DECLARATION OF JUNE ANN MUNFORD

1. My name is June Ann Munford. I am over the age of 18, have personal knowledge of the facts set forth herein, and am competent to testify to the same.

2. I earned a Master of Library and Information Science (MLIS) from the University of Wisconsin-Milwaukee in 2009. I have over ten years of experience in the library/information science field. Beginning in 2004, I have served in various positions in the public library sector including Assistant Librarian, Youth Services Librarian and Library Director. I have attached my Curriculum Vitae as Appendix CV.

3. During my career in the library profession, I have been responsible for materials acquisition for multiple libraries. In that position, I have cataloged, purchased and processed incoming library works. That includes purchasing materials directly from vendors, recording publishing data from the material in question, creating detailed material records for library catalogs and physically preparing that material for circulation. In addition to my experience in acquisitions, I was also responsible for analyzing large collections of library materials, tailoring library records for optimal catalog

search performance and creating lending agreements between libraries during my time as a Library Director.

4. I am fully familiar with the catalog record creation process in the library sector. In preparing a material for public availability, a library catalog record describing that material would be created. These records are typically written in Machine Readable Catalog (herein referred to as "MARC") code and contain information such as a physical description of the material, metadata from the material's publisher, and date of library acquisition. In particular, the 008 field of the MARC record is reserved for denoting the date of creation of the library record itself. As this typically occurs during the process of preparing materials for public access, it is my experience that an item's MARC record indicates the date of an item's public availability.

5. Typically, in creating a MARC record, a librarian would gather various bits of metadata such as book title, publisher and subject headings among others and assign each value to a relevant numerical field. For example, a book's physical description is tracked in field 300 while title/attribution is tracked in field 245. The 008 field of the MARC record is reserved for denoting the creation of the library record itself. As this is the only date reflecting the inclusion of said materials within the library's collection, it is my experience

that an item's 008 field accurately indicates the date of an item's public availability.

6. I have reviewed Exhibit APPLE-1023, 4G, LTE-Advanced Pro and the road to 5G by Erik Dahlman, et al.

7. Attached hereto as Appendix DAHLMAN01 is a true and correct copy of the MARC record for 4G, LTE-Advanced Pro and the road to 5G as held by the Penn State University Library. I secured this record myself from the library's public catalog. The MARC record contained within Appendix DAHLMAN01 accurately describes the title, author, publisher, and ISBN number of 4G, LTE-Advanced Pro and the road to 5G.

8. Attached hereto as Appendix DAHLMAN02 is a true and correct copy of 4G, LTE-Advanced Pro and the road to 5G as held by the Penn State University Library. These scans were secured from the library's collection. In comparing Exhibit APPLE-1023 to Appendix DAHLMAN02, it is my determination that Exhibit APPLE-1023 is a true and correct copy of 4G, LTE-Advanced Pro and the road to 5G by Erik Dahlman, et al.

9. The 008 field of the MARC record in Appendix DAHLMAN01 indicates the date of record creation. The 008 field of Appendix DAHLMAN01 indicates the Penn State University Library first acquired 4G, LTE-Advanced Pro and the road to 5G as of July 25, 2016. Considering this information, it is my determination that 4G, LTE-Advanced Pro and the road to 5G was made available to the public shortly after its initial acquisition as of July 25, 2016.

10. I have been retained on behalf of the Petitioner to provide assistance in the above-illustrated matter in establishing the authenticity and public availability of the documents discussed in this declaration. I am being compensated for my services in this matter at the rate of $200.00 per hour plus reasonable expenses. My statements are objective, and my compensation does not depend on the outcome of this matter.

11. I declare under penalty of perjury that the foregoing is true and correct. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: 4/24/2025



June Ann Munford

APPENDIX CV

June A. Munford
Curriculum Vitae

Education

University of Wisconsin-Milwaukee - MS, Library & Information Science, 2009
Milwaukee, WI

- Coursework included cataloging, metadata, data analysis, library systems, management strategies and collection development.
- Specialized in library advocacy, cataloging and public administration.

Grand Valley State University - BA, English Language & Literature, 2008
Allendale, MI

- Coursework included linguistics, documentation and literary analysis.
- Minor in political science with a focus in local-level economics and government.

Professional Experience

Researcher / Expert Witness, October 2017 – present
Freelance ● Pittsburgh, Pennsylvania & Grand Rapids, Michigan

- Material authentication and public accessibility determination. Declarations of authenticity and/or public accessibility provided upon research completion. Experienced with appeals and deposition process.

- Research provided on topics of public library operations, material publication history, digital database services and legacy web resources.

- Past clients include Alston & Bird, Arnold & Porter, Baker Botts, Fish & Richardson, Erise IP, Irell & Manella, O'Melveny & Myers, Perkins-Coie, Pillsbury Winthrop Shaw Pittman and Slayden Grubert Beard.

Library Director, February 2013 - March 2015
Dowagiac District Library ● Dowagiac, Michigan

- Executive administrator of the Dowagiac District Library. Located in Southwest Michigan, this library has a service area of 13,000, an annual operating budget of over $400,000 and total assets of approximately $1,300,000.

● Developed careful budgeting guidelines to produce a 15% surplus during the 2013-2014 & 2014-2015 fiscal years while being audited.

● Using this budget surplus, oversaw significant library investments including the purchase of property for a future building site, demolition of existing buildings and building renovation projects on the current facility.

● Led the organization and digitization of the library's archival records.

● Served as the public representative for the library, developing business relationships with local school, museum and tribal government entities.

● Developed an objective-based analysis system for measuring library services - including a full collection analysis of the library's 50,000+ circulating items and their records.

November 2010 - January 2013
Librarian & Branch Manager, Anchorage Public Library ● Anchorage, Alaska

● Headed the 2013 Anchorage Reads community reading campaign including event planning, staging public performances and creating marketing materials for mass distribution.

● Co-led the social media department of the library's marketing team, drafting social media guidelines, creating original content and instituting long-term planning via content calendars.

● Developed business relationships with The Boys & Girls Club, Anchorage School District and the US Army to establish summer reading programs for children.

June 2004 - September 2005, September 2006 - October 2013
Library Assistant, Hart Area Public Library
Hart, MI

● Responsible for verifying imported MARC records and original MARC cataloging for the local-level collection as well as the Michigan Electronic Library.

● Handled OCLC Worldcat interlibrary loan requests & fulfillment via ongoing communication with lending libraries.

Professional Involvement

Alaska Library Association - Anchorage Chapter

● Treasurer, 2012

Library Of Michigan
- ● Level VII Certification, 2008
- ● Level II Certification, 2013

Michigan Library Association Annual Conference 2014
- ● New Directors Conference Panel Member

Southwest Michigan Library Cooperative
- ● Represented the Dowagiac District Library, 2013-2015

Professional Development

Library Of Michigan Beginning Workshop, May 2008
Petoskey, MI
- ● Received training in cataloging, local history, collection management, children's literacy and reference service.

Public Library Association Intensive Library Management Training, October 2011
Nashville, TN
- ● Attended a five-day workshop focused on strategic planning, staff management, statistical analysis, collections and cataloging theory.

Alaska Library Association Annual Conference 2012 - Fairbanks, February 2012
Fairbanks, AK
- ● Attended seminars on EBSCO advanced search methods, budgeting, cataloging, database usage and marketing.

Depositions

2019 ● Fish & Richardson

Apple v. Qualcomm (Case No. IPR2018-001281, 39521-00421IP, IPR2018-01282 and 39521-00421IP2)

2019 ● Erise IP

Implicit, LLC v. Netscout Systems, Inc (Case No. 2:18-cv-53-JRG)

2019 ● Perkins-Coie

Adobe Inc. v. RAH Color Technologies LLC (Case No. IPR2019-00627, IPR2019-00628, IPR2019-00629 and IPR2019-00646)

2020 ● O'Melveny & Myers

Maxell, Ltd. v. Apple Inc. (Case No. 5:19-cv-00036-RWS)

2021 ● Pillsbury Winthrop Shaw Pittman LLP

Intel v. SRC (Case No. IPR2020-1449)

2022 ● Perkins-Coie

Realtek v. Future Link (Case No. IPR2021-01182)

2023 ● Fish & Richardson

Neuroderm Ltd. v. Abbvie, Inc (Case No. PGR2022-00040)

2023 ● Fish & Richardson

Nearmap US Inc. v. Pictometry International Corp. (Case No. IPR2022-00735)

2023 ● Fish & Richardson

Samsung Electronics v. MemoryWeb LLC (Case No. 39843-0136PS1)

2023 ● Pillsbury Winthrop Shaw Whitman LLP

Gravel Rating Systems v. Costco Wholesale Corp. (Case No. 4:21-cv-149-ALM)

2024 ● Willkie-Farr

Netflix, Inc. v. VideoLabs. Inc. (Case No. IPR2023-00628)

2024 ● Quinn Emanuel Urquhart & Sullivan

Purdue University v. Wolfspeed (Case No. 1:21-CV-840)

2024 ● Fish & Richardson

Dish Network v. Entropic Communications (Case No. IPR2024-00393)

2025 ● Perkins Coie

Amazon.com, Inc. v. Nokia Technologies Oy (Case No. IPR2024-00572)

2025 ● Fish & Richardson

Dish Network v. Entropic Communications (Case No. IPR2024-00373)

Limited Case and Clientele History

Alston & Bird

- Ericsson

  v. Collision Communications (Case No. IPR2022-01233)

- Nokia

  v. Neptune Subsea, Xtera (Case No. 1:17-cv-01876)

- Universal Electronics Inc

  v. Roku Inc (Case No. IPR2022-00818)

Alavi & Anaipakos PLLC

- Stingray Group Inc. and Stingray Music USA, Inc.

  v. Edwin A. Hernandez-Mondragon and Egla Corp. (Case No. 1:24-cv-21226-RAR)

Arnold & Porter

- Ivantis

  v. Glaukos (Case No. 8:18-cv-00620)

- Samsung

  v. Jawbone (Case No. 2:21-cv-00186)

Benesch Friedlander Coplan & Aronoff

- Voyis

  v. Cathx (Case No. 5:21-cv-00077-RWS)

Buchanan, Ingersoll & Rooney PC

- Google LLC, AT&T Service Inc., T Mobile USA Inc., Cellco Partnership, Ericsson Inc. & Nokia of America Corp.

  v. KT Corp & Pegasus Wireless Innovation LLC (Case No. IPR2025-00293)

Deschert LLP

- Smaxtec Animal Care

  v. ST Reproductive Technologies, LLC (Case No. IPR2024-00885)

Duane Morris

- Kangxi Communication Technologies

  v. Skyworks Solutions Inc. (Case No. IPR2025-00373)

- Microsoft

  v. Edge Networking Systems LLC (Case No. 1:24-cv-00215)

Erise I.P.

- Apple

  v. Ericsson Inc. (Case No. IPR2022-00715)

  v. Future Link Systems (Case No. IPRs 6317804, 6622108, 6807505, and 7917680)

  v. INVT (Case No. 20-1881)

  v. Navblazer LLC (Case No. IPR2020-01253)

  v. Qualcomm (Case No. IPR2018-001281, 39521-00421IP, IPR2018-01282, 39521-00421IP2)

v. Quest Nettech Corp (Case No. 2:19-cv-00118-JRG)

v. Telefonaktiebolaget LM Ericsson (Case No. IPR2022-00275)

v. Theta IP, LLC (Case No. IPR2024-00818)

v. THL Holding Company LLC (Case No. 1:23-cv-00548)

- Fanduel

v. CGT (Case No. 19-1393)

- Garmin

v. Phillips North America LLC (Case No. 2:19-cv-6301-AB-KS)

- Netscout

v. Longhorn HD LLC (Case No. 2:20-cv-00349)

v. Implicit, LLC (Case No. 2:18-cv-53-JRG)

- Sony Interactive Entertainment LLC

v. Bot M8 LLC (Case No. IPR2020-01288)

v. Infernal Technology LLC (Case No. 2:19-CV-00248-JRG)

- Tesla

v. Charge Fusion Technologies LLC (Case No. IPR2025-00032)

- Unified Patents

v. GE Video Compression (Case No. 2:19-cv-248)

Fish & Richardson

- Apple

v. AliveCor (Case No. 3:21-cv-03958)

v. LBS Innovations (Case No. 2:19-cv-00119-JRG-RSP)

v. Koss Corporation (Case No. IPR2021-00305)

v. Masimo (Case No. IPR 50095-0012IP1, 50095-0012IP2, 50095-0013IP1, 50095-0013IP2, 50095-0006IP1, 50095-0135IP1)

v. Neonode (Case No. 21-cv-08872-EMC)

v. Qualcomm (Case No. IPR2018-001281, 39521-00421IP, IPR2018-01282, 39521-00421IP2)

v. Resonant Sys. (Case No. 7:23-cv-00077-ADA)

- AsusTek Computer Inc.
  v. Videolabs, Inc. (Case No. 22-CV-00720-ADA)

- Dell
  v. Neo Wireless (Case No. IPR2022-00616)

- Dish Network

  v. Entropic Communications, LLC (Case No. 2:2023-CV-01043)

  v. Entropic Communications LLC (Case No. IPR2024-00393)

  v. Realtime Adaptive Streaming (Case No. 1:17-CV-02097-RBJ)

  v. TQ Delta LLC (Case No. 18-1798)

- Evapco Dry Cooling

  v. SPG Dry Cooling (Case No. IPR2021-00688)

- Genetec

  v. Sensormatic Electronics (Case No. 1:20-CV-00760)

- Huawei

  v. Bell Northern Research LLC (Case No. IPR2019-01174)

- Kianxis

  v. Blue Yonder (Case No. 3:20-cv-03636)

- LG Electronics

  v. Bell Northern Research LLC (Case No. 3:18-cv-2864-CAB-BLM)

- MediaTek

  v. MOSAID (Case No. IPR2024-00718)

- Metaswitch

  v. Sonus Networks (Case No. IPR2018-01719)

- Microsoft

  v. Throughputer Inc (Case No. IPR2022-00757)

- Mom Enterprises

  v. Ddrops Company (Case No. 1:22-cv-00332)

- MLC Intellectual Property

  v. MicronTech (Case No. 3:14-cv-03657-SI)

- Nearmap Inc

  v. EagleView Technologies (Case No. IPR2022-01009)

- Neuroderm Ltd.

  v. Abbvie, Inc (Case No. PGR2022-00040)

- Posco Co Ltd.

  v. Pascal Drillet et al. (Case No. IPR2025-00370)

- Realtek Semiconductor

  v. Future Link (Case No. IPR2021-01182)

- Quectel
  - v. Koninklijke Philips (Case No. 1:20-cv-01710)

- Samsung
  - v. Aire Technology (Case No. IPR2022-00877)
  - v. Bell Northern Research (Case No. 2:19-cv-00286-JRG)
  - v. Communication Technologies Inc (Case No. IPR2022-01221)
  - v. Jawbone Innovations (Case No. IPR2022-00865)
  - v. MemoryWeb LLC (Case No. IPR2022-00885)
  - v. Dodots Licensing Solutions (Case No. IPR2023-00701)
  - v. Telefonaktiebolaget LM Ericsson (Case No. IPR2021-00615)

- Texas Instruments
  - v. Vantage Micro LLC (Case No. IPR2020-01261)

- Xilinx
  - v. Sentient Sensors LCC (Case No. 1:22-cv-00173)

Goodwin Proctor

- Intel
  - v. Proxense, LLC (Case No. 6:24-cv-00283)

- Samsung Electronics
  - v. Mobile Data Technologies (Case No. 2:24-cv-00435)

- Taiwan Semiconductor Manufacturing Company Ltd.
  - v. Advanced Integrated Circuit Process LLC (Case No. 2:24-cv-00623)

Groombridge Wu

- Nearmap

  v. Eagle View Technologies Inc. and Pictometry International Corp. (Case No. 2:21-cv-00283-TS-DAO)

Irell & Manella

- Curium

Latham & Watkins LLP

- Apple

  v. Advanced Coding Research LLC

Leydig, Voit & Mayer, Ltd.

- Public Availability Consultancy

O'Melveny & Myers

- Apple

  v. Maxell (Case No. 5:19-cv-00036-RWS)

- Disney Media & Entertainment Distribution LLC

  v. Digital Media Technology Holdings, LLC (Case No. 2:22-cv-01642)

- Micron Technology Inc.

  v. Besang Inc (Case No. IPR2023-00900)

- Samsung

  v. Daedalus Prime LLC (1335 Investigation)

Perkins-Coie

- Amazon
  - v. Nokia (Case No. IPR2024-00847 and IPR2024-00848)
  - v. Nokia Technologies Oy (Case No. IPR2024-00572)
- Heru Industries
  - v. The UAB Research Foundation (Case No. IPR2022-01148)
- Intel Corporation
  - v. BeSang Inc. (Case No. IPR2023-00991)
  - v. BitMICRO (Case No. 5:23-cv-00625)
  - v. Telefonaktiebolaget L M Ericsson (Case No. IPR2024-00610)
- Realtek Semiconductor
  - v. Future Link (Case No. IPR2021-01182)
- r-pac
  - v. Adasa, Inc. (Case No. 1:2024cv06102)
- Twitter Inc
  - v. VOIP-Pal.com (Case No. 3:20-cv-02397-JD)
- TCL Industries
  - v. Koninklijke Philips NV (Case No. IPR2021-00495, IPR2021-00496 and IPR2021-00497)
- VusionGroup
  - v. Hanshow (Case No. IPR2024-00857)

Pillsbury Winthrop Shaw Pittman

- Intel

  v. FG SRC LLC (Case No. 6:20-cv-00315)

- Gravel Rating Systems

  v. Costco (Case No. 4:21-cv-149)
  v. Lowe's Home Centers (Case No. 4:21-cv-150)
  v. T-Mobile USA (Case No. 4:21-cv-152)
  v. Kohl's Inc. (Case No. 4:21-cv-258)
  v. Under Armor (Case No. 4:21-cv-356)

Quinn Emanuel

- Exact Sciences Corporation

  v. Geneoscopy, Inc. (Case No. IPR2024-00459 and IPR2024-1330)

- ServiceNow, Inc.

  v. InQuisient, Inc. (Case No. 22-900-CJB)

- Wolfspeed

  v. The Trustees of Purdue University (Case No. 1:2021-cv-00840)

Sheppard, Mullin, Richter & Hampton LLP

- Advanced Micro Devices, Inc and Pensando Systems, Inc.

  v. Concurrent Ventures LLC and Xtreamedge, Inc. (Case No. IPR2025-00478)

- Cadence Design Systems Inc.

  v. Semiconductor Design Technologies (Case No. 3:23-cv-01001)

Shook, Hardy & Bacon LLP

- HP Inc.

  v. Universal Connectivity Technologies (Case No. 1:23-cv-01177)

Stern, Kessler, Goldstein & Fox

- Preliminary Research

Wilmer, Cutler, Pickering, Hale and Dorr

- Apple

  v. Shunock (Case No. 1:23-cv-08598)

- Quest Diagnostics Inc., Haystack Oncology Inc. and Haystack Oncology GmbH

  v. Natera, Inc. (Case No. 1:23-cv-08598)

- Roche Diabetes, Inc.

  v. Trividia Health, Inc. (Case No. 1:24-cv-00668)

Willkie, Farr & Gallagher LLP

- Lenovo, Dell, HP

  v. Universal Connectivity Technologies Inc. (Case No. 2:2023cv00449)

- Neurent

  v. The Foundry, LLC (Case No. IPR2024-00669)

- Netflix, Inc.

  v. VideoLabs, Inc. (Case No. IPR2023-00628)

