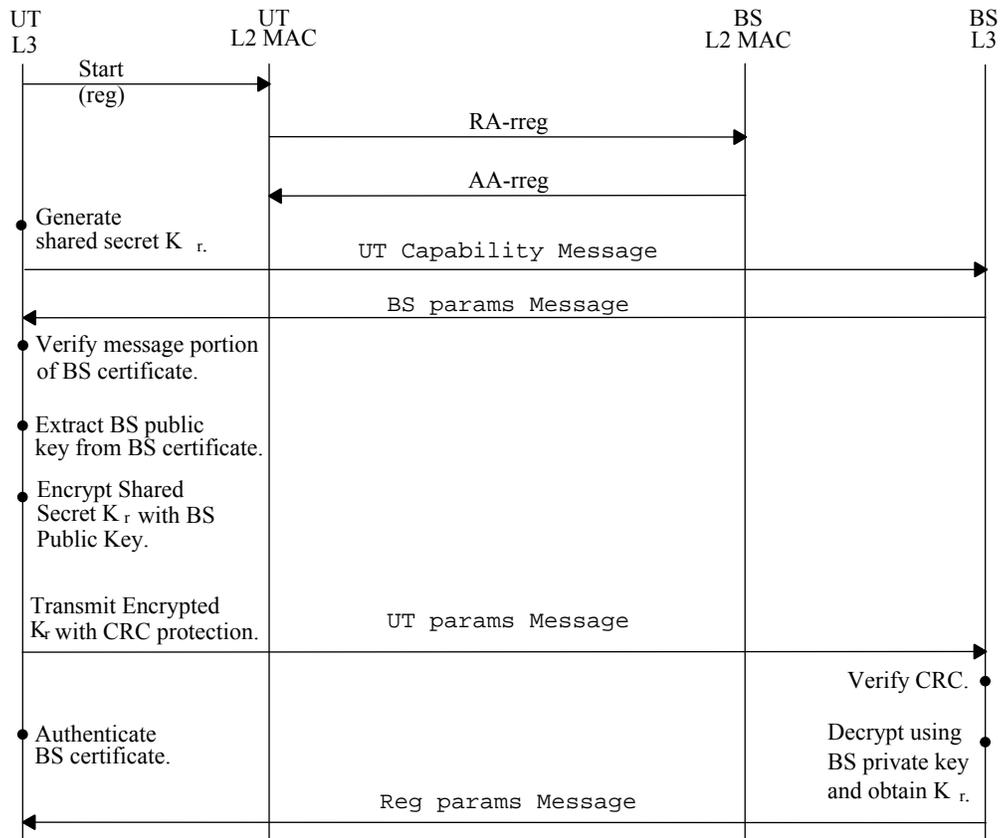


Figure 105: Baseline *i-HAP* Event Diagram

11.2.1 User Terminal Requests for Registration

The UT shall initiate the registration procedure. Before the UT registers, it shall perform the following steps (see Figure 106 - BCH Search and Registration):

- ◆ The UT shall listen to BCH frequency/time synchronization channels to gather frequency and timing information.
- ◆ The UT shall evaluate the BCH payload and acquire minimal parameters about the system.
- ◆ The UT shall send a configuration request (CR) burst to the BS of its choice, and shall expect a configuration message (CM) as response.
- ◆ Based on the contents of the CM burst, the UT shall initiate the registration stream by sending a Request Access (RA) message to the BS and as a response it shall expect an Access Assignment (AA) burst that opens a TCH stream to exchange in-band registration messages.
- ◆ The AA burst shall deliver the ten bit random identifier that was sent in the RA message to the UT (refer to clause 7.3.5.1.2.1, Common message format and field definitions).

The communication between the UT and the BS is as yet neither encrypted nor authenticated.

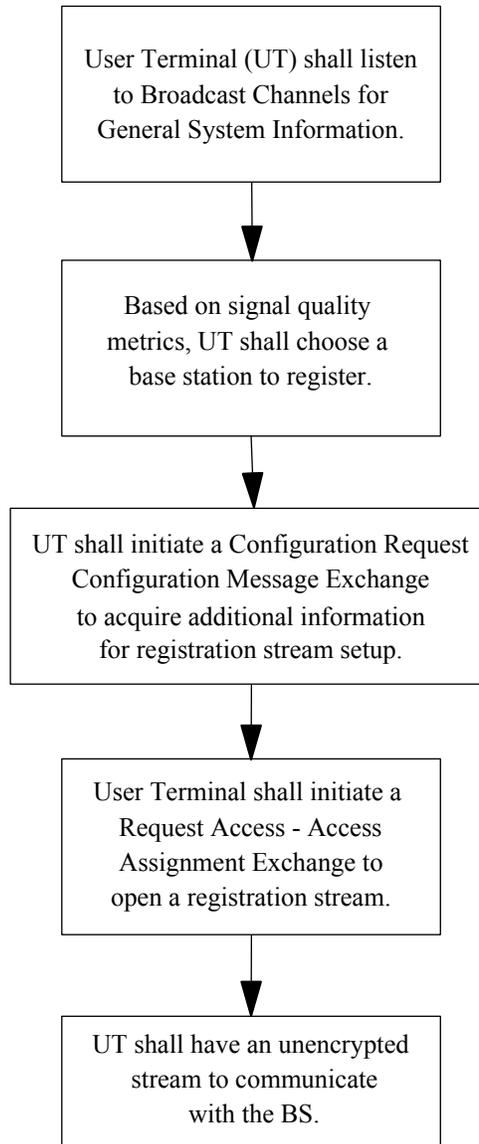


Figure 106: BCH Search and Registration

11.2.2 Base Station Transmits Its Digital Certificate

As shown in Figure 105 - Baseline *i-HAP* Event Diagram, the BS shall send its digital certificate using the BS Params AM message (see clause 9.5.1, Acknowledged Mode (AM) Messages) to the user terminal.

The BS certificate should contain a 297 bit message field followed by a 1024 bit signature field. The message field size shall depend on the public key encryption method¹¹ used for sending the shared secret K_r to the base station and the signature field shall depend on the public key encryption algorithm¹¹ used by the certificate authority to sign the message.

¹¹ Currently, elliptic curve cryptography on curve K-163, see clause 11B.1, Elliptic Curve Domain Parameters (ECDP).

The certificate message fields are defined below (see Figure 107 - *HC-SDMA* BS certificate message fields. Mandatory field values are identified in square brackets. Not to scale. for an illustration). For the fields that represent integer values, little endian format (least significant bit to most significant bit) shall be assumed:

1. Certificate Version: 2 bits. The value shall be 00. Other values are reserved.
2. Signature Generation Method: 5 bits. The first two bits shall define the hash or formatting method to be used and the last three bits shall indicate the public key method employed to generate the signature.

The hash or formatting method subfield shall enable the selection among the following options:

- 00: ISO/IEC 9796 signature scheme¹² (refer to [ISO 9796]).
- 10: Reserved.
- 01: Reserved.
- 11: Reserved.

The public key method subfield shall allow the selection among the following for signature generation:

- 000: RSA-1024.
- 100: RSA-1366.
- 010: RSA-1706.
- 110: RSA-2048.
- xx1: Reserved.

Option '000' shall be supported. The remaining options should be supported.

3. Certificate Authority Identifier: Shall be 8 bits that identify the CA. Note that for verification, the UT shall use the CA identity together with the contents of field (2). Furthermore, the all-zeros CA identifier (00000000) shall be reserved for the test CA.
4. Sharing: 4 bits. If the public key being certified by the certificate is not shared among multiple base stations, then this field shall be equal to '0000', otherwise it shall be equal to '1000'.
 - 0000: Not shared.
 - 1000: Shared.

The remaining choices shall be reserved for defining different levels of public key sharing in the future.

5. Subject: Owner of the certificate, associated with the public key being certified. This field shall consist of 48 bits identifying the BS. This field shall be in big endian format.
6. Validity Period: Two 16 bit fields that shall define the beginning and end of the certificate validity: 5 bits for day, 4 bits for month, and 7 bits for the year encoded in little endian format. For example, March 5, 2004 shall be encoded as: 10100-1100-0010000. The first 16 bit subfield of the 32 bit validity period shall define the beginning day and the last 16 bit subfield shall define the ending day. The certificate shall be valid from 12:00 AM GMT on the beginning day until 12:00 PM GMT on the end day.
7. Public Key Algorithm for BS Encryption: Shall define the public/private key algorithm at the BS.

¹² This method uses a reversible formatting scheme.

- 00: Elliptic Curve Cryptography: Curve K-163 as defined in clause 11B.1, Elliptic Curve Domain Parameters (ECDP). The BS public key field shall require 164 bits and contain the minimal representation of the BS public key.
- 10: Reserved.
- 01: Reserved.
- 11: Reserved.
8. Shared Secret Exchange Method: Shall be 2 bits.
- 00: UT shall determine a shared secret and encrypt it with the BS public key¹³.
- 10: Reserved.
- 01: Reserved.
- 11: Reserved.
9. Certificate Number: Shall be 4 bits to identify multiple certificates with overlapping validity periods held by the BS. This shall be an integer value encoded in little endian format.
10. i-SEC Bulk Encryption Algorithm Choices: Shall be 4 bits. This shall be a bitmap of bulk encryption algorithms supported by the BS. Every supported algorithm shall have a '1' in the relevant bit. The stream cipher method shall be supported and this field shall be set to [1,0,0,0].
- Bit 0: Stream Cipher (see clause 11.4.2, Symmetric Key Stream Cipher Algorithm).
- Bit 1: Reserved.
- Bit 2: Reserved.
- Bit 3: Reserved.
11. i-SEC Bulk Encryption Parameters: Shall be 20 bits. This field shall contain 4 subfields, each 5 bits. The k -th subfield shall contain parameters related to the k -th bulk encryption algorithm presented in Field (10) above. If a bulk encryption algorithm is not supported, then the corresponding subfield shall be set to all zeros. The information in each subfield shall depend on the definition and parameters of the corresponding bulk encryption algorithm. The first subfield shall provide the secret key size used in the stream cipher: if this field is all zeros, then the stream cipher is disabled; otherwise, the secret key size is the value represented in first subfield (in little endian format) plus four in bytes. The resulting value (in bytes) shall be defined as the maximum value of encryption key size (L_{key}) supported by the BS (refer to clause 11.4.2.3, Determination of Encryption Key from Shared Secret). For example, a BS that supports the stream cipher alone, with a maximum encryption secret key size of 56 bits (7 bytes) has:
- ◆ [1,0,0,0] as its Algorithm Choices field, and
 - ◆ [1,1,0,0,0;0,0,0,0,0;0,0,0,0,0;0,0,0,0,0] as its Bulk Encryption Parameters field. Note that each subfield represents an integer as a bitstring encoded with LSB first.
12. Physical Layer Integrity Check Method: Shall have 2 bits.
- 00: No integrity check at physical layer.
- 01: Reserved.
- 10: Reserved.
- 11: Reserved.
13. BS Public Key Minimal Representation: Shall be variable size field that depends on the BS public key encryption infrastructure selected in field (7). Determination of BS public key Q is

¹³ The length of the shared secret shall be twice the length of the size of the polynomial representation of the field in (7) -- i.e., current selection of $GF(2^{163})$ amounts to a shared secret transfer of 326 bits.

outlined in clause 11C.2.1, i-HAP BS Key Pair Generation. For the selection of elliptic curve cryptography on K-163, Q_M (the minimal representation¹⁴of the BS public key) shall be 164 bits.

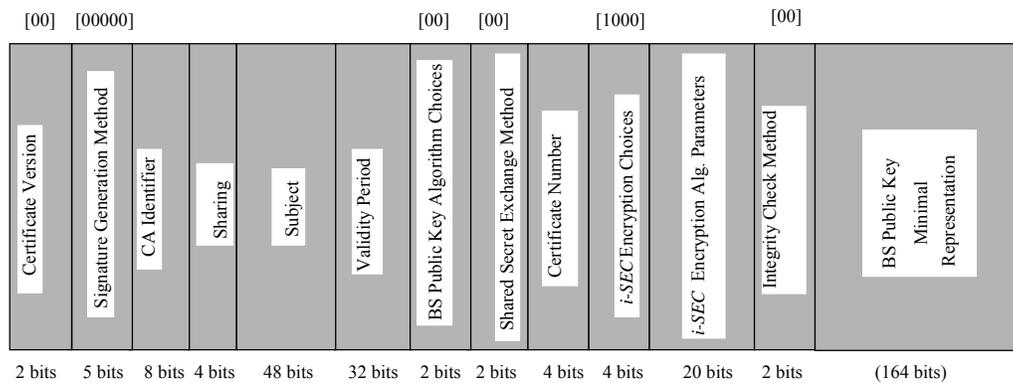


Figure 107: HC-SDMA BS certificate message fields. Mandatory field values are identified in square brackets. Not to scale.

11.2.3 User Terminal Analyzes Base Station Certificate Message

The UT shall examine the message part of the BS certificate and shall extract the BS public key:

- ◆ The UT application software should determine whether it has access to reliable time information¹⁵. If this is affirmative, then the UT shall check the certificate validity date to ensure that the certificate has not expired. If the certificate has expired, then the UT shall scan the BCH channels and attempt to register with a different BS. If the UT does not have access to time information:
 1. The UT may request verification of the AFN supplied by the BS from a centralized time server/certificate authority present in the HC-SDMA network (refer to clause 11.3.3, AFN Verification with a Centralized Time Server). The UT shall possess the capability to verify the signature of the centralized time server (CTS).
 2. The UT may shut down communications.
 3. The UT may ignore the certificate validity date (field (6) of the certificate message, refer to clause 11.2.2, Base Station Transmits Its Digital Certificate).

The behavior among the above three choices should be configurable. The default action for the user terminal should be to verify the AFN (for a UT which possesses the algorithmic capability) and to shut down communications if AFN verification is not possible. (Ignoring the validity period of the certificate lets the user terminal become vulnerable against certain man-in-the-middle attacks.)
- ◆ If the certificate has not expired, then the UT shall extract Q_M (the minimal representation of the BS public key) from the message part of the certificate (see clause 11.2.2, Base Station Transmits

¹⁴ Q_M shall be obtained from the base station public key Q using point compression (see clause 11B.2, Elliptic Curve Arithmetic). The first 163 bits of this field contains the x-coordinate of the BS public key Q , and it is in little endian format -- i.e., the first bit is the coefficient of the lowest order (constant) term in the polynomial representation of the x-coordinate of the public key.

¹⁵ This excludes the time information supplied by a BS, since an unauthorized BS may transmit a false AFN.

Its Digital Certificate). Q_M is 164 bits long.

- ◆ The UT shall obtain Q (the public key of the BS) from Q_M using point decompression (refer to clause 11B.2, Elliptic Curve Arithmetic). Q is a point on the elliptic curve (defined in clause 11B.1, Elliptic Curve Domain Parameters (ECDP) and it is represented by two 163 bit components Q_x and Q_y .
- ◆ The UT shall verify that Q is a point on the elliptic curve, i.e., it satisfies $(Q_y)^2 + Q_x \cdot Q_y = (Q_x)^3 + (Q_x)^2 + 1$ in $GF(2^{163})$.
- ◆ Furthermore, the UT shall verify that Q is not equal to either the point at infinity represented by $(0,0)$, or $(1,0)$, or $(0,1)$.
- ◆ If the point Q does not pass all of verification tests outlined in the preceding list, then the UT shall scan BCH channels and attempt to register with a different BS.
- ◆ User terminals should cache BS public keys and their expiration dates from prior registrations. Such user terminals may compare the acquired BS public key with cached public keys and eliminate key verification and certificate signature verification (see BS certificate field (8) in clause 11.2.6, User Terminal Verifies BS Certificate Signature). Field (4) of the BS certificate provides additional information about the sharing status of the BS keys.

11.2.4 User Terminal Encrypted Shared Secret Transmission

After the UT derives the public key Q of the BS from its minimal representation Q_M in the BS certificate and verifies that Q is a valid point on the elliptic curve, the UT shall generate a shared secret and encrypt it with the BS public key prior to transmission. The BS public key encryption routine is defined in field (7) of the BS certificate. and is an elliptic curve encryption using curve K-163 (refer to clause 11B.1, Elliptic Curve Domain Parameters (ECDP), which sets the shared secret size to 326 bits (refer to clause 11.2.2, Base Station Transmits Its Digital Certificate). Note that BS certificate signature is not verified by the UT yet. Shared secret transmission from a regular user terminal¹⁶ shall proceed as follows:

- ◆ The UT shall generate 326 random bits as the shared secret K_r . The UT shall store this value for bulk encryption using *i-SEC* and for UT authentication using *i-TAP*.
- ◆ The UT shall encrypt K_r using the BS public key Q (refer to clause 11C.2.2, *i-HAP* UT Encryption to obtain N_r (490 bits long). Note that Step 5 in the UT encryption process imposes constraints on N_r .
- ◆ The UT shall append a 16 bit cyclic redundancy check (CRC_{16}) to N_r in order to obtain:
- ◆ $O_r = [N_r, CRC_{16}(N_r)]$.

CRC_{16} is defined in clause 11C.3, Cyclic Redundancy Code for Shared Secret Exchange. O_r is 506 bits long.

- ◆ The UT shall transmit O_r to the BS (see Figure 112 - Encryption in *HC-SDMA* Secure Communication protocol, *i-SEC*) via a UT Params AM message (see Figure 105 - Baseline *i-HAP* Event Diagram).

¹⁶ As opposed to a test terminal defined as a UT with extended test capability. In this context, the encryption shall be turned off.

Because of parameter checks involved in UT encryption (see step 5 in clause 11C.2.2, *i-HAP* UT Encryption), the first 164 bits of N_r (therefore of O_r) cannot be all zeros. A test user terminal may use this property and the linearity of the CRC to disable *i-SEC* bulk encryption by transmitting O_r as the all zeros 506 bit vector.

11.2.5 Base Station Decryption of Encrypted Shared Secret

The base station shall check the CRC in the message O_r , and if the CRC check succeeds, shall decrypt N_r using its private key, and recover the shared secret K_r :

- ◆ The BS shall compare O_r to the all zeros 506 bit vector first. If O_r is all zeros, then the BS shall disable *i-SEC* bulk encryption for this UT. Otherwise, it shall continue with the following.
- ◆ The BS shall partition O_r into two parts: P_r which consists of the first 490 bits of O_r , and Q_r which consists of the last 16 bits of O_r .
- ◆ The BS shall check the CRC in O_r : it determines whether $CRC_{16}(P_r)$ is equal to Q_r . CRC_{16} is defined in clause 11C.3 Cyclic Redundancy Code for Shared Secret Exchange.
- ◆ If the CRC check fails, *i-HAP* shall be terminated and the registration stream shall be closed.
- ◆ If the CRC check is affirmative, then the BS shall check for the validity of the contents of N_r (refer to step 5 in clause 11C.2.2, *i-HAP* UT Encryption for valid values of N_r). If N_r is found to be invalid, then *i-HAP* shall be terminated and the registration stream shall be closed.
- ◆ If the BS finds the contents of N_r to be valid, then it shall decrypt N_r (490 bits long) using its private key (refer to clause 11C.2.3, *i-HAP* BS Decryption) to obtain the shared secret K_r (326 bits long). The shared secret K_r shall be used for bulk encryption and UT authentication.

11.2.6 User Terminal Verifies BS Certificate Signature

Before using the shared secret, the UT shall verify the signature of the certificate authority presented in the BS certificate. This step may be performed while the base station is performing public key decryption to recover the shared secret determined by the UT. The mathematical operations necessary for BS certificate verification are presented in clause 11C.1.3, *i-HAP* Certificate Signature Verification.

If the user terminal determines the signature of the BS certificate to be invalid, then it shall terminate communications with the BS and shall start the BCH search procedure to acquire a different BS for registration.

If the user terminal determines that the BS public key presented to it in the BS certificate is identical to a non-expired key stored in its public key cache and the base station identity is the same as in the certificate, then the UT should accept the certificate without signature verification. The public key cache shall contain public keys from verified but non-expired digital certificates (as determined certificate's validity period). The public key cache shall not contain public keys from expired digital certificates (as determined certificate's validity period).

11.3 HC-SDMA Terminal Authentication Protocol, *i-TAP*

The services within *i-TAP* shall be performed following the initial access to the selected base station and in conjunction with the *i-HAP* protocol message exchange for efficient use of the air interface. The *i-TAP* protocol proceeds as follows:

1. UT shall create¹⁷ the UT authenticator string (A_s).
2. UT shall scramble the UT certificate and UT specific parameters.
3. UT shall transmit the scrambled UT certificate, scrambled UT specific parameters and the authenticator string to the BS.
4. BS shall verify UT certificate authenticity and determine the UT public key.
5. BS shall verify the UT authenticator string and shall either admit the UT to the network or reject network access.

¹⁷ The UT shall compute the SHA-1 hash of the authenticator message A_m first, and shall encrypt the hash value using its private key (using the elliptic curve K-163 from [FIPS]) to obtain the UT authenticator string A_s . UT Authenticator message A_m is defined in 11.3.2, UT Authenticator Message and String Generation.

