

4G

LTE / LTE-Advanced for Mobile Broadband

Erik Dahlman
Stefan Parkvall
Johan Sköld



4G LTE/LTE-Advanced for Mobile Broadband

**Erik Dahlman, Stefan Parkvall, and
Johan Sköld**



AMSTERDAM • BOSTON • HEIDELBERG • LONDON • NEW YORK • OXFORD
PARIS • SAN DIEGO • SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Academic Press is an imprint of Elsevier



Academic Press is an imprint of Elsevier
The Boulevard, Langford Lane, Kidlington, Oxford, OX5 1GB, UK
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA

First published 2011

Copyright © 2011 Erik Dahlman, Stefan Parkvall & Johan Sköld. Published by Elsevier Ltd. All rights reserved

The rights of Erik Dahlman, Stefan Parkvall & Johan Sköld to be identified as the authors of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher. Details on how to seek permission, further information about the Publisher's permissions policies and our arrangement with organizations such as the Copyright Clearance Center and the Copyright Licensing Agency, can be found at our website: www.elsevier.com/permissions

This book and the individual contributions contained in it are protected under copyright by the Publisher (other than as may be noted herein).

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2011921244

ISBN: 978-0-12-385489-6

For information on all Academic Press publications
visit our website at www.elsevierdirect.com

Typeset by MPS Limited, a Macmillan Company, Chennai, India
www.macmillansolutions.com

Printed in the United States of America
Transferred to Digital Printing, 2012

Working together to grow
libraries in developing countries

www.elsevier.com | www.bookaid.org | www.sabre.org

ELSEVIER

BOOK AID
International

Sabre Foundation

Contents

Preface.....	xiii
Acknowledgements.....	xv
Abbreviations and Acronyms.....	xvii
CHAPTER 1 Background of LTE.....	1
1.1 Introduction.....	1
1.2 Evolution of Mobile Systems Before LTE.....	2
1.2.1 The First 3G Standardization	3
1.3 ITU Activities	4
1.3.1 IMT-2000 and IMT-Advanced	4
1.3.2 Spectrum for IMT Systems.....	6
1.4 Drivers for LTE.....	7
1.5 Standardization of LTE.....	8
1.5.1 The Standardization Process	8
1.5.2 The 3GPP Process.....	9
1.5.3 The 3G Evolution to 4G.....	11
CHAPTER 2 High Data Rates in Mobile Communication.....	15
2.1 High Data Rates: Fundamental Constraints.....	15
2.1.1 High Data Rates in Noise-Limited Scenarios	17
2.1.2 Higher Data Rates in Interference-Limited Scenarios	18
2.2 Higher Data Rates Within a Limited Bandwidth: Higher-Order Modulation	19
2.2.1 Higher-Order Modulation in Combination with Channel Coding	20
2.2.2 Variations in Instantaneous Transmit Power.....	20
2.3 Wider Bandwidth Including Multi-Carrier Transmission.....	21
2.3.1 Multi-Carrier Transmission.....	23
CHAPTER 3 OFDM Transmission	27
3.1 Basic Principles of OFDM.....	27
3.2 OFDM Demodulation.....	29
3.3 OFDM Implementation Using IFFT/FFT Processing	30
3.4 Cyclic-Prefix Insertion.....	32
3.5 Frequency-Domain Model of OFDM Transmission.....	34
3.6 Channel Estimation and Reference Symbols.....	35
3.7 Frequency Diversity with OFDM: Importance of Channel Coding	36

3.8	Selection of Basic OFDM Parameters	37
3.8.1	OFDM Subcarrier Spacing	38
3.8.2	Number of Subcarriers	39
3.8.3	Cyclic-Prefix Length	40
3.9	Variations in Instantaneous Transmission Power	40
3.10	OFDM as a User-Multiplexing and Multiple-Access Scheme	41
3.11	Multi-Cell Broadcast/Multicast Transmission and OFDM	43
CHAPTER 4	Wider-Band "Single-Carrier" Transmission	45
4.1	Equalization Against Radio-Channel Frequency Selectivity	45
4.1.1	Time-Domain Linear Equalization	45
4.1.2	Frequency-Domain Equalization	47
4.1.3	Other Equalizer Strategies	49
4.2	Uplink FDMA with Flexible Bandwidth Assignment	50
4.3	DFT-Spread OFDM	52
4.3.1	Basic Principles	52
4.3.2	DFTS-OFDM Receiver	54
4.3.3	User Multiplexing with DFTS-OFDM	55
4.3.4	Distributed DFTS-OFDM	55
CHAPTER 5	Multi-Antenna Techniques	59
5.1	Multi-Antenna Configurations	59
5.2	Benefits of Multi-Antenna Techniques	60
5.3	Multiple Receive Antennas	60
5.4	Multiple Transmit Antennas	65
5.4.1	Transmit-Antenna Diversity	65
5.4.2	Transmitter-Side Beam-Forming	68
5.5	Spatial Multiplexing	71
5.5.1	Basic Principles	71
5.5.2	Precoder-Based Spatial Multiplexing	74
5.5.3	Nonlinear Receiver Processing	75
CHAPTER 6	Scheduling, Link Adaptation, and Hybrid ARQ	79
6.1	Link Adaptation: Power and Rate Control	79
6.2	Channel-Dependent Scheduling	81
6.2.1	Downlink Scheduling	81
6.2.2	Uplink Scheduling	85
6.2.3	Link Adaptation and Channel-Dependent Scheduling in the Frequency Domain	87
6.2.4	Acquiring on Channel-State Information	87
6.2.5	Traffic Behavior and Scheduling	88

37	6.3	Advanced Retransmission Schemes	89
38	6.4	Hybrid ARQ with Soft Combining	90
39	CHAPTER 7	LTE Radio Access: An Overview	95
40	7.1	Basic Principles	96
40	7.1.1	Transmission Scheme	96
41	7.1.2	Channel-Dependent Scheduling and Rate Adaptation	97
43	7.1.3	Inter-Cell Interference Coordination	99
45	7.1.4	Hybrid ARQ with Soft Combining	99
45	7.1.5	Multi-Antenna Transmission	100
45	7.1.6	Spectrum Flexibility	100
47	7.2	LTE Release 9	102
49	7.2.1	Multicast and Broadcast Support	102
50	7.2.2	Positioning	103
52	7.2.3	Dual-Layer Beam-Forming	103
52	7.3	LTE Release 10 and IMT-Advanced	103
54	7.3.1	Carrier Aggregation	104
55	7.3.2	Extended Multi-Antenna Transmission	104
55	7.3.3	Relaying	105
59	7.3.4	Heterogeneous Deployments	105
59	7.4	Terminal Capabilities	105
60	CHAPTER 8	Radio-Interface Architecture	109
60	8.1	Overall System Architecture	109
65	8.1.1	Core Network	109
65	8.1.2	Radio-Access Network	111
68	8.2	Radio Protocol Architecture	111
71	8.2.1	Radio-Link Control	113
71	8.2.2	Medium-Access Control	115
74	8.2.3	Physical Layer	123
75	8.3	Control-Plane Protocols	124
79	8.3.1	State Machine	125
79	CHAPTER 9	Physical Transmission Resources	127
81	9.1	Overall Time-Frequency Structure	127
81	9.2	Normal Subframes and MBSFN Subframes	131
85	9.3	Carrier Aggregation	132
87	9.4	Frequency-Domain Location of LTE Carriers	134
87	9.5	Duplex Schemes	135
88	9.5.1	Frequency-Division Duplex (FDD)	136
	9.5.2	Time-Division Duplex (TDD)	137
	9.5.3	LTE and TD-SCDMA Coexistence	140