DAHLMAN01

PennState University Libraries | Catalog

Bookmarks 0 Course Reserves My Account

Keyword Search... Search Advanced search Start Over

Share Bookmark Report an Issue

MARC View

LEADER 05643cam a2200553li 4500			
001 39316530			
003 SIRSI			
005 20221202150813.0			
006 m o d			
007 cr cnu---unuuu			
008 160725s2016 enk ob 001 0 eng d			
019			a| 957617141 a| 958084370 a| 958350602 a| 958448405 a| 958778823
019			a| MARS
020			a| 9780128046111 q| (electronic bk.)
020			a| 0128046112 q| (electronic bk.)
020			z| 9780128045756
020			z| 0128045752
035			a| (OCoLC)953969117 z| (OCoLC)957617141 z| (OCoLC)958084370 z| (OCoLC)958350602 z| (OCoLC)958448405 z| (OCoLC)958778823
040			a| N$T b\| eng e\| rda e\| pn c\| N$T d| UIU d| N$T d| YDX d| OCLCF d| KNOVL d| OCLCQ d| Z5A d| U3W d| OCLCQ d| WYU d| S2H d| OCLCO d| HNC d| Y@Y d| UtOrBLW
049			a| UPMM
050		4	a| TK5103.48325 b| .D34 2016
072		7	a| TEC x| 009070 2| bisacsh
082	0	4	a| 621.38456 2| 23
100	1		a| Dahlman, Erik, e| author. 4| http://id.loc.gov/vocabulary/relators/aut 0| http://id.loc.gov/authorities/names/nb2007015616
245	1	0	a| 4G, LTE-advanced pro and the road to 5G / c| Erik Dahlman, Stefan Parkvall, Johan Sköld.
250			a| Third edition.
264		1	a| London : b| Academic Press, c| [2016]
264		4	c| ©2016
300			a| 1 online resource
336			a| text b| txt 2| rdacontent
337			a| computer b| c 2| rdamedia
338			a| online resource b| cr 2| rdacarrier
504			a| Includes bibliographical references and index.
505	0		a| Chapter 1. Introduction -- Chapter 2. Spectrum Regulation and Standardization From 3G to 5G -- Chapter 3. LTE Radio Access: An Overview -- Chapter 4. Radio-Interface Architecture -- Chapter 5. Physical Transmission Resources -- Chapter 6. Downlink Physical-Layer Processing -- Chapter 7. Uplink Physical-Layer Processing -- Chapter 8. Retransmission Protocols -- Chapter 9. Scheduling and Rate Adaptation -- Chapter 10. Channel-State Information and Full-Dimension MIMO -- Chapter 11. Access Procedures -- Chapter 12. Carrier Aggregation -- Chapter 13. Multi-Point Coordination and Transmission -- Chapter 14. Heterogeneous Deployments -- Chapter 15. Small-Cell Enhancements and Dynamic TDD -- Chapter 16. Dual Connectivity -- Chapter 17. Unlicensed Spectrum and License-Assisted Access -- Chapter 18. Relaying -- Chapter 19. Multimedia Broadcast/Multicast Services -- Chapter 20. LTE for Massive MTC Applications -- Chapter 21. Device-to-Device Connectivity -- Chapter 22. Spectrum and RF Characteristics -- Chapter 23. 5G Wireless Access -- Chapter 24. New 5G Radio-Access Technology -- Chapter 25. Concluding Remarks -- References -- Index.
520			a| The upcoming 5G specifications from 3GPP, to be available in 2018, will include LTE-Advanced Pro as well as a new 5G radio-access technology. This practical and very successful book, written by engineers working closely with 3GPP, gives insight into the newest technologies and standards adopted by 3GPP, with detailed explanations of the specific solutions chosen and their implementation in LTE, LTE-Advanced, and LTE-Advanced Pro, as well as providing a detailed description of the path to 5G and the associated underlying technologies. This edition has been thoroughly revised and updated to reflect the large extensions to LTE as introduced in 3GPP Releases 12 and 13 and the role of LTE in the upcoming 5G era. New to this edition includes updated content on: 4G and 5G Radio Access Spectrum for 4G and 5G Machine-Type Communication Device-to-Device Communication License-assisted Access Full-dimension MIMO Small-cell enhancements, eIMTA, FDD+TDD aggregation, dual connectivity Requirements on and general structure of 5G wireless access, addressing the existing and new usage scenarios for 5G Technical solutions for the new 5G radio-access technology The authors of this book all work at Ericsson Research and have been deeply involved in 3G and 4G development and standardization. They are leading experts in the field and are today actively contributing to the standardization of 4G and 5G within 3GPP. The leading book on 3GPP specifications for LTE, LTE-Advanced, and LTE-Advanced Pro covering up to and including Release 13, written by Ericsson engineers who are heavily involved in the development of 3GPP specifications Ten new chapters and coverage of all major features introduced with Release 12 and 13 Two completely new chapters on 5G wireless access including a detailed description of the key technology components under development by 3GPP.

ASK A LIBRARIAN

588	0		a| Online resource; title from PDF title page (EBSCO, viewed August 29, 2016).
650		0	a| Long-Term Evolution (Telecommunications) 0| http://id.loc.gov/authorities/subjects/sh2010010510
650		6	a| Technologie d'évolution à long terme (Norme)
650		7	a| TECHNOLOGY & ENGINEERING x| Mechanical. 2| bisacsh
650		7	a| Long-Term Evolution (Telecommunications) 2| fast 0| (OCoLC)fst01762769
655		0	a| Electronic books.
655		4	a| Electronic books.
700	1		a| Parkvall, Stefan, d| 1967- e| author. 4| http://id.loc.gov/vocabulary/relators/aut 0| http://id.loc.gov/authorities/names/no2011103561
700	1		a| Sköld, Johan, e| author. 4| http://id.loc.gov/vocabulary/relators/aut 0| http://id.loc.gov/authorities/names/no2011103562
776	0	8	i| Print version: z| 9780128045756 z| 0128045752 w| (OCoLC)936532306
856	4	0	u| https://ezaccess.libraries.psu.edu/login?url=https://app.knovel.com/hotlink/toc/id:kpGLTEAPR9/4g-lte-advanced?kpromoter=marc
938			a| EBSCOhost b| EBSC n| 1144736
938			a| YBP Library Services b| YANK n| 13089662
949			a| Electronic resource w| ASIS c| 1 i| 39316530-1001 l| ONLINE m| ONLINE r| Y s| Y t| ONLINE u| 11/10/2022 o| .STAFF. Knovel 8066 records from WCM 20221110

PennState

Copyright @2019 The Pennsylvania State University. All rights reserved. Except where otherwise noted, this work is subject to a Creative Commons Attribution 4.0 license. Details and exceptions.

Legal Statements | PSU Hotlines

PENN STATE UNIVERSITY LIBRARIES

Libraries Home
Libraries Intranet (Staff Only)
Accessibility Help
Website Feedback
Policies and Guidelines
Staff Directory

(814) 865-6368

CONNECT WITH PENN STATE UNIVERSITY LIBRARIES

Facebook
Twitter
Instagram

PennState University Libraries | Catalog

Bookmarks 0 Course Reserves My Account

Keyword Search... Search Advanced search Start Over

Share Bookmark Report an Issue



4G, LTE-advanced pro and the road to 5G / Erik Dahlman, Stefan Parkvall, Johan Sköld

Author	Dahlman, Erik
Published	London : Academic Press, [2016]
Copyright Date	©2016
Edition	Third edition.
Physical Description	1 online resource
Additional Creators	Parkvall, Stefan, 1967- and Sköld, Johan

ASK A LIBRARIAN

Access Online

ezaccess.libraries.psu.edu

Google Preview Search inside at Google Books

Contents

Chapter 1. Introduction -- Chapter 2. Spectrum Regulation and Standardization From 3G to 5G -- Chapter 3. LTE Radio Access: An Overview -- Chapter 4. Radio-Interface Architecture -- Chapter 5. Physical Transmission Resources -- Chapter 6. Downlink Physical-Layer Processing -- Chapter 7. Uplink Physical-Layer Processing -- Chapter 8. Retransmission Protocols -- Chapter 9. Scheduling and Rate Adaptation -- Chapter 10. Channel-State Information and Full-Dimension MIMO -- Chapter 11. Access Procedures -- Chapter 12. Carrier Aggregation -- Chapter 13. Multi-Point Coordination and Transmission -- Chapter 14. Heterogeneous Deployments -- Chapter 15. Small-Cell Enhancements and Dynamic TDD -- Chapter 16. Dual Connectivity -- Chapter 17. Unlicensed Spectrum and License-Assisted Access -- Chapter 18. Relaying -- Chapter 19. Multimedia Broadcast/Multicast Services -- Chapter 20. LTE for Massive MTC Applications -- Chapter 21. Device-to-Device Connectivity -- Chapter 22. Spectrum and RF Characteristics -- Chapter 23. 5G Wireless Access -- Chapter 24. New 5G Radio-Access Technology -- Chapter 25. Concluding Remarks -- References -- Index.

Summary

The upcoming 5G specifications from 3GPP, to be available in 2018, will include LTE-Advanced Pro as well as a new 5G radio-access technology. This practical and very successful book, written by engineers working closely with 3GPP, gives insight into the newest technologies and standards adopted by 3GPP, with detailed explanations of the specific solutions chosen and their implementation in LTE, LTE-Advanced, and LTE-Advanced Pro, as well as providing a detailed description of the path to 5G and the associated underlying technologies. This edition has been thoroughly revised and updated to reflect the large extensions to LTE as introduced in 3GPP Releases 12 and 13 and the role of LTE in the upcoming 5G era. New to this edition includes updated content on: 4G and 5G Radio Access Spectrum for 4G and 5G Machine-Type Communication Device-to-Device Communication License-assisted Access Full-dimension MIMO Small-cell enhancements, eIMTA, FDD+TDD aggregation, dual connectivity Requirements on and general structure of 5G wireless access, addressing the existing and new usage scenarios for 5G Technical solutions for the new 5G radio-access technology The authors of this book all work at Ericsson Research and have been deeply involved in 3G and 4G development and standardization. They are leading experts in the field and are today actively contributing to the standardization of 4G and 5G within 3GPP. The leading book on 3GPP specifications for LTE, LTE-Advanced, and LTE-Advanced Pro covering up to and including Release 13, written by Ericsson engineers who are heavily involved in the development of 3GPP specifications Ten new chapters and coverage of all major features introduced with Release 12 and 13 Two completely new chapters on 5G wireless access including a detailed description of the key technology components under development by 3GPP.

Subject(s)

Long-Term Evolution (Telecommunications)
Technologie d'évolution à long terme (Norme)
TECHNOLOGY & ENGINEERING—Mechanical

Genre(s)

Electronic books

ISBN

9780128046111 (electronic bk.)
0128046112 (electronic bk.)

9780128045756
0128045752

Bibliography Note

Includes bibliographical references and index.

View MARC record | catkey: 39316530

PennState

Copyright @2019 The Pennsylvania State University. All rights reserved. Except where otherwise noted, this work is subject to a Creative Commons Attribution 4.0 license. Details and exceptions.

Legal Statements | PSU Hotlines

PENN STATE UNIVERSITY LIBRARIES

Libraries Home
Libraries Intranet (Staff Only)
Accessibility Help
Website Feedback
Policies and Guidelines
Staff Directory

(814) 865-6368

CONNECT WITH PENN STATE UNIVERSITY LIBRARIES

Facebook
Twitter
Instagram

DAHLMAN02

Third Edition

4G

LTE-Advanced Pro and The Road to 5G

Erik Dahlman
Stefan Parkvall
Johan Sköld





Third Edition

4G LTE-Advanced Pro and The Road to 5G

Dahlman
Parkvall
Sköld

ACADEMIC
PRESS

Third Edition

4G

LTE-Advanced Pro and The Road to 5G

Learn the latest developments to LTE-Advanced Pro (Release 13) and the path to 5G

The upcoming 5G specifications from 3GPP, to be available in 2018, will include LTE-Advanced Pro as well as a new 5G radio-access technology. This practical and very successful book, written by engineers working closely with 3GPP, gives insight into the newest technologies and standards adopted by 3GPP, with detailed explanations of the specific solutions chosen and their implementation in LTE, LTE-Advanced, and LTE-Advanced Pro, as well as providing a detailed description of the path to 5G and the associated underlying technologies.

- The leading book on 3GPP specifications for LTE, LTE-Advanced, and LTE-Advanced Pro covering up to and including Release 13, written by Ericsson engineers who are heavily involved in the development of 3GPP specifications
- Ten new chapters and coverage of all major features introduced with Releases 12 and 13
- Two completely new chapters on 5G wireless access including a description of key technology components for the new 5G radio-access technology under development by 3GPP

This edition has been thoroughly revised and updated to reflect the large extensions to LTE as introduced in 3GPP Releases 12 and 13 and the role of LTE in the upcoming 5G era. New to this edition includes updated content on:

- 4G and 5G Radio Access
- Spectrum for 4G and 5G
- Machine-Type Communication
- Device-to-Device Communication
- License-assisted Access
- Full-dimension MIMO
- Small-cell enhancements, eIMTA, FDD + TDD aggregation, dual connectivity
- Requirements on and general structure of 5G wireless access, addressing the existing and new usage scenarios for 5G
- Technical solutions for the new 5G radio-access technology

The authors of this book all work at Ericsson Research and have been deeply involved in 3G and 4G development and standardization. They are leading experts in the field and are today actively contributing to the standardization of 4G and 5G within 3GPP.

ENGINEERING / COMMUNICATIONS

ISBN 978-0-12-804575-6

ELSEVIER

ACADEMIC PRESS

An imprint of Elsevier

elsevier.com

X001AIPVSN

New - 4G, LTE-Advanced Pro and The Road to 5G, Third Edition

9780128045756

4G, LTE-Advanced Pro and The Road to 5G

4G, LTE-Advanced Pro and The Road to 5G

Third Edition

Erik Dahlman

Stefan Parkvall

Johan Sköld



ELSEVIER

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
125 London Wall, London EC2Y 5AS, United Kingdom
525 B Street, Suite 1800, San Diego, CA 92101-4495, United States
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, United States
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, United Kingdom

Copyright © 2016, 2014, 2011 Erik Dahlman, Stefan Parkvall and Johan Sköld.
Published by Elsevier Ltd. All rights reserved.

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 arrangements 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.

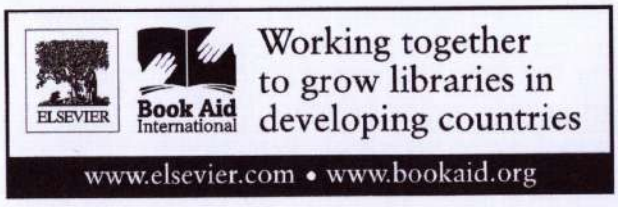
To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions,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

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-0-12-804575-6

For information on all Academic Press publications visit our website at https://www.elsevier.com/



Publisher: Joe Hayton
Acquisition Editor: Tim Pitts
Editorial Project Manager: Charlotte Kent
Production Project Manager: Lisa Jones
Designer: Greg Harris

Typeset by TNQ Books and Journals
Last digit is the print number: 10 9 8 7 6 5 4 3 2

Contents

Preface xv
Acknowledgments xvii
Abbreviations and Acronyms xix

CHAPTER 1 Introduction 1
1.1 1G and 2G—Voice-Centric Technologies 1
1.2 3G and 4G—Mobile Broadband 2
1.3 5G—Beyond Mobile Broadband—Networked Society 4
1.4 Outline 5

CHAPTER 2 Spectrum Regulation and Standardization From 3G to 5G 7
2.1 Overview of Standardization and Regulation 7
2.2 ITU-R Activities From 3G to 5G 8
2.2.1 The Role of ITU-R 8
2.2.2 IMT-2000 and IMT-Advanced 10
2.2.3 IMT-2020 11
2.3 Spectrum for Mobile Systems 13
2.3.1 Spectrum Defined for IMT Systems by the ITU-R 14
2.3.2 Frequency Bands for LTE 15
2.3.3 New Frequency Bands 20
2.4 Spectrum for 5G 21
2.4.1 New Frequency Bands to be Studied by WRC 21
2.4.2 RF Exposure Above 6 GHz 23
2.5 3GPP Standardization 23
2.5.1 The 3GPP Process 24
2.5.2 Specifying 5G in 3GPP 27

CHAPTER 3 LTE Radio Access: An Overview 29
3.1 LTE Basic Technologies 30
3.1.1 Transmission Scheme 31
3.1.2 Channel-Dependent Scheduling and Rate Adaptation 32
3.1.3 Inter-Cell Interference Coordination 34
3.1.4 Multi-Antenna Transmission 35
3.1.5 Spectrum Flexibility 36
3.1.6 Multicast and Broadcast Support 38
3.1.7 Positioning 39
3.2 LTE Evolution 39

3.3 Spectrum Flexibility 41
3.3.1 Carrier Aggregation 41
3.3.2 License-Assisted Access 43
3.4 Multi-Antenna Enhancements 44
3.4.1 Extended Multi-Antenna Transmission 44
3.4.2 Multi-Point Coordination and Transmission 45
3.4.3 Enhanced Control-Channel Structure 46
3.5 Densification, Small Cells, and Heterogeneous Deployments 46
3.5.1 Relaying 47
3.5.2 Heterogeneous Deployments 47
3.5.3 Small-Cell On-Off 48
3.5.4 Dual Connectivity 48
3.5.5 Dynamic TDD 48
3.5.6 WLAN Interworking 49
3.6 Device Enhancements 49
3.7 New Scenarios 50
3.7.1 Device-to-Device Communication 50
3.7.2 Machine-Type Communication 50
3.8 Device Capabilities 51

CHAPTER 4 Radio-Interface Architecture 55
4.1 Overall System Architecture 55
4.1.1 Core Network 55
4.1.2 Radio-Access Network 57
4.2 Radio Protocol Architecture 58
4.2.1 Packet-Data Convergence Protocol 61
4.2.2 Radio-Link Control 61
4.2.3 Medium-Access Control 62
4.2.4 Physical Layer 70
4.3 Control-Plane Protocols 72
4.3.1 State Machine 73

CHAPTER 5 Physical Transmission Resources 75
5.1 Overall Time–Frequency Structure 75
5.2 Normal Subframes and MBSFN Subframes 79
5.3 Antenna Ports 81
5.3.1 Quasi-co-Located Antenna Ports 82
5.4 Duplex Schemes 82
5.4.1 Frequency-Division Duplex 83
5.4.2 Time-Division Duplex 84
5.4.3 LTE and TD-SCDMA Coexistence 88
5.4.4 License-Assisted Access 89
5.5 Carrier Aggregation 89
5.6 Frequency-Domain Location of LTE Carriers 90

CHAPTER 6 Downlink Physical-Layer Processing 93
6.1 Transport-Channel Processing 93
6.1.1 Processing Steps 94
6.1.2 Localized and Distributed Resource Mapping 100
6.2 Downlink Reference Signals 103
6.2.1 Cell-Specific Reference Signals 104
6.2.2 Demodulation Reference Signals 107
6.2.3 CSI Reference Signals 111
6.2.4 Quasi-Colocation Relations 116
6.3 Multi-Antenna Transmission 116
6.3.1 Transmission Modes 117
6.3.2 Transmit Diversity 119
6.3.3 Codebook-Based Precoding 121
6.3.4 Non-Codebook-Based Precoding 125
6.3.5 Downlink MU-MIMO 127
6.4 Downlink L1/L2 Control Signaling 130
6.4.1 Physical Control Format Indicator Channel 132
6.4.2 Physical Hybrid-ARQ Indicator Channel 135
6.4.3 Physical Downlink Control Channel 138
6.4.4 Enhanced Physical Downlink Control Channel 144
6.4.5 Blind Decoding of PDCCHs and EPDCCHs 149
6.4.6 Downlink Scheduling Assignments 153
6.4.7 Uplink Scheduling Grants 163
6.4.8 Power-Control Commands 168
CHAPTER 7 Uplink Physical-Layer Processing 169
7.1 Transport-Channel Processing 169
7.1.1 Processing Steps 169
7.1.2 Mapping to the Physical Resource 172
7.1.3 PUSCH Frequency Hopping 173
7.2 Uplink Reference Signals 176
7.2.1 Demodulation Reference Signals 177
7.2.2 Sounding Reference Signals 185
7.3 Uplink Multi-Antenna Transmission 189
7.3.1 Precoder-Based Multi-Antenna Transmission for PUSCH 189
7.3.2 Uplink Multi-User MIMO 193
7.3.3 PUCCH Transmit Diversity 194
7.4 Uplink L1/L2 Control Signaling 195
7.4.1 Basic PUCCH Structure 196
7.4.2 Uplink Control Signaling on PUCCH 206
7.4.3 Uplink L1/L2 Control Signaling on PUSCH 213

7.5 Uplink Power Control 216
7.5.1 Uplink Power Control: Some Basic Rules 217
7.5.2 Power Control for PUCCH 219
7.5.3 Power Control for PUSCH 221
7.5.4 Power Control for SRS 223
7.6 Uplink Timing Alignment 223

CHAPTER 8 Retransmission Protocols 227
8.1 Hybrid ARQ with Soft Combining 228
8.1.1 Downlink Hybrid ARQ 233
8.1.2 Uplink Hybrid ARQ 234
8.1.3 Hybrid-ARQ Timing 236
8.2 Radio-Link Control 241
8.2.1 Segmentation, Concatenation, and Reassembly of RLC SDUs 242
8.2.2 RLC Retransmission 243
8.2.3 In-Sequence Delivery 244
8.2.4 RLC Operation 244