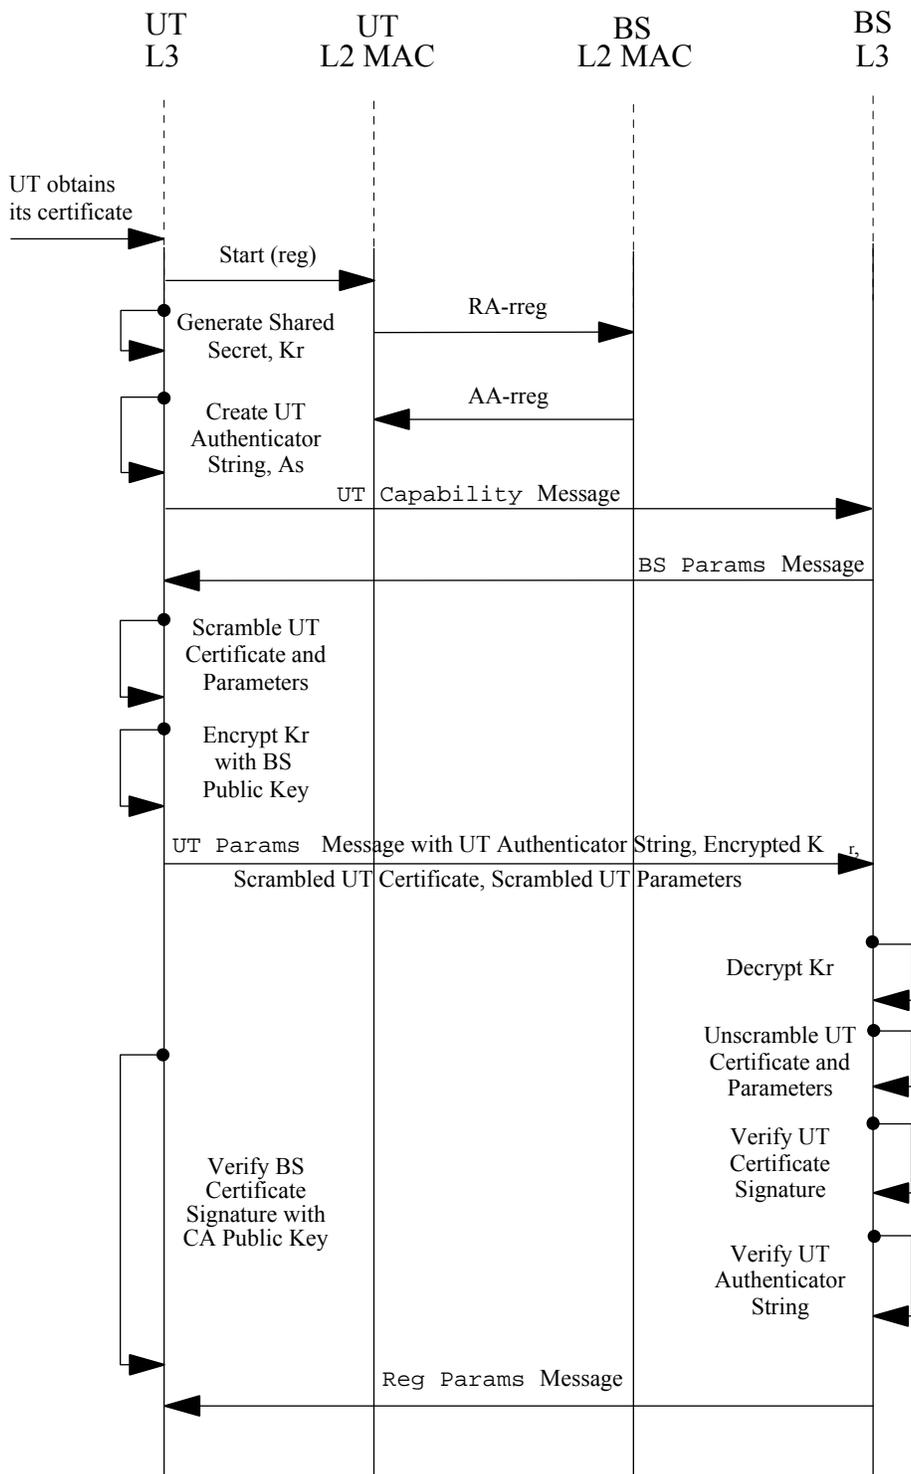


Figure 108: Baseline *i-TAP* Event Diagram

11.3.1 UT Certificates

The *i-TAP* protocol shall conform to Figure 108 - Baseline *i-TAP* Event Diagram. The UT shall send its digital certificate using the UT Params AM message (see clause 9.5.1, Acknowledged Mode (AM) Messages) to the BS.

UTs shall use three types of certificates to identify themselves to Base Stations (BSs). These are:

1. *Identity Certificate* - This certificate is generated by an *HC-SDMA* Certificate Authority (CA) and factory seeded in the UT by the UT manufacturer. It ties the UT's hardware serial number to an elliptic curve public/private key pair using the elliptic curve K-163 defined in clause 11B.1, Elliptic Curve Domain Parameters (ECDP). A UT public key is represented by 164 bits using its minimal representation. This certificate is used in the process of obtaining a Service Authorization Certificate or Network Access Certificate over the air interface (at the BS) from a management entity in the operator's or ISP's network. This certificate shall only be used if the UT does not have a valid Service Authorization or Network Access Certificate. While the Identity Certificate shall be factory seeded in the UT, it may never be used if the Service Authorization or Network Access Certificates are already obtained by other means.
2. *Service Authorization Certificate* - This certificate is optional. This certificate should be used in deployments where Network Access Certificate validity is short and restricted to a duration approximating a typical single network access. When employed, the Service Authorization Certificate shall be acquired by the UT at service authorization time over the air interface. When employed, the Service Authorization Certificate shall be used each time the UT powers up and shall be used to obtain the Network Access Certificate from the local management entity. It ties a service identifier (i.e., IMSI) assigned to the UT with a service type and an elliptic curve public/private key pair using the elliptic curve K-163 defined in clause 11B.1, Elliptic Curve Domain Parameters (ECDP). A UT public key is represented by 164 bits using its minimal representation.
3. *Network Access Certificate* - This certificate is an extension of the Service Authorization Certificate. It shall be used each time a UT establishes a new communications link with a BS except when this link is established to obtain a Service Authorization or Network Access Certificate (in which case either an identity or a Service Authorization Certificate shall be presented). The Network Access Certificate is the nominal credential provided by the UT to register with a base station. It may either persist across network access requests¹⁸, or it may be refreshed at the beginning of each network access or on a periodic basis by means of a Service Authorization Certificate.

11.3.1.1 Certificate Usage Rules

The UT shall present a certificate to the BS in the UT Params message for each EUD link level session (e.g., PPP session) the UT is facilitating. If multiple link level sessions are permitted per UT, then the UT shall have a distinct corresponding Network Access Certificate for each of them. The certificates shall be signed by the CA and shall be unique to each link level EUD session. The UT identification and session identification information certified in the Identity Certificate or Network Access Certificate may be used for authenticated handover switching by Packet Services Switches (PSSs).

Different service providers (i.e., network operators and/or ISPs) may define different service classes and issue different types of Network Access Certificates. UT Identity shall be certified by assigning the

¹⁸ A series of consecutive *HC-SDMA* sessions with one or more base stations.

UT an Identity Certificate at manufacture. This Identity Certificate can be used to securely assign the user terminal a Service Authorization and/or Network Access Certificate at the time it is being authorized for a service provider's service; either directly by the service provider or its representative, or by the end-user as instructed by the service provider.

For deployments in which user terminals support a single EUD session and are not used for distinct classes of service, the Identity Certificate may be used in place of a Network Access Certificate. In such deployments, the BS shall interpret the UT Identity Certificates as UT Network Access Certificates.

The Service Authorization Certificate should be used only in *HC-SDMA* deployments where an automatic certificate download over *HC-SDMA* is possible (See clause 11D.4 i-TAP Automatic Certificate Retrieval) and where Network Access Certificates with a short (for example, a single day) certificate validity period are used. In such deployments, the UT shall present the Service Authorization Certificate to the BS when the UT does not possess a valid Network Access Certificate. In such a scenario, access of the UT shall be restricted to communication with a management entity in the service provider's network. The BS/PDSN shall ensure that the UT communicates only with that entity (e.g., switch the UT PPP session to an isolated network hosting only the management entity).

The UT shall present an authenticator string (A_s) to the BS (see clause 11.3.2 UT Authenticator Message and String Generation). The authenticator string shall be transferred in the UT Params message. Depending on which certificate type the UT uses for authentication, the UT shall choose the appropriate private key to produce the authenticator string (i.e., the identity private key, the service authorization private key, or the network access private key). The BS shall use the corresponding public key, which is contained in the certificate to verify that the UT is in possession of both the certificate and the private key.

In summary, UT hardware authentication is accomplished in a single message transmission (once the UT is authorized for service and obtains a Network Access Certificate by one of several means as explained earlier) -- i.e., using the *HC-SDMA* UT Params message. The Network Access Certificate based authentication shall be used both when the UT initially enters the network for establishing an EUD link-level session and when the UT session is in handover. The PSS shall facilitate handovers based on the UT identity and session reference.

The BS shall verify: (1) the signature of the UT certificate (any of the 3 certificate types); (2) that the UT certificate has not been revoked by checking a local Certificate Revocation List (CRL); and (3) the authenticator string (A_s) with the public key presented in the received UT certificate. If the verification fails, the BS shall revoke the UT's registration. The management system that monitors UT network access should distribute CRLs to revoke stolen private keys or stolen UTs.

The UT should obtain its certificate and its private and public keys either through EUD download or automatically over the initial PPP/IP/TCP management connection.

In the description above, the Identity Certificate is used to prove authenticity of a particular piece of end-user equipment, the UT. The Identity Certificate may alternatively be used to prove authenticity of a different piece of end-user equipment, a dongle or a SIM (subscriber identity module). In this case, a UT may effectively assume the identity of the SIM.

11.3.1.2 UT Certificate Scheme

The *HC-SDMA* UT certificate signature scheme consists of using the RSA algorithm with 1024 bit modulus and the ISO/IEC 9796 signature standard to produce a 1024 bit signature for each of the three certificate types. The certificate signature scheme is defined in clause 11D.1, *i-TAP* UT Certificate

Signatures. Only the signature shall be transmitted in the UT Params message. The message is recoverable from the signature and the signature does not utilize a hash function.

The signature shall be sent scrambled using a linear feedback shift register initialized with a 39 bit scrambling seed (see clause 11.3.4, UT Parameter Scrambling). The seed shall be equal to the 39 most significant bits of the 326 bit *i-HAP* shared secret K_r , presented to the BS in encrypted form with the UT Params message (see clause 11.2.4, User Terminal Encrypted Shared Secret Transmission).

11.3.1.2.1 UT Identity Certificate

The UT Identity Certificate contains a 284 bits message field followed by a 1024 bit signature field. Only the signature shall be transmitted from the UT to the BS in the UT Params message. The CA identifier shall be included in the UT Params message (refer to clause 11.3.5, UT Params Message Transmission). This enables the BS to select the CA public key corresponding to the CA that issued the UT certificate.

The UT Identity Certificate message shall be composed as follows (fields that define integer values shall use the little endian format):

1. Certificate Version Number: 2 bits. The version shall be 00. Other values are reserved.
2. Certificate Authority Identifier: Shall be 8 bits that identify the CA. The CA identifier shall be included in the clear in the UT Params message as the BS needs to select the CA public key corresponding to the CA that issued the UT certificate.
3. Certificate Type: Shall be 2 bits that specify the certificate type. Three types are defined ('01' for Identity, '10' for Service Authorization, '11' for Network Access).
4. Serial Number of Certificate: Shall be 28 bits assigned by the CA that within a certificate validity period uniquely identify any certificate issued by that CA. This shall be in little-endian format.
5. Validity Period: 32 bits. Shall be two 16 bit fields that define the beginning and end of certificate validity: 5 bits for day, 4 bits for month, and 7 bits for year in little-endian format. For example, March 5, 2004 is encoded as: (LSB) 10100-1100-0010000. The year encoded shall be the actual year minus 2000. The beginning day field shall come first, followed by the end day field. The certificate shall be valid from 12:00 AM GMT on the beginning day until 12:00 PM GMT on the end day.
6. UT Serial Number: Shall be 48 bits that identify the UT. This shall be formatted as the IEEE Ethernet MAC address of the UT.
7. UT Identity Public Key: Shall be 164 bits for Q_M ; the minimal representation of the UT identity public key Q . This shall be based on the elliptic curve K-163 (refer to clause 11B.1, Elliptic Curve Domain Parameters (ECDP)). Q_M shall be obtained from the UT identity public key using point compression¹⁹ (see clause 11B.2, Elliptic Curve Arithmetic). Determination of Q is outlined in clause 11D.2.1, *i-TAP* UT Key Pair Generation.

¹⁹ The first 163 bits of this field contains the x-coordinate of the UT public key Q , and it is in little endian format -- i.e., the first bit is the coefficient of the lowest order (constant) term in the polynomial representation of the x-coordinate of public key.

11.3.1.2.2 UT Service Authorization Certificate

The UT Service Authorization Certificate shall contain a 472 bit message field followed by a 1024 bit signature field. Only the signature shall be transmitted from the UT to the BS in the UT Params message. The UT Service Authorization Certificate fields that define integer values shall use the little endian format.

The UT Service Authorization Certificate message shall be composed as follows (the first six fields are identical to the UT Identity Certificate message):

1. Certificate Version Number: Shall be 2 bits. The version shall be 00. Other values are reserved.
2. Certificate Authority Identifier: Shall be 8 bits that identify the CA. The CA identifier shall be included in the clear in the UT Params message as the BS needs to select the CA public key corresponding to the CA that issued the UT certificate.
3. Certificate Type: Shall be 2 bits that specify the certificate type. Three types are defined ('01' for Identity, '10' for Service Authorization, '11' for Network Access).
4. Serial Number of Certificate: Shall be 28 bits assigned by the CA that within a certificate validity period uniquely identify any certificate issued by that CA. This shall be in little-endian format.
5. Validity Period: 32 bits. Shall be two 16 bit fields that define the beginning and end of certificate validity: 5 bits for day, 4 bits for month, and 7 bits for year in little-endian format. For example, March 5, 2004 is encoded as: (LSB) 10100-1100-0010000. The year encoded shall be the actual year minus 2000. The beginning day field shall come first, followed by the end day field. The certificate shall be valid from 12:00 AM GMT on the beginning day until 12:00 PM GMT on the end day.
6. UT Serial Number: Shall be 48 bits that identify the UT that owns the certificate. This shall be formatted as the IEEE Ethernet MAC address of the UT.
7. UT IMSI: Shall be 60 bits. International Mobile Station Identifier shall be assigned to the UT by the service authority (e.g., the network operator).
8. UT Service Category Authorization: Shall be 128 bits that shall be encoded as defined by the service authority and each bit set corresponds to a service category authorized.
9. UT Service Authorization Public Key: Shall be 164 bits for Q_M ; the minimal representation of the UT identity public key Q . This shall be based on the elliptic curve K-163 (refer to clause 11B.1, Elliptic Curve Domain Parameters (ECDP)). Q_M shall be obtained from the UT identity public key using point compression (see clause 11B.2, Elliptic Curve Arithmetic). Determination of Q is outlined in clause 11D.2.1, *i-TAP* UT Key Pair Generation.

In deployments where only the Identity Certificate is used and IMSI is not available, the least significant 48 bits of the IMSI shall be set to the UT Serial Number and the remaining 12 bits shall be set to 0.

11.3.1.2.3 UT Network Access Certificate

The UT Network Access Certificate shall contain a 488 bits message field followed by a 1024 bit signature field. Only the signature shall be transmitted from the UT to the BS in the UT Params message. The UT Network Access Certificate fields that define integer values shall use the little endian format.

The UT Network Access Certificate message shall be composed as follows (the first eight fields are identical to the UT Service Authorization Certificate message):

1. Certificate Version Number: Shall be 2 bits. The version shall be 00. Other values are reserved.
2. Certificate Authority Identifier: Shall be 8 bits that identify the CA. The CA identifier shall be included in the clear in the UT Params message as the BS needs to select the CA public key corresponding to the CA that issued the UT certificate.
3. Certificate Type: Shall be 2 bits that specify the certificate type. Three types are defined ('01' for Identity, '10' for Service Authorization, '11' for Network Access).
4. Serial Number of Certificate: Shall be 28 bits assigned by the CA that within a certificate validity period uniquely identify any certificate issued by that CA. This shall be in little-endian format.
5. Validity Period: 32 bits. Shall be two 16 bit fields that define the beginning and end of certificate validity: 5 bits for day, 4 bits for month, and 7 bits for year in little-endian format. For example, March 5, 2004 is encoded as: (LSB) 10100-1100-0010000. The year encoded shall be the actual year minus 2000. The beginning day field shall come first, followed by the end day field. The certificate shall be valid from 12:00 AM GMT on the beginning day until 12:00 PM GMT on the end day.
6. UT Serial Number: Shall be 48 bits that identify the UT that owns the certificate. This shall be formatted as the IEEE Ethernet MAC address of the UT.
7. UT IMSI: Shall be 60 bits. International Mobile Station Identifier shall be assigned to the UT by the service authority (e.g., the network operator).
8. UT Service Category Authorization: Shall be 128 bits that shall be encoded as defined by the service authority and each bit set corresponds to a service category authorized.
9. UT Session Reference: Shall be 16 bits that shall uniquely identify a link level session (e.g., PPP) within the UT.
10. UT Network Access Public Key: Shall be 164 bits for Q_M ; the minimal representation of the UT identity public key Q . This shall be based on the elliptic curve K-163 (refer to clause 11B.1, Elliptic Curve Domain Parameters (ECDP)). Q_M shall be obtained from the UT identity public key using point compression (see clause 11B.2, Elliptic Curve Arithmetic). Determination of Q is outlined in clause 11D.2.1, *i-TAP* UT Key Pair Generation.

UT session reference shall enable authorization of multiple simultaneous link level sessions (e.g., PPP) per UT. For the k -th session, the UT shall present a certificate that authorizes the UT to establish a session with reference k . UTs that are authorized to establish a single session only, shall set the session reference to 0.

11.3.1.3 Operational Recommendations

Service Authorization Certificates are possibly necessary in different *HC-SDMA* deployments depending on the size of deployment and the number of revoked Network Access Certificates. Network Access Certificates may be revoked due to changes in the service agreement, or if there is a reason to believe that they have been compromised.

As the number of updates to the Network Access Certificate Revocation List (CRL) and the length of that list grows to be unmanageable by the BSs, the validity period of the Network Access Certificate may be reduced to one month, two weeks, one week, three days, or a single day. A shorter certificate validity period makes it possible to drop expired certificates from the CRL and shortens the CRL to the

list of revoked certificates that are not expired. Once the Network Access Certificate becomes valid for network access with handovers over a very short period of time (e.g., one day), then the Service Authorization Certificate should be introduced since it contains sufficient service authorization information to minimize the number of service server look-up requests.

The Network Access Certificate validity period should not need to be uniform across all Network Access Certificates. Short validity and long validity certificates can be combined to enable applications that require mobility over an extended period of time.

11.3.2 UT Authenticator Message and String Generation

Before sending its digital certificate in the UT Params message, the UT shall generate a UT authenticator message A_m and shall encrypt²⁰ the SHA-1 hash of the authenticator message to obtain the UT authenticator string A_s . The UT authenticator message A_m shall be 438 bits long, and it shall consist of the following fields:

1. AFN: Shall be the 48 bit Absolute Frame Number (AFN). See below for more details.
2. UT Serial Number: Shall be 48 bits that shall identify the UT. This shall be formatted as the IEEE Ethernet MAC address of the UT.
3. UT Session Reference: Shall be 16 bits that uniquely identifies the PPP session within the UT. A UT operating without a Network Access Certificate shall set the session reference to zero.
4. K_r : Shall be 326 bits UT/BS shared secret²¹.

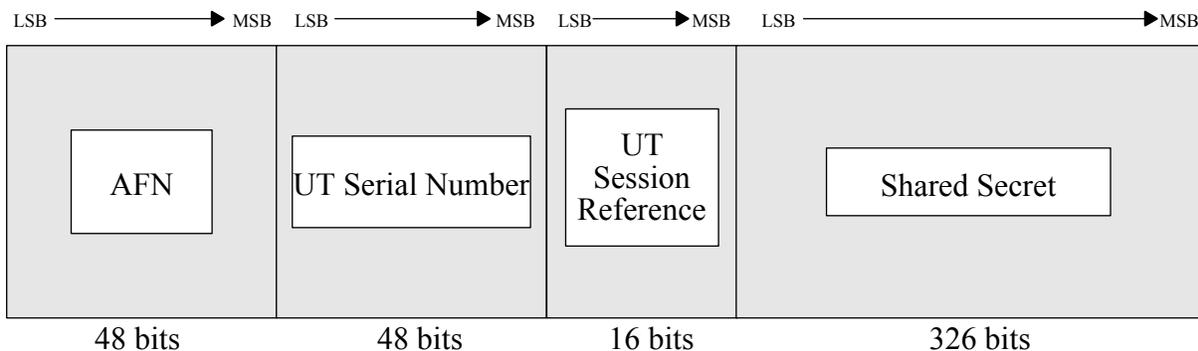


Figure 109: UT Authenticator Message Formatting (not to scale)

The default value of the AFN field in the authenticator message (refer to Figure 109 - UT Authenticator Message Formatting (not to scale)) shall be the AFN on which the registration stream started -- i.e., the

²⁰ Encryption shall be performed with the UT private key corresponding to the certificate in use. The UT public key algorithm for *i-TAP* shall be based on elliptic curve cryptography in the finite field $GF(2^{163})$ which offers the equivalent in security to a 96-bit symmetric key encryption algorithm [FFW].

²¹ The length of the shared secret may vary depending on the BS public key encryption method employed. Currently, the shared secret shall be 326 bits long.

frame in which the RA/AA exchange²² took place. The UT may use its estimate of starting time of the registration stream as the authenticator message AFN in cases where it may choose to pre-encrypt the authenticator message hash to produce the authenticator string. The little endian representation of the AFN shall be used in this context.

The AFN extracted from the BS transmissions is possibly the only source of time for a UT. A UT with BS supplied time information may choose to verify the AFN before it utilizes the AFN derived time to check BS certificate expiration. The *i-TAP* protocol allows a UT to request (from the BS) the verification of the AFN implied in the UT authenticator string (A_s) by a CA (clause 11.3.3, AFN Verification with a Centralized Time Server). This shall be done by including the 8 bit CA identifier, that the UT wishes to verify the AFN, in the UT Params message.

The UT shall use its private key to sign (encrypt) the 160 bit SHA-1 hash of the authenticator message (refer to clause 11D.2.2, *i-TAP* UT Private Key Encryption) to obtain the 326 bit UT authenticator string A_s . The corresponding UT public key shall be included in the UT certificate, and the UT shall transmit its certificate to the BS after scrambling (refer to clause 11.3.4, UT Parameter Scrambling). The UT Params message shall also include the scrambled AFN used by the UT for authentication and time verification purposes and an optional 8-bit AFN verification CA identifier.