CHAPTER 10	Downlink Physical-Layer Processing	143
10.1	Transport-Channel Processing	143
10.1.1	Processing Steps	143
10.1.2	Localized and Distributed Resource Mapping	149
10.2	Downlink Reference Signals	152
10.2.1	Cell-Specific Reference Signals	153
10.2.2	Demodulation Reference Signals	156
10.2.3	CSI Reference Signals	158
10.3	Multi-Antenna Transmission	161
10.3.1	Transmit Diversity	163
10.3.2	Codebook-Based Precoding	165
10.3.3	Non-Codebook-Based Precoding	169
10.3.4	Downlink Multi-User MIMO	171
10.4	Downlink L1/L2 Control Signaling	173
10.4.1	Physical Control Format Indicator Channel	174
10.4.2	Physical Hybrid-ARQ Indicator Channel	177
10.4.3	Physical Downlink Control Channel	179
10.4.4	Downlink Scheduling Assignment	181
10.4.5	Uplink Scheduling Grants	189
10.4.6	Carrier Aggregation and Cross-Carrier Scheduling	193
10.4.7	Power-Control Commands	195
10.4.8	PDCCH Processing	195
10.4.9	Blind Decoding of PDCCHs	199
CHAPTER 11	Uplink Physical-Layer Processing	203
11.1	Transport-Channel Processing	203
11.1.1	Processing Steps	203
11.1.2	Mapping to the Physical Resource	205
11.1.3	PUSCH Frequency Hopping	207
11.2	Uplink Reference Signals	210
11.2.1	Uplink Demodulation Reference Signals	210
11.2.2	Uplink Sounding Reference Signals	217
11.3	Uplink Multi-Antenna Transmission	221
11.3.1	Precoder-Based Multi-Antenna Transmission for PUSCH	221
11.3.2	Uplink Multi-User MIMO	224
11.3.3	PUCCH Transmit Diversity	226
11.4	Uplink L1/L2 Control Signaling	226
11.4.1	Uplink L1/L2 Control Signaling on PUCCH	227
11.4.2	Uplink L1/L2 Control Signaling on PUSCH	242
11.5	Uplink Timing Alignment	245

143			
143			
143			
149			
152			
153			
156			
158			
161			
163			
165			
169			
171			
173			
174			
177			
179			
181			
189			
193			
195			
195			
199			
203			
203			
203			
205			
207			
210			
210			
217			
221			
221			
224			
226			
226			
227			
242			
245			
	CHAPTER 12	Retransmission Protocols	247
	12.1	Hybrid ARQ with Soft Combining	248
	12.1.1	Downlink Hybrid ARQ	251
	12.1.2	Uplink Hybrid ARQ	251
	12.1.3	Hybrid-ARQ Timing	255
	12.2	Radio-Link Control	259
	12.2.1	Segmentation, Concatenation, and Reassembly of RLC SDUs	260
	12.2.2	RLC Retransmission	261
	12.2.3	In-Sequence Delivery	262
	12.2.4	RLC Operation	262
	CHAPTER 13	Power Control, Scheduling, and Interference Handling	265
	13.1	Uplink Power Control	265
	13.1.1	Uplink Power Control – Some Basic Rules	266
	13.1.2	Power Control for PUCCH	267
	13.1.3	Power Control for PUSCH	269
	13.1.4	Power Control for SRS	271
	13.1.5	Power Headroom	271
	13.2	Scheduling and Rate Adaptation	272
	13.2.1	Downlink Scheduling	273
	13.2.2	Uplink Scheduling	275
	13.2.3	Semi-Persistent Scheduling	280
	13.2.4	Scheduling for Half-Duplex FDD	281
	13.2.5	Channel-State Reporting	282
	13.2.6	Discontinuous Reception (DRX) and Component Carrier Deactivation	287
	13.3	Inter-Cell Interference Coordination	290
	13.4	Heterogeneous Network Deployments	293
	13.4.1	Interference Handling in a Heterogeneous Deployment	294
	13.4.2	Interference Coordination in the Case of Home-eNodeB	299
	CHAPTER 14	Access Procedures	301
	14.1	Acquisition and Cell Search	301
	14.1.1	Overview of LTE Cell Search	301
	14.1.2	PSS Structure	303
	14.1.3	SSS Structure	304
	14.2	System Information	304
	14.2.1	MIB and BCH Transmission	305
	14.2.2	System-Information Blocks	308

14.3	Random Access.....	310
14.3.1	Step 1: Random-Access Preamble Transmission.....	312
14.3.2	Step 2: Random-Access Response.....	317
14.3.3	Step 3: Terminal Identification.....	318
14.3.4	Step 4: Contention Resolution.....	319
14.4	Paging.....	319
CHAPTER 15	Multimedia Broadcast/Multicast Services.....	323
15.1	Architecture.....	324
15.2	Overall Channel Structure and Physical-Layer Processing.....	325
15.3	Scheduling of MBMS Services.....	328
CHAPTER 16	Relaying.....	331
16.1	Relays in LTE.....	332
16.2	Overall Architecture.....	333
16.3	Backhaul Design for Inband Relaying.....	334
16.3.1	Access-Link Hybrid-ARQ Operation.....	335
16.3.2	Backhaul-Link Hybrid-ARQ Operation.....	336
16.3.3	Backhaul Downlink Control Signaling.....	337
16.3.4	Reference Signals for the Backhaul Link.....	341
16.3.5	Backhaul-Access Link Timing.....	342
CHAPTER 17	Spectrum and RF Characteristics.....	347
17.1	Spectrum for LTE.....	347
17.1.1	Spectrum Defined for IMT Systems by the ITU-R.....	347
17.1.2	Frequency Bands for LTE.....	348
17.1.3	New Frequency Bands.....	351
17.2	Flexible Spectrum Use.....	352
17.3	Flexible Channel Bandwidth Operation.....	353
17.4	Carrier Aggregation for LTE.....	355
17.5	Multi-Standard Radio Base Stations.....	359
17.6	Overview of RF Requirements for LTE.....	361
17.6.1	Transmitter Characteristics.....	362
17.6.2	Receiver Characteristics.....	363
17.6.3	Regional Requirements.....	363
17.6.4	Band-Specific UE Requirements Through Network Signaling.....	364
17.6.5	Base-Station Classes.....	365
17.7	Output Power Level Requirements.....	366
17.7.1	Base-Station Output Power and Dynamic Range.....	366
17.7.2	UE Output Power and Dynamic Range.....	366

310	17.8 Transmitted Signal Quality	366
312	17.8.1 EVM and Frequency Error	367
317	17.8.2 UE In-Band Emissions	367
318	17.8.3 Base-Station Time Alignment	367
319	17.9 Unwanted Emissions Requirements	367
319	17.9.1 Implementation Aspects	368
	17.9.2 Spectrum Emission Mask	368
323	17.9.3 Adjacent Channel Leakage Ratio	370
324	17.9.4 Spurious Emissions	372
325	17.9.5 Occupied Bandwidth	372
328	17.9.6 Transmitter Intermodulation	372
	17.10 Sensitivity and Dynamic Range	373
	17.11 Receiver Susceptibility to Interfering Signals	373
331	CHAPTER 18 Performance	377
332	18.1 Performance Assessment	377
333	18.1.1 End-User Perspective of Performance	378
334	18.1.2 Operator Perspective	379
335	18.2 Performance in Terms of Peak Data Rates and Latency	379
336	18.3 Performance Evaluation of LTE-Advanced	380
337	18.3.1 Models and Assumptions	381
341	18.3.2 Evaluation Criteria	383
342	18.3.3 Performance Numbers for FDD	384
	18.3.4 Performance Numbers for TDD	385
347	18.4 Conclusion	387
347	CHAPTER 19 Other Wireless Communications Systems	389
347	19.1 HSPA	389
348	19.1.1 Architecture	391
351	19.1.2 Channel-Dependent Scheduling	392
352	19.1.3 Hybrid ARQ with Soft Combining	393
353	19.1.4 Control-Plane Latency Reductions	394
355	19.1.5 Spatial Multiplexing	395
359	19.1.6 Carrier Aggregation	395
361	19.1.7 UTRA TDD	395
362	19.2 GSM/EDGE	396
363	19.2.1 Objectives for GSM/EDGE Evolution	397
363	19.2.2 Dual-Antenna Terminals	398
364	19.2.3 Multi-Carrier EDGE	398
365	19.2.4 Reduced TTI and Fast Feedback	399
366	19.2.5 Improved Modulation and Coding	400
366		