CHAPTER 9 Scheduling and Rate Adaptation 249
9.1 Scheduling Strategies 250
9.2 Downlink Scheduling 254
9.3 Uplink Scheduling 255
9.3.1 Uplink Priority Handling 256
9.3.2 Scheduling Requests 257
9.3.3 Buffer Status Reports 259
9.3.4 Power Headroom Reports 260
9.4 Timing of Scheduling Assignments/Grants 261
9.4.1 Downlink Scheduling Timing 261
9.4.2 Uplink Scheduling Timing 262
9.5 Semi-Persistent Scheduling 263
9.6 Scheduling for Half-Duplex FDD 264
9.7 Discontinuous Reception 266

CHAPTER 10 Channel-State Information and Full-Dimension MIMO 269
10.1 CSI Reports 269
10.2 Periodic and Aperiodic CSI Reports 270
10.2.1 Aperiodic CSI Reports 271
10.2.2 Periodic CSI Reports 273
10.3 Interference Estimation 273
10.4 Channel-Quality Indicator 276
10.5 Rank Indicator and Precoder Matrix Indicator 276

10.6 Full-Dimension MIMO 279
10.6.1 CSI Feedback for Massive Antenna Arrangements 279
10.6.2 CSI Reporting Class A 280
10.6.3 CSI Reporting Class B 281

CHAPTER 11 Access Procedures 285
11.1 Acquisition and Cell Search 285
11.1.1 Overview of LTE Cell Search 285
11.1.2 PSS Structure 288
11.1.3 SSS Structure 288
11.2 System Information 289
11.2.1 MIB and BCH Transmission 290
11.2.2 System-Information Blocks 293
11.3 Random Access 296
11.3.1 Step 1: Random-Access Preamble Transmission 298
11.3.2 Step 2: Random-Access Response 304
11.3.3 Step 3: Device Identification 305
11.3.4 Step 4: Contention Resolution 306
11.4 Paging 307

CHAPTER 12 Carrier Aggregation 309
12.1 Overall Protocol Structure 312
12.2 Primary and Secondary Component Carriers 313
12.3 Self-Scheduling and Cross-Carrier Scheduling 314
12.3.1 Scheduling Timing for Aggregation of FDD Carriers 315
12.3.2 Scheduling Timing for Aggregation of TDD Carriers 315
12.3.3 Scheduling Timing for Aggregation of FDD and TDD Carriers 318
12.4 DRX and Component-Carrier Deactivation 319
12.5 Downlink Control Signaling 322
12.5.1 PCFICH 322
12.5.2 PHICH 322
12.5.3 PDCCH and EPDCCH 323
12.6 Uplink Control Signaling 325
12.6.1 Hybrid-ARQ Acknowledgments on PUCCH 326
12.6.2 CSI Reporting on PUCCH 328
12.6.3 Control Signaling on PUSCH 329
12.7 Timing Advance and Carrier Aggregation 330

CHAPTER 13 Multi-Point Coordination and Transmission 331
13.1 Inter-Cell Interference Coordination 334
13.2 Multi-Point Coordination/Transmission 335
13.2.1 Multi-Point Coordination 336

13.2.2 Multi-Point Transmission 341
13.2.3 Uplink Multi-Point Coordination/Reception 345

CHAPTER 14 Heterogeneous Deployments 347
14.1 Interference Scenarios in Heterogeneous Deployments 348
14.2 Heterogeneous Deployments Using Rel-8 Functionality 352
14.3 Frequency-Domain Partitioning 353
14.4 Time-Domain Partitioning 354
14.5 Shared Cell 358
14.6 Closed Subscriber Groups 361

CHAPTER 15 Small-Cell Enhancements and Dynamic TDD 363
15.1 Small-Cell On/Off 363
15.1.1 Discovery Signal and Associated Measurements 365
15.2 Dynamic TDD and eIMTA 366
15.2.1 Basic Principles of eIMTA 367
15.2.2 Scheduling and Hybrid-ARQ Retransmissions 370
15.2.3 RRM Measurements and CSI Reporting 370
15.2.4 Uplink Power Control 371
15.2.5 Inter-Cell Interference Coordination 371

CHAPTER 16 Dual Connectivity 373
16.1 Architecture 375
16.2 Physical-Layer Impact 378
16.2.1 Timing 378
16.2.2 Power Control 379
16.3 Scheduling in Dual Connectivity 381

CHAPTER 17 Unlicensed Spectrum and License-Assisted Access 383
17.1 Spectrum for LAA 385
17.2 Wi-Fi Basics 386
17.3 Technology Components for LAA 389
17.3.1 Dynamic Frequency Selection 391
17.3.2 Listen Before Talk 391
17.3.3 Frame Structure and Burst Transmission 395
17.3.4 Reference Signals and Discontinuous Transmission 397
17.3.5 Scheduling, Hybrid-ARQ, and Retransmissions 399
17.3.6 Radio Bearer Mapping and QoS Control 399
17.4 Enhancements Beyond Release 13 400

CHAPTER 18 Relaying 403
18.1 Relays in LTE 404
18.2 Overall Architecture 405

18.3 Backhaul Design for Inband Relaying 407
18.3.1 Access-Link Hybrid-ARQ Operation 408
18.3.2 Backhaul-Link Hybrid-ARQ Operation 409
18.3.3 Backhaul Downlink Control Signaling 410
18.3.4 Reference Signals for the Backhaul Link 415
18.3.5 Backhaul–Access Link Timing 415

CHAPTER 19 Multimedia Broadcast/Multicast Services 421
19.1 Architecture 423
19.2 MBSFN Channel Structure and Physical-Layer Processing 425
19.3 Scheduling of MBSFN Services 428
19.4 Single-Cell Point to Multipoint Transmission 431

CHAPTER 20 LTE for Massive MTC Applications 433
20.1 Introduction 433
20.2 MTC Enhancements for LTE Release 12 434
20.2.1 Data-Rate Capability and UE Category 0 434
20.2.2 Type-B Half-Duplex Operation 435
20.2.3 Possibility for Devices With a Single Receive Antenna 436
20.2.4 Power-Saving Mode 436
20.3 MTC Enhancements for LTE Release 13: eMTC 436
20.3.1 Narrow–Band Operation 437
20.3.2 Coverage Enhancements by Means of Repetition 439
20.3.3 Downlink Transmission: PDSCH and MPDCCH 441
20.3.4 Uplink Transmission: PUSCH and PUCCH 446
20.3.5 Synchronization Signals and BCH 450
20.3.6 System-Information Blocks 452
20.3.7 Random Access 453
20.3.8 Extended DRX 454
20.4 Narrow-Band Internet of Things 455
20.4.1 Background 455
20.4.2 NB-IoT Deployment Modes 456
20.4.3 Downlink Data Transmission 456
20.4.4 Uplink Transmission 457
20.4.5 NB-IoT System Information 459

CHAPTER 21 Device-to-Device Connectivity 461
21.1 Overview 461
21.1.1 Sidelink Transmission 462
21.1.2 In-Coverage vs. Out-of-Coverage Sidelink Connectivity 463
21.1.3 Sidelink Synchronization 464
21.1.4 Configuration for Sidelink Connectivity 466

21.1.5 Architecture for Sidelink 466
21.1.6 Sidelink Channel Structure 467
21.2 Sidelink Communication 468
21.2.1 Resource Pools and Assignment/Selection of Transmission Resources 468
21.2.2 Physical Sidelink Control Channel Periods 469
21.2.3 Sidelink Control Information/Physical Sidelink Control Channel Transmission 469
21.2.4 Sidelink Shared Channel/Physical Sidelink Shared Channel Transmission 472
21.2.5 Sidelink Control Information Content 474
21.2.6 Scheduling Grants and DCI Format 5 475
21.2.7 Reception Resource Pools 476
21.3 Sidelink Discovery 477
21.3.1 Resource Pools and Selection/Assignment of Transmission Resources 477
21.3.2 Discovery Transmission 479
21.3.3 Reception Resource Pools 479
21.4 Sidelink Synchronization 479
21.4.1 Sidelink Identity and Structure of the Sidelink Synchronization Signal 480
21.4.2 The Sidelink Broadcast Channel and Sidelink Master Information Block 482
21.4.3 SyncRef UE Selection 482
21.4.4 Transmission of Sidelink Synchronization Signals 483
21.5 Device-to-Device Extensions in LTE Release 13 486
21.5.1 Out-of-Coverage Discovery 486
21.5.2 Layer-3 Relaying 486

CHAPTER 22 Spectrum and RF Characteristics 487
22.1 Flexible Spectrum Use 487
22.2 Flexible Channel Bandwidth Operation 488
22.3 Carrier Aggregation for LTE 491
22.4 Operation in Noncontiguous Spectrum 495
22.5 Multi-Standard Radio Base Stations 496
22.6 Overview of RF Requirements for LTE 500
22.6.1 Transmitter Characteristics 501
22.6.2 Receiver Characteristics 502
22.6.3 Regional Requirements 503
22.6.4 Band-Specific Terminal Requirements Through Network Signaling 503
22.6.5 Base-Station Classes 504

22.7 Output Power Level Requirements 505
22.7.1 Base-Station Output Power and Dynamic Range 505
22.7.2 Terminal Output Power and Dynamic Range 506
22.8 Transmitted Signal Quality 506
22.8.1 EVM and Frequency Error 507
22.8.2 Terminal In-Band Emissions 507
22.8.3 Base-Station Time Alignment 507
22.9 Unwanted Emissions Requirements 507
22.9.1 Implementation Aspects 508
22.9.2 Spectrum Emission Mask 508
22.9.3 Adjacent Channel Leakage Ratio 510
22.9.4 Spurious Emissions 512
22.9.5 Occupied Bandwidth 513
22.9.6 Transmitter Intermodulation 513
22.10 Sensitivity and Dynamic Range 513
22.11 Receiver Susceptibility to Interfering Signals 514
22.12 Multiband-Capable Base Stations 516
22.13 RF Requirements for Relays 520
22.14 RF Requirements for License-Assisted Access 521
22.14.1 Regulatory Requirements for the Unlicensed 5 GHz Band 521
22.14.2 Specific BS RF Requirements for LAA Operation 522
22.14.3 Specific Terminal RF Requirements for LAA Operation 524
22.15 RF Requirements for BS With Active Antenna Systems 524

CHAPTER 23 5G Wireless Access 527
23.1 What Is 5G? 527
23.1.1 Data Rates 528
23.1.2 Latency 528
23.1.3 Extreme Reliability 529
23.1.4 Low-Cost Devices With Very Long Battery Life 529
23.1.5 Network Energy Efficiency 529
23.2 5G and IMT-2020 530
23.2.1 Usage Scenarios for IMT-2020 531
23.2.2 Capabilities of IMT-2020 533
23.2.3 Studies of 5G in Regional and Operator Groups 536
23.3 One Versus Multiple Technologies: "Network Slicing" 538
23.4 5G Spectrum 539
23.4.1 Expansion Into Higher-Frequency Bands 539
23.4.2 Licensed Versus Unlicensed Spectrum 542

23.5 LTE Evolution Versus New 5G Technology 542
23.6 Frequency Bands for 5G Initial Deployments 544
23.7 5G Technical Specification 544

CHAPTER 24 New 5G Radio-Access Technology 547
24.1 5G: Some General Design Principles 547
24.1.1 Radio-Access Evolution and Forward Compatibility 547
24.1.2 Ultra-Lean Design: Minimize "Always-On" Transmissions 548
24.1.3 Stay in the Box 550
24.1.4 Avoid Strict Timing Relations 551
24.2 5G: Key Technology Components 551
24.2.1 Waveform 551
24.2.2 Flexible Duplex 557
24.2.3 Frame Structure 560
24.2.4 Channel Coding 562
24.2.5 Multi-Antenna Transmission and Beam-Forming 563
24.2.6 Multi-Site Connectivity and Tight Interworking 565
24.2.7 System-Access Functionality 566
24.2.8 Scheduled and Contention-Based Transmissions 569
24.2.9 New Types of Wireless Links 570

CHAPTER 25 Concluding Remarks 575

References 577
Index 581

Preface

LTE has become the most successful mobile wireless broadband technology, serving over one billion users as of the beginning of 2016 and handling a wide range of applications. Compared to the analog voice-only systems 25 years ago, the difference is dramatic. Although LTE is still at a relatively early stage of deployment, the industry is already well on the road toward the next generation of mobile communication, commonly referred to as the fifth generation or 5G. Mobile broadband is, and will continue to be, an important part of future cellular communication, but future wireless networks are to a large extent also about a significantly wider range of use cases and a correspondingly wider range of requirements.

This book describes LTE, developed in 3GPP (*Third-Generation Partnership Project*) and providing true fourth-generation (4G) broadband mobile access, as well as the new radio-access technology 3GPP is currently working on. Together, these two technologies will provide 5G wireless access.

Chapter 1 provides a brief introduction, followed by a description of the standardization process and relevant organizations such as the aforementioned 3GPP and ITU in Chapter 2. The frequency bands available for mobile communication are also be covered, together with a discussion on the process for finding new frequency bands.

An overview of LTE and its evolution is found in Chapter 3. This chapter can be read on its own to get a high-level understanding of LTE and how the LTE specifications evolved over time. To underline the significant increase in capabilities brought by the LTE evolution, 3GPP introduced the names LTE-Advanced and LTE-Advanced Pro for some of the releases.

Chapters 4–11 cover the basic LTE structure, starting with the overall protocol structure in Chapter 4 and followed by a detailed description of the physical layer in Chapters 5–7. The remaining Chapters 8–11, cover connection setup and various transmission procedures, including multi-antenna support.

Some of the major enhancements to LTE introduced over time is covered in Chapters 12–21, including carrier aggregation, unlicensed spectrum, machine-type communication, and device-to-device communication. Relaying, heterogeneous deployments, broadcast/multicast services, and dual connectivity multi-site coordination are other examples of enhancements covered in these chapters.

Radio frequency (RF) requirements, taking into account spectrum flexibility and multi-standard radio equipment, is the topic of Chapter 22.

Chapters 23 and 24 cover the new radio access about to be standardized as part of 5G. A closer look on the requirements and how they are defined is the topic of Chapter 23, while Chapter 24 digs into the technical realization.

Finally, Chapter 25 concludes the book and the discussion on 5G radio access.

CHAPTER

3 LTE RADIO ACCESS: AN OVERVIEW

The work on LTE was initiated in late 2004 with the overall aim of providing a new radio-access technology focusing on packet-switched data only. The first phase of the 3GPP work on LTE was to define a set of performance and capability targets for LTE [6]. This included targets on peak data rates, user/system throughput, spectral efficiency, and control/user-plane latency. In addition, requirements were also set on spectrum flexibility, as well as on interaction/compatibility with other 3GPP radio-access technologies (GSM, WCDMA/HSPA, and TD-SCDMA).

Once the targets were set, 3GPP studies on the feasibility of different technical solutions considered for LTE were followed by development of detailed specifications. The first release of the LTE specifications, release 8, was completed in 2008 and commercial network operation began in late 2009. Release 8 has been followed by additional LTE releases, introducing additional functionality and capabilities in different areas, as illustrated in Figure 3.1. Of particular interest are release 10, being the first release of LTE-Advanced, and release 13, finalized late 2015 and the first release of LTE-Advanced Pro. Currently, as of this writing, 3GPP is working on LTE release 14.

In parallel to the development of LTE, there has also been an evolution of the overall 3GPP network architecture, termed *system architecture evolution* (SAE), including both the radio-access network and the core network. Requirements were also set on the architecture evolution, leading to a new flat radio-access-network architecture with a single type of node, the *eNodeB*,[1] as well as a new core-network architecture. An excellent description of the LTE-associated core-network architecture, the *evolved packet core* (EPC), can be found in [5].

The remaining part of this chapter provides an overview of LTE. The most important technology components of LTE release 8—including transmission schemes, scheduling, multi-antenna support, and spectrum flexibility—are presented, as well as the additional features and enhancements introduced in later releases up to and including release 13. The chapter can either be read on its own to get a high-level overview of LTE, or as an introduction to the subsequent chapters.

[1] eNodeB is a 3GPP term that can roughly be seen as being equivalent to a base station, see further Chapter 4.

4G, LTE-Advanced Pro and The Road to 5G. http://dx.doi.org/10.1016/B978-0-12-804575-6.00003-0
Copyright © 2016 Erik Dahlman, Stefan Parkvall and Johan Sköld. Published by Elsevier Ltd. All rights reserved.

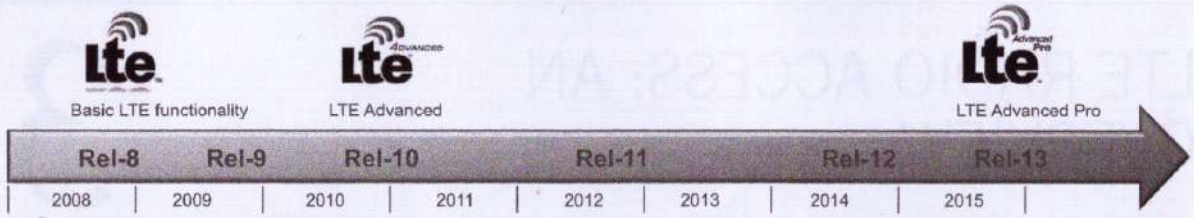


FIGURE 3.1

LTE and its evolution.

The following chapters, Chapters 4–22, provide a detailed description of the LTE radio-access technology. Chapter 4 provides an overview of the LTE protocol structure, including RLC, MAC, and the physical layer, explaining the logical and physical channels, and the related data flow. The time–frequency structure on which LTE is based is covered in Chapter 5, together with a brief overview of the LTE antenna-port concept. This is followed by a detailed description of the LTE physical-layer functionality for downlink and uplink transmission in Chapters 6 and 7, respectively. Chapter 8 contains a description of LTE retransmission mechanisms, followed by a discussion on the mechanisms available in LTE to support advanced scheduling and link adaptation in Chapter 9. Channel-state reporting to support scheduling, including handling of large antenna arrays, is covered in Chapter 10. Access procedures, necessary for a device (or a terminal, in 3GPP known as a *user equipment*, UE) to connect to the network, are the topic of Chapter 11.

The following chapters focus on some of the enhancements incorporated into LTE from release 10 onwards, starting with carrier aggregation (CA) in Chapter 12. Multi-point coordination/transmission is discussed in Chapter 13, followed by a discussion of heterogeneous deployments based on LTE in Chapter 14. Chapter 15 covers small-cell enhancements and dynamic TDD, Chapter 16 dual connectivity, and Chapter 17 operation in unlicensed spectrum. Relaying and broadcast/multicast are the topics of Chapters 18 and 19, respectively. Chapter 20 describes enhancements for improved support of machine-type communication while Chapter 21 focuses on direct device-to-device communication. The LTE part of the book is concluded with a discussion on the definition of radio-frequency (RF) requirements, taking into account the spectrum flexibility.

Finally, Chapters 23 and 24 provide an overview of the new 5G radio-access technology currently under discussion in 3GPP and Chapter 25 concludes the discussion on 5G radio access.

3.1 LTE BASIC TECHNOLOGIES

Release 8 is the first LTE release and forms the basis for the following releases. Due to time limitations for the release 8 work, some smaller features originally planned for release 8 was postponed for release 9, which thus can be seen as part of the basic LTE framework. In the following sections, the basic LTE technologies in release 8/9 are described.

3.1.1 TRANSMISSION SCHEME

The LTE downlink transmission scheme is based on conventional *orthogonal frequency-division multiplexing* (OFDM) [53], an attractive transmission scheme for several reasons. Due to the relatively long OFDM symbol time in combination with a cyclic prefix, OFDM provides a high degree of robustness against channel-frequency selectivity. Although signal corruption due to a frequency-selective channel can, in principle, be handled by equalization at the receiver side, the complexity of such equalization starts to become unattractively high for implementation in a device at larger bandwidths and especially in combination with advanced multi-antenna transmission schemes such as spatial multiplexing. Therefore, OFDM is an attractive choice for LTE for which a wide bandwidth and support for advanced multi-antenna transmission were key requirements.

OFDM also provides some additional benefits relevant for LTE:

- OFDM provides access to the frequency domain, thereby enabling an additional degree of freedom to the channel-dependent scheduler compared to time-domain-only scheduling used in major 3G systems.
- Flexible transmission bandwidth to support operation in spectrum allocations of different sizes is straightforward with OFDM, at least from a baseband perspective, by varying the number of OFDM subcarriers used for transmission. Note, however, that support of a flexible transmission bandwidth also requires flexible RF filtering, and so on, for which the exact transmission scheme is to a large extent irrelevant. Nevertheless, maintaining the same baseband-processing structure, regardless of the bandwidth, eases device development and implementation.
- Broadcast/multicast transmission, where the same information is transmitted from multiple base stations, is straightforward with OFDM as described in Chapter 19.