11.3.3 AFN Verification with a Centralized Time Server

An HC-SDMA deployment may support AFN verification. A UT may request AFN verification. The BS may support AFN verification in which case AFN verification shall be accomplished by communication with a centralized time server. The selection of the centralized server shall be based on the UT request (refer to the last field described in clause 11.3.4, UT Parameter Scrambling). If requesting AFN verification, the UT shall possess the public key of the centralized time server²³.

If AFN verification is supported, then the BS shall maintain a cryptographically secure link with the centralized time server.

The AFN verification shall conform to Figure 110 - AFN Verification Procedure:

1. The BS shall check the validity of the UT identity by first verifying the authenticator string A_s . If the check fails, then the BS shall not grant AFN verification service.
2. The BS shall forward the UT authenticator message A_m , authenticator string A_s and the digital certificate of the UT to the centralized time server upon verification of the UT identity by checking the validity of the UT certificate and verifying the authenticator string. The identity of the centralized time server shall be determined²⁴ by the UT.
3. The centralized time server should check the validity of the UT identity by verifying the authenticator string A_s (refer to clause 11D.2.3 *i-TAP* BS Verification of UT Signature), after

²² The UT shall recover the 10 LSBs of the AFN from the CM burst. The BS Params message shall supply the 38 MSBs of the AFN of the frame belonging to the RA-rreg/AA-rreg exchange that started the registration stream. This shall allow the UT to have complete AFN synchronization with the BS.

²³ The time server may be the same entity as the certificate authority that generates BS certificates. The time server shall perform RSA public key encryption operations exactly as a CA that signs BS certificates. In addition, the centralized time server may need to verify the UT authenticator string (an elliptic curve operation, refer to clause 11D.1.3, *i-TAP* UT Certificate Signature Verification).

²⁴ The UT shall choose a centralized time server whose RSA public key is known to the UT.

- verifying²⁵ the signature on the UT digital certificate (refer to clause 11D.1.3, *i-TAP* UT Certificate Signature Verification) and shall not grant time verification service if this check fails.
4. Optionally, the centralized time server should check the validity of the UT certificate from a Certificate Revocation List. The centralized time server shall not grant time verification service if this check fails.
 5. The centralized time server shall verify that the AFN presented in the UT authenticator message A_m is within T_{allow} seconds of the time perceived at the centralized time server. The centralized time server shall not grant time verification service if this check fails. T_{allow} is a system variable that is set to 60 by default.
 6. The centralized time server shall append the value of the 48 bit AFN perceived by the centralized time server AFN_{TTS} (in little endian format) to the 326 bit authenticator string A_s from the UT to obtain the 374 bit centralized time server message M_{TTS} : $M_{\text{TTS}} = [A_s, AFN_{\text{TTS}}]$.
 7. The centralized server shall generate a 1024 bit digital signature A_{ts} calculated by the RSA algorithm using its RSA private key and the 374 bit centralized time server message M_{TTS} . The digital certificate generation process is identical to that required in *i-HAP*, refer to clause 11C.1.2, *i-HAP* Certificate Signature Generation for details. Generation of RSA parameters for the centralized time server shall be identical to generation of CA parameters for *i-HAP*, refer to clause 11C.1.1, *i-HAP* Certificate Authority (CA) Key Pair Generation. The centralized time server message M_{TTS} shall be fully recoverable from the 1024 bit signature A_{ts} because of the formatting method employed, refer to [ISO 9796] for details.
 8. The centralized time server shall transmit a message to the BS to confirm the authentication of the UT (if the optional steps 3 and 4 are implemented) and the 1024 bit RSA digital signature A_{ts} that verifies the centralized time server's value for AFN.
 9. The BS shall communicate the centralized time server response to the UT in the L3 AFN Verify AM message. The L3 AFN Verify AM message (see clause 9.5.1.3.13, AFN Verify Message) includes the 8 bit Centralized Time Server Identifier. The UT should abort its session with the BS if the UT does not receive the centralized time server response within T_{abort} seconds where T_{abort} is a UT configuration parameter.
 10. Upon timely receipt of the centralized time server response, the UT shall obtain the digital signature A_{ts} from the L3 AFN Verify AM message, shall extract A_s and AFN_{TTS} from A_{ts} , and verify the signature A_{ts} using the RSA public key of the centralized time server. Furthermore, the UT shall compare its own value for the authenticator string with that extracted from A_{ts} . The RSA signature verification operation shall be identical to that performed by a UT during *i-HAP* (refer to clause 11C.1.3, *i-HAP* Certificate Signature Verification). For extracting A_s and AFN_{TTS} from A_{ts} refer to [ISO 9796]. If all the checks are successful, then the UT shall accept the AFN_{TTS} and make its certificate timeliness decisions accordingly.

²⁵ Although this appears redundant, it is included as a precaution to prevent potential cooperative attacks.

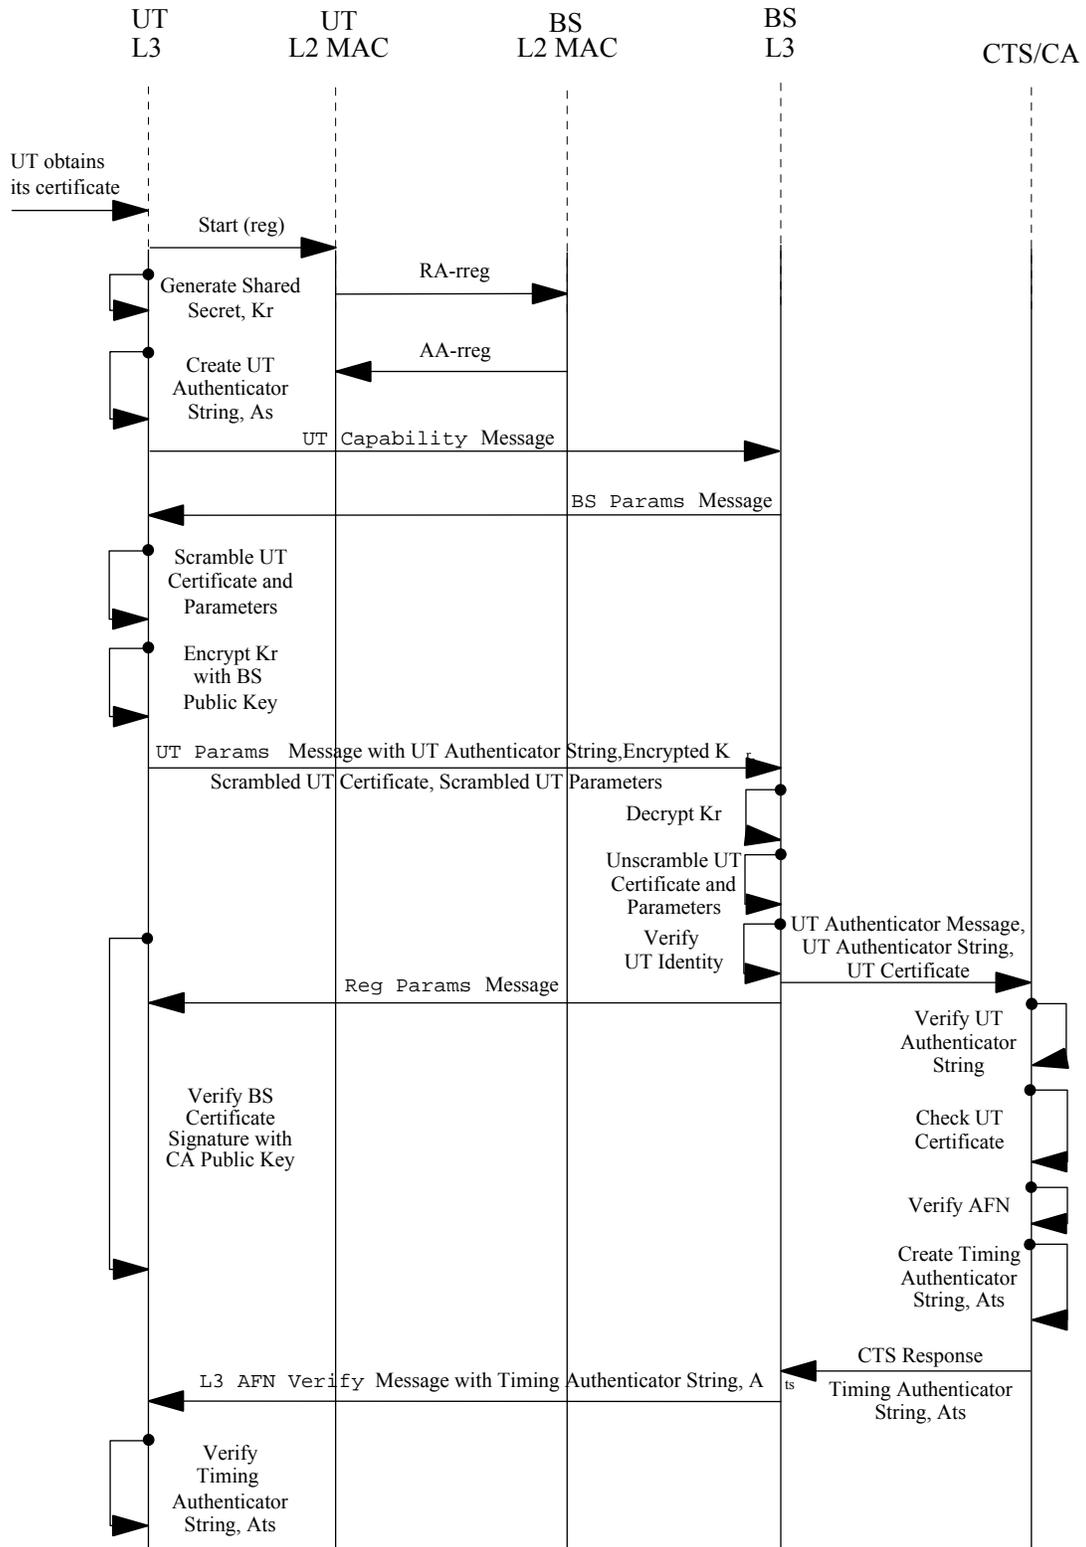


Figure 110: AFN Verification Procedure

11.3.4 UT Parameter Scrambling

The digital certificate and certain registration parameters and preferences of the UT shall be grouped together and scrambled with a linear feedback shift register (LFSR).

The UT shall assemble the following 1088 bit message for scrambling (refer to Figure 111 - i-TAP UT parameters before scrambling):

1. The 48 bit AFN used in the UT authenticator message.
2. The 8 bits plaintext CA identifier of the certificate authority that created the UT certificate.
3. 1024 bit UT certificate signature.
4. The 8 bits plaintext identifier of the centralized time server that the UT wants to verify the AFN. All zeros identifier shall be used if the UT does not require AFN verification.

This message shall be scrambled as specified in clause 11D.3 i-TAP Scrambling Algorithm using a 39 bit LFSR. The LFSR shall be initialized with the 39 most significant bits of the shared secret K_r , and shall produce 1088 scrambling bits.

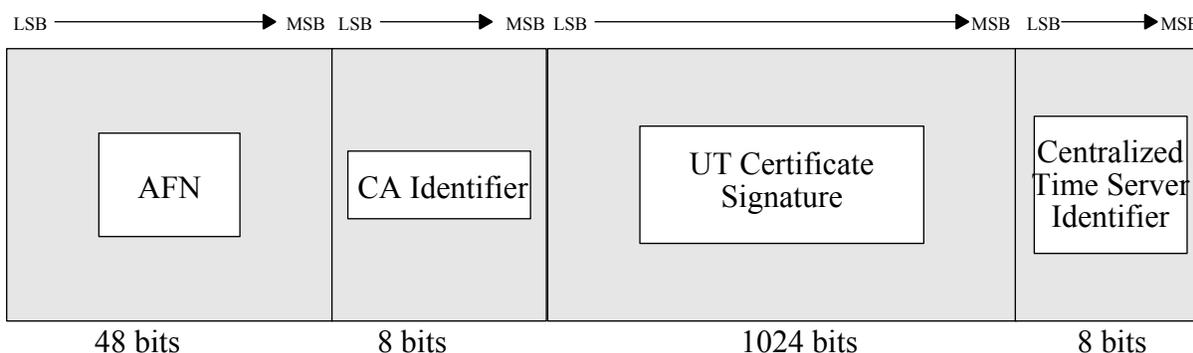


Figure 111: i-TAP UT parameters before scrambling

11.3.5 UT Params Message Transmission

In the UT Params message the BS shall receive from the UT three terms of importance to the i-TAP protocol. These are:

1. The message O_r shall be composed of a 506 bit long string $O_r = [N_r, CRC_{16}(N_r)]$ where N_r shall be the encrypted UT/BS shared secret K_r (refer to clause 11.2.4, User Terminal Encrypted Shared Secret Transmission).
2. The 326 bit authenticator string shall be derived from the UT authenticator message (see clause 11.3.2, UT Authenticator Message and String Generation).
3. Scrambled UT parameters (refer to clause 11.3.4, UT Parameter Scrambling). This shall be 1088 bits long.

11.3.6 BS Analysis of UT Certificate and Authenticator String

Following the successful recovery of the shared secret (refer to clause 11.2.5, Base Station Decryption of Encrypted Shared Secret), the BS shall process the received UT certificate signature S' as follows:

1. The BS shall descramble the relevant part of the UT Params message as specified in clause 11D.3 i-TAP Scrambling Algorithm.
2. The BS shall obtain the public key (n, e) for the CA from its memory corresponding to the CA identifier field obtained from the UT Params message.
3. BS shall decrypt the received CA signature S' by using the CA public key e : $V = (S')^e \bmod n$. Due to the specific value of e ($e = 2^{16} + 1$), this operation requires only sixteen modular squarings and one modular multiplication.
4. The BS shall analyze the format of V according to the verification process for ISO/IEC 9796 formatting scheme (see [ISO 9796]).
5. If the verification process fails, the certificate shall be rejected and the registration stream shall be closed.
6. If the ISO/IEC 9796 format verification succeeds, then the BS shall recover the message M from the analysis of V .

Following the successful recovery of the UT certificate message M , the BS shall perform the following steps:

1. The BS shall check the certificate validity date to ensure that the certificate has not expired. If the certificate has expired, the BS shall reject the UT (refer to clause 11.3.7, Failed UT Authentication).
2. If the certificate has not expired, then the BS shall extract Q_M (the minimal representation of the UT public key) from the message part of the certificate (refer to clause 11.3.1, UT Certificates). Q_M shall be 164 bits long.
3. The BS shall obtain Q (the public key of the UT) from Q_M using point decompression function (refer to clause 11B.2, Elliptic Curve Arithmetic). Q shall be represented by two 163 bit components Q_x and Q_y .
4. BS shall verify that Q is a point on the elliptic curve, i.e., it satisfies: $(Q_y)^2 + Q_x \cdot Q_y = (Q_x)^3 + (Q_x)^2 + 1$ in $GF(2^{163})$.
5. Furthermore, the BS shall verify that Q is not equal to either the point at infinity represented by $(0,0)$, or $(1,0)$, or $(0,1)$.
6. If the point Q does not pass all of the tests in this list, the BS shall reject the UT (see clause 11.3.7, Failed UT Authentication).

Following the successful analysis of the UT certificate, the BS shall analyze the authenticator string using the verification procedure described in clause 11D.2.3, i-TAP BS Verification of UT Signature. If the verification fails, then the BS shall reject the UT (see clause 11.3.7, Failed UT Authentication).

11.3.7 Failed UT Authentication

In case that any of the steps in clause 11.3.6 BS Analysis of UT Certificate and Authenticator String lead to a rejection of the UT, the BS shall send a rejection code in the Reg Params message. Table 101 - Connection Rejection Values provides the details.

11.4 HC-SDMA Secure Communications Protocol, *i*-SEC

i-SEC protocol is responsible for the encryption of TCH traffic streams that carry AM data and control messages. The BS certificate provides information about the capabilities of the base station (refer to clause 11.2.2, Base Station Transmits Its Digital Certificate) in terms of supported algorithms and secret key sizes. *i*-SEC shall employ the stream cipher algorithm explained in this section and shall use a portion²⁶ of the shared secret K_r established during *i*-HAP.

In addition, the BS shall have two thresholds: `regThreshBS` that limits the number of TCH streams that use the same shared secret and `streamThreshBS` that limits the number of bursts per TCH stream. The UT shall have its own thresholds `regThreshUT` and `streamThreshUT` for the same purpose. Refer to see clause 11.4.2.2, Stream Counter Test and clause 11.4.2.6, Burst Counter Test for details on the use of these thresholds. The UT shall be responsible for closing TCH streams or initiating new registrations based on the updated values of `regThreshUT` and `streamThreshUT` gracefully. The BS shall be responsible for terminating TCH streams and registrations based on its thresholds `streamThreshBS` and `regThreshBS`.

The *i*-SEC encryption key K shall consist of the L_{key} lowest significant bytes of the shared secret K_r (326 bit long shared secret which shall be exchanged with the current BS public key encryption algorithm). The largest value of the parameter L_{key} supported by a BS shall be presented the first subfield of field (12) of the BS certificate (refer to clause 11.2.2, Base Station Transmits Its Digital Certificate). L_{key} shall have a maximum value of 35 bytes, allowing for encryption keys of up to 280 bits. The actual length of the *i*-SEC encryption key K is negotiated between the UT and the BS. The remaining 46 bits of the shared secret shall be reserved.

11.4.1 TCH Streams

i-SEC stream encryption shall consist of XORing the RMU header and data octets that are going to be transmitted by the physical layer with keystream generated by the *i*-SEC cipher on a per burst basis. The CRC shall be computed over the encrypted bits and the CRC itself shall not be encrypted. See clause 3.3.3, Broadcast Burst for further information on burst formats.

11.4.2 Symmetric Key Stream Cipher Algorithm

i-SEC shall employ the symmetric key encryption algorithm illustrated in Figure 112 - Encryption in HC-SDMA Secure Communication protocol, *i*-SEC. The following sections provide the details.

²⁶ A disjoint portion of the shared secret shall be used for scrambling in *i*-TAP.

11.4.2.1 Initialization Vector Selection

The *i-SEC* encryption key K shall be combined with a 32 bit initialization vector (IV) to initialize a stream cipher state at the beginning of a stream. This IV shall be different for each TCH stream over the expected lifetime of the shared secret key.

In the *HC-SDMA* system, a stream begins with a burst containing a RA message from the UT. The IV for a stream shall be determined from the absolute frame number (AFN), slot number (SN) and relative carrier number (RCN) of the burst containing the RA message that initiates the stream:

1. The least significant 26 bits of the IV shall be:

$$IV_k = AFN_k ; k=0,1,\dots,25 .$$

where AFN_0 is the LSB of the AFN during which the burst containing the RA was transmitted.

2. Slot Number (SN) shall be represented as two bits $SN = \{SN_0, SN_1\}$, taking values {0,0}, {1,0}, {0,1} for slots 0, 1, 2 respectively. The slot number shall contribute two bits to the IV:

$$IV_{26} = SN_0 \text{ and } IV_{27} = SN_1 .$$

SN shall refer to the burst in which the RA was transmitted.

3. Relative Carrier Number (RCN) shall be the relative carrier number (the carrier number relative to the first carrier, starting at 0. See the BCH *offset* field description in clause 7.3.4.2, Configuration Message (CM) for detailed information) on which the burst with the RA was transmitted. The four least significant bits of the RCN shall contribute to the IV:

$$IV_{28} = RCN_0 ; IV_{29} = RCN_1 ; IV_{30} = RCN_2 ; IV_{31} = RCN_3$$

where RCN_0 shall be the LSB of RCN that the burst containing the RA was transmitted.

11.4.2.2 Stream Counter Test

The UT shall initialize a counter entitled `streamCount` to 0 after *i-HAP*. This counter shall be incremented at the beginning of every TCH stream and then compared to a threshold entitled `regThreshUT`. The UT application software shall set the value for the registration termination threshold `regThreshUT`. If the value of `streamCount` is equal to `regThreshUT`, then the UT shall re-register with the BS after the closure of the current stream in order to exchange a fresh shared secret using the *i-HAP* protocol. If the value of `regThreshUT` is zero, then the UT shall not force a re-registration in order to update the shared secret.

The BS shall perform the same operations independently by using its own counter `streamCount`, and its threshold `regThreshBS`. A nominal value for `regThreshBS` and `regThreshUT` should be 10,000 streams. A UT that needs to open a new stream every second may exhaust its shared secret after approximately three hours. Furthermore, with the nominal values for `streamThreshBS` and `streamThreshUT` (refer to clause 11.4.2.6, Burst Counter Test), the nominal value of 10,000 stream threshold amounts to a maximum of 5 Gigabytes of over the air traffic to be encrypted with the stream cipher using the secret key derived from the same shared secret (but with different initialization vectors).

11.4.2.3 Determination of Encryption Key from Shared Secret

While the shared secret K_r shall be 326 bits long²⁷, at most 280 bits (35 bytes) of the key shall be actually utilized for symmetric key encryption. The length (in bytes) of *i*-SEC encryption key (K) shall be defined as L_{key} . The first subfield of Field (12) of the BS certificate (refer to clause 11.2.2, Base Station Transmits Its Digital Certificate) shall specify the maximum value of L_{key} supported by the BS. The value of L_{key} shall be less than 36.