19.2.6	Higher Symbol Rates	400
19.2.7	Voice Service over Adaptive Multi-User Channels	401
19.3	CDMA2000 and HRPD/1x EV-DO.....	402
19.3.1	CDMA2000 1x.....	403
19.3.2	1x EV-DO Rev 0.....	403
19.3.3	1x EV-DO Rev A.....	404
19.3.4	1x EV-DO Rev B.....	405
19.3.5	1x EV-DO Rev C.....	405
19.4	IEEE 802.16e, Mobile WiMAX and 802.16 m.....	405
19.4.1	IEEE 802.16e and Mobile WiMAX	406
19.4.2	IEEE 802.16 m – WiMAX for IMT-Advanced.....	408
19.5	Summary.....	409
CHAPTER 20	Final Thoughts	411
20.1	Where to Go in the Future?	411
20.1.1	Advanced Multi-Cell Coordination.....	412
20.1.2	Network Energy Efficiency	412
20.1.3	Machine-Type Communication.....	413
20.1.4	New ways of Using Spectrum	414
20.1.5	Direct Device-to-Device Communication.....	414
20.2	Concluding Remarks	415
References		417
Index		423

Preface

During the past years, there has been a quickly rising interest in radio access technologies for providing mobile as well as nomadic and fixed services for voice, video, and data. The difference in design, implementation, and use between telecom and datacom technologies is also becoming more blurred. One example is cellular technologies from the telecom world being used for broadband data and wireless LAN from the datacom world being used for voice-over IP.

Today, the most widespread radio access technology for mobile communication is digital cellular, with the number of users passing 5 billion by 2010, which is more than half of the world's population. It has emerged from early deployments of an expensive voice service for a few car-borne users, to today's widespread use of mobile-communication devices that provide a range of mobile services and often include camera, MP3 player, and PDA functions. With this widespread use and increasing interest in mobile communication, a continuing evolution ahead is foreseen.

This book describes LTE, developed in 3GPP (*Third Generation Partnership Project*) and providing true 4G broadband mobile access, starting from the first version in release 8 and through the continuing evolution to release 10, the latest version of LTE. Release 10, also known as LTE-Advanced, is of particular interest as it is the major technology approved by the ITU as fulfilling the IMT-Advanced requirements. The description in this book is based on LTE release 10 and thus provides a complete description of the LTE-Advanced radio access from the bottom up.

Chapter 1 gives the background to LTE and its evolution, looking also at the different standards bodies and organizations involved in the process of defining 4G. It also gives a discussion of the reasons and driving forces behind the evolution.

Chapters 2–6 provide a deeper insight into some of the technologies that are part of LTE and its evolution. Because of its generic nature, these chapters can be used as a background not only for LTE as described in this book, but also for readers who want to understand the technology behind other systems, such as WCDMA/HSPA, WiMAX, and CDMA2000.

Chapters 7–17 constitute the main part of the book. As a start, an introductory technical overview of LTE is given, where the most important technology components are introduced based on the generic technologies described in previous chapters. The following chapters provide a detailed description of the protocol structure, the downlink and uplink transmission schemes, and the associated mechanisms for scheduling, retransmission and interference handling. Broadcast operation and relaying are also described. This is followed by a discussion of the spectrum flexibility and the associated requirements from an RF perspective.

Finally, in Chapters 18–20, an assessment is made on LTE. Through an overview of similar technologies developed in other standards bodies, it will be clear that the technologies adopted for the evolution in 3GPP are implemented in many other systems as well. Finally, looking into the future, it will be seen that the evolution does not stop with LTE-Advanced but that new features are continuously added to LTE in order to meet future requirements.

Access Procedures

14

The previous chapters have described the LTE uplink and downlink transmission schemes. However, prior to transmission of data, the terminal needs to connect to the network. This chapter describes the procedures necessary for a terminal to be able to access an LTE-based network.

14.1 ACQUISITION AND CELL SEARCH

Before an LTE terminal can communicate with an LTE network it has to do the following:

- Find and acquire synchronization to a cell within the network.
- Receive and decode the information, also referred to as the *cell system information*, needed to communicate with and operate properly within the cell.

The first of these steps, often simply referred to as *cell search*, is discussed in this section. The next section then discusses, in more detail, the means by which the network provides the cell system information.

Once the system information has been correctly decoded, the terminal can access the cell by means of the random-access procedure as described in Section 14.3.

14.1.1 Overview of LTE Cell Search

A terminal does not only need to carry out cell search at power-up – that is, when initially accessing the system. Rather, to support mobility, terminals need to continuously search for, synchronize to, and estimate the reception quality of neighboring cells. The reception quality of the neighboring cells, in relation to the reception quality of the current cell, is then evaluated to conclude if a *handover* (for terminals in RRC_CONNECTED) or *cell reselection* (for terminals in RRC_IDLE) should be carried out.

LTE cell search consists of the following basic parts:

- Acquisition of frequency and symbol synchronization to a cell.
- Acquisition of frame timing of the cell – that is, determine the start of the downlink frame.
- Determination of the physical-layer cell identity of the cell.

As already mentioned (e.g. in Chapter 10), there are 504 different physical-layer cell identities defined for LTE, where each cell identity corresponds to one specific downlink reference-signal sequence. The set of physical-layer cell identities is further divided into 168 *cell-identity groups*, with three cell identities within each group.




FIGURE 14.1

Time-domain position of PSS and SSS for FDD and TDD.

To assist the cell search, two special signals are transmitted on each downlink component carrier, the *Primary Synchronization Signal* (PSS) and the *Secondary Synchronization Signal* (SSS). Although having the same detailed structure, the time-domain positions of the synchronization signals within the frame differ somewhat depending on whether the cell is operating in FDD or TDD:

- In the case of FDD (upper part of Figure 14.1), the PSS is transmitted within the last symbol of the first slot of subframes 0 and 5, while the SSS is transmitted within the second last symbol of the same slot – that is, just prior to the PSS.
- In the case of TDD (lower part of Figure 14.1), the PSS is transmitted within the third symbol of subframes 1 and 6 – that is, within the DwPTS – while the SSS is transmitted in the last symbol of subframes 0 and 5 – that is, three symbols ahead of the PSS.

There are several reasons for the difference in the positions of the synchronization signals for FDD and TDD. For example, the difference allows for the detection of the duplex scheme used on a carrier if this is not known in advance.

Within one cell, the two PSSs within a frame are identical. Furthermore, the PSS of a cell can take three different values depending on the physical-layer cell identity of the cell. More specifically, the three cell identities within a cell-identity group always correspond to different PSS. Thus, once the terminal has detected and identified the PSS of the cell, it has found the following:

- Five-millisecond timing of the cell and thus also the position of the SSS, which has a fixed offset relative to the PSS.
- The cell identity within the cell-identity group. However, the terminal has not yet determined the cell-identity group itself – that is, the number of possible cell identities has been reduced from 504 to 168.

Thus, from the SSS, the position of which is known once the PSS has been detected, the terminal should find the following:

- frame timing (two different alternatives given the found position of the PSS); and
- the cell-identity group (168 alternatives).




FIGURE 14.2

Definition and structure of PSS.

Furthermore, it should be possible for a terminal to do this by the reception of one single SSS. The reason is that, for example, in the case when the terminal is searching for cells on other carriers, the search window may not be sufficiently large to cover more than one SSS.

To enable this, each SSS can take 168 different values corresponding to the 168 different cell-identity groups. Furthermore, the set of values valid for the two SSSs within a frame (SSS₁ in subframe 0 and SSS₂ in subframe 5) are different, implying that, from the detection of a single SSS, the terminal can determine whether SSS₁ or SSS₂ has been detected and thus determine frame timing.