The LTE uplink is also based on OFDM transmission. However, to enable high-power amplifier efficiency on the device side, different means are taken to reduce the *cubic metric* [10] of uplink transmissions. Cubic metric is a measure of the amount of additional back-off needed for a certain signal waveform, relative to the back-off needed for some reference waveform. It captures similar properties as the more commonly known peak-to-average ratio but better represents the actual back-off needed in an implementation. Low cubic metric is achieved by preceding the OFDM modulator by a DFT precoder, leading to *DFT-spread OFDM* (DFTS-OFDM), see Chapter 7. Often the term DFTS-OFDM is used to describe the LTE uplink transmission scheme in general. However, it should be understood that DFTS-OFDM is only used for uplink data transmission. As described in more detail in later chapters, other means are used to achieve a low cubic metric for other types of uplink transmissions. Thus, the LTE uplink transmission scheme should be described as OFDM with different techniques, including DFT precoding for data transmission, being used to reduce the cubic metric of the transmitted signal.

The use of OFDM-based transmission for the LTE uplink allows for orthogonal separation of uplink transmissions also in the frequency domain. Orthogonal separation is in many cases

beneficial as it avoids interference between uplink transmissions from different devices within the cell (*intra-cell interference*). Allocating a very large instantaneous bandwidth for transmission from a single device is not an efficient strategy in situations where the data rate is mainly limited by the available device transmit power rather than the bandwidth. In such situations a device can instead be allocated only a part of the total available bandwidth and other devices within the cell can be scheduled to transmit in parallel on the remaining part of the spectrum. In other words, the LTE uplink transmission scheme allows for both *time-division multiple access* (TDMA) and *frequency-division multiple access* (FDMA) to separate users.

3.1.2 CHANNEL-DEPENDENT SCHEDULING AND RATE ADAPTATION

One key characteristic of mobile radio communication is the large and typically rapid variations in the instantaneous channel conditions stemming from frequency-selective fading, distance-dependent path loss, and random interference variations due to transmissions in other cells and by other terminals. Instead trying to combat these variations through, for example, power control, LTE tries to exploit these variations through *channel-dependent scheduling* where the time–frequency resources are dynamically shared between users. Dynamic sharing of resources across the users is well matched to the rapidly varying resource requirements posed by packet-data communication and also enables several of the other key technologies on which LTE is based.

The scheduler controls, for each time instant, to which users the different parts of the shared resource should be assigned and the data rate to be used for each transmission. Thus, *rate adaptation* (i.e., trying to dynamically adjust the data rate to match the instantaneous channel conditions) can be seen as a part of the scheduling functionality.

However, even if the rate adaptation successfully selects an appropriate data rate, there is a certain likelihood of transmission errors. To handle transmission errors, *fast hybrid-ARQ with soft combining* is used in LTE to allow the device to rapidly request retransmissions of erroneously received data blocks and to provide a tool for implicit rate adaptation. Retransmissions can be rapidly requested after each packet transmission, thereby minimizing the impact on end-user performance from erroneously received packets. Incremental redundancy is used as the soft combining strategy, and the receiver buffers the soft bits to be able to perform soft combining between transmission attempts.

The scheduler is a key element and to a large extent determines the overall system performance, especially in a highly loaded network. Both downlink and uplink transmissions are subject to tight scheduling in LTE. A substantial gain in system capacity can be achieved if the channel conditions are taken into account in the scheduling decision, so-called *channel-dependent scheduling*, where transmission are directed to user with momentarily favorable channel conditions. Due to the use of OFDM in both the downlink and uplink transmission directions, the scheduler has access to both the time and frequency domains. In other words, the scheduler can, for each time instant and frequency region, select the user with the best channel conditions, as illustrated in Figure 3.2.

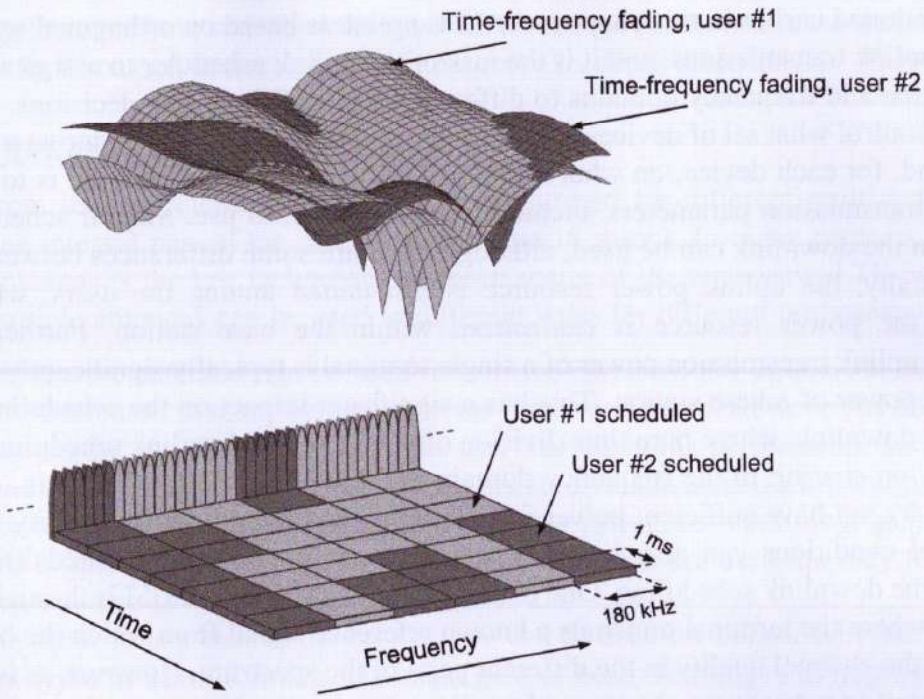


FIGURE 3.2

Downlink channel-dependent scheduling in time and frequency domains.

The possibility of channel-dependent scheduling in the frequency domain is particularly useful at low device speeds—in other words, when the channel is varying slowly in time. Channel-dependent scheduling relies on channel-quality variations between users to obtain a gain in system capacity. For delay-sensitive services, a time-domain-only scheduler may, due to the delay constraints, be forced to schedule a particular user, despite the channel quality not being at its peak. In such situations, exploiting channel-quality variations also in the frequency domain will help to improve the overall performance of the system. For LTE, scheduling decisions can be taken as often as once every 1 ms, and the granularity in the frequency domain is 180 kHz. This allows for relatively rapid channel variations in both the time and frequency domains to be tracked and utilized by the scheduler.

To support downlink scheduling, a device may provide the network with *channel-state* reports indicating the instantaneous downlink channel quality in both the time and frequency domains. The channel state is typically obtained by measuring on *reference signals* transmitted in the downlink. Based on the channel-state reports, also referred to as *channel-state information* (CSI), the downlink scheduler can assign resources for downlink transmission to different devices, taking the channel quality into account in the scheduling decision. In principle, a scheduled device can be assigned an arbitrary combination of 180 kHz wide *resource blocks* in each 1 ms scheduling interval.

As mentioned earlier in this chapter, the LTE uplink is based on orthogonal separation of different uplink transmissions, and it is the task of the uplink scheduler to assign resources in both the time and frequency domains to different devices. Scheduling decisions, taken once per 1 ms, control what set of devices are allowed to transmit within a cell during a given time interval and, for each device, on what frequency resources the transmission is to take place and what transmission parameters, including the data rate, to use. Similar scheduling strategies as in the downlink can be used, although there are some differences between the two. Fundamentally, the uplink power resource is *distributed* among the users, while in the downlink the power resource is *centralized* within the base station. Furthermore, the maximum uplink transmission power of a single terminal is typically significantly lower than the output power of a base station. This has a significant impact on the scheduling strategy. Unlike the downlink, where pure time division often can be used, uplink scheduling typically has to rely on sharing in the frequency domain in addition to the time domain, as a single terminal may not have sufficient power for efficiently utilizing the link capacity.

Channel conditions can also be taken into account in the uplink scheduling process, similar to the downlink scheduling. One possibility to acquire uplink CSI is through so-called *sounding*, where the terminal transmits a known reference signal from which the base station can assess the channel quality in the different parts of the spectrum. However, as is discussed in more detail in subsequent chapters, obtaining information about the uplink channel conditions may not be feasible or desirable in all situations. Therefore, different means to obtain *uplink diversity* are important as a complement in situations where uplink channel-dependent scheduling is not suitable.

3.1.3 INTER-CELL INTERFERENCE COORDINATION

LTE is designed to operate with a frequency reuse of one implying that the same carrier frequency can be used at neighboring transmission points. In particular, the basic control channels of LTE are designed to operate properly with the relatively low signal-to-interference ratio that may be experienced in a reuse-one deployment.

Fundamentally, having access to all available frequency resources at each transmission point is always beneficial. However, system efficiency and end-user quality are further improved if transmissions from neighboring transmission points can be coordinated in such a way that the most severe interference situations can be avoided.

Already the first release of LTE included explicit support for such coordination, in the release-8 context referred to as *inter-cell interference coordination* (ICIC), see Chapter 13. More specifically, the release-8 specifications defined a set of messages that can be exchanged between eNodeBs using the so-called *X2 interface*, see Chapter 4. These messages provide information about the interference situation experienced by the eNodeB issuing the message and can be used by a neighboring eNodeB receiving the message as input to its scheduling process, thereby providing a means for at least partly coordinating the transmissions and controlling the interference between cells of different eNodeBs. Especially severe interference situations may occur in so-called *heterogeneous network deployments* consisting of

overlapping layers of base stations with large differences in the downlink transmission power. This is briefly discussed in Section 3.5, with a more in-depth discussion in Chapter 14.

3.1.4 MULTI-ANTENNA TRANSMISSION

Already from its first release, LTE has included support for different multi-antenna techniques as an integral part of the radio-interface specifications. In many respects, the use of multiple antennas is the key technology to reach many of the aggressive LTE performance targets. Multiple antennas can be used in different ways for different purposes:

- Multiple receive antennas can be used for receive diversity. For uplink transmissions, they have been used in many cellular systems for several years. However, as dual receive antennas are the baseline for all LTE devices,[2] the downlink performance is also improved. The simplest way of using multiple receive antennas is classical receive diversity to collect additional energy and suppress fading, but additional gains can be achieved in interference-limited scenarios if the antennas are used not only to provide diversity, but also to suppress interference.
- Multiple transmit antennas at the base station can be used for transmit diversity and different types of beam-forming. The main goal of beam-forming is to improve the received signal-to-interference-and-noise ratio (SINR) and, eventually, improve system capacity and coverage.
- *Spatial multiplexing*, sometimes referred to as multiple input, multiple output (MIMO) or more specifically single-user MIMO (SU-MIMO) using multiple antennas at both the transmitter and receiver, is supported by LTE. Spatial multiplexing results in an increased data rate, channel conditions permitting, in bandwidth-limited scenarios by creating several parallel "channels." Alternatively, by combining the spatial properties with the appropriate interference-suppressing receiver processing, multiple devices can transmit on the same time–frequency resource in order to improve the overall cell capacity. In 3GPP this is referred to as *multiuser MIMO*.

In general, the different multi-antenna techniques are beneficial in different scenarios. As an example, at relatively low SINR, such as at high load or at the cell edge, spatial multiplexing provides relatively limited benefits. Instead, in such scenarios multiple antennas at the transmitter side should be used to raise the SINR by means of beam-forming. On the other hand, in scenarios where there already is a relatively high SINR, for example, in small cells, raising the signal quality further provides relatively minor gains as the achievable data rates are then mainly bandwidth limited rather than SINR limited. In such scenarios, spatial multiplexing should instead be used to fully exploit the good channel conditions. The multi-antenna scheme used is under control of the base station, which therefore can select a suitable scheme for each transmission.

[2]For low-end MTC devices single-antenna operation is the baseline.

Release 9 enhanced the support for combining spatial multiplexing with beam-forming. Although the combination of beam-forming and spatial multiplexing was already possible in release 8, it was restricted to so-called *codebook-based precoding* (see Chapter 6). In release 9, the support for spatial multiplexing combined with so-called *non-codebook-based precoding* was introduced, allowing for improved flexibility in deploying various multi-antenna transmission schemes.

Up to four layers downlink spatial multiplexing was already supported from the first release of LTE. Later releases further extended the LTE multi-antenna capabilities as described in Section 3.4.

3.1.5 SPECTRUM FLEXIBILITY

A high degree of spectrum flexibility is one of the main characteristics of the LTE radio-access technology. The aim of this spectrum flexibility is to allow for the deployment of LTE radio access in difference frequency bands with various characteristics, including different duplex arrangements and different sizes of the available spectrum. Chapter 22 outlines further details of how spectrum flexibility is achieved in LTE.

3.1.5.1 Flexibility in duplex arrangements

One important part of the LTE requirements in terms of spectrum flexibility is the possibility to deploy LTE-based radio access in both paired *and* unpaired spectra. Therefore, LTE supports both frequency- and time-division-based duplex arrangements. *Frequency-division duplex* (FDD), as illustrated on the left in Figure 3.3, implies that downlink and uplink transmission take place in different, sufficiently separated, frequency bands. *Time-division duplex* (TDD), as illustrated on the right in Figure 3.3, implies that downlink and uplink

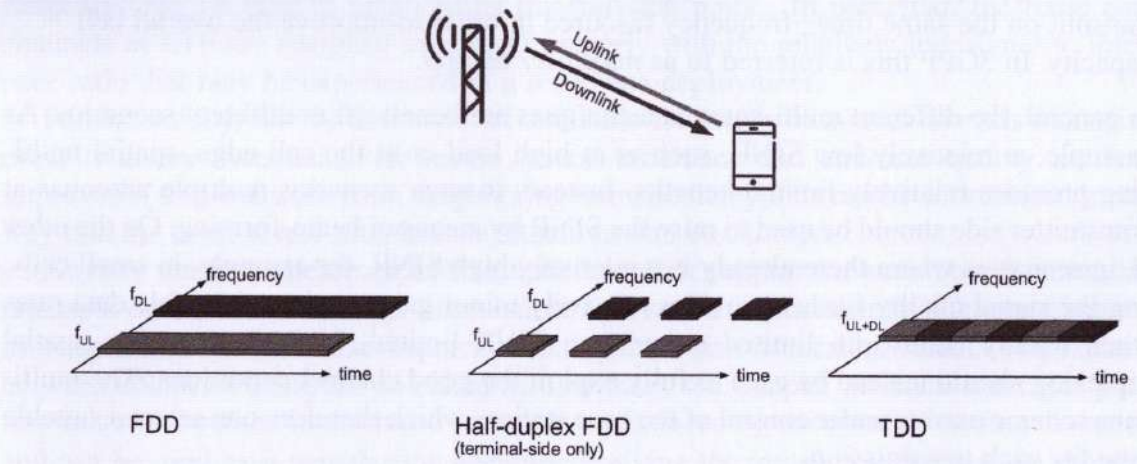


FIGURE 3.3

Frequency- and time-division duplex.

transmission take place in different, nonoverlapping time slots. Thus, TDD can operate in unpaired spectrum, whereas FDD requires paired spectrum. The required flexibility and resulting requirements to support LTE operation in different paired and unpaired frequency arrangements are further discussed in Chapter 22.

Operation in both paired and unpaired spectrum has been supported by 3GPP radio-access technologies even before the introduction of LTE by means of FDD-based WCDMA/HSPA in combination with TDD-based TD-SCDMA radio. However, this was then achieved by means of, at least in the details, relatively different radio-access technologies leading to additional effort and complexity when developing and implementing dual-mode devices capable of both FDD and TDD operation. LTE, on the other hand, supports both FDD and TDD *within a single radio-access technology*, leading to a minimum of deviation between FDD and TDD for LTE-based radio access. As a consequence of this, the overview of the LTE radio access provided in the following chapters is, to a large extent, valid for both FDD and TDD. In the case of differences between FDD and TDD, these differences are explicitly indicated. Furthermore, the TDD mode, also known as TD-LTE, is designed with coexistence between LTE (TDD) and TD-SCDMA in mind to simplify a gradual migration from TD-SCDMA to LTE.

LTE also supports *half-duplex* FDD at the device (illustrated in the middle of Figure 3.3). In half-duplex FDD, transmission and reception *at a specific device* are separated in both frequency and time. The base station still uses full-duplex FDD as it simultaneously may schedule *different* devices in uplink and downlink; this is similar to, for example, GSM operation. The main benefit with half-duplex FDD is the reduced device complexity as no duplex filter is needed in the device. This is especially beneficial in the case of multiband devices which otherwise would need multiple sets of duplex filters.

3.1.5.2 Bandwidth flexibility

An important characteristic of LTE is the support for a range of different transmission bandwidths on both downlink and uplink. The main reason for this is that the amount of spectrum available for LTE deployment may vary significantly between different frequency bands and also depending on the exact situation of the operator. Furthermore, the possibility of operating in different spectrum allocations gives the possibility for gradual migration of spectrum from other radio-access technologies to LTE.

LTE supports operation in a wide range of spectrum allocations, achieved by a flexible transmission bandwidth being part of the LTE specifications. To efficiently support very high data rates when spectrum is available, a wide transmission bandwidth is necessary. However, a sufficiently large amount of spectrum may not always be available, either due to the band of operation or due to a gradual migration from another radio-access technology, in which case LTE can be operated with a narrower transmission bandwidth. In such cases, the maximum achievable data rates will be reduced accordingly. As discussed in the following, the spectrum flexibility is further improved in later releases of LTE.

The LTE physical-layer specifications [24–27] are bandwidth agnostic and do not make any particular assumption on the supported transmission bandwidths beyond a minimum value. As is seen in the following, the basic radio-access specification, including the physical-layer and protocol specifications, allows for any transmission bandwidth ranging from roughly 1 MHz up to around 20 MHz. At the same time, at an initial stage, RF requirements are only specified for a limited subset of transmission bandwidths, corresponding to what is predicted to be relevant spectrum-allocation sizes and relevant migration scenarios. Thus, in practice the LTE radio-access technology supports a limited set of transmission bandwidths, but additional transmission bandwidths can easily be introduced by updating only the RF specifications.

3.1.6 MULTICAST AND BROADCAST SUPPORT

In situations where a large number of users want to receive the same information, for example, a TV news clip, information about the local weather conditions, or stock-market information, separate transmission of the information to each user may not be the most efficient approach. Instead, transmitting the same information once to all interested users may be a better choice. This is known as *broadcast*, or *multimedia broadcast–multicast services* (MBMS) in 3GPP, implying transmission of the same information to multiple receivers. In many cases the same information is of interest over a large area in which case identical signals can be transmitted from multiple cell sites with identical coding and modulation and with timing and frequency synchronized across the sites. From a device perspective, the signal will appear exactly as a signal transmitted from a single cell site and subject to multipath propagation, see Figure 3.4. Due to the OFDM robustness to multipath propagation, such multicell transmission, in 3GPP also referred to as *multicast/broadcast single-frequency network* (MBSFN) transmission, will then not only improve the received signal strength, but also eliminate the inter-cell interference. Thus, with OFDM, multicell broadcast/multicast throughput may eventually be limited by noise only and can then, in the case of small cells, reach extremely high values.

It should be noted that the use of MBSFN transmission for multicell broadcast/multicast assumes the use of tight synchronization and time alignment of the signals transmitted from different cell sites.

MBSFN support was part of the LTE work from the beginning, but due to time limitations, the support was not completed until the second release of LTE, release 9. There are also

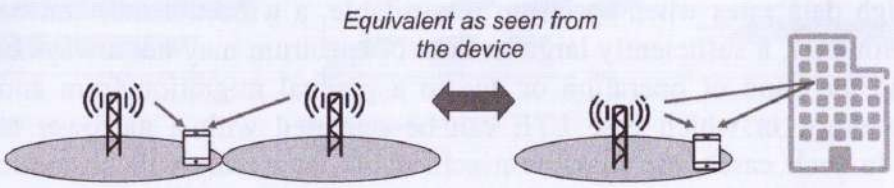


FIGURE 3.4

Equivalence between broadcast and multipath propagation from a device perspective.

enhancements for supporting multicast/broadcast services in a single cell in release 13, known as *single-cell point to multi-point* (SC-PTM).