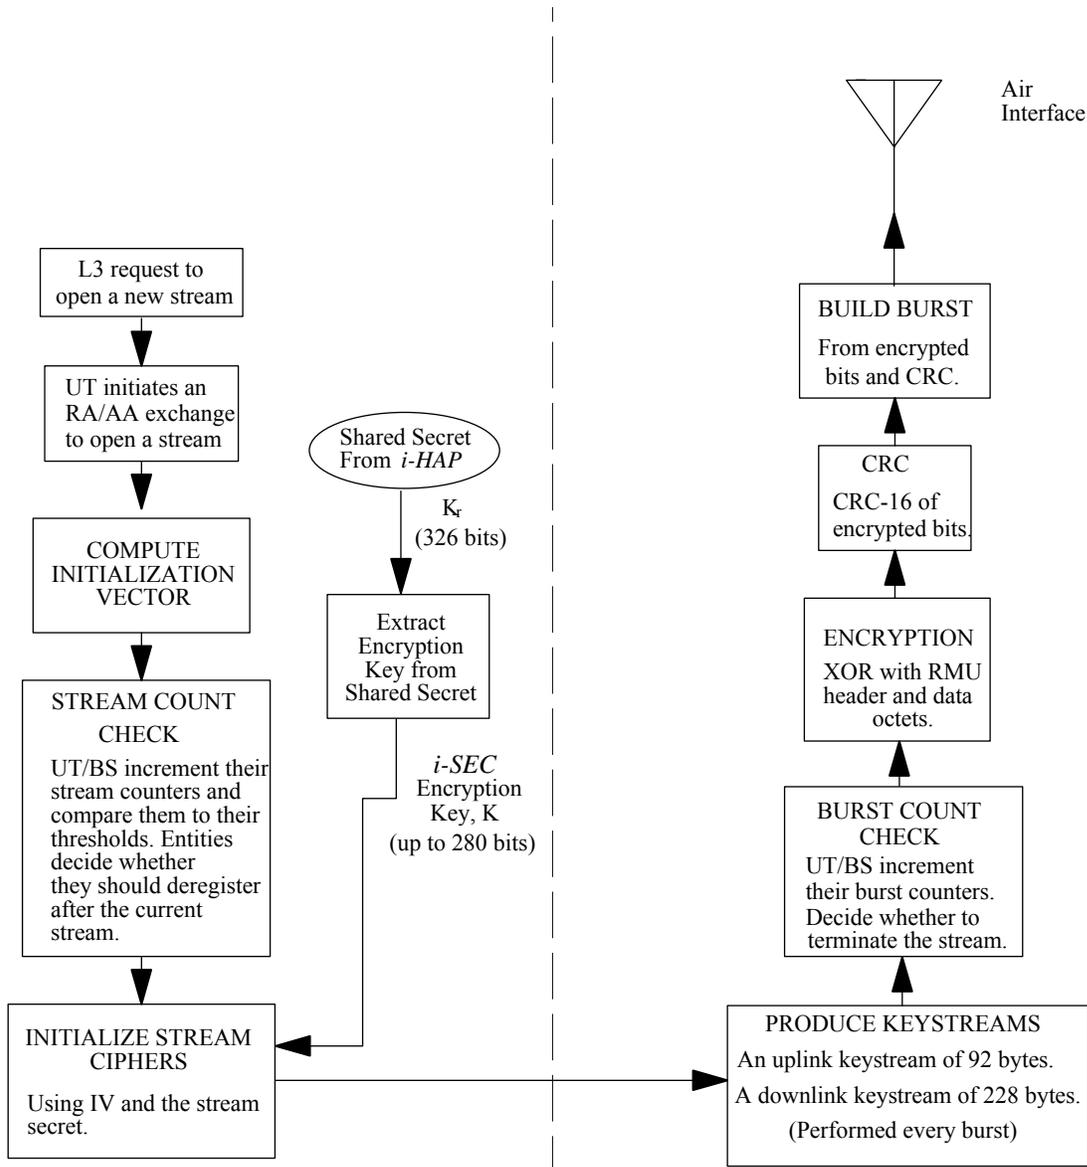


Figure 112: Encryption in *HC-SDMA* Secure Communication protocol, *i*-SEC

²⁷ When the BS public key encryption is set to elliptic curve using K-163, and two elements of $GF(2^{163})$ are used to transport the shared secret.

The first ($L_{\text{key}} \times 8$) least significant bits of the shared secret K_r shall be defined as the encryption key, K :

$$K_n = K_{r,n}; n = 0, 1, \dots, L_{\text{key}} \times 8 - 1.$$

where K_n is the n -th bit of the i -SEC encryption key K , and $K_{r,n}$ is the n -th bit of the i -SEC encryption key K_r . For maximally secure US deployments, L_{key} may be set to 35¹. The actual key size used by i -SEC shall be negotiated during registration. The BS certificate (see clause 11.2.2, Base Station Transmits Its Digital Certificate) shall set the maximum length on the key size that the BS supports. The UT shall send the key length to be used for the registration in the UT Params message (see clause 9.5.1, Acknowledged Mode (AM) Messages). Test UTs may use a method to bypass i -SEC encryption as explained in clause 11.2.4, User Terminal Encrypted Shared Secret Transmission.

11.4.2.4 Stream Cipher State Initialization

Each stream shall maintain a single encryption engine which is composed of two 8-bit counters i and j , and a state vector S that consists of 256, 8-bit states $S[0], S[1], \dots, S[255]$. The stream cipher shall be initialized per TCH stream during the Request Access / Access Assignment (RA/AA) frame. The initialization procedure shall distribute the i -SEC encryption key (K) bits and the initialization vector bits (IV) into the 256 element state vector S of the cipher as follows:

1. Initialization the state vector S and counters:

$$S[k] = k, k = 0, 1, \dots, 255; i = 0; j = 0.$$

2. Insertion of the i -SEC encryption key K and the initialization vector IV into the state vector of the cipher. Perform the following sequence of operations for 256 times:

- ◆ Compute the following pointer variables: $m = i \bmod 8$, and $n = i \bmod (2 \cdot L_{\text{key}})$, in which L_{key} denotes the i -SEC encryption key size.

- ◆ Form the interleaved byte w :

$$w = [IV_{3+4 \times m}, K_{3+4 \times n}, IV_{2+4 \times m}, K_{2+4 \times n}, IV_{1+4 \times m}, K_{1+4 \times n}, IV_{4 \times m}, K_{4 \times n}]$$

in which: IV_0 is the LSB of IV and IV_{31} is the MSB of IV ; K_0 is the LSB of K and $K_{8 \times L_{\text{key}} - 1}$ is the MSB of K ; the leftmost bit in the above expression for w (i.e., $IV_{3+4 \times m}$) is the MSB of w .

- ◆ Update j : $j = (j + S[i] + w) \bmod 256$, in which is w the interleaved byte from the previous step.

- ◆ Swap $S[i]$ and $S[j]$.

- ◆ Increment i : $i = (i + 1) \bmod 256$.

3. Key diffusion: The cipher is operated without providing keystream.

- ◆ Initialize counters: $i = 0; j = 0$.

- ◆ Repeat the following set of operations 256 times:

$$\text{Update } j: j = (j + S[i]) \bmod 256.$$

$$\text{Swap } S[i] \text{ and } S[j].$$

$$\text{Increment } i: i = (i + 1) \bmod 256.$$

- ◆ Finally, reinitialize counters: $i = 0; j = 0$.

11.4.2.5 Keystream Generation

On every frame following the RA/AA frame, 320 keystream bytes shall be generated for uplink and downlink encryption. The first 92 keystream bytes shall be allocated for uplink encryption and the remaining 228 keystream bytes are allocated for downlink encryption.

A keystream byte P shall be generated by performing the following set of operations:

- ◆ $i = (i + 1) \bmod 256$.
- ◆ $j = (j + S[i]) \bmod 256$.
- ◆ $t = (S[i] + S[j]) \bmod 256$.
- ◆ Keystream byte P : $P = S[t]$.
- ◆ Swap $S[i]$ and $S[j]$.

A total of 320 keystream bytes shall be generated per frame for each stream and counters i and j shall never be reset during the duration of a stream. This stream shall be defined as: $P[0], P[1], \dots, P[319]$. The keystream bytes shall be partitioned into two groups:

1. Uplink Keystream Bytes ($UP[k]$): The first 92 bytes of the keystream shall be allocated for uplink encryption: $UP[k] = P[k]$; $k = 0, 1, \dots, 91$.
2. Downlink Keystream Bytes ($DP[k]$): The last 228 bytes of the keystream shall be allocated for downlink encryption: $DP[k] = P[k + 92]$; $k = 0, 1, \dots, 227$.

During each frame, a constant number of stream cipher keystream bytes shall be generated independent of the actual burst sizes. The number of keystream bytes to generate for uplink (92 bytes) and downlink (228 bytes) are chosen to exceed the size of the largest uplink and downlink bursts in the *HC-SDMA* standard (refer to Figure 15 - Block Diagram of Error Control Coding Scheme and Tables 24 and Table 25.):

- ◆ The maximum uplink burst is 623 bits, including 16 bits of CRC. The CRC shall not be encrypted in *i-SEC*.
- ◆ The maximum downlink burst is 1817 bits, including 16 bits of CRC. The CRC shall not be encrypted in *i-SEC*.

11.4.2.6 Burst Counter Test

The UT shall increment its counter `burstCount` after generating the keystream for a burst. This counter shall be set to zero at every stream start. After incrementing `burstCount`, the UT shall compare the value of the counter to a threshold entitled `streamThreshUT`. The UT application software shall set the value for the stream termination threshold `streamThreshUT`. If `streamThreshUT` is non-zero, and the value of `burstCount` is equal to `streamThreshUT`, then the UT shall close the current stream and shall start a new stream with a new stream secret and IV. If the value of `streamThreshUT` is zero, then the UT shall not force a stream closure.

The BS shall perform the same operations independently by using its own counter `burstCount`, and its own threshold `streamThreshBS`. A nominal value for `streamThreshBS` and `streamThreshUT` shall be 20,000 bursts. At a combined uplink/downlink approximate data rate of 400 kilobits/sec (single burst/frame, 200 frames/sec), this results in a stream closure approximately

after 5 megabytes of over the air traffic before the initialization vector is changed and the stream cipher is reinitialized. This selection also limits the stream duration to 100 seconds regardless of the data rate.

11.4.2.7 Encryption and Decryption Using *i*-SEC Keystream

i-SEC encryption shall be performed by XORing the RMU header and payload (see clause 7.3.7.2 RMU Mux Function and Figure 30 - AM RMU Definition for definitions) by the *i*-SEC keystream before transmission. Therefore, *i*-SEC encryption is not a bandwidth expanding transformation. The 16 bit CRC shall be calculated over the encrypted bits. The CRC shall not be encrypted.

i-SEC decryption shall be performed by XORing the received encrypted RMU header and payload by the *i*-SEC keystream. The 16 bit CRC check shall be performed on the received encrypted bits before decryption.

The uplink RMU header and payload shall be XORed with the uplink keystream, and the downlink header and payload shall be XORed with the downlink keystream. In streams using the basic RMU header, the RMU header and payload octets shall be aligned such that the last payload octet falls exactly at the end of the relevant keystream buffer. In streams using the minimized RMU header, the last burst bit before the CRC bits shall be aligned with the end of the relevant keystream buffer. During this process, some of the keystream bytes or bits are possibly left unused and they shall be discarded.

For example, consider an uplink burst payload that contains 28 bits of basic RMU header H and 3 payload octets $\{A, B, C\}$ with the basic RMU header H represented by a bit vector $[h_0, h_1, \dots, h_{27}]$ where h_0 is transmitted first and the header is partitioned into four bytes:

$$H_0 = [0, 0, 0, 0, h_0, h_1, h_2, h_3],$$

$$H_1 = [h_4, h_5, h_6, h_7, h_8, h_9, h_{10}, h_{11}], H_2 = [h_{12}, h_{13}, h_{14}, h_{15}, h_{16}, h_{17}, h_{18}, h_{19}], \text{ and}$$

$$H_3 = [h_{20}, h_{21}, h_{22}, h_{23}, h_{24}, h_{25}, h_{26}, h_{27}]:$$

- ◆ The *i*-SEC cipher generates 320 keystream bytes and allocates the first 92 bytes for uplink encryption: $\{UP[0], UP[1], \dots, UP[91]\}$.
- ◆ Payload octet C is XORed with $UP[91]$ to produce C_e , payload octet B is XORed with $UP[90]$ to produce B_e , and payload octet A is XORed with $UP[89]$ to produce A_e .
- ◆ Header octet H_0 is XORed with $UP[85]$ to produce G_0 , octet H_1 is XORed with $UP[86]$ to produce G_1 , octet H_2 is XORed with $UP[87]$ to produce G_2 , octet H_3 is XORed with $UP[88]$ to produce G_3 .
- ◆ Let $G_0 = [g_0, g_1, g_2, g_3, g_4, g_5, g_6, g_7]$. The encrypted RMU header and the payload is represented by the bit vector $X = [g_4, g_5, g_6, g_7, G_1, G_2, G_3, A_e, B_e, C_e]$.
- ◆ 16 bit CRC of X is computed and appended after the payload bits to prepare the burst content Y defined as $Y = [X, CRC_{16}(X)]$. Note that the CRC shall be not encrypted.
- ◆ The unused uplink keystream bytes $\{UP[0], UP[1], \dots, UP[84]\}$ is discarded. These discarded keystream bytes serve as additional key diffusion steps among frames.

APPENDIX 11A REFERENCES

The following are references used within this chapter:

- ◆ [BSS] I. Blake, G. Seroussi and N. Smart, *Elliptic Curves in Cryptography*, Cambridge University Press, Cambridge, United Kingdom, 1999.

- ◆ [FFW] J. Fegghi, J. Fegghi, and P. Williams, *Digital Certificates: Applied Internet Security*, page 52, Addison Wesley Longman Inc., 1999.
- ◆ [FIPS] FIPS PUB 186-2: Federal Standards Publication: *Digital Signature Standard (DSS)*, January 27, 2000.
- ◆ [ISO 9796] ISO/IEC 9796 First Edition 1991-09-15: *Information Technology – Security Techniques – Digital Signature Scheme Giving Message Recovery*, September 15, 1991.
- ◆ [MENEZES] A.J. Menezes, P. van Oorschot, and S. Vanstone, *Handbook of Applied Cryptography*, CRC Press Series on Discrete Mathematics and Its Applications, 1996. Online version is available at: <http://www.cacr.math.uwaterloo.ca/hac>.
- ◆ [PETERSON] Roger L. Peterson, Rodger E. Ziemer, and David E. Borth, *Introduction to Spread-Spectrum Communications*, page 119, Prentice Hall, 1995.
- ◆ [SCHNEIER] B. Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*, John Wiley and Sons, New York, Second Edition, 1996.
- ◆ [WICKER] S.B. Wicker, *Error Control Systems for Digital Communication and Storage*, Prentice-Hall, New Jersey, 1995.

APPENDIX 11B ELLIPTIC CURVE PARAMETERS AND ARITHMETIC

An elliptic curve (EC) is defined by a set of parameters called the EC domain parameters (ECDP). These parameters are: a finite field size and representation, an elliptic curve equation, and a base point that satisfies the elliptic curve equation in the selected finite field. The set of all pairs from the selected finite field ($GF(2^m)$ in our application) that satisfy the elliptic curve equation form a group under the operation entitled point addition (refer to clause 11B.2, Elliptic Curve Arithmetic). The elliptic group is represented as $E\{GF(2^m)\}$. A hypothetical point entitled as *the* point at infinity O , is also included in $E\{GF(2^m)\}$ as the identity element.

11B.1 Elliptic Curve Domain Parameters (ECDP)

The elliptic curve domain parameters (ECDP) used for *i-HAP* and *i-TAP* public key encryption are defined below (Curve K-163 in [FIPS]):

1. Field size, 2^{163} .
2. Field representation (polynomial form): $t^{163} + t^7 + t^6 + t^3 + 1$.
3. Elliptic Curve: $y^2 + x \cdot y = x^3 + x^2 + 1$.
4. Base point of prime order $G = [G_x, G_y]$, in which the components of G are presented in polynomial form in $GF(2^{163})$. In hexadecimal notation:

$$G_x = 2FE13C0537BBC11ACAA07D793DE4E6D5E5C94EEE8 .$$

$$G_y = 289070FB05D38FF58321F2E800536D538CCDAA3D9 .$$

5. The order²⁸ of the base point G (in decimal base):

$$r_G = 5846006549323611672814741753598448348329118574063.$$

²⁸ The smallest positive integer r_G , such that the point multiplication $[r_G]G = O$, the point at infinity. Refer to clause 11B.2, Elliptic Curve Arithmetic for point multiplication.

The integers G_x , G_y , and r_G are defined above in big-endian format to comply with the standard document [FIPS].

11B.2 Elliptic Curve Arithmetic

Using the ECDP, we define the arithmetic operations on the elements of $E\{GF(2^m)\}$. These definitions are used with $m = 163$.

- ◆ **Point Negation:** Given a point $P = [P_x, P_y]$ on the elliptic curve, the ‘negative’ of P is denoted as $[-1]P$ and computed as $[-1]P = [P_x, P_x + P_y]$. The addition $P_x + P_y$ is performed using $GF(2^m)$ arithmetic.
- ◆ **Point Addition:** Given two points $P = [P_x, P_y]$ and $Q = [Q_x, Q_y]$ on the elliptic curve, their point addition results in the point R defined as: $R = [R_x, R_y] = P + Q$, provided that P is not equal to $[-1]Q$. The computation for the components of R is illustrated in [BSS]. If P is equal to $[-1]Q$, then the result R is defined as the point at infinity, O . If Q is equal to O , then $R = P$.
- ◆ **Point Multiplication:** For any positive integer $k > 1$, the multiplication of a point on the elliptic curve P with k is denoted as $[k]P$. This multiplication is defined as a recursive point addition: $[k]P = [k - 1]P + P$, with the definition $[0]P = O$.
- ◆ **Point Compression:** $Q = [Q_x, Q_y]$ with components in $GF(2^m)$ can be represented using $2 \cdot m$ bits. This storage requirement can be reduced to $(m + 1)$ bits by using the minimal representation of Q , defined as Q_M : (1) First m bits of Q_M are set to the polynomial representation of Q_x ; (2) The $(m + 1)$ -th bit of Q_M is determined as the coefficient of the n -th order term in the polynomial representation of Q_y where n is the highest power term in the polynomial representation of Q_x . Q_M uniquely describes Q provided that Q is not equal to the point at infinity, O . Refer to [BSS] for details.
- ◆ **Point Decompression:** This is the inverse of the point compression function. Given Q_x , the elliptic curve is quadratic in Q_y : (1) Solve for one of the roots of the elliptic curve equation for Q_y using the half trace function (see [BSS], page 26) and denote this value as z ; (2) Let n be the highest order term in the polynomial representation of Q_x . If the coefficient of the n -th order term in the polynomial representation of z has a coefficient equal to the $(m + 1)$ -th bit of Q_M , then we accept z as the y -coordinate of Q ; $Q_y = Q_x + z$ where addition takes place in $GF(2^m)$.

APPENDIX 11C. I-HAP ALGORITHMS

This appendix describes the base station certificate signature generation and verification by the CA, and CA public/private key pair generation methods. This description is used to explain UT encryption and BS decryption operations.

11C.1 i-HAP Certificate Signatures

This section describes the key pair generation process for the certificate authority (CA), the process of generating a certificate signature, and the process for certificate signature verification performed by the UT using the public key of the CA.

The *HC-SDMA* signature scheme (currently) employs the RSA algorithm. Signature generation method shall be defined in the message portion of the certificate (see clause 11.2.2, Base Station Transmits Its Digital Certificate): RSA algorithm with a 1024 bit modulus shall be the baseline public key algorithm together with the ISO/IEC 9796 formatting scheme (see [ISO 9796]) which does not require a hash function. The signature shall be 1024 bits long. The message field may vary depending on the length of the minimal representation of the BS public key; with elliptic curve cryptography using curve K-163, the message field shall be 297 bits.

The BS certificate shall be 1321 bits long and obtained by concatenating the message and the signature fields. Although the message field is redundant (since the message can be recovered from the signature provided the ISO/IEC 9796 formatting is used instead of a hash function), it provides a mechanism to initiate a speculative shared secret exchange by the UT.

The *HC-SDMA* privacy and authentication protocol does not specify a specific mechanism on the generation and delivery of the base station certificates. The BS certificate message field may be composed by the certificate authority. In addition, the BS public/private key pairs may be generated by the certificate authority and transmitted to the BS by a secure communication link.

The BS key pair should not be used more than three days. This shall be enforced by the certificate authority: the CA should not sign certificates with validity period more than 3 days. In addition, the CA should not sign more than 16 certificates from a base station during a period of 3 days.

The *HC-SDMA* protocol allows the certificate authority to generate BS key pairs and sent to the BS using a secure link. *HC-SDMA* privacy and authentication protocol does not dictate but recommends the certificate authority to send different key pairs to different base stations. The BSs that receive their certificates from the CA shall be authenticated by the CA but this requirement is beyond the scope of *HC-SDMA* protocols.

11C.1.1 *i-HAP* Certificate Authority (CA) Key Pair Generation

This section describes CA public and private key generation based on the RSA algorithm. The modulus n used in the RSA algorithm is 1024 bits. This is approximately equivalent to 96 bit security provided by a symmetric key encryption algorithm [FFW]. The CA public key pair shall be generated as follows:

- ◆ The certificate authority shall randomly choose two large prime numbers p and q of 512 bits.
- ◆ The CA shall compute the RSA modulus $n = p \cdot q$. The RSA modulus²⁹ is represented by 1024 bits.
- ◆ The CA shall compute $\phi = (p - 1) \cdot (q - 1)$ and shall destroy the values p and q .
- ◆ The CA shall set the variable $e = 2^{16} + 1$, check³⁰ whether e and ϕ are relatively prime, i.e., $\text{gcf}(e, \phi) = 1$. If this is not the case, then the CA shall restart the key generation process.
- ◆ The CA shall compute³¹ the unique integer d , such that $1 < d < \phi$ and $e \cdot d = 1 \pmod{\phi}$ □□
- ◆ The CA public key shall be the combination (n, e) and the private key is d . The public key shall be distributed to the users of the *HC-SDMA* system. A UT shall have the public key of the CA prior to registration.
- ◆ If there exist multiple CAs, then each CA shall generate its keys independently through the

²⁹ The most significant bit of n shall be non-zero.

³⁰ Note that e is chosen as a prime with a simple binary expansion in order to expedite signature verification at the UT.

³¹ This may be accomplished using the extended Euclidean algorithm.

procedure outlined above, and shall not share their private parameters $\{p, q\}$.