Once the terminal has acquired frame timing and the physical-layer cell identity, it has identified the cell-specific reference signal. The behavior is slightly different depending on whether it is an initial cell search or cell search for the purpose of neighboring cell measurements:

- In the case of initial cell search – that is, the terminal state is in RRC_IDLE mode – the reference signal will be used for channel estimation and subsequent decoding of the BCH transport channel to obtain the most basic set of system information.
- In the case of mobility measurements – that is, the terminal is in RRC_CONNECTED mode – the terminal will measure the received power of the reference signal. If the measurement fulfills a configurable condition, it will trigger sending of a *reference signal received power* (RSRP) measurement report to the network. Based on the measurement report, the network will conclude whether a handover should take place. The RSRP reports can also be used for component carrier management,¹ for example whether an additional component carrier should be configured or if the primary component carrier should be reconfigured.

14.1.2 PSS Structure

On a more detailed level, the three PSSs are three length-63 Zadoff-Chu (ZC) sequences extended with five zeros at the edges and mapped to the center 73 subcarriers (center six resource blocks) as illustrated in Figure 14.2. It should be noted though that the center subcarrier is actually not transmitted as it coincides with the DC subcarrier. Thus, only 62 elements of the length-63 ZC sequences are actually transmitted (element X_{32}^{PSS} is not transmitted).

¹As discussed in Chapter 9, the specifications use the terms primary and secondary cells instead of primary and secondary component carriers.




FIGURE 14.3

Definition and structure of SSS.

The PSS thus occupies 72 resource elements (not including the DC carrier) in subframes 0 and 5 (FDD) and subframes 1 and 6 (TDD). These resource elements are then not available for transmission of DL-SCH.

14.1.3 SSS Structure

Similar to PSS, the SSS occupies the center 72 resource elements (not including the DC carrier) in subframes 0 and 5 (for both FDD and TDD). As described above, the SSS should be designed so that:

- The two SSS (SSS_1 in subframe 0 and SSS_2 in subframe 5) take their values from sets of 168 possible values corresponding to the 168 different cell-identity groups.
- The set of values applicable for SSS_2 is different from the set of values applicable for SSS_1 to allow for frame-timing detection from the reception of a single SSS.

The structure of the two SSS is illustrated in Figure 14.3. SSS_1 is based on the frequency interleaving of two length-31 m -sequences X and Y , each of which can take 31 different values (actually 31 different shifts of the same m -sequence). Within a cell, SSS_2 is based on exactly the same two sequences as SSS_1 . However, the two sequences have been swapped in the frequency domain, as outlined in Figure 14.3. The set of valid combinations of X and Y for SSS_1 has then been selected so that a swapping of the two sequences in the frequency domain is not a valid combination for SSS_1 . Thus, the above requirements are fulfilled:

- The set of valid combinations of X and Y for SSS_1 (as well as for SSS_2) are 168, allowing for detection of the physical-layer cell identity.
- As the sequences X and Y are swapped between SSS_1 and SSS_2 , frame timing can be found.

14.2 SYSTEM INFORMATION

By means of the basic cell-search procedure described in Section 14.1, a terminal synchronizes to a cell, acquires the physical-layer identity of the cell, and detects the cell frame timing. Once this has

been achieved, the terminal has to acquire the cell *system information*. This is information that is repeatedly broadcast by the network and which needs to be acquired by terminals in order for them to be able to access and, in general, operate properly within the network and within a specific cell. The system information includes, among other things, information about the downlink and uplink cell bandwidths, the uplink/downlink configuration in the case of TDD, detailed parameters related to random-access transmission and uplink power control, etc.

In LTE, system information is delivered by two different mechanisms relying on two different transport channels:

- A limited amount of system information, corresponding to the so-called *Master-Information Block* (MIB), is transmitted using the BCH.
- The main part of the system information, corresponding to different so-called *System-Information Blocks* (SIBs), is transmitted using the downlink shared channel (DL-SCH).

It should be noted that system information in both the MIB and the SIBs corresponds to the BCCH logical channel. Thus, as also illustrated in Figure 8.7 in Chapter 8, BCCH can be mapped to both BCH and DL-SCH depending on the exact BCCH information.

14.2.1 MIB and BCH Transmission

As mentioned above, the MIB transmitted using BCH consists of a very limited amount of system information, mainly such information that is absolutely needed for a terminal to be able to read the remaining system information provided using DL-SCH. More specifically, the MIB includes the following information:


- *Information about the downlink cell bandwidth.* Four bits are available within the MIB to indicate the downlink bandwidth. Thus, up to 16 different bandwidths, measured in number of resource blocks, can be defined for each frequency band.
- *Information about the PHICH configuration of the cell.* As mentioned in Section 10.4.2, the terminal must know the PHICH configuration to be able to receive the L1/L2 control signaling on PDCCH which, in turn, is needed to receive DL-SCH. Thus, information about the PHICH configuration (three bits) is included in the MIB – that is, transmitted using BCH, which can be received and decoded without first receiving any PDCCH.
- *The System Frame Number (SFN)* or, more exactly, all bits except the two least significant bits of the SFN are included in the MIB. As described below, the terminal can indirectly acquire the two least significant bits of the SFN from the BCH decoding.

BCH physical-layer processing, such as channel coding and resource mapping, differs quite substantially from the corresponding processing and mapping for DL-SCH outlined in Chapter 10.

As can be seen in Figure 14.4, one BCH transport block, corresponding to the MIB, is transmitted every 40 ms. The BCH Transmissions Time Interval (TTI) thus equals 40 ms.

The BCH relies on a 16-bit CRC, in contrast to a 24-bit CRC used for all other downlink transport channels. The reason for the shorter BCH CRC is to reduce the relative CRC overhead, having the very small BCH transport-block size in mind.

BCH channel coding is based on the same rate-1/3 tail-biting convolutional code as is used for the PDCCH control channel. The reason for using convolutional coding for BCH, rather than the Turbo



subframes 0 and 5
le for transmission

the DC carrier) in
e designed so that:

m sets of 168 pos-

icable for SSS₁ to

the frequency inter-
nt values (actually
ctly the same two
y domain, as out-
en selected so that
on for SSS₁. Thus,

168, allowing for

n be found.

synchronizes to a
ing. Once this has




FIGURE 14.4

Channel coding and subframe mapping for the BCH transport channel.

code used for all other transport channels, is the small size of the BCH transport block. With such small blocks, tail-biting convolutional coding actually outperforms Turbo coding. The channel coding is followed by rate matching, in practice repetition of the coded bits, and bit-level scrambling. QPSK modulation is then applied to the coded and scrambled BCH transport block.

BCH multi-antenna transmission is limited to transmit diversity – that is, SFBC in the case of two antenna ports and combined SFBC/FSTD in the case of four antenna ports. Actually, as mentioned in Chapter 10, if two antenna ports are available within the cell, SFBC *must* be used for BCH. Similarly, if four antenna ports are available, combined SFBC/FSTD *must* be used. Thus, by blindly detecting what transmit-diversity scheme is used for BCH, a terminal can indirectly determine the number of cell-specific antenna ports within the cell and also the transmit-diversity scheme used for the L1/L2 control signaling.

As can also be seen from Figure 14.4, the coded BCH transport block is mapped to the first subframe of each frame in four consecutive frames. However, as can be seen in Figure 14.5 and in contrast to other downlink transport channels, the BCH is not mapped on a resource-block basis. Instead, the BCH is transmitted within the first four OFDM symbols of the second slot of subframe 0 and only over the 72 center subcarriers.² Thus, in the case of FDD, BCH follows immediately after the PSS and SSS in subframe 0. The corresponding resource elements are then not available for DL-SCH transmission.

The reason for limiting the BCH transmission to the 72 center subcarriers, regardless of the cell bandwidth, is that a terminal may not know the downlink cell bandwidth when receiving BCH. Thus,

²Not including the DC carrier.