3.1.7 POSITIONING

Positioning, as the name implies, refers to functionality in the radio-access network to determine the location of individual devices. Determining the position of a device can, in principle, be done by including a GPS receiver in the device. Although this is a quite common feature, not all devices include the necessary GPS receiver and there may also be cases when the GPS service is not available. LTE release 9 therefore introduced positioning support inherent in the radio-access network. By letting the device measure and report to the network the relative time of arrival of special reference signals transmitted regularly from different cell sites, the location of the device can be determined by the network.

3.2 LTE EVOLUTION

Releases 8 and 9 form the foundation of LTE, providing a highly capable mobile-broadband standard. However, to meet new requirements and expectations, the releases following the basic ones provides additional enhancements and features in different areas. Figure 3.5 illustrates some of the major areas in which LTE has been enhanced with details provided in the following. Table 3.1 captures the major enhancements per release.

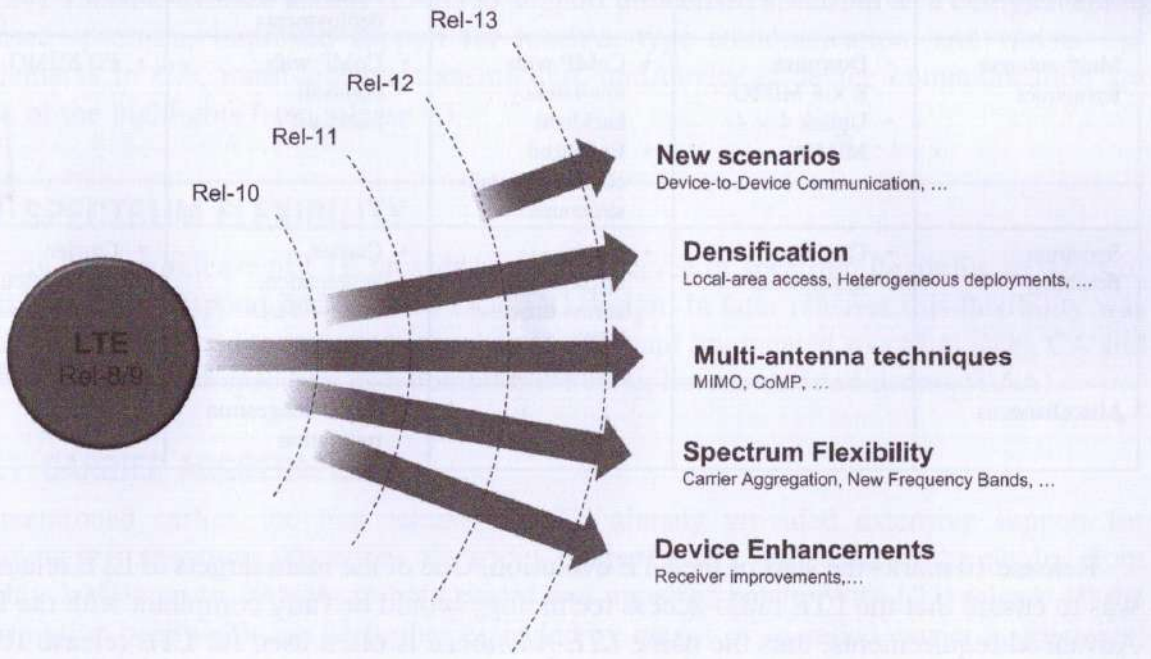


FIGURE 3.5

LTE evolution.

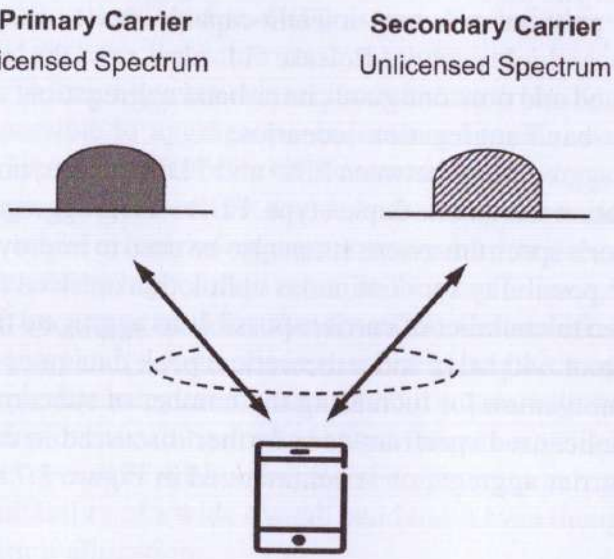


FIGURE 3.8

License-assisted access.

small-cell deployments is the target. Fair sharing of the spectrum resources with other systems, in particular Wi-Fi, is an important characteristic of LAA which therefore incudes a listen-before-talk mechanism.

Chapter 17 provides an in-depth discussion of license-assisted access.

3.4 MULTI-ANTENNA ENHANCEMENTS

Multi-antenna support has been enhanced over the different releases, increasing the number of transmission layers in the downlink to 8 and introducing uplink spatial multiplexing. Full-dimension MIMO (FD-MIMO) and two-dimensional beam-forming are other enhancements, as is the introduction of CoMP transmission.

3.4.1 EXTENDED MULTI-ANTENNA TRANSMISSION

In release 10, downlink spatial multiplexing was expanded to support up to eight transmission layers. This can be seen as an extension of the release-9 dual-layer beam-forming to support up to eight antenna ports and eight corresponding layers. Together with the support for carrier aggregation this enables downlink data rates up to 3 Gbit/s in 100 MHz of spectrum in release 10, increased to 25 Gbit/s in release 13 using 32 carriers, 8 layers spatial multiplexing and 256QAM.

Uplink spatial multiplexing of up to four layers was also introduced as part of LTE release 10. Together with the possibility for uplink carrier aggregation this allows for uplink data rates up to 1.5 Gbit/s in 100 MHz of spectrum. Uplink spatial multiplexing consists of a

codebook-based scheme under control of the base station, which means that the structure can also be used for uplink transmitter-side beam-forming.

An important consequence of the multi-antenna extensions in LTE release 10 was the introduction of an enhanced downlink *reference-signal structure* that more extensively separated the function of channel estimation and the function of acquiring CSI. The aim of this was to better enable novel antenna arrangements and new features such as more elaborate multi-point coordination/transmission in a flexible way.

In release-13, improved support for a large number of antennas was introduced, in particular in terms of more extensive feedback of CSI. The larger degrees of freedom can be used for, for example, beam-forming in both elevation and azimuth and massive multiuser MIMO where several spatially separated devices are simultaneously served using the same time–frequency resource. These enhancements are sometimes termed full-dimension MIMO and form a step into massive MIMO with a very large number of steerable antenna elements.

The multi-antenna support is described as part of the general downlink processing in Chapter 6, uplink processing in Chapter 7, and the supporting CSI reporting mechanism in Chapter 10.

3.4.2 MULTI-POINT COORDINATION AND TRANSMISSION

As discussed earlier, the first release of LTE included specific support for coordination between transmission points, referred to as ICIC, as a means to at least partly control the interference between cells. However, the support for such coordination was significantly expanded as part of LTE release 11 including the possibility for much more dynamic coordination between transmission points.

In contrast to release-8 ICIC, which was limited to the definition of certain X2 messages to assist coordination between cells, the release 11 activities focused on radio-interface features and device functionality to assist different coordination means, including the support for channel-state feedback for multiple transmission points. Jointly these features and functionality go under the name CoMP transmission/reception. Refinement to the reference-signal structure was also an important part of the CoMP support, as was the enhanced control-channel structure introduced as part of release 11, see later.

Support for CoMP includes *multi-point coordination*, that is, when transmission to a device is carried out from one specific transmission point but where scheduling and link adaptation is coordinated between the transmission points, as well as *multi-point transmission* in which case transmission to a device can be carried out from multiple transmission points either in such a way that that transmission can switch dynamically between different transmission points (*Dynamic Point Selection*) or be carried out jointly from multiple transmission points (*Joint Transmission*), see Figure 3.9.

A similar distinction can be made for uplink where one can distinguish between (uplink) multi-point coordination and multi-point *reception*. In general, though uplink CoMP is

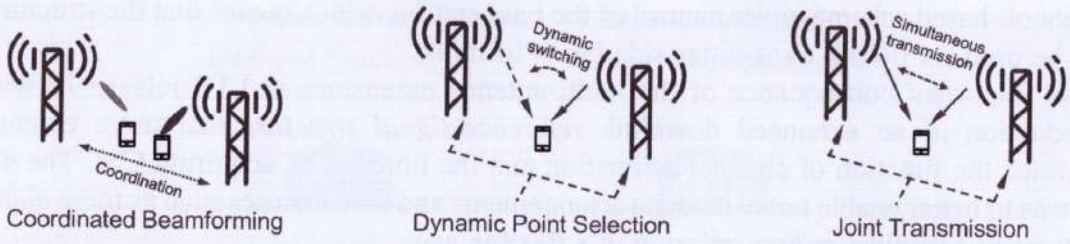


FIGURE 3.9

Different types of CoMP.

mainly a network implementation issue and has very little impact on the device and very little visibility in the radio-interface specifications.

The CoMP work in release 11 assumed "ideal" backhaul, in practice implying centralized baseband processing connected to the antenna sites using low-latency fiber connections. Extensions to relaxed backhaul scenarios with noncentralized baseband processing were introduced in release 12. These enhancements mainly consisted of defining new X2 messages for exchanging information about so-called CoMP hypotheses, essentially a potential resource allocation, and the associated gain/cost.

CoMP is described in more detail in Chapter 13.

3.4.3 ENHANCED CONTROL-CHANNEL STRUCTURE

In release 11, a new complementary control-channel structure was introduced in order to support ICIC and to exploit the additional flexibility of the new reference-signal structure not only for data transmission, which was the case in release 10, but also for control signaling. The new control-channel structure can thus be seen as a prerequisite for many CoMP schemes, although it is also beneficial for beam-forming and frequency-domain interference coordination. It is also used to support narrow-band operation for MTC enhancements in releases 12 and 13. A description of the enhanced control-channel structure is found in Chapter 6.

3.5 DENSIFICATION, SMALL CELLS, AND HETEROGENEOUS DEPLOYMENTS

Small cells and dense deployment has been in focus for several releases as means to provide very high capacity and data rates. Relaying, small-cell on/off, dynamic TDD, and heterogeneous deployments are some examples of enhancements over the releases. LAA, discussed in the area of spectrum flexibility, can also be seen as primarily an enhancement for small cells.

MME by means of the *S1 control-plane part*, S1-c. One eNodeB can be connected to multiple MMEs/S-GWs for the purpose of load sharing and redundancy.

The *X2 interface*, connecting eNodeBs to each other, is mainly used to support active-mode mobility. This interface may also be used for multi-cell *radio-resource management* (RRM) functions such as *inter-cell interference coordination* (ICIC) discussed in Chapter 13. The X2 interface is also used to support lossless mobility between neighboring cells by means of packet forwarding.

The interface between the eNodeB to the device is known as the *Uu interface*. Unless dual connectivity is used, see Chapter 16, a device is connected to a single eNodeB as a time. There is also a *PC5 interface* defined for direct device-to-device communication, see Chapter 21.

4.2 RADIO PROTOCOL ARCHITECTURE

With the overall network architecture in mind, the RAN protocol architecture for the user and control planes can be discussed. Figure 4.3 illustrates the RAN protocol architecture (the MME is, as discussed in the previous section, not part of the RAN but is included in the figure for completeness). As seen in the figure, many of the protocol entities are common to the user and control planes. Therefore, although this section mainly describes the protocol architecture from a user-plane perspective, the description is in many respects also applicable to the control plane. Control-plane-specific aspects are discussed in Section 4.3.

The LTE RAN provides one or more *radio bearers* to which IP packets are mapped according to their QoS requirements. A general overview of the LTE (user-plane) protocol architecture for the downlink is illustrated in Figure 4.4. As will become clear in the subsequent discussion, not all the entities illustrated in Figure 4.4 are applicable in all situations. For example, neither MAC scheduling nor hybrid ARQ with soft combining is used for broadcast of the basic system information. The LTE protocol structure related to uplink

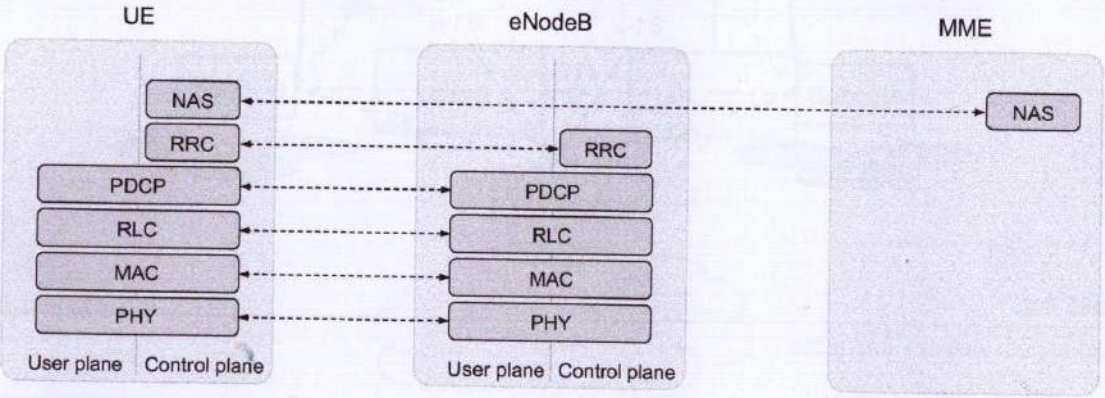


FIGURE 4.3

Overall RAN protocol architecture.

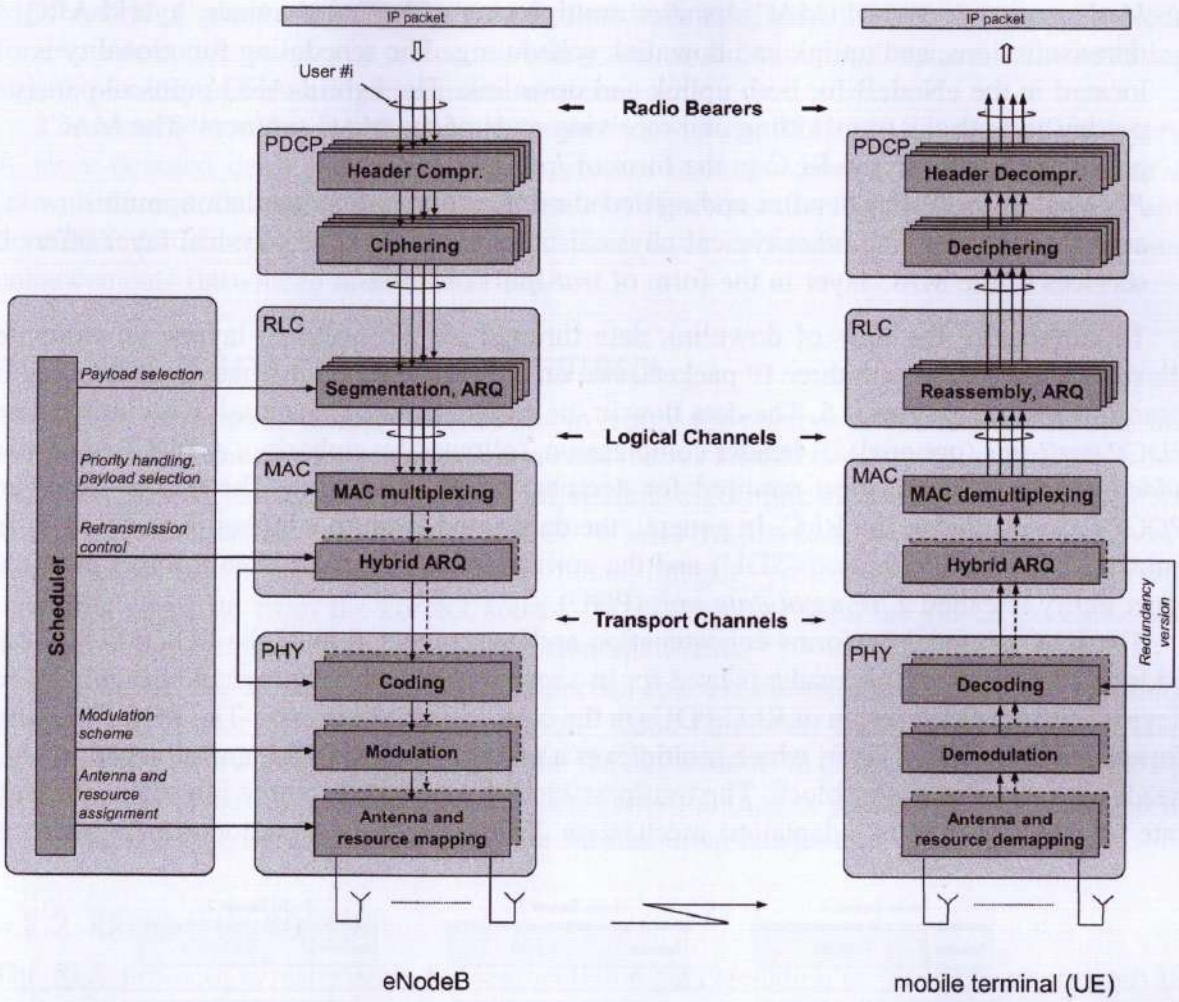


FIGURE 4.4

LTE protocol architecture (downlink).

transmissions is similar to the downlink structure in Figure 4.4, although there are some differences with respect to, for example, transport-format selection.

The different protocol entities of the RAN are summarized and described in more detail in the following sections:

- *Packet data convergence protocol* (PDCP) performs IP header compression, ciphering, and integrity protection. It also handles in-sequence delivery and duplicate removal in case of handover. There is one PDCP entity per radio bearer configured for a device.
- *Radio-link control* (RLC) is responsible for segmentation/concatenation, retransmission handling, duplicate detection, and in-sequence delivery to higher layers. The RLC provides services to the PDCP. There is one RLC entity per radio bearer configured for a device.

- *Medium-access control* (MAC) handles multiplexing of logical channels, hybrid-ARQ retransmissions, and uplink and downlink scheduling. The scheduling functionality is located in the eNodeB for both uplink and downlink. The hybrid-ARQ protocol part is present in both the transmitting and receiving ends of the MAC protocol. The MAC provides services to the RLC in the form of *logical channels*.
- *Physical layer* (PHY) handles coding/decoding, modulation/demodulation, multi-antenna mapping, and other typical physical-layer functions. The physical layer offers services to the MAC layer in the form of *transport channels*.

To summarize the flow of downlink data through all the protocol layers, an example illustration for a case with three IP packets, two on one radio bearer and one on another radio bearer, is given in Figure 4.5. The data flow in the case of uplink transmission is similar. The PDCP performs (optional) IP-header compression, followed by ciphering. A PDCP header is added, carrying information required for deciphering in the device. The output from the PDCP is forwarded to the RLC. In general, the data entity from/to a higher protocol layer is known as a *service data unit* (SDU) and the corresponding entity to/from a lower protocol layer entity is called a *protocol data unit* (PDU).

The RLC protocol performs concatenation and/or segmentation of the PDCP SDUs and adds an RLC header. The header is used for in-sequence delivery (per logical channel) in the device and for identification of RLC PDUs in the case of retransmissions. The RLC PDUs are forwarded to the MAC layer, which multiplexes a number of RLC PDUs and attaches a MAC header to form a transport block. The transport-block size depends on the instantaneous data rate selected by the link-adaptation mechanism. Thus, the link adaptation affects both the

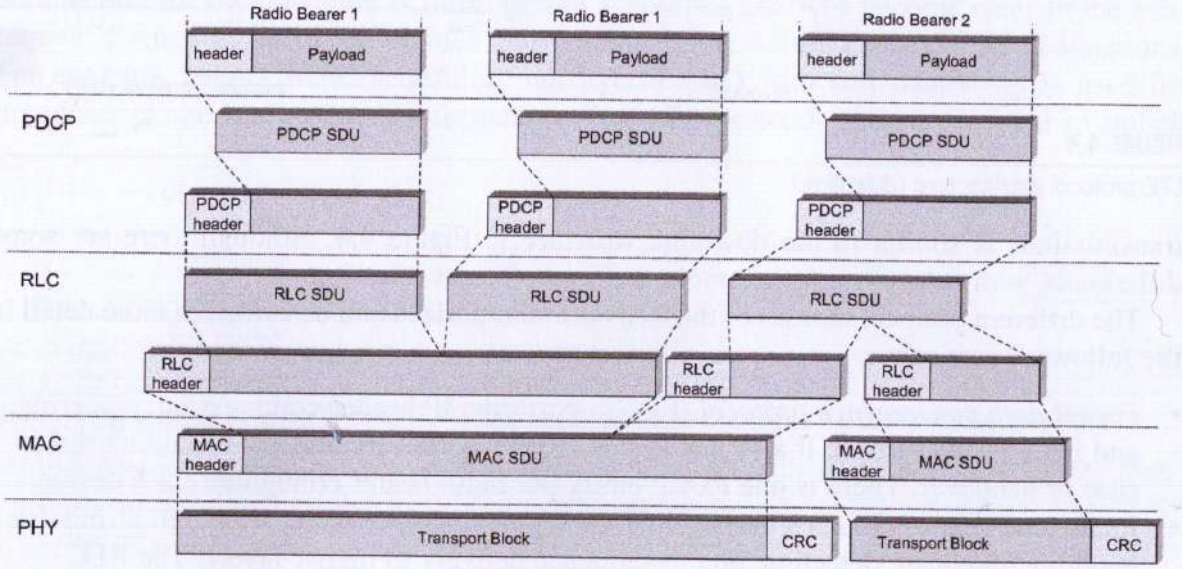


FIGURE 4.5

Example of LTE data flow.