11C.1.2 *i-HAP* Certificate Signature Generation

The CA shall generate its signature for the BS certificate by processing the BS message M (currently 297 bits, see clause 11.2.2, Base Station Transmits Its Digital Certificate):

- ◆ The CA shall format M according to ISO/IEC 9796 formatting scheme (see [ISO 9796]) to obtain a 1024 bit integer Y . This transformation is reversible.
- ◆ The certificate signature shall be computed as $S = Y^d \bmod n$. n is the RSA modulus and is public. However, d shall be private to the CA. The signature S shall be 1024 bits long.
- ◆ The certificate $C = [M, S]$ shall be the concatenation of the message and signature. C shall be 1321 bits long³².

The ISO/IEC 9796 formatting does not utilize a hash function and therefore signature generation and verification steps are simpler than signature schemes that utilize hash functions. On the other hand, this lets the CA signature scheme become more vulnerable to chosen plaintext attacks and necessitates a BS authentication scheme at the CA. However, since the CA should not sign more than 16 certificates during a period of 3 days, the potential of chosen plaintext attacks from unauthorized base stations are kept to a minimum.

11C.1.3 *i-HAP* Certificate Signature Verification

The CA shall serve the certificate to the appropriate base station. The BS shall deliver its certificate to a UT when the UT starts registration. After receipt, the UT shall check the signature of the CA:

- ◆ UT shall receive the BS certificate $C' = [M', S']$ from the BS. The UT may perform a speculative shared secret transmission at this point.
- ◆ UT shall obtain the public key (n, e) for the CA from its memory by analyzing the CA identifier field in M' .
- ◆ UT shall decrypt the received CA signature S' by using the CA public key e :
- ◆ $V = (S')^e \bmod n$.

Due to the specific value of e , this operation can be computed with only 16 modular squarings and one modular multiplication. The following details the process of verifying the certificate when ISO/IEC 9796 formatting is employed (see field (2) of the BS certificate detailed in clause 11.2.2, Base Station Transmits Its Digital Certificate).

- ◆ The UT shall analyze the format of V according to the verification process for ISO/IEC 9796 formatting scheme (refer to [ISO 9796]).
- ◆ If the verification process fails, the certificate shall be rejected. If a speculative shared secret transmission was made, then the on-going *i-HAP* protocol shall be aborted.
- ◆ If the ISO/IEC 9796 format verification succeeds, the UT shall recover its estimate of the message M'' from the analysis of V . The UT shall compare the message component of the received

³² Since the ISO/IEC 9796 formatting is reversible, it is possible to recover M from S at the UT. Therefore, the *HC-SDMA* BS certificate defined as the concatenation of M and S are redundant by 297 bits. However, this redundancy is allowed in order to let the UT perform speculative shared secret transmissions without verifying the certificate. With the availability of faster UT processors, it is possible to eliminate the message part from the certificate during the evolution of the *HC-SDMA* protocol.

certificate M' and M'' . If these are identical, then the UT shall accept the authenticity of the certificate. Otherwise, the certificate shall be rejected, and any on-going speculative *i-HAP* protocol attempt shall be aborted.

11C.2 *i-HAP Base Station Public Key Encryption Algorithm*

This section specifies algorithms for the generation and use of public/private key pairs for the base stations in the *HC-SDMA* system. Each BS shall possess a private and public key pair. Public keys shall be delivered to any UT that requests them.

The base station public key algorithm for *HC-SDMA* is based on elliptic curve (EC) cryptography in the finite field $GF(2^{163})$ which offers the equivalent in security to a 96-bit symmetric key-encryption algorithm [FFW].

11C.2.1 *i-HAP BS Key Pair Generation*

The following steps describe key pair generation for a base station using ECDP (defined in clause 11B.1, Elliptic Curve Domain Parameters (ECDP):

1. The base station shall generate a 162 bit random integer d .
2. The base station shall compute the point multiplication (refer to clause 11B.2, Elliptic Curve Arithmetic) $Q = [d]G$ in the elliptic group where G is the base point of the elliptic group $E\{GF(2^{163})\}$.
3. The point $Q = [Q_x, Q_y]$ shall be checked for validity: (a) Q shall be on the elliptic curve -- i.e., $(Q_y)^2 + Q_x \cdot Q_y = (Q_x)^3 + (Q_x)^2 + 1$ in $GF(2^{163})$; (b) Q shall not be equal to the point at infinity O represented as $(0,0)$; (c) Q shall not be equal to point $(0,1)$; (d) Q shall not be equal to $(1,0)$. If any of these tests fail, then the key generation process restarts.
4. The base station's public key shall be $Q = [Q_x, Q_y]$ combined with ECDP, and its private key shall be d .

The *HC-SDMA* protocol allows the certificate authority to generate BS key pairs and sent to the BS using a secure link. In this case, the above steps shall be performed by the certificate authority.

11C.2.2 *i-HAP UT Encryption*

1. The UT shall perform the following operations to extract the public key from the BS certificate. The BS public key shall be used to encrypt the shared secret K_r generated by the UT. After verifying the BS certificate, the UT shall extract the minimal representation of the BS public key Q_M from the message part of the certificate (refer to clause 11.2.2, Base Station Transmits Its Digital Certificate).
2. The UT shall determine the x-coordinate of $Q = [Q_x, Q_y]$: Q_x is set to the first 163 bits of Q_M . The UT shall determine Q_y using Q_x and the last bit of Q_M using point decompression (refer to clause 11B.2, Elliptic Curve Arithmetic).
3. The UT shall generate a random 162 bit random bit integer a .

4. The UT shall compute the following two points on the elliptic curve using point multiplication (refer to clause 11B.2, Elliptic Curve Arithmetic): $S = [a]G$ and $V = [a]Q$. G is the base point of the elliptic group and Q is the BS public key.
5. The two elliptic curve points S and V shall be tested against the following requirements: (a) Points shall be on the elliptic curve; (b) Points shall not be equal to the point at infinity O represented as $(0,0)$; (c) Points shall not be equal to point $(0,1)$; (d) Points shall not be equal to $(1,0)$. If any of these tests fail, then the encryption process shall return to Step 3 to generate a new a .
6. The UT shall determine the minimal representation of S defined as S_M using point compression (refer to clause 11B.2, Elliptic Curve Arithmetic). S_M requires 164 bits of storage.
7. The UT shall compute the following two $GF(2^{163})$ multiplications: $U_x = V_x \cdot M_x$, and $U_y = V_y \cdot M_y$; where M is the 326 bit message to be encrypted, M_x denotes the least significant 163 bits of M , and M_y denotes the remaining bits of M such that $M = [M_x, M_y]$. V_x is the x -component of the point V , and V_y is the y -component of the point V .
8. The UT shall transmit $N_r = [S_M, U_x, U_y]$ to the BS. The length of N_r is 490 bits. The components shall be transmitted in little endian format -- i.e., the coefficient of the lowest order polynomial term of a finite field element is transmitted first.

11C.2.3 *i*-HAP BS Decryption

A base station that receives the message from the user terminal that contains $N_r = [S_M, U_x, U_y]$ shall perform decryption as follows:

1. The base station shall determine the point S from its minimal representation S_M using point decompression (refer to clause 11B.2, Elliptic Curve Arithmetic).
2. The base station shall compute the point multiplication (refer to clause 11B.2, Elliptic Curve Arithmetic): $H = [d]S = [H_x, H_y]$, using the BS private key d .
3. The base station shall perform the following $GF(2^{163})$ division: $M_x = U_x / H_x$ where H_x is the x -component of point H and shall recover the 163 bit message component M_x .
4. The base station shall perform the following $GF(2^{163})$ division: $M_y = U_y / H_y$ where H_y is the y -component of point H and shall recover the 163 bit message component M_y .
5. The final message M shall be assembled as: $M = [M_x, M_y]$.

11C.3 *Cyclic Redundancy Code for Shared Secret Exchange*

Error detection for shared secret exchange shall be performed by using a cyclic-redundancy code (CRC). This code adds sixteen CRC bits to the end of a message and defined as CRC_{16} . The generator polynomial for CRC_{16} is (also known as CRC-ANSI):

$$g_{16}(x) = x^{16} + x^{15} + x^2 + 1 = (x + 1) \cdot (x^{15} + x + 1)$$

Encoding and decoding the CRC bits are described in clause 5 of [WICKER].

APPENDIX 11D. *i-TAP* ALGORITHMS

This section describes signature generation for UT certificates by the CA, key pair generation for the CA and a UT. Using these primitives, this section also describes the UT authenticator string generation (A_s) from the hash of the authenticator message (A_m), and the BS verification of the UT authenticator string.

11D.1 *i-TAP* UT Certificate Signatures

The *HC-SDMA* signature scheme (currently) employs the RSA algorithm. Signature generation method is defined based on the RSA algorithm with a 1024 bit modulus as the public key algorithm together with the the ISO/IEC 9796 formatting scheme (see [ISO 9796]) which does not require a hash function. The signature is 1024 bits long. The message field may vary depending on the UT certificate type (see clause 11.3.1 UT Certificates).

The *HC-SDMA* Terminal Authentication Protocol does not dictate a specific mechanism on the generation and delivery of UT certificates. The degree of control and management of UT certificates is at the discretion of the operating entities involved. The UT certificate message field may be composed by the certificate authority. In addition, the UT public/private key pairs may be generated by the certificate authority and transmitted to the UT by a secure communication link.

11D.1.1 *i-TAP* UT Certificate Authority (CA) Key Pair Generation

This section describes CA public and private key generation based on the RSA algorithm. The modulus n used in the RSA algorithm is 1024 bits. This is approximately equivalent to 96 bit security provided by a symmetric key encryption algorithm [FFW]. The CA public key pair is generated as follows:

1. The certificate authority shall randomly choose two large prime numbers p and q of 512 bits.
2. The CA shall compute the RSA modulus $n = p \cdot q$. The RSA modulus³³ n is represented by 1024 bits.
3. The CA shall compute $\phi = (p - 1) \cdot (q - 1)$ and destroy the values p and q .
4. The CA shall set the variable $e = 2^{16} + 1$, check³⁴ whether e and ϕ are relatively prime -- i.e., $\text{gcf}(e, \phi) = 1$. If this is not the case, then the CA shall restart the key generation process. Note that e is chosen as a prime with a simple binary expansion in order to expedite signature verification at the BS.
5. The CA shall compute³⁵ the unique integer d , such that $1 < d < \phi$ and $e \cdot d = 1 \pmod{\phi}$. This may be accomplished using the extended Euclidean algorithm.
6. The CA public key shall be the combination (n, e) and the private key is d . The public key shall be distributed to the users of the *HC-SDMA* system. A BS shall have the public key of the CA prior to registration.
7. If there exist multiple CAs, then each CA shall not share their private parameters and shall maintain independence.

³³ The most significant bit of n shall be non-zero.

³⁴ Note that e is chosen as a prime with a simple binary expansion in order to expedite signature verification at the UT.

³⁵ This may be accomplished using the extended Euclidean algorithm.

11D.1.2 *i-TAP* UT Certificate Signature Generation

The CA shall generate its signature for the UT certificate by processing the UT message M (the size depends on the type of the UT certificate, refer to clause 11.3.1 UT Certificates):

1. The CA shall format M according to ISO/IEC 9796 formatting scheme (refer to [ISO 9796]) to obtain a 1024 bit integer Y . This transformation is reversible.
2. The certificate signature shall be computed as $S = Y^d \bmod n$. n is the RSA modulus and it is public. However, d shall be private to the CA. The signature S is 1024 bits long.
3. The certificate $C = [M, S]$ shall be the concatenation of the message and signature.
4. Since the ISO/IEC 9796 formatting is reversible, it is possible to recover M from S at the BS. Therefore, the *HC-SDMA* UT certificate defined as the concatenation of M and S are redundant. Only the signature shall be transmitted in the UT Params message.

The ISO/IEC 9796 formatting does not utilize a hash function and therefore signature generation and verification steps are simpler than signature schemes that utilize hash functions.

11D.1.3 *i-TAP* UT Certificate Signature Verification

The CA shall serve the certificate to the appropriate UT. The UT shall deliver its certificate to a BS during UT registration. After receipt, the BS shall check the signature of the CA:

1. BS shall receive the UT certificate signature S' from the UT and the CA identifier.
2. BS shall obtain the public key (n, e) for the CA from its memory corresponding to the CA identifier field obtained from the UT.
3. BS shall decrypt the received CA signature S' by using the CA public key e : $V = (S')^e \bmod n$. Due to the specific value of e , this operation requires only 16 modular squarings and one modular multiplication.
4. The BS shall analyze the format of V according to the verification process for ISO/IEC 9796 formatting scheme (see [ISO 9796]).
5. If the verification process fails, the certificate shall be rejected.
6. If the ISO/IEC 9796 format verification succeeds, the BS shall recover the message M from the analysis of V . The BS shall use the UT public key contained in M to verify the UT authenticator string A_s sent in the same message.

11D.2 *i-TAP* UT Public Key Encryption Algorithm

This section³⁶ specifies algorithms for the generation and use of public/private key pairs for UTs in the *HC-SDMA* system. Each UT shall possess a private and public key pair with each of its certificates.

The UT public key algorithm for *i-TAP* is based on elliptic curve (EC) cryptography in the finite field $GF(2^{163})$ which offers the equivalent in security to a 96-bit symmetric key encryption algorithm [FFW].

³⁶ The generation and use of public/private key pairs for a centralized time server are identical to the descriptions presented in this section. The role of the centralized time server is to verify the time perceived at a UT. Refer to clause 11.3.3, AFN Verification with a Centralized Time Server.

11D.2.1 *i-TAP* UT Key Pair Generation

The following steps describe key pair generation for a UT using parameters in Elliptic Curve Domain Parameters (ECDP) earlier in this Appendix:

1. Generate a 162 bit random integer d .
2. Compute the point multiplication (refer to clause 11B.2, Elliptic Curve Arithmetic) $Q = [d]G$ in the elliptic group $E\{GF(2^m)\}$ in which G is defined as the base point (refer to clause 11B.1, Elliptic Curve Domain Parameters (ECDP)).
3. The point $Q = [Q_x, Q_y]$ shall be checked for validity:
 - a. Q shall be on the elliptic curve, i.e., $(Q_y)^2 + Q_x \cdot Q_y = (Q_x)^3 + (Q_x)^2 + 1$ in $GF(2^{163})$;
 - b. Q shall not be equal to the point at infinity represented as $(0,0)$;
 - c. Q shall not be equal to point $(0,1)$;
 - d. Q shall not be equal to $(1,0)$.

If any of these tests fail, then the key generation process shall restart.
4. The public key of the UT shall be combined with ECDP, and its private key shall be the 162 bit random integer d .

The *HC-SDMA* protocol allows the certificate authority to generate UT key pairs and sent to the UT using a secure link. In this case, the above steps shall be performed by the certificate authority.

11D.2.2 *i-TAP* UT Private Key Encryption

Let r_G denote the order of the base point G in the elliptic group $E\{GF(2^m)\}$. The value of r_G is defined in clause 11B.1, Elliptic Curve Domain Parameters (ECDP) and can be expressed as:

$$r_G = 2^{162} + 2422733316269705388402159$$

Note that r_G can be represented by 163 unsigned bits and the most significant bit of this representation is equal to one.

Let A_m denote the authenticator message defined in clause 11.3.2 UT Authenticator Message and String Generation. Furthermore, let $h_{SHA-1}(A_m)$ denote the 160 bit SHA-1 hash value of the authenticator message. The UT shall proceed as follows (all integer arithmetic below assumes little endian representation of integers):

1. The UT shall generate a random 162 bit random integer k . Note that $k < r_G$.
2. The UT shall compute the following the following point on the elliptic curve using point multiplication (refer to clause 11B.2, Elliptic Curve Arithmetic): $P = [k]G$. G is the base point of the elliptic group. Let the point P be represented by the ordered pair (P_x, P_y) in which each component is a 163 bit integer.
3. The UT shall compute $r = P_x \bmod r_G$. If $r = 0$, then the algorithm starts again with step 1. r is represented using 163 unsigned bits.
4. The UT shall compute³⁷ k_{inv} which is defined implicitly as: $k_{inv} \cdot k = 1 \bmod r_G$

³⁷ This can be accomplished by the extended Euclidean algorithm.

5. The UT shall compute $s = k_{\text{inv}} \cdot (h_{\text{SHA-1}}(A_m) + d \cdot r) \bmod r_G$ in which d is the private key of the UT. If $s = 0$, then the algorithm starts again with step 1. s is represented using 163 unsigned bits. The integer arithmetic is performed by treating the SHA-1 output string $h_{\text{SHA-1}}(A_m)$ as an integer in little endian format and appending three zeros on the right hand side (MSB).
6. The UT authenticator string A_s shall be defined as concatenated pair $A_s = [r, s]$. The length of the authenticator string shall be 326 bits. The components shall be transmitted in little endian format (i.e., if $A_s = [r, s]$ then LSB of A_s is $A_{s,0} = r_0$, in which r_0 is the LSB of r ; likewise $A_{s,163} = s_0$).

11D.2.3 *i-TAP* BS Verification of UT Signature

The base station shall generate a copy of the 438 bit authenticator message \hat{A}_m based on its information³⁸ on the UT (refer to Figure 109: UT Authenticator Message Formatting). Let $h_{\text{SHA-1}}(\hat{A}_m)$ denote the 160 bit string obtained by the SHA-1 hash of the BS copy of the authenticator message. The BS shall receive³⁹ the authenticator string \hat{A}_s from the user terminal ($\hat{A}_s = [\hat{r}, \hat{s}]$) and shall perform signature verification as follows (all integer arithmetic below assumes little endian representation of integers):

1. The BS shall verify that \hat{r} and \hat{s} are integers in the interval $[1, r_G - 1]$. r_G denotes the order of the base point G in the elliptic group $E\{\text{GF}(2^m)\}$.
2. The BS shall compute w which satisfies: $w \cdot \hat{s} = 1 \bmod r_G$.
3. The BS shall compute $u_1 = (h_{\text{SHA-1}}(\hat{A}_m)) \cdot w \bmod r_G$ and $u_2 = \hat{r} \cdot w \bmod r_G$. The integer arithmetic shall be performed by treating the SHA-1 output string $h_{\text{SHA-1}}(\hat{A}_m)$ as an integer in little endian format and appending three zeros on the right hand side (MSB).
4. The BS shall compute the point Z defined as the addition of two points $[u_1]G$ and $[u_2]Q$: $Z = [u_1]G + [u_2]Q$.
 G is the base point of the elliptic group (refer to clause 11B.1, Elliptic Curve Domain Parameters (ECDP)); Q is the public key of the UT obtained through point decomposition (refer to clause 11B.2, Elliptic Curve Arithmetic) from the UT certificate (refer to clause 11.3.1 UT Certificates).
5. If Z is equal to the point at infinity represented by the pair $(0,0)$, then the BS shall reject the signature. Otherwise, the BS shall convert the x -coordinate of Z , denoted here as Z_x to an integer and shall compute $v = Z_x \bmod r_G$.
6. The BS shall accept the signature if and only if $v = \hat{r} \bmod r_G$.

11D.3 *i-TAP* Scrambling Algorithm

This section specifies algorithms for Linear Feedback Shift Register (LFSR) scrambling used for UT certificate signature and UT parameter transmission. UT certificate signature and certain UT specific parameters shall be scrambled before they are encapsulated in the UT Params message for over the air

³⁸ Nominally, \hat{A}_m is equal to A_m

³⁹ Nominally, \hat{A}_s is equal to A_s ; and therefore, $\hat{r} = r$, $\hat{s} = s$.

transmission. This is intended for concealing the UT identity from unauthorized listeners. UT data confidentiality is provided by the *i-SEC* protocol.

The LFSR employed shall be generated by the following primitive polynomial:

$$g(x) = 1 + x^4 + x^{39}$$

The pseudo-random sequence generated by the sequence generator that is defined by $g(x)$ has a period of $2^{39} - 1$ bits. Figure 113: *i*TAP Linear Feedback Shift Register illustrates an implementation of the LFSR. The memory elements S_0 to S_{38} shall be initialized with the 39 most significant bits of the shared secret K_r , i.e., S_0 shall assume the value of the most significant bit of K_r , S_1 shall assume the value of the next significant bit of K_r , etc. If all of the memory elements are equal to zero after initialization, then S_0 shall be reset to 1.

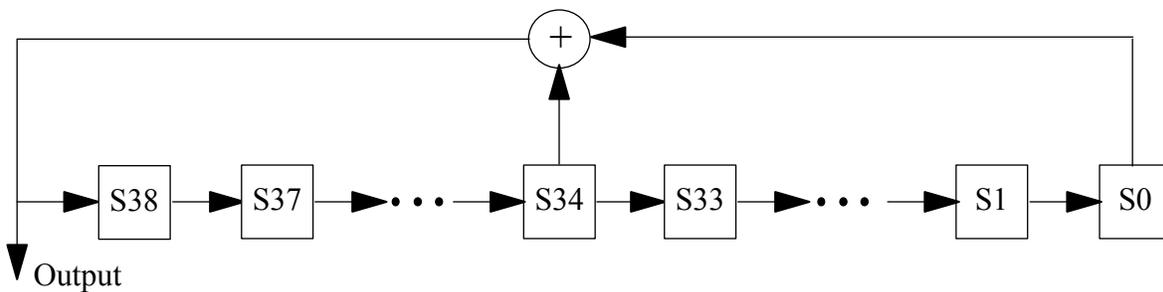


Figure 113: *i*-TAP Linear Feedback Shift Register

The UT shall produce 1088 output bits using the LFSR after the initialization. These output bits shall be used to scramble the UT certificate signature and certain UT specific parameters outlined in clause 11.3.4, UT Parameter Scrambling. The scrambling shall be performed by computing the XOR of a data bit with an output bit from the LFSR. The BS shall extract the scrambled data bits from the UT Params message (refer to clause 11.3.5, UT Params Message Transmission), shall initialize its own LFSR by the 39 most significant bits of the shared secret K_r , shall produce 1088 output bits from its LFSR, and shall perform descrambling by computing the XOR of a scrambled bit with an output bit from its LFSR.

11D.4 *i*-TAP Automatic Certificate Retrieval

UTs which are IP enabled (i.e., also include a native TCP/IP stack) may also implement the *HC-SDMA* automatic certificate download feature. IP enabled UTs shall establish a link-level connection (i.e., PPP) autonomously (i.e., without a request by an EUD) for purposes of IP level communication between the UT and a management server.