FIGURE 14.5

Detailed resource mapping for the BCH transport channel.

when first receiving BCH of a cell, the terminal can assume a cell bandwidth equal to the minimum possible downlink bandwidth – that is, six resource blocks corresponding to 72 subcarriers. From the decoded MIB, the terminal is then informed about the actual downlink cell bandwidth and can adjust the receiver bandwidth accordingly.

Clearly, the total number of resource elements to which the coded BCH is mapped is very large compared to the size of the BCH transport block, implying extensive repetition coding or, equivalently, massive processing gain for the BCH transmission. Such large processing gain is needed as it should be possible to receive and correctly decode the BCH also by terminals in neighboring cells, implying potentially very low receiver *Signal-to-Interference-and-Noise Ratio* (SINR) when decoding the BCH. At the same time, many terminals will receive BCH in much better channel conditions. Such terminals then do not need to receive the full set of four subframes over which a BCH transport block is transmitted to acquire sufficient energy for correct decoding of the transport block. Instead, already by receiving only a few or perhaps only a single subframe, the BCH transport block may be decodable.

From the initial cell search, the terminal has found only the cell frame timing. Thus, when receiving BCH, the terminal does not know to what set of four subframes a certain BCH transport block is mapped. Instead, a terminal must try to decode the BCH at four possible timing positions. [Depending on which decoding is successful, indicated by a correct CRC check, the terminal can implicitly determine 40 ms timing or, equivalently, the two least significant bits of the SFN.³] This is the reason why these bits do not need to be explicitly included in the MIB.

³BCH scrambling is defined with 40 ms periodicity, hence even if the terminal successfully decodes the BCH after observing only a single transmission instant, it can determine the 40 ms timing.

14.2.2 System-Information Blocks

As already mentioned, the MIB on the BCH only includes a very limited part of the system information. The main part of the system information is instead included in different *System-Information Blocks* (SIBs) that are transmitted using the DL-SCH. The presence of system information on DL-SCH in a subframe is indicated by the transmission of a corresponding PDCCH marked with a special *System-Information RNTI* (SI-RNTI). Similar to the PDCCH providing the scheduling assignment for “normal” DL-SCH transmission, this PDCCH also indicates the transport format and physical resource (set of resource blocks) used for the system-information transmission.

LTE defines a number of different SIBs characterized by the type of information that is included within them:

- SIB1 includes information mainly related to whether a terminal is allowed to camp on the cell. This includes information about the operator/operators of the cell, if there are restrictions with regards to what users may access the cell, etc. SIB1 also includes information about the allocation of subframes to uplink/downlink and configuration of the special subframe in the case of TDD. Finally, SIB1 includes information about the time-domain scheduling of the remaining SIBs (SIB2 and beyond).
- SIB2 includes information that terminals need in order to be able to access the cell. This includes information about the uplink cell bandwidth, random-access parameters, and parameters related to uplink power control.
- SIB3 mainly includes information related to cell reselection.
- SIB4–SIB8 include neighboring-cell-related information, including information related to neighboring cells on the same carrier, neighboring cells on different carriers, and neighboring non-LTE cells, such as WCDMA/HSPA, GSM, and CDMA2000 cells.
- SIB9 contains the name of the home-eNodeB.
- SIB10–SIB12 contain public warning messages, for example earthquake information.
- SIB13 contains information necessary for MBMS reception (see also Chapter 15).

Not all the SIBs need to be present. For example, SIB9 is not relevant for an operator-deployed node and SIB13 is not necessary if MBMS is not provided in the cell.

Similar to the MIB, the SIBs are broadcasted repeatedly. How often a certain SIB needs to be transmitted depends on how quickly terminals need to acquire the corresponding system information when entering the cell. In general, a lower-order SIB is more time critical and is thus transmitted more often compared to a higher-order SIB. SIB1 is transmitted every 80 ms, whereas the transmission period for the higher-order SIBs is flexible and can be different for different networks.

The SIBs represent the basic system information to be transmitted. The different SIBs are then mapped to different *System-Information messages* (SIs), which correspond to the actual transport blocks to be transmitted on DL-SCH. SIB1 is always mapped, by itself, on to the first system-information message SI-1,⁴ whereas the remaining SIBs may be group-wise multiplexed on to the same SI subject to the following constraints:

- The SIBs mapped to the same SI must obviously have the same transmission period. Thus, as an example, two SIBs with a transmission period of 320 ms can be mapped to the same SI, whereas an SIB with a transmission period of 160 ms must be mapped to a different SI.

⁴Strictly speaking, as SIB1 is not multiplexed with any other SIBs, it is not even said to be mapped to an SI. Rather, SIB1 in itself directly corresponds to the transport block.




FIGURE 14.6

Example of mapping of SIBs to SIs.

- The total number of information bits that is mapped to a single SI must not exceed what is possible to transmit within a transport block.

It should be noted that the transmission period for a given SIB might be different in different networks. For example, different operators may have different requirements concerning the period when different types of neighboring-cell information needs to be transmitted. Furthermore, the amount of information that can fit into a single transport block very much depends on the exact deployment situation, such as cell bandwidth, cell size, and so on.

Thus, in general, the SIB-to-SI mapping for SIBs beyond SIB1 is flexible and may be different for different networks or even within a network. An example of SIB-to-SI mapping is illustrated in Figure 14.6. In this case, SIB2 is mapped to SI-2 with a transmission period of 160 ms. SIB3 and SIB4 are multiplexed into SI-3 with a transmission period of 320 ms, whereas SIB5, which also requires a transmission period of 320 ms, is mapped to a separate SI (SI-4). Finally, SIB6, SIB7, and SIB8 are multiplexed into SI-5 with a transmission period of 640 ms. Information about the detailed SIB-to-SI mapping, as well as the transmission period of the different SIs, is provided in SIB1.

Regarding the more detailed transmission of the different system-information messages there is a difference between the transmission of SI-1, corresponding to SIB1, and the transmission of the remaining SIs.

The transmission of SI-1 has only a limited flexibility. More specifically, SI-1 is always transmitted within subframe 5. However, the bandwidth or, in general, the set of resource blocks over which SI-1 is transmitted, as well as other aspects of the transport format, may vary and is signaled on the associated PDCCH.

For the remaining SIs, the scheduling on DL-SCH is more flexible in the sense that each SI can, in principle, be transmitted in any subframe within *time windows* with well-defined starting points and durations. The starting point and duration of the time window of each SI are provided in SIB-1. It should be noted that an SI does not need to be transmitted on consecutive subframes within the time window, as is illustrated in Figure 14.7. Within the time window, the presence of system information in a subframe is indicated by the SI-RNTI on PDCCH, which also provides the frequency-domain scheduling as well as other parameters related to the system-information transmission.

Different SIs have different non-overlapping time windows. Thus, a terminal knows what SI is being received without the need for any specific identifier for each SI.

In the case of a relatively small SI and a relatively large system bandwidth, a single subframe may be sufficient for the transmission of the SI. In other cases, multiple subframes may be needed for the transmission of a single SI. In the latter case, instead of segmenting each SI into sufficiently small




FIGURE 14.7

Transmission window for the transmission of an SI.

blocks that are separately channel coded and transmitted in separate subframes, the complete SI is channel coded and mapped to multiple, not necessarily consecutive, subframes.

Similar to the case of the BCH, terminals that are experiencing good channel conditions may then be able to decode the complete SI after receiving only a subset of the subframes to which the coded SI is mapped, while terminals in bad positions need to receive more subframes for proper decoding of the SI. This approach has two benefits:

- Similar to BCH decoding, terminals in good positions need to receive fewer subframes, implying the possibility for reduced terminal power consumption.
- The use of larger code blocks in combination with Turbo coding leads to improved channel-coding gain.