MAC and RLC processing. Finally, the physical layer attaches a CRC to the transport block for error-detection purposes, performs coding and modulation, and transmits the resulting signal, possibly using multiple transmit antennas.

The remainder of the chapter contains an overview of the RLC, MAC, and physical layers. A more detailed description of the LTE physical-layer processing is given in Chapter 6 (downlink) and Chapter 7 (uplink), followed by descriptions of some specific uplink and downlink radio-interface functions and procedures in the subsequent chapters as well as some enhancements introduced after release 8/9.

4.2.1 PACKET-DATA CONVERGENCE PROTOCOL

The PDCP performs IP header compression to reduce the number of bits to transmit over the radio interface. The header-compression mechanism is based on robust header compression (ROHC) [31], a standardized header-compression algorithm also used for several other mobile-communication technologies. PDCP is also responsible for ciphering to protect against eavesdropping and, for the control plane, integrity protection to ensure that control messages originate from the correct source. At the receiver side, the PDCP performs the corresponding deciphering and decompression operations.

In addition, the PDCP also plays an important role for intra-eNodeB handover, handling in-sequence delivery and duplicate removal.[3] Upon handover, undelivered downlink data packets will be forwarded by the PDCP from the old eNodeB to the new eNodeB. The PDCP entity in the device will also handle retransmission of all uplink packets not yet delivered to the eNodeB as the hybrid-ARQ buffers are flushed upon handover.

4.2.2 RADIO-LINK CONTROL

The RLC protocol is responsible for segmentation/concatenation of (header-compressed) IP packets, also known as RLC SDUs, from the PDCP into suitably sized RLC PDUs. It also handles retransmission of erroneously received PDUs, as well as removal of duplicated PDUs. Finally, the RLC ensures in-sequence delivery of SDUs to upper layers. Depending on the type of service, the RLC can be configured in different modes to perform some or all of these functions.

Segmentation and concatenation, one of the main RLC functions, is illustrated in Figure 4.6. Depending on the scheduler decision, a certain amount of data is selected for transmission from the RLC SDU buffer, and the SDUs are segmented/concatenated to create the RLC PDU. Thus, for LTE the RLC PDU size varies *dynamically*. For high data rates, a large PDU size results in a smaller relative overhead, while for low data rates, a small PDU size is required as the payload would otherwise be too large. Hence, as the LTE data rates may

[3]Reordering is done in a similar way as the RLC reordering handling out-of-sequence PDUs from the hybrid-ARQ entity, see Chapter 8.

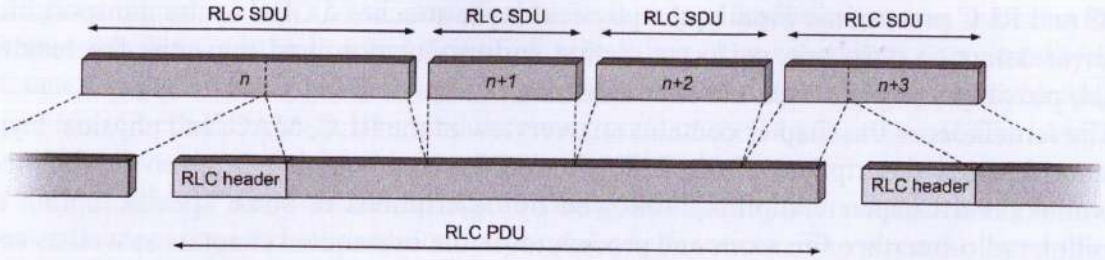


FIGURE 4.6

RLC segmentation and concatenation.

range from a few kbit/s up to several Gbit/s, dynamic PDU sizes are motivated for LTE in contrast to earlier mobile-communication technologies, which typically used a fixed PDU size. Since the RLC, scheduler, and rate-adaptation mechanisms are all located in the eNodeB, dynamic PDU sizes are easily supported for LTE. In each RLC PDU, a header is included, containing, among other things, a sequence number used for in-sequence delivery and retransmission handling.

The RLC retransmission mechanism is also responsible for providing error-free delivery of data to higher layers. To accomplish this, a retransmission protocol operates between the RLC entities in the receiver and transmitter. By monitoring the sequence numbers of the incoming PDUs, the receiving RLC can identify missing PDUs. Status reports are then fed back to the transmitting RLC entity, requesting retransmission of missing PDUs. Based on the received status report, the RLC entity at the transmitter can take appropriate action and retransmit the missing PDUs if needed.

Although the RLC is capable of handling transmission errors due to noise, unpredictable channel variations, and so on, error-free delivery is in most cases handled by the MAC-based hybrid-ARQ protocol. The use of a retransmission mechanism in the RLC may therefore seem superfluous at first. However, as is discussed in Section 4.2.3.3, this is not the case, and the use of both RLC- and MAC-based retransmission mechanisms is in fact well motivated by the differences in the feedback signaling.

The details of RLC are further described in Chapter 8.

4.2.3 MEDIUM-ACCESS CONTROL

The MAC layer handles logical-channel multiplexing, hybrid-ARQ retransmissions, and uplink and downlink scheduling. It is also responsible for multiplexing/demultiplexing data across multiple component carriers when carrier aggregation is used, see Chapter 12, and clear-channel assessment for license-assisted access, see Chapter 17.

4.2.3.1 Logical Channels and Transport Channels

The MAC provides services to the RLC in the form of *logical channels*. A logical channel is defined by the *type* of information it carries and is generally classified as a *control channel*,

used for transmission of control and configuration information necessary for operating an LTE system, or as a *traffic channel*, used for the user data. The set of logical-channel types specified for LTE includes:

- The *broadcast control channel* (BCCH), used for transmission of *system information* from the network to all devices in a cell. Prior to accessing the system, a device needs to acquire the system information to find out how the system is configured and, in general, how to behave properly within a cell.
- The *paging control channel* (PCCH), used for paging of devices whose location on a cell level is not known to the network. The paging message therefore needs to be transmitted in multiple cells.
- The *common control channel* (CCCH), used for transmission of control information in conjunction with random access.
- The *dedicated control channel* (DCCH), used for transmission of control information to/from a device. This channel is used for individual configuration of devices such as different handover messages.
- The *dedicated traffic channel* (DTCH), used for transmission of user data to/from a device. This is the logical channel type used for transmission of all uplink and non-MBSFN downlink user data.
- The *multicast control channel* (MCCH), used for transmission of control information required for reception of the MTCH (see later).
- The *single-cell multicast control channel* (SC-MCCH), used for transmission of control information for single-cell MTCH reception (see later).
- The *multicast traffic channel* (MTCH), used for downlink transmission of MBMS across multiple cells.
- The *single-cell multicast traffic channel* (SC-MTCH), used for downlink transmission of MBMS in a single cell.
- The *sidelink broadcast control channel* (SBCCH), used for sidelink synchronization.
- The *sidelink traffic channel* (STCH), used for sidelink communication (sidelink is the link for direct device-to-device communication, compare uplink and downlink for communication between a device and an eNodeB), see Chapter 21 for a description of device-to-device communication.

From the physical layer, the MAC layer uses services in the form of *transport channels*. A transport channel is defined by *how* and *with what characteristics* the information is transmitted over the radio interface. Data on a transport channel is organized into *transport blocks*. In each *transmission time interval* (TTI), at most one transport block of dynamic size is transmitted over the radio interface to/from a device in the absence of spatial multiplexing. In the case of spatial multiplexing (MIMO), there can be up to two transport blocks per TTI.

Associated with each transport block is a *transport format* (TF), specifying *how* the transport block is to be transmitted over the radio interface. The transport format includes

information about the transport-block size, the modulation-and-coding scheme, and the antenna mapping. By varying the transport format, the MAC layer can thus realize different data rates. Rate control is also known as *transport-format selection*.

The following transport-channel types are defined for LTE:

- The *broadcast channel* (BCH) has a fixed transport format, provided by the specifications. It is used for transmission of parts of the BCCH system information, more specifically the so-called *master information block* (MIB), as described in Chapter 11.
- The *paging channel* (PCH) is used for transmission of paging information from the PCCH logical channel. The PCH supports *discontinuous reception* (DRX) to allow the device to save battery power by waking up to receive the PCH only at predefined time instants. The LTE paging mechanism is also described in Chapter 11.
- The *downlink shared channel* (DL-SCH) is the main transport channel used for transmission of downlink data in LTE. It supports key LTE features such as dynamic rate adaptation and channel-dependent scheduling in the time and frequency domains, hybrid ARQ with soft combining, and spatial multiplexing. It also supports DRX to reduce device power consumption while still providing an always-on experience. The DL-SCH is also used for transmission of the parts of the BCCH system information not mapped to the BCH. There can be multiple DL-SCHs in a cell, one per device[4] scheduled in this TTI, and, in some subframes, one DL-SCH carrying system information.
- The *multicast channel* (MCH) is used to support MBMS. It is characterized by a semi-static transport format and semi-static scheduling. In the case of multi-cell transmission using MBSFN, the scheduling and transport format configuration are coordinated among the transmission points involved in the MBSFN transmission. MBSFN transmission is described in Chapter 19.
- The *uplink shared channel* (UL-SCH) is the uplink counterpart to the DL-SCH—that is, the uplink transport channel used for transmission of uplink data.
- The *sidelink shared channel* (SL-SCH) is the transport channel used for sidelink communication as described in Chapter 21.
- The *sidelink broadcast channel* (SL-BCH) is used for sidelink synchronization, see Chapter 21.
- The *sidelink discovery channel* (SL-DCH) is used in the sidelink discovery process as described in Chapter 21.

In addition, the *random-access channel* (RACH) is also defined as a transport channel, although it does not carry transport blocks. Furthermore, the introduction of NB-IoT in release 13, see Chapter 20, resulted in a set of channels optimized for narrowband operation.

Part of the MAC functionality is multiplexing of different logical channels and mapping of the logical channels to the appropriate transport channels. The mapping between logical-

[4]For carrier aggregation, a device may receive multiple DL-SCHs, one per component carrier.

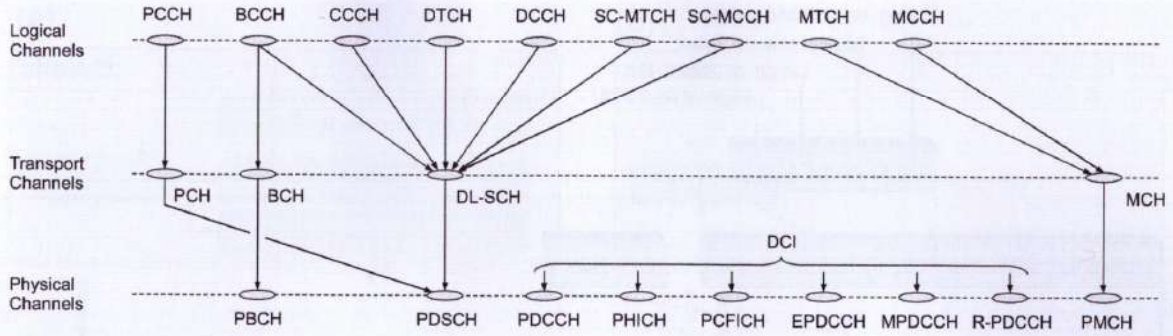


FIGURE 4.7

Downlink channel mapping.

channel types and transport channel types is given in Figure 4.7 for the downlink and Figure 4.8 for the uplink. The figures clearly indicate how DL-SCH and UL-SCH are the main downlink and uplink transport channels, respectively. In the figures, the corresponding physical channels, described later, are also included and the mapping between transport channels and physical channels is illustrated. For details on the mapping of sidelink channels see Chapter 21.

To support priority handling, multiple logical channels, where each logical channel has its own RLC entity, can be multiplexed into one transport channel by the MAC layer. At the receiver, the MAC layer handles the corresponding demultiplexing and forwards the RLC PDUs to their respective RLC entity for in-sequence delivery and the other functions handled by the RLC. To support the demultiplexing at the receiver, a MAC header, as shown in

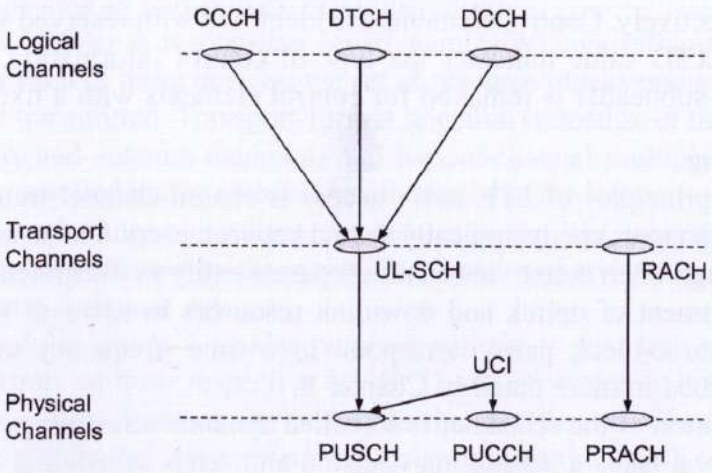


FIGURE 4.8

Uplink channel mapping.

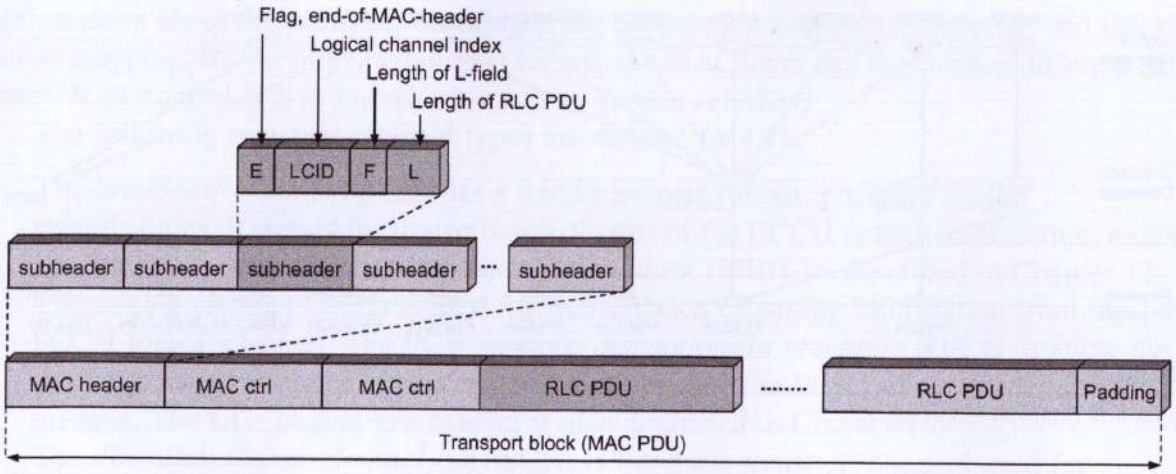


FIGURE 4.9

MAC header and SDU multiplexing.

Figure 4.9, is used. To each RLC PDU, there is an associated subheader in the MAC header. The subheader contains the identity of the logical channel (LCID) from which the RLC PDU originated and the length of the PDU in bytes. There is also a flag indicating whether this is the last subheader or not. One or several RLC PDUs, together with the MAC header and, if necessary, padding to meet the scheduled transport-block size, form one transport block which is forwarded to the physical layer.

In addition to multiplexing of different logical channels, the MAC layer can also insert the so-called *MAC control elements* into the transport blocks to be transmitted over the transport channels. A MAC control element is used for inband control signaling—for example, timing-advance commands and random-access response, as described in Sections 7.6 and 11.3, respectively. Control elements are identified with reserved values in the LCID field, where the LCID value indicates the type of control information. Furthermore, the length field in the subheader is removed for control elements with a fixed length.

4.2.3.2 Scheduling

One of the basic principles of LTE radio access is shared-channel transmission—that is, time–frequency resources are dynamically shared between users. The *scheduler* is part of the MAC layer (although often better viewed as a separate entity as illustrated in Figure 4.4) and controls the assignment of uplink and downlink resources in terms of so-called *resource-block pairs*. Resource-block pairs correspond to a time–frequency unit of 1 ms times 180 kHz, as described in more detail in Chapter 9.

The basic operation of the scheduler is so-called *dynamic* scheduling, where the eNodeB in each 1 ms interval takes a scheduling decision and sends scheduling information to the selected set of devices. However, there is also a possibility for semi-persistent scheduling

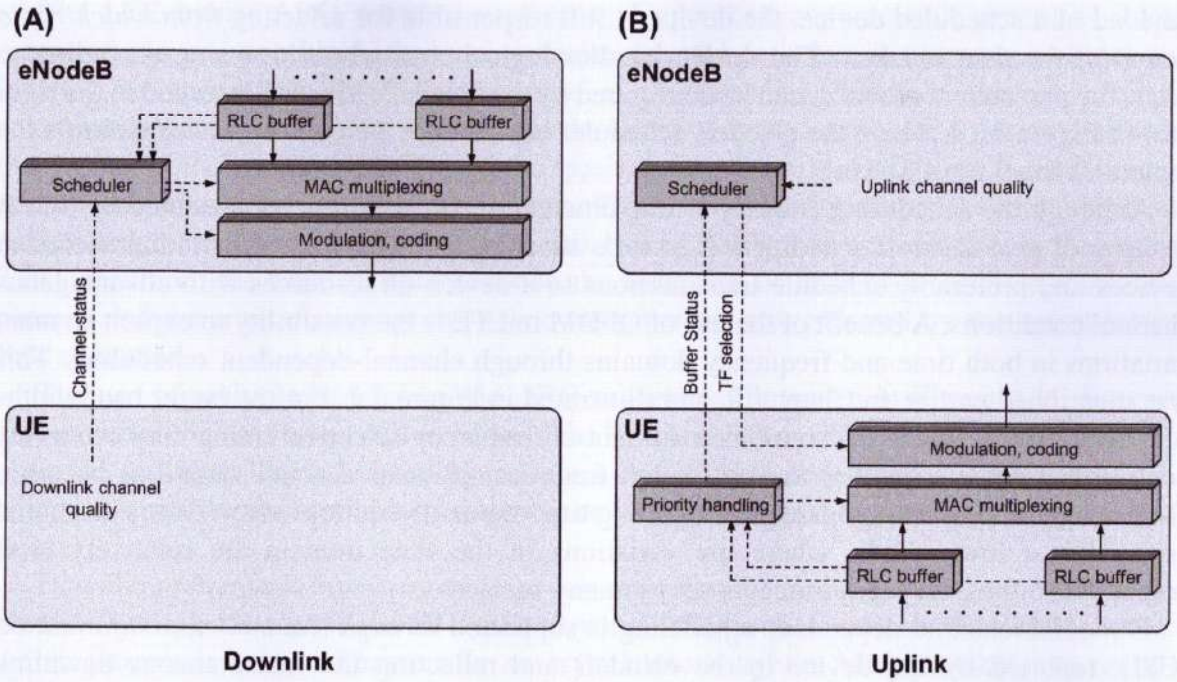


FIGURE 4.10

TF selection in (A) downlink and (B) uplink.

where a semi-static scheduling pattern is signaled in advance to reduce the control-signaling overhead.

Uplink and downlink scheduling are separated in LTE, and uplink and downlink scheduling decisions can be taken independently of each other (within the limits set by the uplink/downlink split in the case of half-duplex FDD operation).