An IP enabled UT may support the *HC-SDMA* automatic certificate download feature by using TCP/IP to download a certificate. A BS shall restrict link-level access for link-level (e.g., PPP) sessions originating at the UT. The BS shall do that by monitoring the Network Access Identifier (NAI) used by the UT to gain access. The NAI presented by the UT is of the form <UT-Serial-Number>@<manufacturer-domain>. A BS shall compare that information with the information presented in the certificate.

A BS shall distinguish between native UT link-level access and EUD link-level access by the type of certificate presented by the UT. A Network Access Certificate shall always signal EUD level access. A Service Authorization Certificate shall always signal UT level access. An Identity Certificate shall signal UT level access except for early deployment cases in which only the Identity Certificate is used and where a BS can be configured to allow EUD level access to UTs authenticated only via an Identity Certificate.

The *HC-SDMA* automatic certificate download feature shall use native TCP/IP to establish an LDAP binding between the UT and an LDAP management directory server. The LDAP management directory server shall have a fixed host and domain name on the service providers network and shall be accessed by the UT as LDAPSrvr@<service-provider-domain>. The certificate shall be downloaded using an LDAP search operation.

The *HC-SDMA* automatic certificate download feature enables the UT to automatically obtain a Service Authorization Certificate or a Network Access Certificate. In either case, the network operator's LDAP directory shall be configured to accept the LDAP binding of a UT. That can be accomplished by either manual directory server configuration or by automatic configuration through a service provisioning system.

In summary, when using the *HC-SDMA* automatic certificate download feature, the UT shall present either an Identity certificate or a Service Authorization Certificate to the BS. The BS shall authenticate the UT using the authentication procedure defined in clause 11.3, *HC-SDMA* Terminal Authentication Protocol, *i-TAP*. After a successful UT authentication, the BS shall verify that the link-level Network Access Identifier (NAI) used by the UT to gain network access at the Network Access Server (NAS) is of the predefined form <UT-Serial-Number>@<manufacturer-domain> where the UT serial number and the manufacturer domain correspond to the information presented in the certificate. Figure 114 - UT Automatic Certificate Download illustrates automatic certificate download.

The message exchange in *Figure 11.3* assumes that the PSS is a PDSN and that the LDAP directory server is not colocated with the PDSN. If the PDSN is colocated with the LDAP directory server it can act as an LNS as well and therefore the PDSN and the LNS may be the same physical entity.

A person wishing to subscribe to an *HC-SDMA* based service offered by a service provider should call the service provider. The service provider should use the customer care system to create a new subscriber record that should be either manually or automatically downloaded to the LDAP directory server used by the service provider. Alternatively, the new subscriber may access a web site that interfaces the customer self-care system and initiates the service provisioning.

The information provisioned into the LDAP directory shall include the UT make and serial number. The provisioning procedure shall also include the generation of the certificate and public/private key as defined in Figure 114 - UT Automatic Certificate Download.

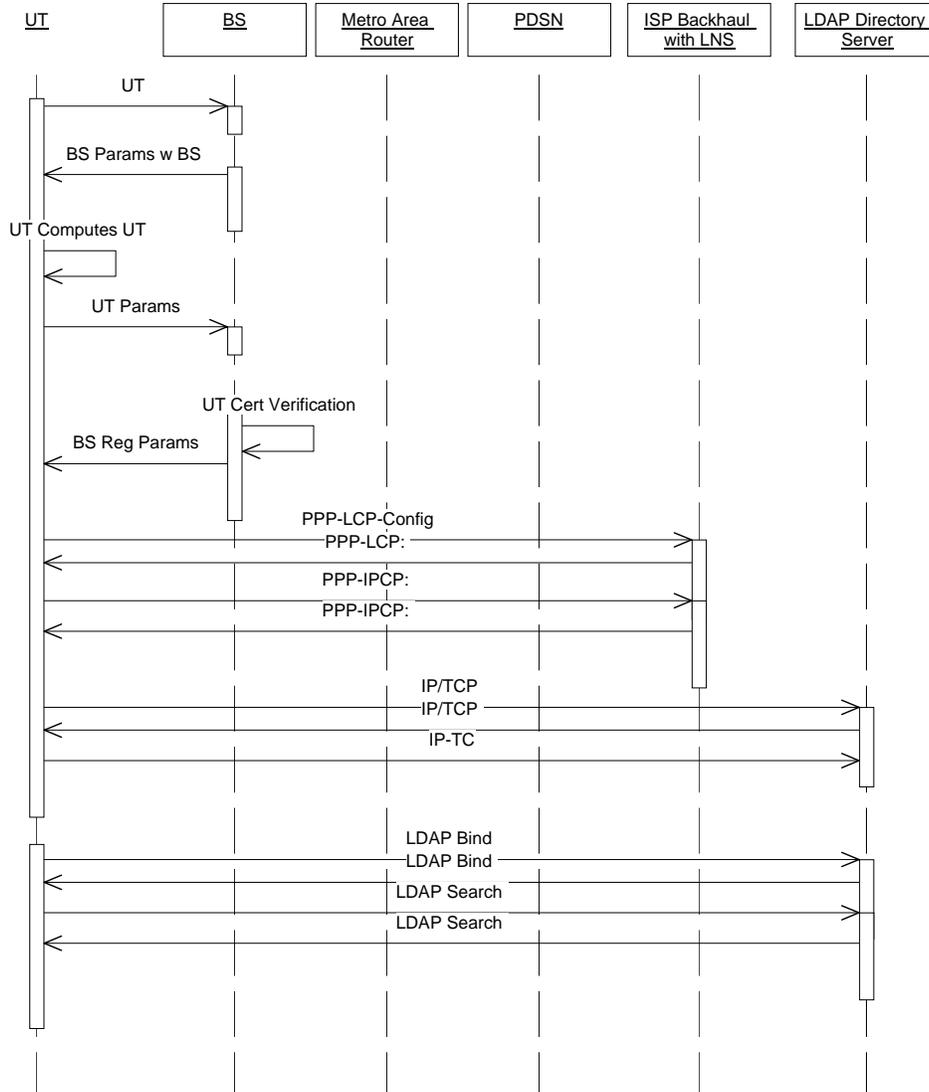


Figure 114: UT Automatic Certificate Download

Once the UT attempts access to the service providers network, the Network Access Server (NAS) -- i.e., the LNS -- shall access a co-located RADIUS server which has access to the LDAP directory information. Once the UT has gained access to the LDAP directory, it shall get provisioned with the certificate and possibly other service and operator specific provisioning data, for example, a BS CRL.

12 PERFORMANCE REQUIREMENTS AND TEST PROCEDURES

12.1 Introduction

This clause contains the performance specifications for various transmission channels in *HC-SDMA*. In addition, this clause specifies the input operating conditions and the measurement procedures under which the specifications are to be met.

Clause 12.2, Frequency (F) Synchronization Channel Receiver Requirements describes the performance requirements for the Frequency Synchronization (F) channel. The requirements for the F channel are specified for both AWGN and fading channels.

Clauses 12.3, Timing (T) Channel Receiver Requirements, and 12.4, Broadcast (B) Channel Receiver Requirements, describe the performance specification for the Timing (T) and Broadcast (B) channels in AWGN and fading channels, respectively.

Clause 12.5, BCH Acquisition Requirements describes specifications for BCH acquisition. BCH acquisition involves acquiring the F burst, followed by the T burst and then the B burst. The time taken to acquire each of the broadcast bursts with an associated detect rate is specified.

Clause 12.6, Paging Channel Receiver Requirements describes paging channel performance specifications. The rate of detection for different slot-diversity schemes are specified for both AWGN and fading channels.

Clause 12.7, Configuration Channel Receiver Requirements presents the FER specifications for the Configuration Channel (CCH).

Clauses 12.8, Access Channel (RACH) Receiver Requirements, and 12.9, Traffic Channel Receiver Requirements, present the FER specifications for the Access (RACH) and Traffic (TCH) Channels, respectively.

Clause 12.10, Fast Associated Control Channel Receiver Requirements, presents the Decode Error Rate (DER) specifications for the Fast Associated Control Channel (FACCH).

12.1.1 Definitions

12.1.1.1 input signal power: Input signal power shall be measured at the antenna connector, and the measurement shall be within the channel bandwidth centered around the specified carrier frequency. Averaging shall take place between the start of the first useful symbol and the end of the last useful symbol of a slot (ramp-up, ramp-down, and guard symbols excluded). Unless otherwise specified, the measurement procedure requires a root-raised cosine filter with roll-off factor $\alpha = 0.25$ and symbol period $T_s = 2 \mu\text{sec}$.

12.1.1.2 maximum doppler frequency: Maximum doppler frequency is defined as the maximum frequency component in the Doppler spectrum which characterizes the fading channel. The fading envelope is characterized by a Rayleigh distribution, while the power spectral density is defined by Clark's spectrum. The maximum Doppler frequency is defined in Hz for an operating carrier frequency of 1.9 GHz.

12.2 Frequency (F) Synchronization Channel Receiver Requirements

This section describes the tests to validate the frequency synchronization channel receiver performance under AWGN, fading, and noise-only conditions.

12.2.1 Performance in Additive White Gaussian Noise

12.2.1.1 Definition

The *frequency synchronization channel receiver performance* under the AWGN condition is determined by the detection rate and the distribution of the frequency and timing estimation errors at the specified input signal power. *Detection rate* is defined as the rate at which the presence of an F burst is detected at

the specified input signal power. *Frequency estimation error* is defined as the difference between the known frequency and the estimated frequency offset of the tones. *Timing estimation error* is defined as the difference between the estimated frame timing and a known reference timing.

12.2.1.2 Performance Specifications

12.2.1.2.1 Input Specifications⁴⁰: The input frequency offset shall be within ± 10 KHz and the input signal power shall be -116 dBm.

12.2.1.2.2 Output Specifications: The detection rate shall be greater than 0.85. The allowed frequency offset estimation error shall be as specified in Table 108 - F Burst Frequency Offset Specification in an AWGN Channel

Table 108: F Burst Frequency Offset Specification in an AWGN Channel

Percentile	Frequency Offset Estimation Error (Hz)
99	≤ 1000

The allowed timing alignment estimation error shall be as specified in Table 109 - F Burst Timing Alignment Specification in an AWGN Channel below.

Table 109: F Burst Timing Alignment Specification in an AWGN Channel

Percentile	Timing Alignment Estimation Error (μ s)
95	≤ 64

12.2.1.3 Measurement Procedure

1. Connect the UT under test and signal generators as shown in Figure 115 – User Terminal Frequency Synchronization Channel Receiver Test Setup for the AWGN Condition.
2. Synchronize the signal generator and the UT so the frequency offset between the transmitter and the receiver is zero.
3. Configure the signal generator such that it generates two frequency synchronization tones.
4. Detune the signal generator so that there is a frequency offset, less than 10 kHz, between the UT and signal generator.
5. Adjust the signal generator so the signal power at the input of the receiver is -116 dBm. Then measure the detection rate, frequency offset, and time alignment estimation errors for at least 10,000 bursts. The timing alignment error is computed using an accurate knowledge of the signal generator’s frame timing at the user terminal.

⁴⁰ *General guideline:* Throughout the document, most the input specifications are presented as a range. The measurement procedure must span the entire input specifications range. Bursts stored in a signal generator or generated by a base station must adhere to these input specifications. In case of timing and frequency offsets introduced at the input, the generated burst should have random timing and frequency offsets within the specified range and it should change from burst to burst.

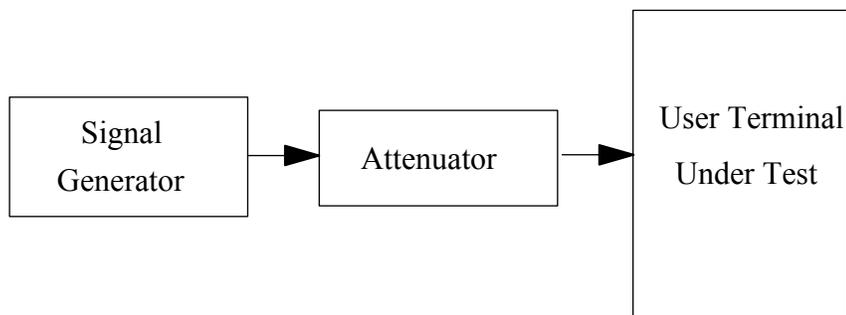


Figure 115: User Terminal Frequency Synchronization Channel Receiver Test Setup for the AWGN Condition

12.2.2 Performance in Fading

12.2.2.1 Definition

The frequency synchronization channel receiver performance under fading conditions is determined by the detection rate and the distributions of the frequency offset estimation error and timing estimation error at the specified input signal power.

12.2.2.2 Performance Specifications

12.2.2.2.1 Input Specifications

The input frequency offset shall be within ± 10 KHz and the input signal power shall be -114 dBm⁴¹. The maximum Doppler frequency shall be set to 30 Hz.

12.2.2.2.2 Output Specifications

The detection rate shall be greater than 0.88. The allowed frequency offset estimation error shall be as specified in Table 110 - F Burst Frequency Offset Specification in a Fading Channel.

Table 110: F Burst Frequency Offset Specification in a Fading Channel

Percentile	Frequency Offset Estimation Error (Hz)
99	≤ 1000

The allowed timing alignment estimation error shall be as specified in Table 111- F Burst Timing Alignment Specification in a Fading Channel.

Table 111: F Burst Timing Alignment Specification in a Fading Channel

Percentile	Timing Alignment Estimation Error (s)
95	≤ 64

⁴¹ This specified signal power is the long term average of the signal power across thousands of frames. This is true for all fading specifications in this document.

12.2.2.3 Measurement Procedure

1. Connect the UT under test, fading generator and signal generators as shown in Figure 116 - User Terminal Frequency Synchronization Channel Receiver Test Setup for Fading Condition.
2. Synchronize the signal generators and the UT so the frequency offset between the transmitter and the receiver is zero.
3. Configure the signal generators such that each of them generates one of the two frequency synchronization tones.
4. Configure the fading generator to generate independent fading for a maximum Doppler frequency of 30 Hz.
5. Detune both the signal generators so that there is a frequency offset, less than 10 kHz, between the UT and signal generators.
6. Adjust the signal generator so the average signal power at the input of the receiver is -114 dBm. Then measure the detection rate, frequency offset and timing alignment estimation errors using at least 10,000 bursts.

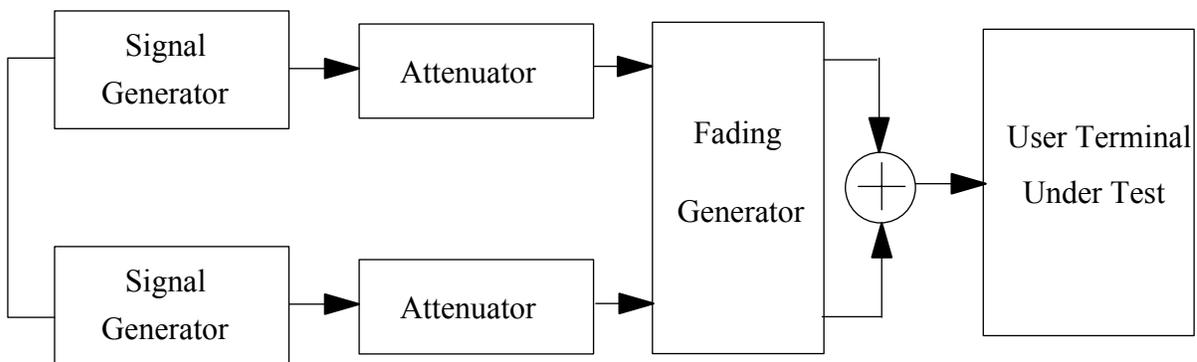


Figure 116: User Terminal Frequency Synchronization Channel Receiver Test Setup for Fading Condition

12.2.3 Performance in Noise-only Condition

12.2.3.1 Definition

The frequency synchronization channel receiver performance under noise-only conditions is determined by the false alarm rate. False alarm rate is defined as the percentage of time the detector falsely detects the presence of a burst when the burst is actually not present.

12.2.3.2 Performance Specifications

12.2.3.2.1 Input Specifications

There shall be no input signal in this test procedure.

12.2.3.2.2 Output Specifications

The false alarm rate shall be less than 0.0055.

12.2.3.3 Measurement Procedure

1. Turn off the signal and fading generator in the set up shown in Figure 116 - User Terminal Frequency Synchronization Channel Receiver Test Setup for Fading Condition.
2. Measure the false alarm rate using at least 100,000 bursts.

12.3 Timing (T) Channel Receiver Requirements

This section describes the tests to validate the timing channel receiver performance under AWGN and fading conditions.

12.3.1 Performance in Additive White Gaussian Noise

12.3.1.1 Definition

The timing channel receiver performance under the AWGN condition is determined by the decode rate and timing estimation error at the specified input signal power. Decode rate is a measure of correctly decoding the color code of the BS.

12.3.1.2 Performance Specifications

12.3.1.2.1 Input Specifications

The input frequency offset shall be within ± 0.9 KHz, the input timing misalignment shall be within $\pm 64 \mu s$ and the input signal power at the UT shall be -119 dBm.

12.3.1.2.2 Output Specifications

The decode rate shall be as specified in Table 112 - T Burst Decode Rate Specification in an AWGN Channel.

Table 112: T Burst Decode Rate Specification in an AWGN Channel

Input Signal Power (dBm)	Decode Rate
- 119	≥ 0.97

The allowed timing alignment estimation error at an input signal power of -119 dBm shall be as specified in Table 113 - T Burst Timing Alignment Specification.

Table 113: T Burst Timing Alignment Specification in an AWGN Channel

Percentile	Timing Alignment Estimation Error (μs)
99	≤ 2

12.3.1.3 Measurement Procedure

1. Connect the UT under test and a BS with multiple antennas as shown in Figure 117 - User Terminal and Base Station Test Setup for the AWGN Condition.
2. Synchronize the BS and the UT so that the frequency offset between the transmitter and the receiver is within ± 0.9 KHz and the timing misalignment is within ± 64 μ s.
3. Configure the BS to transmit T burst corresponding to any BSCC. Also the transmit weights in the base station are different for each repetition.
4. Adjust the BS transmit power to maintain the desired signal power at the input of the receiver and then measure the detection rate. Also measure the timing alignment error for at least 10,000 bursts.

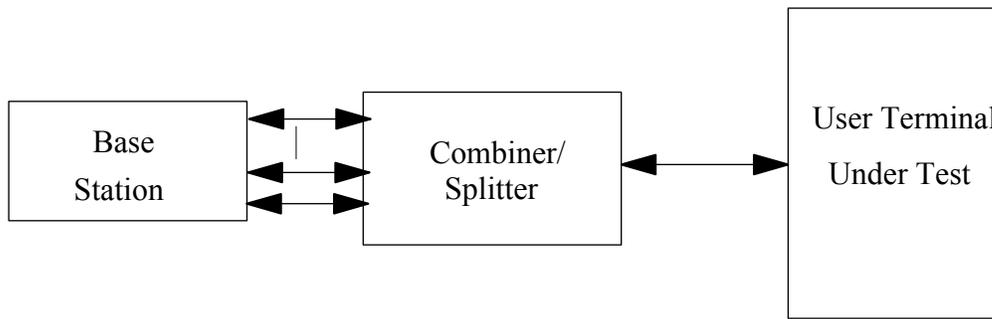


Figure 117: User Terminal and Base Station Test Setup for the AWGN Condition

12.3.2 Performance in Fading

12.3.2.1 Definition

The timing channel receiver performance under fading condition is determined by the decode rate and timing estimation error at the specified input signal power.

12.3.2.2 Performance Specifications

12.3.2.2.1 Input Specifications

The input frequency offset shall be within ± 0.9 KHz, input timing misalignment shall be within ± 64 μ s and the input signal power shall be -114 dBm. The maximum Doppler frequency shall be set to 30 Hz.

12.3.2.2.2 Output Specifications

The decode rate shall be as specified in Table 114 - T Burst Decode Rate Specification in a Fading Channel.

Table 114: T Burst Decode Rate Specification in a Fading Channel

Input Signal Power (dBm)	Decode Rate
-114	≥ 0.96

The allowed timing alignment estimation error shall be as specified in Table 115 - T Burst Timing Alignment Specification in a Fading Channel.

Table 115: T Burst Timing Alignment Specification in a Fading Channel

Percentile	Timing Alignment Estimation Error (μs)
95	≤ 2

12.3.2.3 Measurement Procedure

1. Connect the UT under test and a BS with multiple antennas as shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. Synchronize the BS and the UT so that the frequency offset between the transmitter and the receiver is within ± 0.9 KHz and the timing misalignment is within ± 64 μs .
3. Configure the BS to transmit T burst corresponding to any BSCC. Also the transmit weights in the base station are different for each repetition.
4. Configure the fading generator to generate independent fading for a maximum Doppler frequency of 30 Hz.
5. Adjust the transmit signal power to maintain a desired signal power at the input of the receiver and then measure the detection rate. Also measure the timing alignment error using at least 10,000 bursts.

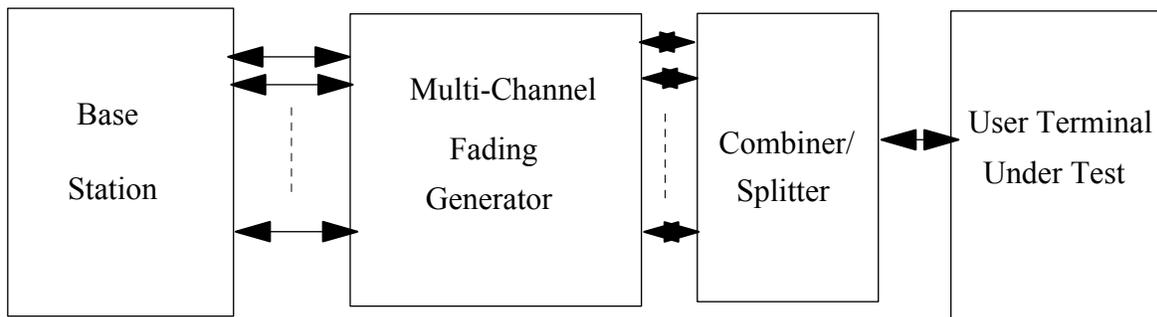


Figure 118: User Terminal and Base Station Test Setup for Fading Condition

12.4 Broadcast (B) Channel Receiver Requirements

This section describes the tests to validate the broadcast channel receiver performance under AWGN, fading, and noise-only conditions.