Strictly speaking the single transport block containing the SI is not transmitted over multiple subframes. Rather, the subsequent SI transmissions are seen as autonomous hybrid-ARQ retransmissions of the first SI transmission – that is, retransmissions taking place without any explicit feedback signaling provided on the uplink.

For terminals capable of carrier aggregation, the system information for the primary component carrier is obtained as described above. For secondary component carriers, the terminal does not need to read the system-information blocks but assumes that the information obtained for the primary component carrier also holds for the secondary component carriers. System information specific for the secondary component carrier is provided through dedicated RRC signaling as part of the procedure to configure an additional secondary component carrier. Using dedicated signaling instead of reading the system information on the secondary component carrier enables faster activation of secondary component carriers as the terminal otherwise would have to wait until the relevant system information had been transmitted.

14.3 RANDOM ACCESS

A fundamental requirement for any cellular system is the possibility for the terminal to request a connection setup, commonly referred to as *random access*. In LTE, random access is used for several purposes, including:

- for initial access when establishing a radio link (moving from RRC_IDLE to RRC_CONNECTED; see Chapter 8 for a discussion on different terminal states);
- to re-establish a radio link after radio-link failure;




FIGURE 14.8

Overview of the random-access procedure.

- for handover when uplink synchronization needs to be established to the new cell;
- to establish uplink synchronization if uplink or downlink data arrives when the terminal is in RRC_CONNECTED and the uplink is not synchronized;
- for the purpose of positioning using positioning methods based on uplink measurements;
- as a scheduling request if no dedicated scheduling-request resources have been configured on PUCCH (see Chapter 13 for a discussion on uplink scheduling procedures).

Acquisition of uplink timing is a main objective for all the cases above; when establishing an initial radio link (that is, when moving from RRC_IDLE to RRC_CONNECTED), the random-access procedure also serves the purpose of assigning a unique identity, the C-RNTI, to the terminal.

A terminal may perform random access on its primary component carrier only.⁵ Either a contention-based or a contention-free scheme can be used.

Contention-based random access uses a four-step procedure, illustrated in Figure 14.8, with the following steps:

1. The transmission of a random-access preamble, allowing the eNodeB to estimate the transmission timing of the terminal. Uplink synchronization is necessary as the terminal otherwise cannot transmit any uplink data.
2. The network transmits a timing advance command to adjust the terminal transmit timing, based on the timing estimate obtained in the first step. In addition to establishing uplink synchronization,

⁵The primary component carrier is UE specific, as already discussed; hence, from an eNodeB perspective, random access may occur on multiple component carriers.




FIGURE 14.9

Preamble subsets.

the second step also assigns uplink resources to the terminal to be used in the third step in the random-access procedure.

3. The transmission of the mobile-terminal identity to the network using the UL-SCH similar to normal scheduled data. The exact content of this signaling depends on the state of the terminal, in particular whether it is previously known to the network or not.
4. The final step consists of transmission of a contention-resolution message from the network to the terminal on the DL-SCH. This step also resolves any contention due to multiple terminals trying to access the system using the same random-access resource.

Only the first step uses physical-layer processing specifically designed for random access. The subsequent three steps utilize the same physical-layer processing as used for normal uplink and downlink data transmission. In the following, each of these steps is described in more detail.

Contention-free random access can only be used for re-establishing uplink synchronization upon downlink data arrival, handover, and positioning. Only the first two steps of the procedure above are used as there is no need for contention resolution in a contention-free scheme.

14.3.1 Step 1: Random-Access Preamble Transmission

The first step in the random-access procedure is the transmission of a random-access preamble. The main purpose of the preamble transmission is to indicate to the base station the presence of a random-access attempt and to allow the base station to estimate the delay between the eNodeB and the terminal. The delay estimate will be used in the second step to adjust the uplink timing.

The time-frequency resource on which the random-access preamble is transmitted is known as the *Physical Random-Access Channel* (PRACH). The network broadcasts information to all terminals in which time-frequency resource random-access preamble transmission is allowed (that is, the PRACH resources, in SIB-2). As part of the first step of the random-access procedure, the terminal selects one preamble to transmit on the PRACH.

In each cell, there are 64 preamble sequences available. Two subsets of the 64 sequences are defined as illustrated in Figure 14.9, where the set of sequences in each subset is signaled as part of the system information. When performing a (contention-based) random-access attempt, the terminal selects at random one sequence in one of the subsets. As long as no other terminal is performing a random-access attempt using the same sequence at the same time instant, no collisions will occur and the attempt will, with a high likelihood, be detected by the eNodeB.

The subset to select the preamble sequence from is given by the amount of data the terminal would like to (and from a power perspective can) transmit on the UL-SCH in the third random-access




FIGURE 14.10

Principal illustration of random-access preamble transmission.

step. Hence, from the preamble the terminal used, the eNodeB, will get some guidance on the amount of uplink resources to be granted to the terminal.

If the terminal has been requested to perform a contention-free random access, for example for handover to a new cell, the preamble to use is explicitly indicated from the eNodeB. To avoid collisions, the eNodeB should preferably select the contention-free preamble from sequences outside the two subsets used for contention-based random access.

14.3.1.1 PRACH Time-Frequency Resources

In the frequency domain, the PRACH resource, illustrated in Figure 14.10, has a bandwidth corresponding to six resource blocks (1.08 MHz). This nicely matches the smallest uplink cell bandwidth of six resource blocks in which LTE can operate. Hence, the same random-access preamble structure can be used, regardless of the transmission bandwidth in the cell.

In the time domain, the length of the preamble region depends on configured preamble, as will be discussed further below. The basic random-access resource is 1 ms in duration, but it is also possible to configure longer preambles. Also, note that the eNodeB uplink scheduler in principle can reserve an arbitrary long-random-access region by simply avoiding scheduling terminals in multiple subsequent subframes.

Typically, the eNodeB avoids scheduling any uplink transmissions in the time-frequency resources used for random access, resulting in the random-access preamble being *orthogonal* to user data. This avoids interference between UL-SCH transmissions and random-access attempts from different terminals. However, from a specification perspective, nothing prevents the uplink scheduler from scheduling transmissions in the random-access region. Hybrid-ARQ retransmissions are examples of this; synchronous non-adaptive hybrid-ARQ retransmissions may overlap with the random-access region and it is up to the implementation to handle this, either by moving the retransmissions in the frequency domain as discussed in Chapter 12 or by handling the interference at the eNodeB receiver.

For FDD, there is at most one random-access region per subframe – that is, multiple random-access attempts are not multiplexed in the frequency domain. From a delay perspective, it is better to spread out the random-access opportunities in the time domain to minimize the average waiting time before a random-access attempt can be initialized.

For TDD, multiple random-access regions can be configured in a single subframe. The reason is the smaller number of uplink subframes per radio frame in TDD. To maintain the same




FIGURE 14.11

Different preamble formats.

random-access capacity as in FDD, frequency-domain multiplexing is sometimes necessary. The number of random-access regions is configurable and can vary from one per 20 ms to one per 1 ms for FDD; for TDD up to six attempts per 10 ms radio frame can be configured.

14.3.1.2 Preamble Structure and Sequence Selection

The preamble consists of two parts:

- Preamble sequence.
- Cyclic prefix.

Furthermore, the preamble transmission uses a guard period to handle the timing uncertainty. Prior to starting the random-access procedure, the terminal has obtained downlink synchronization from the cell-search procedure. However, as uplink synchronization has not yet been established prior to random access, there is an uncertainty in the uplink timing⁶ as the location of the terminal in the cell is not known. The uplink timing uncertainty is proportional to the cell size and amounts to $6.7 \mu\text{s}/\text{km}$. To account for the timing uncertainty and to avoid interference with subsequent subframes not used for random access, a guard time is used as part of the preamble transmission – that is, the length of the actual preamble is shorter than 1 ms.