The downlink scheduler is responsible for (dynamically) controlling which device(s) to transmit to and, for each of these devices, the set of resource blocks upon which the device's DL-SCH should be transmitted. Transport-format selection (selection of transport-block size, modulation scheme, and antenna mapping) and logical-channel multiplexing for downlink transmissions are controlled by the eNodeB, as illustrated in the left part of Figure 4.10. As a consequence of the scheduler controlling the data rate, the RLC segmentation and MAC multiplexing will also be affected by the scheduling decision. The outputs from the downlink scheduler can be seen in Figure 4.4.

The uplink scheduler serves a similar purpose, namely to (dynamically) control which devices are to transmit on their respective UL-SCH and on which uplink time–frequency resources (including component carrier). Despite the fact that the eNodeB scheduler determines the TF for the device, it is important to point out that the uplink scheduling decision is taken *per device* and not per radio bearer. Thus, although the eNodeB scheduler controls the

payload of a scheduled device, the device is still responsible for selecting *from which radio bearer(s)* the data is taken. The device handles logical-channel multiplexing according to rules, the parameters of which can be configured by the eNodeB. This is illustrated in the right part of Figure 4.10, where the eNodeB scheduler controls the TF and the device controls the logical-channel multiplexing.

Although the scheduling strategy is implementation specific and not specified by 3GPP, the overall goal of most schedulers is to take advantage of the channel variations between devices and preferably schedule transmissions to a device on resources with advantageous channel conditions. A benefit of the use of OFDM in LTE is the possibility to exploit channel variations in both time and frequency domains through channel-dependent scheduling. This was mentioned earlier, in Chapter 9, and illustrated in Figure 3.2. For the larger bandwidths supported by LTE, where a significant amount of frequency-selective fading may occur, the possibility for the scheduler to also exploit frequency-domain channel variations becomes increasingly important compared to exploiting time-domain variations only. This is beneficial especially at low speeds, where the variations in the time domain are relatively slow compared to the delay requirements set by many services.

Downlink channel-dependent scheduling is supported through *channel-state information* (CSI), reported by the device to the eNodeB and reflecting the instantaneous downlink channel quality in the time and frequency domains, as well as information necessary to determine the appropriate antenna processing in the case of spatial multiplexing. In the uplink, the channel-state information necessary for uplink channel-dependent scheduling can be based on a *sounding reference signal* transmitted from each device for which the eNodeB wants to estimate the uplink channel quality. To aid the uplink scheduler in its decisions, the device can transmit buffer-status information to the eNodeB using a MAC message. This information can only be transmitted if the device has been given a valid scheduling grant. For situations when this is not the case, an indicator that the device needs uplink resources is provided as part of the uplink L1/L2 control-signaling structure (see Chapter 7).

4.2.3.3 Hybrid ARQ with Soft Combining

Hybrid ARQ with soft combining provides robustness against transmission errors. As hybrid-ARQ retransmissions are fast, many services allow for one or multiple retransmissions, thereby forming an implicit (closed-loop) rate-control mechanism. The hybrid-ARQ protocol is part of the MAC layer, while the actual soft combining is handled by the physical layer.[5]

Hybrid ARQ is not applicable for all types of traffic. For example, broadcast transmissions, where the same information is intended for multiple devices, typically do not rely on hybrid ARQ.[6] Hence, hybrid ARQ is only supported for the DL-SCH and the UL-SCH, although its usage is optional.

[5]The soft combining is done before or as part of the channel decoding, which clearly is a physical-layer functionality.
[6]Autonomous retransmissions, where the transmitter repeats the same information several times without the receivers transmitting an acknowledgment, is possible and sometimes used to obtain a soft combining gain.

The LTE hybrid-ARQ protocol uses multiple parallel stop-and-wait processes. Upon reception of a transport block, the receiver makes an attempt to decode the transport block and informs the transmitter about the outcome of the decoding operation through a single acknowledgment bit indicating whether the decoding was successful or if a retransmission of the transport block is required. Clearly, the receiver must know to which hybrid-ARQ process a received acknowledgment is associated. This is solved by using the timing of the acknowledgment for association with a certain hybrid-ARQ process. Note that, in the case of TDD operation, the time relation between the reception of data in a certain hybrid-ARQ process and the transmission of the acknowledgment is also affected by the uplink/downlink allocation.

The use of multiple parallel hybrid-ARQ processes, illustrated in Figure 4.11, for each user can result in data being delivered from the hybrid-ARQ mechanism out of sequence. For example, transport block 5 in the figure was successfully decoded before transport block 1 which required retransmissions. In-sequence delivery of data is therefore ensured by the RLC layer.

Downlink retransmissions may occur at any time after the initial transmission—that is, the protocol is asynchronous—and an explicit hybrid-ARQ process number is used to indicate which process is being addressed. In an asynchronous hybrid-ARQ protocol, the retransmissions are in principle scheduled similarly to the initial transmissions. Uplink retransmissions, on the other hand, are based on a synchronous protocol, the retransmission occurs at a predefined time after the initial transmission and the process number can be implicitly derived. In a synchronous protocol the time instant for the retransmissions is fixed once the initial transmission has been scheduled, which must be accounted for in the scheduling operation. However, note that the scheduler knows from the hybrid-ARQ entity in the eNodeB whether a retransmission needs to be scheduled or not.

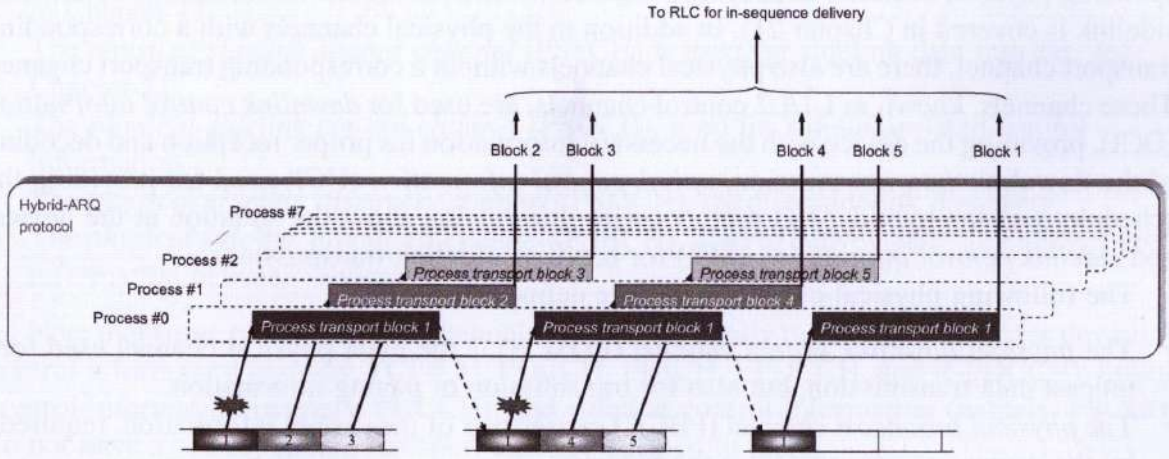


FIGURE 4.11

Multiple parallel hybrid-ARQ processes.

The hybrid-ARQ mechanism will rapidly correct transmission errors due to noise or unpredictable channel variations. As discussed earlier in this chapter, the RLC is also capable of requesting retransmissions, which at first sight may seem unnecessary. However, the reason for having two retransmission mechanisms on top of each other can be seen in the feedback signaling—hybrid ARQ provides fast retransmissions but due to errors in the feedback the residual error rate is typically too high, for example, for good TCP performance, while RLC ensures (almost) error-free data delivery but slower retransmissions than the hybrid-ARQ protocol. Hence, the combination of hybrid ARQ and RLC provides an attractive combination of small round-trip time and reliable data delivery. Furthermore, as the RLC and hybrid ARQ are located in the same node, tight interaction between the two is possible, as discussed in Chapter 8.

4.2.4 PHYSICAL LAYER

The physical layer is responsible for coding, physical-layer hybrid-ARQ processing, modulation, multi-antenna processing, and mapping of the signal to the appropriate physical time–frequency resources. It also handles mapping of transport channels to physical channels, as shown in Figures 4.7 and 4.8.

As mentioned in the introduction, the physical layer provides services to the MAC layer in the form of transport channels. Data transmission in downlink, uplink, and sidelink use the DL-SCH, UL-SCH, and SL-SCH transport-channel types, respectively. There is at most one or, in the case of spatial multiplexing,[7] two transport blocks per TTI on a DL-SCH, UL-SCH, or SL-SCH. In the case of carrier aggregation, there is one DL-SCH (or UL-SCH) per component carrier seen by the device.

A *physical channel* corresponds to the set of time–frequency resources used for transmission of a particular transport channel and each transport channel is mapped to a corresponding physical channel, as shown in Figures 4.7 and 4.8 for the downlink and uplink (the sidelink is covered in Chapter 21). In addition to the physical channels with a corresponding transport channel, there are also physical channels without a corresponding transport channel. These channels, known as L1/L2 control channels, are used for *downlink control information* (DCI), providing the device with the necessary information for proper reception and decoding of the downlink data transmission, *uplink control information* (UCI) used for providing the scheduler and the hybrid-ARQ protocol with information about the situation at the device, and *sidelink control information* (SCI) for handling sidelink transmissions.

The following physical-channel types are defined for LTE:

- The *physical downlink shared channel* (PDSCH) is the main physical channel used for unicast data transmission, but also for transmission of paging information.
- The *physical broadcast channel* (PBCH) carries part of the system information, required by the device in order to access the network.

[7]There is no spatial multiplexing defined for sidelink communication.

- The *physical multicast channel* (PMCH) is used for MBSFN transmission.
- The *physical downlink control channel* (PDCCH) is used for downlink control information, mainly scheduling decisions, required for reception of PDSCH, and for scheduling grants enabling transmission on the PUSCH.
- The *enhanced physical downlink control channel* (EPDCCH) was introduced in release 11. It essentially serves the same purpose as the PDCCH but allows for transmission of the control information in a more flexible way.
- The *MTC physical downlink control channel* (MPDCCH) was introduced in release 13 as part of the improved support for massive machine-type communication, see Chapter 20. In essence it is a variant of the EPDCCH.
- The *relay physical downlink control channel* (R-PDCCH) was introduced in release 10 and is used to carry L1/L2 control signaling on the donor-eNodeB-to-relay link.
- The *physical hybrid-ARQ indicator channel* (PHICH) carries the hybrid-ARQ acknowledgment to indicate to the device whether a transport block should be retransmitted or not.
- The *physical control format indicator channel* (PCFICH) is a channel providing the devices with information necessary to decode the set of PDCCHs. There is only one PCFICH per component carrier.
- The *physical uplink shared channel* (PUSCH) is the uplink counterpart to the PDSCH. There is at most one PUSCH per uplink component carrier per device.
- The *physical uplink control channel* (PUCCH) is used by the device to send hybrid-ARQ acknowledgments, indicating to the eNodeB whether the downlink transport block(s) was successfully received or not, to send channel-state reports aiding downlink channel-dependent scheduling, and for requesting resources to transmit uplink data upon. There is at most one PUCCH per device.
- The *physical random-access channel* (PRACH) is used for random access, as described in Chapter 11.
- The *physical sidelink shared channel* (PSSCH) is used for sidelink data transfer, see Chapter 21.
- The *physical sidelink control channel* (PSCCH), used for sidelink-related control information.
- The *physical sidelink discovery channel* (PSDCH), used for sidelink discovery.
- The *physical sidelink broadcast channel* (PSBCH), used to convey basic sidelink-related information between devices.

Note that some of the physical channels, more specifically the channels used for downlink control information (namely, PCFICH, PDCCH, PHICH, EPDCCH, and R-PDCCH), uplink control information (namely, PUCCH), and sidelink control information (namely, PSCCH) do not have a corresponding transport channel mapped to them.

The remaining downlink transport channels are based on the same general physical-layer processing as the DL-SCH, although with some restrictions in the set of features used. This

is especially true for PCH and MCH transport channels. For the broadcast of system information on the BCH, a device must be able to receive this information channel as one of the first steps prior to accessing the system. Consequently, the transmission format must be known to the devices a priori, and there is no dynamic control of any of the transmission parameters from the MAC layer in this case. The BCH is also mapped to the physical resource (the OFDM time–frequency grid) in a different way, as described in more detail in Chapter 11.

For transmission of paging messages on the PCH, dynamic adaptation of the transmission parameters can, to some extent, be used. In general, the processing in this case is similar to the generic DL-SCH processing. The MAC can control modulation, the amount of resources, and the antenna mapping. However, as an uplink has not yet been established when a device is paged, hybrid ARQ cannot be used as there is no possibility for the device to transmit a hybrid-ARQ acknowledgment.

The MCH is used for MBMS transmissions, typically with single-frequency network operation by transmitting from multiple cells on the same resources with the same format at the same time. Hence, the scheduling of MCH transmissions must be coordinated between the cells involved and dynamic selection of transmission parameters by the MAC is not possible.

4.3 CONTROL-PLANE PROTOCOLS

The control-plane protocols are, among other things, responsible for connection setup, mobility, and security. Control messages transmitted from the network to the devices can originate either from the MME, located in the core network, or from the *radio resource control* (RRC), located in the eNodeB.

NAS control-plane functionality, handled by the MME, includes EPS bearer management, authentication, security, and different idle-mode procedures such as paging. It is also responsible for assigning an IP address to a device. For a detailed discussion about the NAS control-plane functionality, see [5].

The RRC is located in the eNodeB and is responsible for handling the RAN-related procedures, including:

- Broadcast of system information necessary for the device to be able to communicate with a cell. Acquisition of system information is described in Chapter 11.
- Transmission of paging messages originating from the MME to notify the device about incoming connection requests. Paging, discussed further in Chapter 11, is used in the RRC_IDLE state (described later) when the device is not connected to a particular cell. Indication of system-information updates is another use of the paging mechanism, as is public warning systems.
- Connection management, including setting up bearers and mobility within LTE. This includes establishing an RRC context—that is, configuring the parameters necessary for communication between the device and the RAN.
- Mobility functions such as cell (re)selection.

CHAPTER

PHYSICAL TRANSMISSION RESOURCES

5

In Chapter 4, the overall LTE architecture was discussed, including an overview of the different protocol layers. Prior to discussing the detailed LTE downlink and uplink transmission schemes, a description of the basic time–frequency transmission resource of LTE is provided in this chapter. An overview of the concept of *antenna ports* is also provided.

5.1 OVERALL TIME–FREQUENCY STRUCTURE

OFDM is the basic transmission scheme for both the downlink and uplink transmission directions in LTE although, for the uplink, specific means are employed to reduce the cubic metric[1] of the transmitted signal, thereby allowing for improved efficiency for the device transmitter power amplifier. Thus, for uplink user data and higher-layer control signaling corresponding to transmission of the PUSCH physical channel, DFT precoding is applied before OFDM modulation, leading to *DFT-spread OFDM* or *DFTS-OFDM*, see further Chapter 7. As will be described in Chapter 7, for other uplink transmissions, such as the transmission of *L1/L2 control signaling* and different types of *reference-signal transmissions*, other measures are taken to limit the cubic metric of the transmitted signal.

The LTE OFDM subcarrier spacing equals 15 kHz for both downlink and uplink. The selection of the subcarrier spacing in an OFDM-based system needs to carefully balance overhead from the cyclic prefix against sensitivity to Doppler spread/shift and other types of frequency errors and inaccuracies. The choice of 15 kHz for the LTE subcarrier spacing was found to offer a good balance between these different constraints.

Assuming an FFT-based transmitter/receiver implementation, 15 kHz subcarrier spacing corresponds to a sampling rate $f_s = 15000\, N_{\text{FFT}}$, where N_{FFT} is the FFT size. It is important to understand though that the LTE specifications do not in any way mandate the use of FFT-based transmitter/receiver implementations and even less so a particular FFT size or sampling rate. Nevertheless, FFT-based implementations of OFDM are common practice and an

[1]The cubic metric is a better measure than the peak-to-average ratio of the amount of additional back-off needed for a certain signal waveform, relative to the back-off needed for some reference waveform [10].

4G, LTE-Advanced Pro and The Road to 5G. http://dx.doi.org/10.1016/B978-0-12-804575-6.00005-4
Copyright © 2016 Erik Dahlman, Stefan Parkvall and Johan Sköld. Published by Elsevier Ltd. All rights reserved.

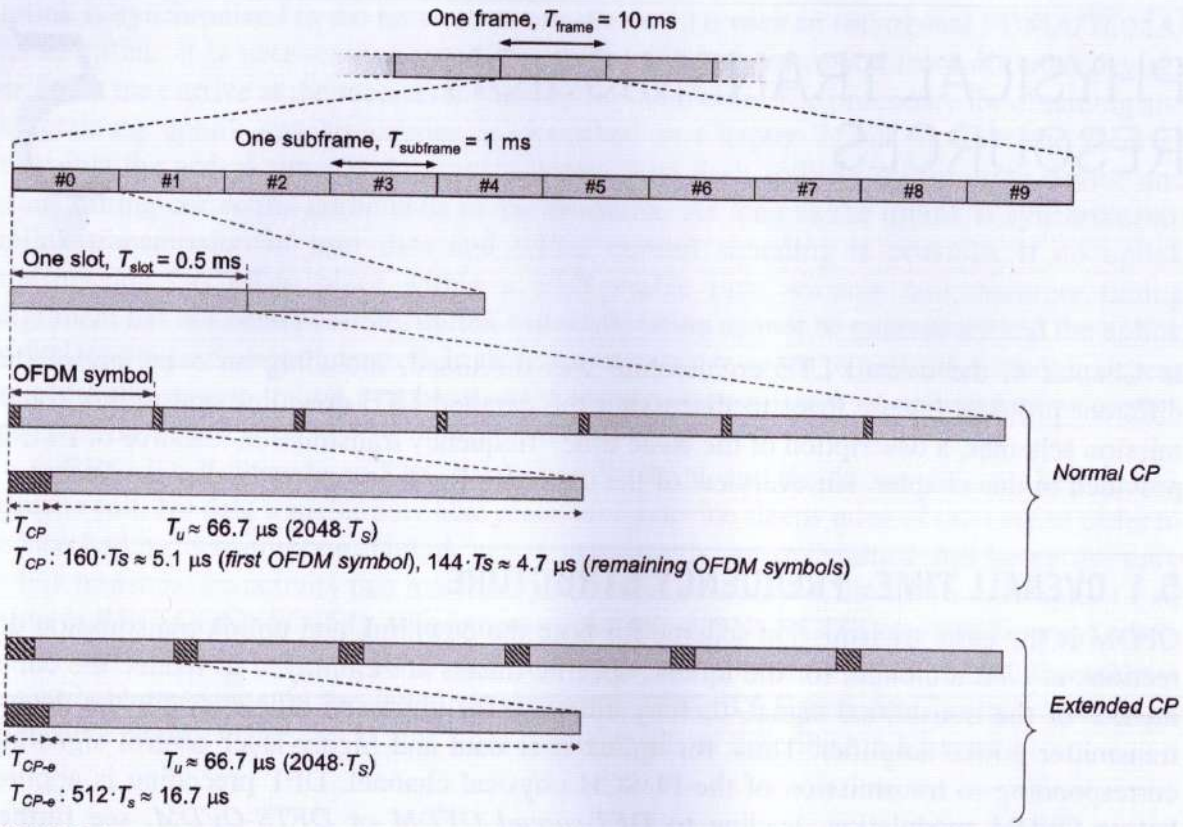


FIGURE 5.1

LTE time-domain structure.

FFT size of 2048, with a corresponding sampling rate of 30.72 MHz, is suitable for the wider LTE carrier bandwidths, such as bandwidths of the order of 15 MHz and above. However, for smaller carrier bandwidths, a smaller FFT size and a correspondingly lower sampling rate can very well be used.

In addition to the 15 kHz subcarrier spacing, a *reduced subcarrier spacing* of 7.5 kHz, with a corresponding OFDM symbol time that is twice as long, is also defined for LTE. The introduction of the reduced subcarrier spacing specifically targeted MBSFN-based multicast/broadcast transmissions (see Chapter 19). However, currently the 7.5 kHz subcarrier numerology is only partly implemented in the LTE specifications. The remaining discussions within this and the following chapters will assume the 15 kHz subcarrier spacing unless explicitly stated otherwise.

In the time domain, LTE transmissions are organized into *frames* of length 10 ms, each of which is divided into ten equally sized *subframes* of length 1 ms, as illustrated in Figure 5.1.

Each subframe consists of two equally sized *slots* of length $T_{slot} = 0.5$ ms, with each slot consisting of a number of OFDM symbols including cyclic prefix.[2] To provide consistent and exact timing definitions, different time intervals within the LTE specifications are defined as multiples of a basic time unit $T_s = 1/(15000 \cdot 2048)$ s. The basic time unit T_s can thus be seen as the sampling time of an FFT-based transmitter/receiver implementation with an FFT size equal to 2048. The time intervals outlined in Figure 5.1 can thus also be expressed as $T_{frame} = 307200\ T_s$, $T_{subframe} = 30720\ T_s$, and $T_{slot} = 15360\ T_s$ for the frame, subframe, and slot durations, respectively.