12.4.1 Performance in Additive White Gaussian Noise

12.4.1.1 Definition

The broadcast channel receiver performance under the AWGN condition is determined by the detection rate, timing estimation error and frequency estimation error at the specified input signal power. Detection rate is quantified by correctly detecting the presence of the burst and decoding the transmitted B burst information bits.

12.4.1.2 Performance Specifications

12.4.1.2.1 Input Specifications

The input frequency offset shall be within ± 0.7 KHz, the input timing misalignment shall be within ± 4 μ s and the input signal power shall be -117 dBm.

12.4.1.2.2 Output Specifications

The detection rate shall be as specified in Table 116 - B Burst Decode Rate Specification in an AWGN Channel.

Table 116: B Burst Decode Rate Specification in an AWGN Channel

Input Signal Power (dBm)	Detection Rate
- 117	≥ 0.97

The allowed timing alignment estimation error shall be as specified in Table 117 - B Burst Timing Alignment Specification in an AWGN Channel.

Table 117: B Burst Timing Alignment Specification in an AWGN Channel

Percentile	Timing Alignment Estimation Error (μ s)
95	≤ 1

The allowed frequency offset estimation error shall be as specified in Table 118 - B Burst Frequency Offset Specification in an AWGN Channel.

Table 118: B Burst Frequency Offset Specification in an AWGN Channel

Percentile	Frequency Offset Estimation Error (Hz)
90	≤ 200

12.4.1.3 Measurement Procedure

1. Connect the UT under test and a BS with multiple antennas as shown in Figure 117 - User Terminal and Base Station Test Setup for the AWGN Condition.

2. Synchronize the BS and the UT so that the frequency offset between the transmitter and the receiver is within ± 0.9 KHz and the timing misalignment is within ± 4 μ s.
3. Configure the BS to transmit the B burst scrambled by any BSCC and carry one of the 64 possible information sequences. Also the transmit weights in the base station are different for each repetition.
4. Adjust the signal power to maintain a desired signal power at the input of the receiver and then measure the detection rate. Also measure the timing alignment and frequency offset errors using at least 10,000 bursts.

12.4.2 Performance in Fading

12.4.2.1 Definition

The broadcast channel receiver performance under fading conditions is determined by the detection rate, timing estimation error and frequency estimation error at the specified input signal power. Detection rate is quantified by correctly detecting the presence of the burst and decoding the transmitted B-burst information bits.

12.4.2.2 Performance Specifications

12.4.2.2.1 Input Specifications

The input frequency offset shall be as within ± 0.7 KHz, the input timing misalignment shall be as within ± 4 μ s, and the input signal power shall be as -114 dBm. The maximum Doppler frequency shall be set to 30 Hz.

12.4.2.2.2 Output Specifications

The detection rate shall be as specified in Table 119 - B Burst Decode Rate Specification in a Fading Channel.

Table 119: B Burst Decode Rate Specification in a Fading Channel

Input Signal Power (dBm)	Detection Rate
- 114	≥ 0.96

The allowed timing alignment estimation error shall be as specified in Table 120 - B Burst Timing Alignment Specification in a Fading Channel.

Table 120: B Burst Timing Alignment Specification in a Fading Channel

Percentile	Timing Alignment Estimation Error (μ s)
90	≤ 1

The allowed frequency offset estimation error shall be as specified in Table 121 - B Burst Frequency Offset Specification in a Fading Channel.

Table 121: B Burst Frequency Offset Specification in a Fading Channel

Percentile	Frequency Offset Estimation Error (Hz)
90	≤ 200

12.4.2.3 Measurement Procedure

1. Connect the UT under test, fading generator and a BS with multiple antenna signal as shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. Synchronize the BS and the UT so that the frequency offset between the transmitter and the receiver is within ± 0.9 KHz and the timing misalignment is within ± 4 μ s.
3. Configure the BS to transmit the B burst scrambled by any BSCC and carry one of 64 possible information sequences. Also the transmit weights in the base station are different for each repetition.
4. Configure the fading generator to generate independent fading for a maximum Doppler frequency of 30 Hz.
5. Adjust the transmit signal power to maintain a desired signal power at the input of the receiver and then measure the detection rate. Also measure the timing alignment and frequency offset errors using at least 10,000 bursts.

12.4.3 Performance in Noise-only Condition**12.4.3.1 Definition**

The broadcast channel receiver performance under noise-only condition is determined by the false alarm rate. False alarm rate is defined as the percentage of time the detector falsely detects the presence of a burst when the burst is actually not present.

12.4.3.2 Performance Specifications**12.4.3.2.1 Input Specifications**

There shall be no input signal in this test procedure.

12.4.3.2.2 Output Specifications

The false alarm rate shall be less than 2×10^{-5} .

12.4.3.3 Measurement Procedure

1. Turn off the BS shown in the set up shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. Measure the false alarm rate using at least 10,000,000 bursts.

12.4.4 Performance in a Multi-basestation environment

12.4.4.1 Definition

The broadcast channel receiver performance in a multi base station environment is determined by the false detect rate. AWGN channel is assumed. False detect rate is defined as the rate with which the detector at the user terminal falsely detects the presence of a B burst from one BS, when it actually received a B burst from some other BS.

12.4.4.2 Performance Specifications

12.4.4.2.1 Input Specifications

The input signal power shall be in the range -120 to -70 dBm.

12.4.4.2.2 Output Specifications

The false detect rate shall be less than 2×10^{-5} .

12.4.4.3 Measurement Procedure

1. Use the set up shown in Figure 117 - User Terminal and Base Station Test Setup for the AWGN Condition.
2. Configure the base station to transmit a B burst with a specific BSCC (e.g., BSCC = 1) and adhering to the input specifications.
3. The user terminal is configured to receive the B burst with a different BSCC (e.g., BSCC = 0).
4. Measure the rate of falsely detecting the B burst using at least 10,000,000 bursts.
5. It is recommended to randomly choose BSCC combinations for each measurement.

12.5 BCH Acquisition Requirements

This section describes the specifications for the time taken by the UT to acquire the broadcast channel. A UT acquires the F burst followed by T and B bursts.

12.5.1 F-Burst Acquisition

12.5.1.1 Definition

F-Burst acquisition is the time taken to successfully detect an F-burst and provide estimates for frequency and timing alignment errors within a specified limit at a specified input signal power and input frequency offset in a fading channel.

12.5.1.2 Specifications

The F-burst shall be detected with 90% probability within 205 ms from the beginning of the search process at an input signal power of -114 dBm, when the input frequency offset is less than ± 10 KHz, and when the maximum Doppler frequency is set to 30 Hz. The estimated frequency offset error from the F-burst at the end of this time interval shall be within ± 200 Hz and the timing alignment error shall be within ± 64 μ s with 90% probability.

12.5.2 T-Burst Acquisition

12.5.2.1 Definition

T-Burst acquisition is the time taken to successfully decode a T-burst, after an F-burst is detected, and provide timing alignment error within a specified limit at a specified input signal power in a fading channel.

12.5.2.2 Specifications

The T-burst shall be detected with 95% probability within 10 ms after an F-burst is detected at an input signal power of -114 dBm. The maximum Doppler frequency shall be set to 30 Hz. The timing alignment error at the end of this time interval shall be within $\pm 1 \mu\text{s}$ with 90% probability.

12.5.3 B-Burst Acquisition

12.5.3.1 Definition

B-Burst acquisition is the time taken to successfully decode a B-burst after a T-burst is detected, and provide frequency and timing alignment errors within a specified limit at a specified input signal power in a fading channel.

12.5.3.2 Specifications

The B-burst shall be detected with 95% probability within 80 ms, after a T-burst is detected at an input signal power of -114 dBm. The maximum Doppler frequency shall be set to 30 Hz. The frequency offset error at the end of this time interval shall be within $\pm 100 \text{ Hz}$ and the timing alignment error shall be within $\pm 1 \mu\text{s}$ with 90% probability.

12.5.4 Measurement Procedure

1. Connect the UT, BS and the fading generator as shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. The BS is configured to transmit all the broadcast bursts (F, T, & B) according to the super frame format for all possible BSCC configurations. Also the transmit weights in the base station are different for each repetition in T and B bursts.
3. The UT is configured to receive all the BCH bursts in sequence.
4. Configure the fading generator to generate independent fading for a maximum Doppler frequency of 30 Hz.
5. Adjust the transmitter power level to set a specified signal power at the input of the receiver.
6. Measure the time taken to acquire each broadcast burst, rate of detection and the corresponding frequency offset and/or timing alignment errors.

12.5.5 Overall BCH Acquisition

12.5.5.1 Definition

Overall BCH acquisition is defined as the rate with which the BS color code (BSCC) corresponding to the selected BS is detected correctly and the information bits (BS load and BCH power) are correctly decoded for the chosen BSCC. Acquisition must be achieved within a specified time for a specified

input signal power in an AWGN or fading channel when the input frequency offset between the UT and the transmitter is within a specified range.

12.5.5.2 Specifications for AWGN Channel

The input frequency offset error shall be within ± 10 KHz. The allowed time to acquire and the acquisition rate shall be as specified in Table 122 - BCH Acquisition Specifications in an AWGN Channel.

Table 122: BCH Acquisition Specifications in an AWGN Channel

Input Signal Power (dBm)	Time to Acquire (ms)	Acquire Rate
-116	≤ 400	≥ 0.99
-118	≤ 400	≥ 0.95

12.5.5.3 Measurement Procedure

1. Connect a UT and a BS as shown in the set up shown in Figure 117 - User Terminal and Base Station Test Setup for the AWGN Condition.
2. BS is configured to transmit the superframe structure with F, T, and B bursts continuously.
3. The frequency offset between the UT and the BS shall be within ± 10 KHz.
4. Adjust the input signal power to the above specification and measure the acquire rate and time to acquire.

12.5.5.4 Specifications for Fading Channel

The input frequency offset shall be less than or equal to ± 10 KHz. The channel shall be uncorrelated Rayleigh fading across the antennas with a maximum Doppler frequency of 30 Hz. The allowed time to acquire and the acquisition rate shall be as specified in Table 123 - BCH Acquisition Specifications in a Fading Channel.

Table 123: BCH Acquisition Specifications in a Fading Channel

Input Signal Power (dBm)	Time to Acquire (ms)	Acquire Rate
-114	≤ 400	≥ 0.99
-116	≤ 400	≥ 0.85

12.5.5.5 Measurement Procedure

1. Connect a BS with multiple antennas to the UT through fading generator as shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. BS is configured to transmit the superframe structure with F, T, and B bursts continuously. Also the transmit weights in the base station are different for each repetition in the T and B bursts.
3. The frequency offset between the UT and the BS shall be within ± 10 KHz.
4. Fading generator is set to 30 Hz maximum Doppler frequency and uncorrelated fading across the antennas.

5. Adjust the input signal power at the UT to the above spec and measure the acquire rate and time to acquire.

12.5.6 Overall False Alarm Rate Specification

12.5.6.1 Definition

False Alarm Rate is defined as the rate with which the detector incorrectly acquires a BCH when the BCH is not present.

12.5.6.2 Specification

False alarm rate shall be less than 2×10^{-5} .

12.5.6.3 Measurement Procedure

Same procedure as in clause 12.4.3.3, Measurement Procedure.

12.6 Paging Channel Receiver Requirements

This section describes the tests to validate the paging channel receiver performance under AWGN, fading, and noise-only conditions. The paging channel has options to use one-slot or two-slot diversity schemes. The receiver shall be able to receive a page listening to just one downlink frame.

12.6.1 Performance in Additive White Gaussian Noise

12.6.1.1 Definition

The paging channel receiver performance under the AWGN condition is determined by the detection rate at the specified input signal power.

12.6.1.2 Performance Specifications

12.6.1.2.1 Input Specifications

The input frequency offset shall be within ± 400 Hz, the input timing misalignment shall be within $\pm 8 \mu s$, and the input signal power shall be either -117 dBm (one-slot) or -118 dBm (two-slot).

12.6.1.2.2 Output Specifications

The detection rate for one-slot and two-slot diversity schemes and the corresponding input signal powers shall be as specified in Table 124 - Paging Specifications in an AWGN Channel.

Table 124: Paging Specifications in an AWGN Channel

Detection Rate	Input Signal Power (dBm)	
	One-Slot	Two-Slot
≥ 0.92	-117	-118

12.6.1.3 Measurement Procedure

1. Connect the UT under test to a signal generator through an attenuator as shown in Figure 119 - User Terminal Test Setup.
2. Configure the signal generator to transmit paging bursts with one-slot or two-slot diversity schemes.
3. Adjust the signal generator to set the desired signal power at the input of the receiver and measure the detection rate using at least 10,000 bursts.

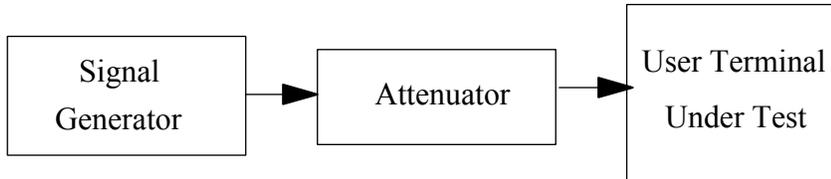


Figure 119: User Terminal Test Setup

12.6.2 Performance in Fading

12.6.2.1 Definition

The paging channel receiver performance under fading conditions is determined by the detection rate at the specified input signal power.

12.6.2.2 Performance Specifications

12.6.2.2.1 Input Specifications

The input frequency offset shall be within ± 400 Hz, the input timing misalignment shall be within $\pm 8 \mu s$, and the input signal power shall be either -109 dBm (one-slot) or -115 dBm (two-slot). The maximum Doppler frequency shall be set to 30 Hz.

12.6.2.2.2 Output Specifications

The detection rate for one-slot and two-slot diversity schemes and the corresponding input signal powers shall be as specified in Table 125 - Paging Specifications in a Fading Channel.

Table 125: Paging Specifications in a Fading Channel

Detection Rate	Input Signal Power (dBm)	
	One-Slot	Two-Slot
≥ 0.90	-109	-115

12.6.2.3 Measurement Procedure

1. Connect the UT under test, fading generator and the BS as shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. Configure the BS to transmit paging bursts with one-slot or two-slot diversity schemes.

3. Configure the fading generator to generate independent fading for a maximum Doppler frequency of 30 Hz.
4. Adjust the BS transmit power so that the desired signal power is maintained at the input of the UT receiver and measure the detection rate using at least 10,000 bursts.

12.6.3 Performance in Noise-only Condition

12.6.3.1 Definition

The paging channel receiver performance under noise-only condition is determined by the false alarm rate.

12.6.3.2 Performance Specifications

12.6.3.2.1 Input Specifications

There shall be no input signal in this test procedure.

12.6.3.2.2 Output Specifications

The false alarm rate shall be less than 2×10^{-5} .

12.6.3.3 Measurement Procedure

1. Turn off the BS in the set up shown in Figure 118 - User Terminal and Base Station Test Setup for Fading Condition.
2. Measure the false alarm rate using at least 10,000,000 bursts.

12.6.4 Performance under Multiuser Paging

12.6.4.1 Definition

The paging channel receiver performance under multiuser paging is determined by the false detect rate. AWGN channel is assumed. False detect rate is defined as the rate with which the detector at the user terminal falsely detects the presence of a page, when it receives a page intended for another user. The unintended page can be in the same subslot or in the adjacent subslot of the same time slot or in both subslots.

12.6.4.2 Performance Specifications

12.6.4.2.1 Input Specifications

The input signal power shall be in the range -120 to -70 dBm.

12.6.4.2.2 Output Specifications

The false detect rate shall be less than 2×10^{-5} .

12.6.4.3 Measurement Procedure

1. Use the set up shown in Figure 119 - User Terminal Test Setup.

2. Configure the signal generator to send a page with a specific PID and adhering to the input specifications.
3. The user terminal is configured to receive the page burst with a different PID.
4. Measure the rate of falsely detecting the page burst using at least 10,000,000 bursts.
5. It is recommended to randomly choose a PID for each measurement.

12.7 Configuration Channel Receiver Requirements

This section describes the tests to validate the configuration channel receiver performance under the AWGN condition. Configuration channel carries two types of bursts: one in the uplink and one in the downlink. The uplink burst is the Configuration Request (CR) burst and the downlink burst is the Configuration Message (CM) burst. The CR burst has a special structure defined in clause 3, Slot and Frame Structure. The CM burst is a standard downlink burst.

12.7.1 Performance in Additive White Gaussian Noise

12.7.1.1 Definition

The configuration channel receiver performance under the AWGN condition is determined by the frame error rate (FER) at the specified input signal power.

12.7.1.2 Uplink Performance Specifications (CR Burst)

12.7.1.2.1 Input Specifications

The input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be between 0 and 110 μ s.

12.7.1.2.2 Output Specifications

The measured uplink frame error rates (FER)⁴² and the corresponding input signal powers shall be as specified in Table 126 - CR Sensitivity Specifications.

Table 126: CR Sensitivity Specifications

Input Signal Power (dBm)	FER
-110.2	≤ 0.100
-109.0	≤ 0.010
-108.1	≤ 0.001

⁴² FER is defined as the ratio of number of frames in error to the total number of received frames.

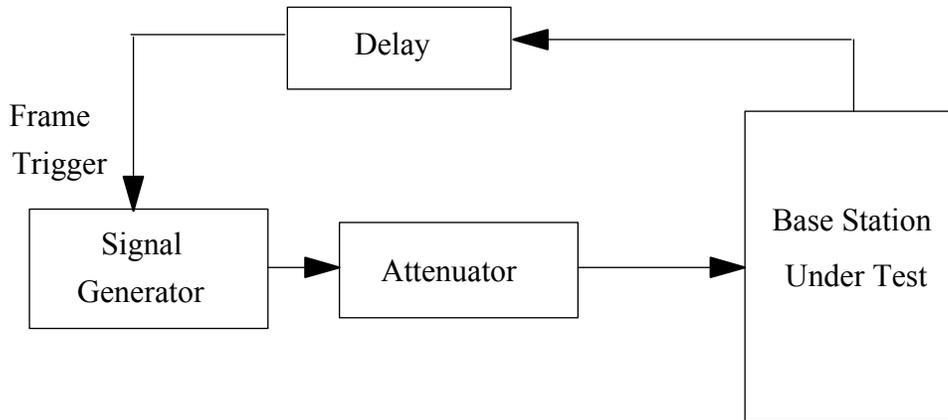


Figure 120: Base Station Test Setup

12.7.1.3 Uplink Measurement Procedure

1. Configure the BS under test to function in single-antenna mode.
2. Connect the BS under test and a signal generator as shown in Figure 120 - Base Station Test Setup.
3. Connect a delay box to the trigger from the BS to the signal generator to synchronize the frame timing.
4. The BS is configured to receive CR bursts.
5. Configure the signal generator to transmit CR bursts with a frequency offset satisfying the specifications. Introduce delay to the received signal according to the specification. The delay in the received signal is introduced by suitably delaying the frame trigger from the BS.
6. Adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 126 - CR Sensitivity Specifications above and then measure the corresponding FER using at least 1,000,000 bursts.

12.7.1.4 Downlink Performance Specifications (CM Burst)

12.7.1.4.1 Input Specifications

The input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within ± 4 μ s.

12.7.1.4.2 Output Specifications

The measured downlink frame error rates and the corresponding input signal powers shall be as specified in Table 127 - CM Sensitivity Specifications.

Table 127: CM Sensitivity Specifications

Input Signal Power (dBm)	FER
-111.0	≤ 0.100
-110.0	≤ 0.010
-109.2	≤ 0.001

12.7.1.5 Downlink Measurement Procedure

1. Connect the UT under test and a signal generator as shown in Figure 119 - User Terminal Test Setup.
2. Set the UT to receive CM bursts.
3. Configure the signal generator to transmit CM bursts with the frequency offset and timing misalignment introduced at the input satisfying the input specifications.
4. Adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 127 - CM Sensitivity Specifications above and then measure the corresponding FER using at least 1,000,000 bursts.

12.8 Access Channel (RACH) Receiver Requirements

This section describes the receiver performance tests to validate the uplink and downlink access channels under the AWGN condition. The uplink burst is the Request Access (RA) burst and the downlink burst is the Access Assignment (AA) burst. RA and AA bursts are standard uplink and downlink bursts, respectively.

12.8.1 Performance in Additive White Gaussian Noise**12.8.1.1 Definition**

The access channel receiver performance under the AWGN condition is determined by the frame error rate (FER) at the specified input signal power.

12.8.1.2 Uplink Performance Specifications (RA Burst)**12.8.1.2.1 Input Specifications**

For base stations supporting ranges upto 12.75 km, the input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within ± 4 μ s.

For base stations supporting ranges upto 15.0 km, the input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be between -4 and $+ 19$ μ s.

12.8.1.2.2 Output Specifications

The measured uplink frame error rates and the corresponding input signal powers shall be as specified in Table 128 - RA Sensitivity Specifications.

Table 128: RA Sensitivity Specifications

Input Signal Power (dBm)	FER
-112.9	≤ 0.100
-111.4	≤ 0.010
-110.6	≤ 0.001

12.8.1.3 Uplink Measurement Procedure

1. Configure the BS under test to function in single-antenna mode.
2. Connect the BS under test and a signal generator as shown in Figure 120 - Base Station Test Setup.
3. Set the BS to receive RA bursts.
4. Configure the signal generator to transmit RA bursts with the input frequency and timing misalignment satisfying the input specifications.
5. Adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 128 - RA Sensitivity Specifications above and then measure the corresponding FER using at least 1,000,000 bursts.

12.8.1.4 Downlink Performance Specifications (AA Burst)**12.8.1.4.1 Input Specifications**

The input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within ± 4 μ s.

12.8.1.4.2 Output Specifications

The measured downlink frame error rates and the corresponding input signal powers shall be as specified in Table 129 - AA Sensitivity Specifications.