Including a cyclic prefix as part of the preamble is beneficial as it allows for frequency-domain processing at the base station (discussed further below), which can be advantageous from a complexity perspective. Preferably, the length of the cyclic prefix is approximately equal to the length

⁶The start of an uplink frame at the terminal is defined relative to the start of a downlink frame received at the terminal.

of the guard period. With a preamble sequence length of approximately 0.8 ms, there is 0.1 ms cyclic prefix and 0.1 ms guard time. This allows for cell sizes up to 15 km and is the typical random-access configuration, configuration 0 in Figure 14.11. To handle larger cells, where the timing uncertainty is larger, preamble configurations 1–3 can be used. Some of these configurations also support a longer preamble sequence to increase the preamble energy at the detector, which can be beneficial in larger cells. The preamble configuration used in a cell is signaled as part of the system information. Finally, note that guard times larger than those in Figure 14.11 can easily be created by not scheduling any uplink transmissions in the subframe following the random-access resource.

The preamble formats in Figure 14.11 are applicable to both FDD and TDD. However, for TDD, there is an additional fourth preamble configuration for random access. In this configuration, the random-access preamble is transmitted in the UpPTS field of the special subframe instead of in a normal subframe. Since this field is at most two OFDM symbols long, the preamble and the possible guard time are substantially shorter than the preamble formats described above. Hence, format 4 is applicable to very small cells only. The location of the UpPTS, next to the downlink-to-uplink switch for TDD, also implies that the interference from distant base stations may interfere with this short random-access format, which limits its usage to small cells and certain deployment scenarios.

14.3.1.3 PRACH Power Setting

The basis for setting the transmission power of the random-access preamble is a downlink path-loss estimate obtained from measuring the cell-specific reference signals on the primary downlink component carrier. From this path-loss estimate, the initial PRACH transmission power is obtained by adding a configurable offset.

The LTE random-access mechanism allows power ramping where the actual PRACH transmission power is increased for each unsuccessful random-access attempt. For the first attempt, the PRACH transmission power is set to the initial PRACH power. In most cases, this is sufficient for the random-access attempts to be successful. However, if the random-access attempt fails (random-access failures are detected at the second of four random-access steps, as described in the following sections), the PRACH transmission power for the next attempt is increased by a configurable step size to increase the likelihood of the next attempt being successful.

Since the random-access preamble is orthogonal to the user data, the need for power ramping to control intra-cell interference is smaller than in other systems with non-orthogonal random access and in many cases the transmission power is set such that the first random-access attempt with a high likelihood is successful. This is beneficial from a delay perspective.

14.3.1.4 Preamble Sequence Generation

The preamble sequences are generated from cyclic shifts of root Zadoff–Chu sequences [74]. Zadoff–Chu sequences are also used for creating the uplink reference signals as described in Chapter 11, where the structure of those sequences is described. From each root Zadoff–Chu sequence, $X_{ZC}^{(u)}(k)$, $[N_{ZC}/N_{CS}]$ cyclically shifted⁷ sequences are obtained by cyclic shifts of N_{CS} each, where N_{ZC} is the length of the root Zadoff–Chu sequence. The generation of the random-access preamble is illustrated

⁷The cyclic shift is in the time domain. Similar to the uplink reference signals and control signaling, this can equivalently be described as a phase rotation.



FIGURE 14.12

Random-access preamble generation.

in Figure 14.12. Although the figure illustrates generation in the time domain, frequency-domain generation can equally well be used in an implementation.

Cyclically shifted Zadoff–Chu sequences possess several attractive properties. The amplitude of the sequences is constant, which ensures efficient power amplifier utilization and maintains the low PAR properties of the single-carrier uplink. The sequences also have ideal cyclic auto-correlation, which is important for obtaining an accurate timing estimation at the eNodeB. Finally, the cross-correlation between different preambles based on cyclic shifts of the same Zadoff–Chu root sequence is zero at the receiver as long as the cyclic shift N_{CS} used when generating the preambles is larger than the maximum round-trip propagation time in the cell plus the maximum delay spread of the channel. Therefore, due to the ideal cross-correlation property, there is no intra-cell interference from multiple random-access attempts using preambles derived from the same Zadoff–Chu root sequence.

To handle different cell sizes, the cyclic shift N_{CS} is signaled as part of the system information. Thus, in smaller cells, a small cyclic shift can be configured, resulting in a larger number of cyclically shifted sequences being generated from each root sequence. For cell sizes below 1.5 km, all 64 preambles can be generated from a single root sequence. In larger cells, a larger cyclic shift needs to be configured and to generate the 64 preamble sequences, multiple root Zadoff–Chu sequences must be used in the cell. Although the larger number of root sequences is not a problem in itself, the zero cross-correlation property only holds between shifts of the *same* root sequence and from an interference perspective it is therefore beneficial to use as few root sequences as possible.

Reception of the random-access preamble is discussed further below, but in principle it is based on correlation of the received signal with the root Zadoff–Chu sequences. One disadvantage of Zadoff–Chu sequences is the difficulties in separating a frequency offset from the distance-dependent delay. A frequency offset results in an additional correlation peak in the time domain, a correlation peak that corresponds to a spurious terminal-to-base-station distance. In addition, the true correlation peak is attenuated. At low-frequency offsets, this effect is small and the detection performance is hardly affected. However, at high Doppler frequencies, the spurious correlation peak can be larger than the true peak. This results in erroneous detection; the correct preamble may not be detected or the delay estimate may be incorrect.

To avoid the ambiguities from spurious correlation peaks, the set of preamble sequences generated from each root sequence can be restricted. Restrictions imply that only some of the sequences that can be generated from a root sequence are used to define random-access preambles. Whether restriction should be applied or not to the preamble generation is signaled as part of the system information. The location of the spurious correlation peak relative to the “true” peak depends on the root sequence and hence different restrictions have to be applied to different root sequences. The restrictions to apply are broadcast as part of the system information in the cell.




FIGURE 14.13

Random-access preamble detection in the frequency domain.

14.3.1.5 Preamble Detection

The base-station processing is implementation specific, but due to the cyclic prefix included in the preamble, low-complexity frequency-domain processing is possible. An example hereof is shown in Figure 14.13. Samples taken in a time-domain window are collected and converted into the frequency-domain representation using an FFT. The window length is 0.8 ms, which is equal to the length of the Zadoff-Chu sequence without a cyclic prefix. This allows handling timing uncertainties up to 0.1 ms and matches the guard time defined for the basic preamble configuration.

The output of the FFT, representing the received signal in the frequency domain, is multiplied by the complex-conjugate frequency-domain representation of the root Zadoff-Chu sequence and the result is fed through an IFFT. By observing the IFFT outputs, it is possible to detect which of the shifts of the root Zadoff-Chu sequence has been transmitted and its delay. Basically, a peak of the IFFT output in interval i corresponds to the i th cyclically shifted sequence and the delay is given by the position of the peak within the interval. This frequency-domain implementation is computationally efficient and allows simultaneous detection of multiple random-access attempts using different cyclically shifted sequences generated from the same root Zadoff-Chu sequence; in the case of multiple attempts there will simply be a peak in each of the corresponding intervals.

14.3.2 Step 2: Random-Access Response

In response to the detected random-access attempt, the eNodeB will, as the second step of the random-access procedure, transmit a message on the DL-SCH, containing:

- The index of the random-access preamble sequences the network detected and for which the response is valid.
- The timing correction calculated by the random-access preamble receiver.

- A scheduling grant, indicating resources the terminal will use for the transmission of the message in the third step.
- A temporary identity, the TC-RNTI, used for further communication between the terminal and the network.

If the network detects multiple random-access attempts (from different terminals), the individual response messages of multiple terminals can be combined in a single transmission. Therefore, the response message is scheduled on the DL-SCH and indicated on a PDCCH using an identity reserved for random-access response, the RA-RNTI. All terminals that have transmitted a preamble monitor the L1/L2 control channels for random-access response within a configurable time window. The timing of the response message is not fixed in the specification in order to be able to respond to sufficiently many simultaneous accesses. It also provides some flexibility in the base-station implementation. If the terminal does not detect a random-access response within the time window, the attempt will be declared as failed and the procedure will repeat from the first step again, possibly with an increased preamble transmission power.