Table 129: AA Sensitivity Specifications

Input Signal Power (dBm)	FER
-111.0	≤ 0.100
-110.0	≤ 0.010
-109.2	≤ 0.001

12.8.1.5 Downlink Measurement Procedure

1. Connect the UT under test and a signal generator as shown in Figure 119 - User Terminal Test Setup.
2. Set the UT to receive AA bursts.
3. Configure the signal generator to transmit AA bursts with the input frequency offset and timing misalignment satisfying the input specifications.

4. Adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 129 - AA Sensitivity Specifications above and then measure the corresponding FER using at least 1,000,000 bursts.

12.9 Traffic Channel Receiver Requirements

This section describes the receiver performance tests to validate the uplink and downlink traffic channels under the AWGN condition.

12.9.1 Performance in Additive White Gaussian Noise

12.9.1.1 Definition

The traffic channel receiver performance under the AWGN condition is determined by the input signal power at a given frame error rate (FER).

12.9.1.2 Uplink Performance Specifications

12.9.1.2.1 Input Specifications

For base stations supporting ranges up to 12.75 km, the input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within ± 4 μ s.

For base stations supporting ranges up to 15.0 km, the input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be between -4 and $+ 19$ μ s.

12.9.1.2.2 Output Specifications

The input signal power for the uplink traffic on all mod classes at given FERs shall be less than the values specified in Table 130 - Uplink TCH Sensitivity Specifications.

Table 130: Uplink TCH Sensitivity Specifications

Mod Class	Input Power (dBm) at FER \leq 0.100	Input Power (dBm) at FER \leq 0.01	Input Power (dBm) at FER \leq 0.001
0	-109.6	-108.6	-107.8
1	-108.0	-107.0	-106.3
2	-106.3	-105.3	-104.6
3	-103.1	-102.4	-101.9
4	-100.9	-100.2	-99.6
5	-98.7	-97.9	-97.3
6	-96.7	-95.9	-95.4
7	-95.3	-94.6	-94.0

12.9.1.3 Uplink Measurement Procedure

1. Configure the base-station under test to function in single-antenna mode.

2. Connect the BS under test and a signal generator as shown in Figure 120 - Base Station Test Setup.
3. Set the BS to receive the correct modulation class.
4. Configure the signal generator to transmit traffic bursts of appropriate modulation class with the input frequency offset and timing misalignment satisfying the input specifications.
5. For each modulation class, adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 130 - Uplink TCH Sensitivity Specifications above and then measure the corresponding FER using at least 1,000,000 bursts.

12.9.1.4 Downlink Performance Specifications

12.9.1.4.1 Input Specifications

The input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within $\pm 4 \mu\text{s}$.

12.9.1.4.2 Output Specifications

The input signal power for the uplink traffic on all mod classes at given FERs shall be less than the values specified in Table 131 - Downlink TCH Sensitivity Specifications.

Table 131: Downlink TCH Sensitivity Specifications

Mod Class	Input Power (dBm) at FER ≤ 0.100	Input Power (dBm) at FER ≤ 0.01	Input Power (dBm) at FER ≤ 0.001
0	-108.5	-107.5	-106.7
1	-106.7	-105.7	-105.0
2	-105.2	-104.2	-103.5
3	-102.0	-101.3	-100.8
4	-99.8	-99.1	-98.5
5	-97.7	-96.9	-96.3
6	-95.6	-94.8	-94.3
7	-94.2	-93.5	-93.0
8	-92.3	-91.6	-91.2

12.9.1.5 Downlink Measurement Procedure

1. Connect the UT under test and a signal generator as shown in Figure 119 - User Terminal Test Setup.
2. Set the UT to receive the correct modulation class.
3. Configure the signal generator to transmit traffic bursts of appropriate modulation class with the input frequency offset and timing misalignment satisfying the input specifications.
4. For each modulation class, adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 131 - Downlink TCH Sensitivity Specifications above and then measure the corresponding FER using at least 1,000,000 bursts.

12.10 Fast Associated Control Channel Receiver Requirements

This section describes the tests to validate the fast access control channel in AWGN condition.

12.10.1 Performance in Additive White Gaussian Noise

12.10.1.1 Definition

The FACCH receiver performance under the AWGN condition is determined by the decode error rate (DER) at the specified input signal power. DER is defined as the rate of decoding the FACCH incorrectly.

12.10.1.2 Uplink Performance Specifications

12.10.1.2.1 Input Specifications

For base stations supporting ranges up to 12.75 km, the input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within ± 4 μ s.

For base stations supporting ranges up to 15 km, the input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be between -4 and $+ 19$ μ s.

12.10.1.2.2 Output Specifications

The decode error rates (DER)⁴³ shall be as specified in Table 132 - Uplink FACCH Sensitivity Specifications.

Table 132: Uplink FACCH Sensitivity Specifications

Input Signal Power (dBm)	DER
-111.3	≤ 0.100
-108.8	≤ 0.010
-107.4	≤ 0.001

12.10.1.3 Uplink Measurement Procedure

1. Configure the base-station under test to function in single-antenna mode.
2. Connect the BS under test and a signal generator as shown in Figure 120 - Base Station Test Setup.
3. Set the BS to receive uplink traffic burst and FACCH.
4. Configure the signal generator to transmit uplink traffic bursts with all possible FACCH messages with input frequency offset and timing misalignment satisfying the input specifications.
5. Adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 132 - Uplink FACCH Sensitivity Specifications above and then measure the corresponding DER using at least 1,000,000 bursts.

⁴³ DER is defined as the ratio of number of frames with the FACCH message in error to the total number of received frames.

12.10.1.4 Downlink Performance Specifications

12.10.1.4.1 Input Specifications

The input frequency offset range shall be within ± 400 Hz and the input timing misalignment shall be within ± 4 μ s.

12.10.1.4.2 Output Specifications

The decode error rates shall be as specified in Table 133 - Downlink FACCH Sensitivity Specifications.

Table 133: Downlink FACCH Sensitivity Specifications

Input Signal Power (dBm)	DER
-110.5	≤ 0.100
-108.0	≤ 0.010
-106.6	≤ 0.001

12.10.1.5 Downlink Measurement Procedure

1. Connect the UT under test and a signal generator as shown in Figure 119 - User Terminal Test Setup.
2. Set the UT to receive downlink traffic bursts and FACCH.
3. Configure the signal generator to transmit downlink traffic bursts with all possible FACCH messages with the input frequency offset and timing misalignment satisfying the input specifications.
4. Adjust the signal generator to maintain the specified signal power at the input of the receiver specified in Table 133 - Downlink FACCH Sensitivity Specifications above and then measure the corresponding DER using at least 1,000,000 bursts.

Annex A (informative)

A GLOSSARY

This Annex contains a glossary of technical terms used in the *HC-SDMA* Protocol Specification.

- A.1 Absolute Frame Number (AFN)** The representation of absolute time used by the *HC-SDMA* system.
- A.2 Allocation Carrier Number:** Identifier for a carrier in a basic network carrier allocation, where the lowest carrier is numbered zero. See also *Carrier Numbers*.
- A.3 Antenna Array Spatial Processing:** The technique of employing an antenna array together with signal processing to mitigate the effects of interference and fading.
- A.4 Authentication:** Verifying the identity of the originator from a received message.
- A.5 Base Station Color Code (BSCC):** An identification that distinguishes a BS from neighboring BSs.
- A.6 Basic HC-SDMA Network:** A contiguous set of TDD spectrum, up to 10 MHz, with an *HC-SDMA* control channel.
- A.7 Broadcast Carrier:** A particular carrier in a basic *HC-SDMA* network carrier allocation. A broadcast carrier provides basic network information via the broadcast (BCH) and control (CCH) channels.
- A.8 BS Carrier Set:** A subset of the basic network carrier allocation that is used by a particular BS.
- A.9 Bulk Encryption:** See *Symmetric Key Encryption*.
- A.10 Burst:** The unit of radio transmission at Layer 1 occupying part of a single timeslot duration.
- A.11 Byte Stream:** Sequence of bytes from L2 to L3; L3 regards data in a structured view (packets), but L2 regards data in a flat view (byte streams).
- A.12 CA:** See *Certificate Authority*.
- A.13 Carrier:** A radio frequency on which an *HC-SDMA* channel is carried.
- A.14 Carrier Allocation:** The set of carriers over which a basic *HC-SDMA* network operates.
- A.15 Carrier Numbers:** *HC-SDMA* carrier numbers are indexed from 0 to $N_f - 1$ with lower indices indicating lower frequencies. The carriers are spaced contiguously on 625 kHz intervals.
- A.16 Carrier Allocation:** The set of carriers over which a basic *HC-SDMA* network operates.
- A.17 Certificate Authority** A trusted entity that generates a *Digital Certificate* using its private key.
- A.18 Chosen Plaintext Attack:** The cryptanalyst not only has access to the ciphertext and the associated plaintext, but also chooses the plaintext that gets encrypted. More powerful than a *Known Plaintext Attack*.
- A.19 Cipher:** A mathematical function used for encryption and decryption.
- A.20 Cipherstream:** Pseudo-random sequence produced by a stream cipher.
- A.21 Ciphertext:** An encrypted message.

- A.22 Connection Control Block (CCB)**The state associated with an instance of a L3 CM state machine.
- A.23 Connection:** Another term for *session*.
- A.24 Conventional Channel:** A particular carrier and a pair of uplink and downlink timeslots.
- A.25 Conventional Channel ID**The conventional channel ID identifies a given conventional channel by a canonical mapping to frequency carriers and timeslots. See Table 75: Mapping of Frequency and Timeslots to Conventional Channels.
- A.26 CRC:** A Cyclic Redundancy Check (CRC) is an error-control code that is used to detect (not correct) errors.
- A.27 Cryptanalysis:** The art and science of breaking ciphertext.
- A.28 Decryption:** The process of converting ciphertext into plaintext.
- A.29 Digital Certificate:** A message appended by the digital signature of the certificate authority.
- A.30 Digital Signature:** Data that vouches for the origin and integrity of a message. The originator (see *Certificate Authority*) uses its *Private Key* to sign the message and the recipient uses the corresponding *Public Key* to verify the origin and integrity of the message.
- A.31 Downlink:** Part or all of the communication path from the back-end network towards the UT and on to the EUD.
- A.32 DSSI:** The *Desired Signal Strength Indicator* (DSSI) is an estimate of the power of the square-root raised cosine filtered burst for a given active stream.
- A.33 ECDP:** Elliptic Curve Domain Parameters.
- A.34 Encryption:** The process of disguising a message to hide its content.
- A.35 End User Device (EUD):** The device (e.g., PC, laptop, or PDA) terminating a network session and responsible for providing the end-user with an interface to viewing and managing the session
- A.36 Frame:** A group of contiguous uplink and downlink timeslots plus guard times. The frame period is 5 milliseconds.
- A.37 Frame Resource:** A single frequency timeslot pair within a frame.
- A.38 Hash Function:** A function that takes a variable length input string (called a pre-image) and converts it to a fixed length (generally shorter) output string called a hash value.
- A.39 HC-SDMA Packet**Unit of data provided by L3 to L2 to transport on a stream; the *Packet Data Conversion Layer* (PDCL) converts L4 packets into HC-SDMA packets and the vice-versa.
- A.40 i-HAP protocol:** HC-SDMA Handshake and Authentication protocol.
- A.41 i-SEC protocol:** HC-SDMA Secure Communications protocol.
- A.42 i-TAP protocol:** HC-SDMA Terminal Authentication protocol.
- A.43 Incoming:** From the air-interface toward the network. For the BS, this is the uplink; for the UT, this is the downlink.
- A.44 Integrity Check:** A method to verify that a received message has not been modified in transit.
- A.45 IV:** The *Initialization Vector* (IV) is a 32-bit parameter derived from the *Absolute Frame Number* (AFN), the *Relative Carrier Number* (RCN) , and *Slot Number* (SN) . It is used to initialize a stream cipher per stream while using the same encryption key.

- A.46 Key Diffusion:** Operating a cipher without providing the keystream.
- A.47 Key Pair Generation:** Generation of a public and a private key pair for a *Public Key Encryption* algorithm.
- A.48 Keystream:** An output stream from a stream cipher, or from a block cipher in a stream-like cipher mode.
- A.49 Known Plaintext Attack**The cryptanalyst not only has access to the ciphertext of several messages, but also to the associated plaintext. See also *Chosen Plaintext Attack*.
- A.50 L0:** RF specifications for communication over the air link including output power levels, transmit frequency and timing error, pulse shaping, in band and out of band spurious emissions, receiver sensitivity and selectivity, dynamic range, etc.
- A.51 L1:** Specifications for channelization, burst structure, training data, modulation and FEC, timing, etc.
- A.52 L2:** The protocol layer consisting of the RLC and MAC sublayers.
- A.53 L2 MAC**The protocol layer that maintains physical channels and determines how to multiplex logical channels onto physical channels.
- A.54 L2 RLC:** The *L2 Radio Link Control (RLC)* sublayer of the *HC-SDMA* Layer 2 protocol provides a set of data link service for *L3* peer-to-peer messages.
- A.55 L2TP Network Server (LNS):** A server within the back end network that terminates the end-user PPP sessions originating on the EUD's, providing authentication, authorization, and accounting data collection.
- A.56 L3:** Specifications for creating and maintaining logical sessions including registration management, session management, resource control, mobility control, packet fragmentation, slot aggregation, in-band messaging, etc.
- A.57 L3 CM**The *L3 Connection Management (CM)* component provides *L4* with a virtual connection between the BS and UT across the air interface.
- A.58 L3 MAC:** The protocol layer that manages allocation of traffic channels to registrations.
- A.59 L3 MMC:** The *L3 Mobility and Management Control (MMC)* component provides management and control for mobility.
- A.60 L3 RM:** The *L3 Registration Management (RM)* component maps the virtual connection across the air interface provided by the *L3 CM* to the physical connections (streams) provided by the *L2 MAC* layer.
- A.61 L3 RRC:** The *L3 Radio Resource Control (RRC)* component maintains the radio frequency connectivity of individual streams by performing power control and modulation class selection.
- A.62 L4:** The protocol layer above *L3* that manages interaction with the network.
- A.63 L4 Packet**The unit of data provided to the *HC-SDMA* system by the network. *L4* Packets are transported by the *HC-SDMA* air interface. Also called a *network packet*.
- A.64 Layer 0:** See *L0*.
- A.65 Layer 1:** See *L1*.
- A.66 Layer 2:** See *L2*.

- A.67 Layer 2 Medium Access Control:** See *L2 MAC*.
- A.68 Layer 3:** See *L3*.
- A.69 Layer 3 Medium Access Control:** See *L3 CM*.
- A.70 Layer 4:** See *L4*.
- A.71 LDAP:** The Lightweight Directory Access Protocol.
- A.72 Logical Channel:** A conceptual communications pathway between the UT and BS that can be mapped to any of a number of physical channels.
- A.73 Logical Connection:** See *Session*.
- A.74 Message Field to Value Conversion:** Nominally, when a message field `msgfield` corresponds to a scalar value, the message field is mapped to an unsigned integer `uint(msgfield)`. For example, `uint(10110) = 22`.
- A.75 Modulation Classes (or ModClasses):** A combination of a signal constellation and FEC scheme.
- A.76 Net:** See *L4*.
- A.77 Network Packet:** See *L4 Packet*.
- A.78 Network Time:** The synchronized time of the BSs that serves as a reference for the fixed times at which the BS receives and sends bursts.
- A.79 Octet:** 8 consecutive bits of data.
- A.80 Operating Temperature Range:** The operating temperature range as specified by the base station manufacturer.
- A.81 Outgoing:** From the network to the air-interface. For the BS, this is the downlink; for the UT, this is the uplink.
- A.82 Packet Data Conversion Layer (PDCL):** The sub-component of L3 RM that converts data into a form suitable for transfer across the air interface on the sending side and undoes the conversion on the receiving side. That is, *L4 Packet* to/from *HC-SDMA Packet conversion*.
- A.83 Packet Services Switch:** A switch within the back-end wired network responsible for switching and forwarding end-user sessions to session termination equipment, as well as providing seamless network-level handovers as user terminals handover from one BS to another.
- A.84 PCH Resource Set:** The set of implied conventional channels a UT uses to listen for pages from the BS. There is one PCH logical channel for every conventional channel in the PCH resource set
- A.85 Physical Channel:** The set of resources utilized to support an individual communications pathway. These resources include the conventional channel and the spatial channel.
- A.86 Plaintext:** A message that is neither encrypted nor scrambled. Also called *cleartext*.
- A.87 Private Key:** Decryption key used in a *Public Key Encryption* method.
- A.88 Public Key Encryption:** Encryption algorithms where the key used for encryption is different from the key used for decryption. Furthermore, the decryption key cannot be calculated from the encryption key. These algorithms are called public key because the encryption key can be made public.
- A.89 Public Key:** Encryption key used in a *Public Key Encryption* method.
- A.90 RACH Resource Set:** The set of conventional channels a UT uses for sending RA messages to the BS. There is one RACH logical channel for every conventional channel in the RACH resource set.

- A.91 Registration:** The sequence of communication between UT and BS that allows or rejects permission for a UT to use an *HC-SDMA* network for data communication. One result is the establishment of a registration identifier (RID) that is unique to the UT connection within the BS registration domain. Registration can also refer to the period of time for which the RID remains valid, or to the RID itself.
- A.92 Registration Control Block (RCB):** The state associated with an instance of a L3 RM state machine.
- A.93 Relative Carrier Number (RCN):** This identifies an RF carrier relative to the lowest RF carrier in the carrier set of a BS. The UT identifies the RF frequency of relative carrier 0 from the RF frequency of the BCH carrier and the BCH Offset in the CM burst.
- A.94 Relative Frame Number (RFN):** Frame counter used by individual streams. Every stream has its own RFN. See clause 7.3.5.1.2.3.7, Field definitions for subtype-specific fields.
- A.95 Replay Attack:** Recording messages and using them later in an attempt to subvert a protocol.
- A.96 RLC Message Unit (RMU):** The basic message unit of outgoing traffic from the RLC.
- A.97 RSA:** A public key encryption algorithm named after its inventors (Rivest, Shamir, and Adleman). This algorithm is based on the intractability of the integer factorization problem and is the most common method for computing and verifying a *Digital Signature*.
- A.98 RSSI:** The *Received Signal Strength Indicator* (RSSI) is an estimate of the power of a received signal.
- A.99 Service Access Point (SAP):** The interface point between protocol layers.
- A.100 Service Data Unit (SDU):** Message unit of incoming L3 traffic that can be sent by RLC in acknowledged or unacknowledged mode under the direction of MAC and L3. Also known as an *L3 packet*.
- A.101 Session:** The period of time encompassing communication between the end user device and termination equipment in the network. A session may persist across multiple registrations and may be handed off among base stations.
- A.102 SHA-1:** Secure Hash Algorithm (SHA) which produces a 160-bit value from a message.
- A.103 Shared Secret:** The 326 bits that are exchanged between the UT and the BS as part of the *i-HAP protocol*. A subset of these bits is used to initialize ciphers to encrypt data and control messages.
- A.104 Signal to Interference and Noise Ratio (SINR):** The ratio of the power of the desired signal to the power of the interference and noise. SINR is expressed in dB.
- A.105 Signal to Noise Ratio (SNR):** The ratio of the power of the desired signal to the power of the noise. SNR is expressed in dB.
- A.106 Slot Number (SN):** Slot Number.
- A.107 Spatial Channel:** One of multiple data communication channels that share the same conventional channel and are prevented from interfering with one another through the use of *Antenna Array Spatial Processing*.
- A.108 Stream:** A sequence of radio bursts between a UT and BS with an access grant via RACH. The sequence ends when the BS and UT stop listening to each other.
- A.109 Stream Aggregation (Channel Aggregation):** The ability of an *HC-SDMA* registration to communicate on multiple streams simultaneously for the purpose of increasing the per-user bandwidth.

- A.110 Stream Cipher:** A *Symmetric Key Encryption* method which operates on plaintext a single bit at a time.
- A.111 Symbol Error Rate (SER):** The fraction of symbols which are incorrectly demodulated for a data stream.
- A.112 Symmetric Key Encryption** Encryption algorithms where the encryption key can be calculated from the decryption key and vice-versa. See also *Public Key Encryption*.
- A.113 Time-Server Authenticator String:** Encrypted *SHA-1* hash value of the *UT Authenticator Message*. The encryption is performed with the *Private Key* of the *Trusted Time Server*.
- A.114 Timing Advance:** The process followed by a UT of advancing its transmissions and delaying its reception proportionally to the UT-BS separation such that, regardless of the UT's range, the bursts to/from the UT are sent/received by/at the BS at fixed times.
- A.115 Trit:** A digit in a ternary representation of a natural number –e.g., $14 = 112$ consists of three *trits*.
- A.116 Trusted Time Server** An entity that is dedicated to certifying time for user terminals without access to a reliable time source.
- A.117 Uplink:** Part of or the entire communication path from the EUD to the UT towards the BS and on into the back end network.
- A.118 User Data Rate:** Data rate available to network packet streams over the air interface after subtraction of TDMA framing, spatio-temporal training, error control coding, and other air interface overhead.
- A.119 User Terminal (UT):** The device responsible for executing procedures for reception, transmission, and messaging within *HC-SDMA* protocol stack L0-L3 and providing the EUD with an interface into the network
- A.120 Useful Symbol:** All symbols in a time slot with the exception of the rampup, rampdown and guard symbols.
- A.121 UT Authenticator Message** A 438-bit string that contains the *Absolute Frame Number (AFN)*, the *UT Serial Number*, the *UT Session Reference*, and *i-HAP protocol Shared Secret*. See Figure 109: *UT Authenticator Message Formatting* (not to scale).
- A.122 UT Authenticator String:** Encrypted *SHA-1* hash value of the *UT Authenticator Message*. The encryption is performed with the *Private Key* of the user terminal.
- A.123 UT IMSI:** The *International Mobile Station Identifier (IMSI)* is a 60-bit string that is assigned to the UT by the service authority.
- A.124 UT Tracking:** Analysis of the activity of a user terminal and estimation of its physical location.