As long as the terminals that performed random access in the same resource used different preambles, no collision will occur and from the downlink signaling it is clear to which terminal(s) the information is related. However, there is a certain probability of contention – that is, multiple terminals using the same random-access preamble at the same time. In this case, multiple terminals will react upon the same downlink response message and a collision occurs. Resolving these collisions is part of the subsequent steps, as discussed below. Contention is also one of the reasons why hybrid ARQ is not used for transmission of the random-access response. A terminal receiving a random-access response intended for another terminal will have incorrect uplink timing. If hybrid ARQ were used, the timing of the hybrid-ARQ acknowledgement for such a terminal would be incorrect and may disturb uplink control signaling from other users.

Upon reception of the random-access response in the second step, the terminal will adjust its uplink transmission timing and continue to the third step. If contention-free random access using a dedicated preamble is used, then this is the last step of the random-access procedure as there is no need to handle contention in this case. Furthermore, the terminal already has a unique identity allocated in the form of a C-RNTI.

14.3.3 Step 3: Terminal Identification

After the second step, the uplink of the terminal is time synchronized. However, before user data can be transmitted to/from the terminal, a unique identity within the cell, the C-RNTI, must be assigned to the terminal. Depending on the terminal state, there may also be a need for additional message exchange for setting up the connection.

In the third step, the terminal transmits the necessary messages to the eNodeB using the UL-SCH resources assigned in the random-access response in the second step. Transmitting the uplink message in the same manner as scheduled uplink data instead of attaching it to the preamble in the first step is beneficial for several reasons. First, the amount of information transmitted in the absence of uplink synchronization should be minimized, as the need for a large guard time makes such transmissions relatively costly. Secondly, the use of the “normal” uplink transmission scheme for message transmission allows the grant size and modulation scheme to be adjusted to, for example, different radio conditions. Finally, it allows for hybrid ARQ with soft combining for the uplink message. The latter is

an important aspect, especially in coverage-limited scenarios, as it allows for the use of one or several retransmissions to collect sufficient energy for the uplink signaling to ensure a sufficiently high probability of successful transmission. Note that RLC retransmissions are not used for the uplink RRC signaling in step 3.

An important part of the uplink message is the inclusion of a terminal identity, as this identity is used as part of the contention-resolution mechanism in the fourth step. If the terminal is in the RRC_CONNECTED state – that is, connected to a known cell and therefore has a C-RNTI assigned – this C-RNTI is used as the terminal identity in the uplink message.⁸ Otherwise, a core-network terminal identifier is used and the eNodeB needs to involve the core network prior to responding to the uplink message in step 3.

UE-specific scrambling is used for transmission on UL-SCH, as described in Chapter 11. However, as the terminal may not yet have been allocated its final identity, the scrambling cannot be based on the C-RNTI. Instead, a temporary identity is used (TC-RNTI).

14.3.4 Step 4: Contention Resolution

The last step in the random-access procedure consists of a downlink message for contention resolution. Note that, from the second step, multiple terminals performing simultaneous random-access attempts using the same preamble sequence in the first step listen to the same response message in the second step and therefore have the same temporary identifier. Hence, in the fourth step, each terminal receiving the downlink message will compare the identity in the message with the identity transmitted in the third step. Only a terminal which observes a match between the identity received in the fourth step and the identity transmitted as part of the third step will declare the random-access procedure successful. If the terminal has not yet been assigned a C-RNTI, the TC-RNTI from the second step is promoted to the C-RNTI; otherwise the terminal keeps its already assigned C-RNTI.

The contention-resolution message is transmitted on the DL-SCH, using the temporary identity from the second step for addressing the terminal on the L1/L2 control channel. Since uplink synchronization has already been established, hybrid ARQ is applied to the downlink signaling in this step. Terminals with a match between the identity they transmitted in the third step and the message received in the fourth step will also transmit a hybrid-ARQ acknowledgement in the uplink.

Terminals that do not find a match between the identity received in the fourth step and the respective identity transmitted as part of the third step are considered to have failed the random-access procedure and need to restart the procedure from the first step. Obviously, no hybrid-ARQ feedback is transmitted from these terminals. Furthermore, a terminal that has not received the downlink message in step 4 within a certain time from the transmission of the uplink message in step 3 will declare the random-access procedure as failed and need to restart from the first step.

14.4 PAGING

Paging is used for network-initiated connection setup when the terminal is in RRC_IDLE. In LTE, the same mechanism as for “normal” downlink data transmission on the DL-SCH is used and the mobile terminal monitors the L1/L2 control signaling for downlink scheduling assignments related to paging.

⁸The terminal identity is included as a MAC control element on the UL-SCH.

Table 14.1 Paging Cycles and Paging Subframes

		Number of Paging Subframes per Paging Cycle							
		1/32	1/16	1/8	1/4	1/2	1	2	4
Paging subframes in a paging frame	FDD	9	9	9	9	9	9	4, 9	0, 4, 5, 9
	TDD	0	0	0	0	0	0	0, 5	0, 1, 5, 6


**FIGURE 14.14**

Illustration of paging cycles.

Since the location of the terminal typically is not known on a cell level, the paging message is typically transmitted across multiple cells in the so-called *tracking area* (the tracking area is controlled by the MME; see [9] for a discussion on tracking areas).

An efficient paging procedure should allow the terminal to sleep with no receiver processing most of the time and to briefly wake up at predefined time intervals to monitor paging information from the network. Therefore, a paging cycle is defined, allowing the terminal to sleep most of the time and only briefly wake up to monitor the L1/L2 control signaling. If the terminal detects a group identity used for paging (the P-RNTI) when it wakes up, it will process the corresponding downlink paging message transmitted on the PCH. The paging message includes the identity of the terminal(s) being paged, and a terminal not finding its identity will discard the received information and sleep according to the DRX cycle. Obviously, as the uplink timing is unknown during the DRX cycles, no hybrid-ARQ acknowledgements can be transmitted and consequently hybrid ARQ with soft combining is not used for paging messages.

The network configures in which subframes a terminal should wake up and listen for paging. Typically, the configuration is cell specific, although there is a possibility to complement the setting by UE-specific configuration. In which frame a given terminal should wake up and search for the P-RNTI on a PDCCH is determined by an equation taking as input the identity of the terminal as well as a cell-specific and (optionally) a UE-specific paging cycle. The identity used is the so-called IMSI, an identity coupled to the subscription, as an idle mode terminal does not have a C-RNTI allocated, and the paging cycle for a terminal can range from once per 256 up to once per 32 frames. The subframe within a frame to monitor for paging is also derived from the IMSI. Since different terminals

have different IMSI, they will compute different paging instances. Hence, from a network perspective, paging may be transmitted more often than once per 32 frames, although not all terminals can be paged at all paging occasions as they are distributed across the possible paging instances as shown in Figure 14.14.

Paging messages can only be transmitted in some subframes, ranging from one subframe per 32 frames up to a very high paging capacity with paging in four subframes in every frame. The configurations are shown in Table 14.1. Note that, from a network perspective, the cost of a short paging cycle is minimal as resources not used for paging can be used for normal data transmission and are not wasted. However, from a terminal perspective, a short paging cycle increases the power consumption as the terminal needs to wake up frequently to monitor the paging instants.

In addition to initiating connection to terminals being in RRC_IDLE, paging can also be used to inform terminals in RRC_IDLE as well as RRC_CONNECTED about changes of system information. A terminal being paged for this reason knows that the system information will change and therefore needs to acquire the update system information as described in Section 14.2.

page is typi-
controlled by

essing most
nation from
he time and
oup identity
link paging
nal(s) being
eep accord-
, no hybrid-
ining is not

for paging.
t the setting
arch for the
inal as well
called IMSI,
I allocated,
es. The sub-
nt terminals