On a higher level, each frame is identified by a *System Frame Number* (SFN). The SFN is used to define different transmission cycles that have a period longer than one frame, for example, paging cycles, see Chapter 11. The SFN period equals 1024, thus the SFN repeats itself after 1024 frames or 10.24 s.

The 15 kHz LTE subcarrier spacing corresponds to a useful symbol time $T_u = 2048\ T_s$ or approximately 66.7 μs. The overall OFDM symbol time is then the sum of the useful symbol time and the cyclic prefix length T_{CP}. As illustrated in Figure 5.1, LTE defines two cyclic prefix lengths, the *normal* cyclic prefix and an *extended* cyclic prefix, corresponding to seven and six OFDM symbols per slot, respectively. The exact cyclic prefix lengths, expressed in the basic time unit T_s, are given in Figure 5.1. It can be noted that, in the case of the normal cyclic prefix, the cyclic prefix length for the first OFDM symbol of a slot is somewhat larger compared to the remaining OFDM symbols. The reason for this is simply to fill the entire 0.5 ms slot, as the number of basic time units T_s per slot (15360) is not divisible by seven.

The reasons for defining two cyclic prefix lengths for LTE are twofold:

- A longer cyclic prefix, although less efficient from a cyclic prefix-overhead point of view, may be beneficial in specific environments with extensive delay spread, for example, in very large cells. It is important to have in mind, though, that a longer cyclic prefix is not necessarily beneficial in the case of large cells, even if the delay spread is very extensive in such cases. If, in large cells, link performance is limited by noise rather than by signal corruption due to residual time dispersion not covered by the cyclic prefix, the additional robustness to radio-channel time dispersion, due to the use of a longer cyclic prefix, may not justify the corresponding additional energy overhead of a longer cyclic prefix.
- As is discussed in Chapter 19, the cyclic prefix in the case of MBSFN-based multicast/broadcast transmission should not only cover the main part of the actual channel time dispersion, but also the timing difference between the transmissions received from the cells involved in the MBSFN transmission. In the case of MBSFN operation, the extended cyclic prefix is therefore used.

It should be noted that different cyclic prefix lengths may be used for different subframes within a frame in case of MBSFN subframes. As discussed further in Chapter 19, MBSFN-

[2]This is valid for the "normal" downlink subframes. As described in Section 5.4, the *special subframe* present in case of TDD operation is not divided into two slots but rather into three fields.

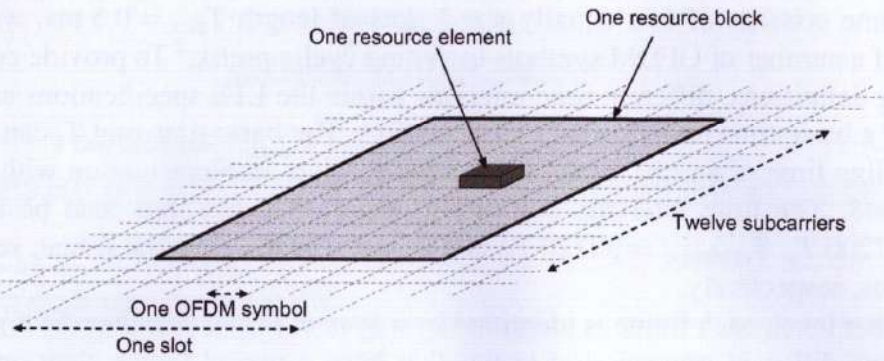


FIGURE 5.2

The LTE physical time–frequency resource.

based multicast/broadcast transmission is always confined to a limited set of subframes, in which case the use of the extended cyclic prefix, with its associated additional cyclic prefix overhead, may only be applied to these subframes.[3]

A *resource element*, consisting of one subcarrier during one OFDM symbol, is the smallest physical resource in LTE. Furthermore, as illustrated in Figure 5.2, resource elements are grouped into *resource blocks*, where each resource block consists of 12 consecutive subcarriers in the frequency domain and one 0.5 ms slot in the time domain. Each resource block thus consists of $7 \cdot 12 = 84$ resource elements in the case of a normal cyclic prefix and $6 \cdot 12 = 72$ resource elements in the case of an extended cyclic prefix.

Although resource blocks are defined over one slot, the basic time-domain unit for dynamic scheduling in LTE is one subframe, consisting of two consecutive slots. The reason for defining the resource blocks over one slot is that *distributed downlink transmission* (described in Chapter 10) and *uplink frequency hopping* (described in Chapter 7) are defined on a slot basis. The minimum scheduling unit, consisting of two time-consecutive resource blocks within one subframe (one resource block per slot), can be referred to as a *resource-block pair.*

The LTE physical-layer specifications allow for a carrier to consist of any number of resource blocks in the frequency domain, ranging from a minimum of six resource blocks up to a maximum of 110 resource blocks. This corresponds to an overall transmission bandwidth ranging from roughly 1 MHz up to in the order of 20 MHz with very fine granularity and thus allows for a very high degree of LTE bandwidth flexibility, at least from a physical-layer-specification point of view. However, as mentioned in Chapter 3, LTE radio-frequency requirements are, at least initially, only specified for a limited set of transmission bandwidths, corresponding to a limited set of possible values for the number

[3]The extended cyclic prefix is then actually applied only to the so-called *MBSFN part* of the MBSFN subframes; see Section 5.2.

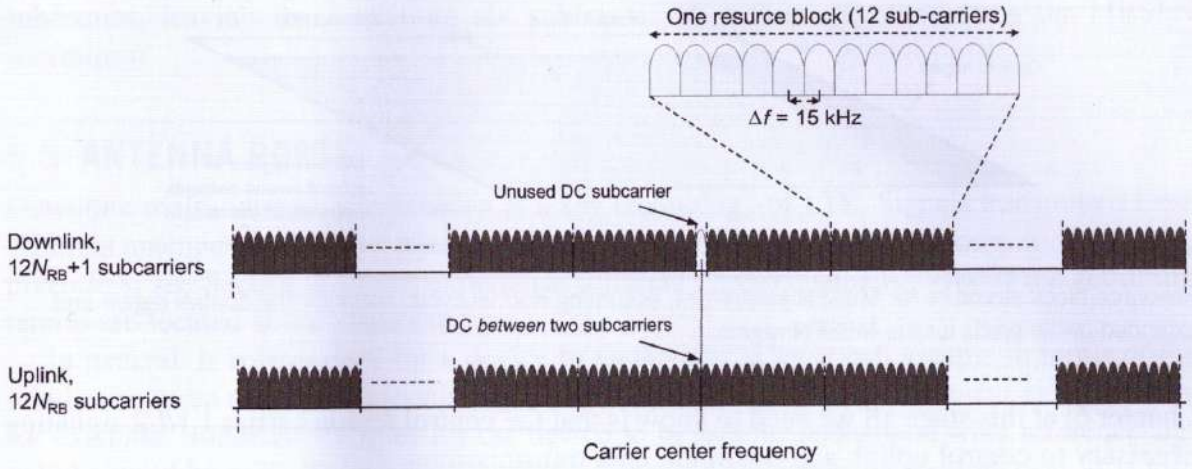


FIGURE 5.3

Frequency-domain structure for LTE.

of resource blocks within a carrier. Also note that, in LTE release 10 and later, the total bandwidth of the transmitted signal can be significantly larger than 20 MHz, up to 100 MHz in release 10 and 640 MHz in release 13, by aggregating multiple carriers, see Section 5.5.

The resource-block definition above applies to both the downlink and uplink transmission directions. However, there is a minor difference between the downlink and uplink in terms of where the carrier center frequency is located in relation to the subcarriers.

In the downlink (upper part of Figure 5.3), there is an unused *DC subcarrier* that coincides with the carrier center frequency. The reason why the DC subcarrier is not used for downlink transmission is that it may be subject to disproportionately high interference due to, for example, local-oscillator leakage.

On the other hand, in the uplink (lower part of Figure 5.3), no unused DC subcarrier is defined and the center frequency of an uplink carrier is located *between* two uplink subcarriers. The presence of an unused DC-carrier in the center of the spectrum would have prevented the assignment of the entire cell bandwidth to a single device and while still retaining the assumption of mapping to consecutive inputs of the OFDM modulator, something that is needed to retain the low-cubic-metric property of the DFTS-OFDM modulation used for uplink data transmission.

5.2 NORMAL SUBFRAMES AND MBSFN SUBFRAMES

In LTE, each downlink subframe (and the DwPTS in the case of TDD; see Section 5.4.2 for a discussion of the TDD frame structure) is normally divided into a *control region*, consisting of the first few OFDM symbols, and a *data region*, consisting of the remaining part of the subframe. The usage of the resource elements in the two regions is discussed in detail in

different. BCH transmission is described in Chapter 11 as part of the discussion on LTE system information.

6.1.1 PROCESSING STEPS

The different steps of the DL-SCH physical-layer processing are outlined in Figure 6.1. In the case of carrier aggregation—that is transmission on multiple component carriers in parallel to the same device—the transmissions on the different carriers correspond to separate transport channels with separate and essentially independent physical-layer processing. The transport-channel processing outlined in Figure 6.1 and the discussion as follows is thus valid also in the case of carrier aggregation.

Within each *transmission time interval* (TTI), corresponding to one subframe of length 1 ms, up to two transport blocks of dynamic size are delivered to the physical layer and transmitted over the radio interface for each component carrier. The number of transport

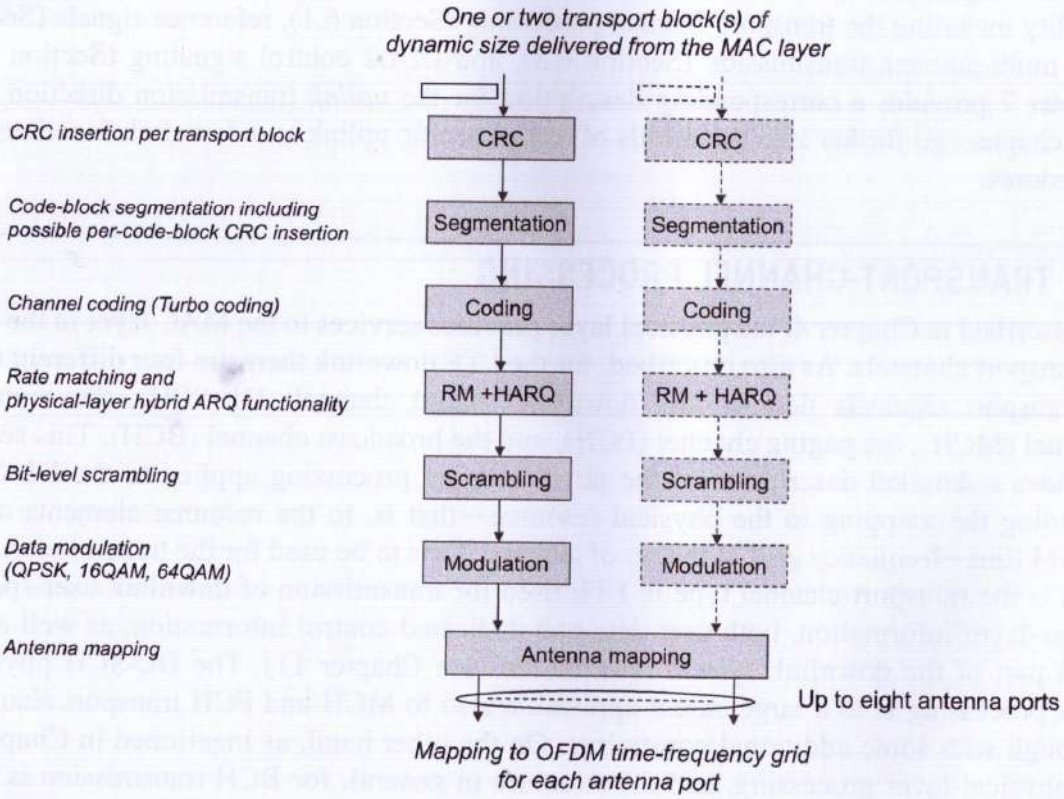


FIGURE 6.1

Physical-layer processing for the downlink shared channel (DL-SCH).

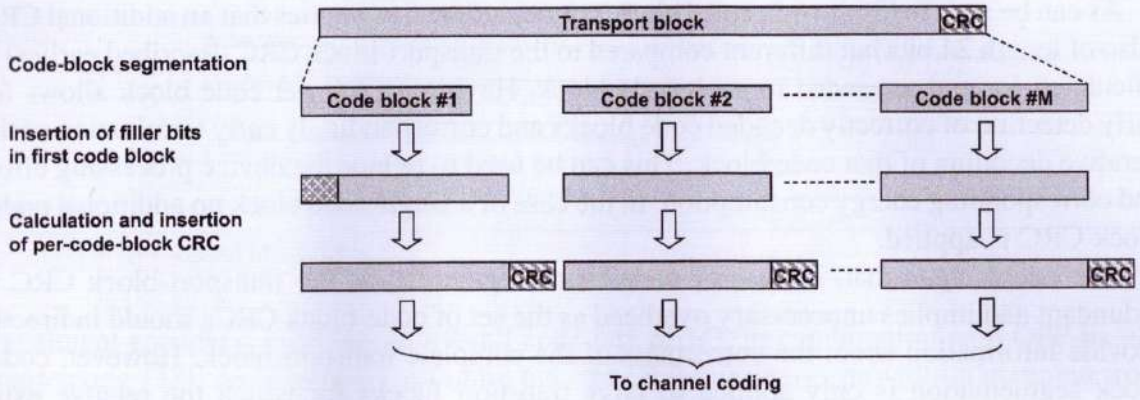


FIGURE 6.2

Code-block segmentation and per-code-block CRC insertion.

blocks transmitted within a TTI depends on the configuration of the multi-antenna transmission scheme (see Section 6.3):

- In the case of no *spatial multiplexing* there is at most a single transport block in a TTI.
- In the case of spatial multiplexing, with transmission on multiple *layers* in parallel to the same device, there are two transport blocks within a TTI.[1]

6.1.1.1 CRC Insertion per Transport Block

In the first step of the physical-layer processing, a 24-bit CRC is calculated for and appended to each transport block. The CRC allows for receiver-side detection of errors in the decoded transport block. The corresponding error indication can, for example, be used by the downlink hybrid-ARQ protocol as a trigger for requesting retransmissions.

6.1.1.2 Code-Block Segmentation and per-Code-Block CRC Insertion

The LTE Turbo-coder internal interleaver is only defined for a limited number of code-block sizes, with a maximum block size of 6144 bits. If the transport block, including the transport-block CRC, exceeds this maximum code-block size, *code-block segmentation* is applied before the Turbo coding as illustrated in Figure 6.2. Code-block segmentation implies that the transport block is segmented into smaller *code blocks*, the sizes of which should match the set of code-block sizes supported by the Turbo coder.

In order to ensure that a transport block of arbitrary size can be segmented into code blocks that match the set of available code-block sizes, the LTE specification includes the possibility to insert "dummy" *filler bits* at the head of the first code block. However, the set of transport-block sizes currently defined for LTE has been selected so that filler bits are not needed.

[1]This is true for initial transmissions. In the case of hybrid-ARQ retransmissions there may also be cases when a single transport block is transmitted over multiple layers as discussed, for example, in Section 6.3.

As can be seen in Figure 6.2, code-block segmentation also implies that an additional CRC (also of length 24 bits but different compared to the transport-block CRC described earlier) is calculated for and appended to each code block. Having a CRC per code block allows for early detection of correctly decoded code blocks and correspondingly early termination of the iterative decoding of that code block. This can be used to reduce the device processing effort and corresponding energy consumption. In the case of a single code block no additional code-block CRC is applied.

One could argue that, in case of code-block segmentation, the transport-block CRC is redundant and implies unnecessary overhead as the set of code-block CRCs should indirectly provide information about the correctness of the complete transport block. However, code-block segmentation is only applied to large transport blocks for which the relative extra overhead due to the additional transport-block CRC is small. The transport-block CRC also adds additional error-detection capabilities and thus further reduces the risk for undetected errors in the decoded transport block.

Information about the transport-block size is provided to the device as part of the scheduling assignment transmitted on the physical downlink control channel/enhanced physical downlink control channel (PDCCH/EPDCCH), as described in Section 6.4. Based on this information, the device can determine the code-block size and number of code blocks. The device receiver can thus, based on the information provided in the scheduling assignment, straightforwardly undo the code-block segmentation and recover the decoded transport blocks.

6.1.1.3 Channel Coding

Channel coding for DL-SCH (as well as for PCH and MCH) is based on Turbo coding [17], with encoding according to Figure 6.3. The encoding consists of two rate-1/2, eight-state

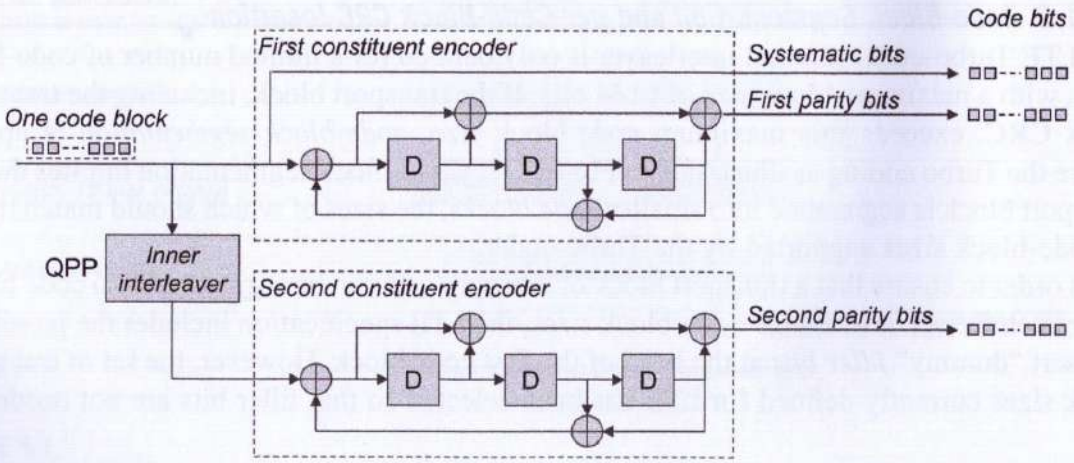


FIGURE 6.3

LTE Turbo encoder.

6.1.2 LOCALIZED AND DISTRIBUTED RESOURCE MAPPING

As already discussed in Chapter 3, when deciding what set of resource blocks to use for transmission to a specific device, the network may take the downlink channel conditions in both the time and frequency domains into account. Such time–frequency-domain channel-dependent scheduling, taking channel variations—for example, due to frequency-selective fading—into account, may significantly improve system performance in terms of achievable data rates and overall cell throughput.

However, in some cases downlink channel-dependent scheduling is not suitable to use or is not practically possible:

- For low-rate services such as voice, the feedback signaling associated with channel-dependent scheduling may lead to extensive relative overhead.
- At high mobility (high device speed), it may be difficult or even practically impossible to track the instantaneous channel conditions to the accuracy required for channel-dependent scheduling to be efficient.

In such situations, an alternative means to handle radio-channel frequency selectivity is to achieve frequency diversity by distributing a downlink transmission in the frequency domain.

One way to distribute a downlink transmission in the frequency domain, and thereby achieve frequency diversity, is to assign multiple nonfrequency-contiguous resource blocks for the transmission to a device. LTE allows for such *distributed resource-block allocation* by means of *resource allocation types 0 and 1* (see Section 6.4.6.1). However, although sufficient in many cases, distributed resource-block allocation by means of these resource-allocation types has certain drawbacks:

- For both types of resource allocations, the minimum size of the allocated resource can be as large as four resource-block pairs and may thus not be suitable when resource allocations of smaller sizes are needed.
- In general, both these resource-allocation methods are associated with a relatively large control-signaling overhead for the scheduling assignment, see Section 6.4.6

In contrast, *resource-allocation type 2* (Section 6.4.6.1) always allows for the allocation of a single resource-block pair and is also associated with a relatively small control-signaling overhead. However, resource allocation type 2 only allows for the allocation of resource blocks that are contiguous in the frequency domain. In addition, regardless of the type of resource allocation, frequency diversity by means of distributed resource-block allocation will only be achieved in the case of resource allocations larger than one resource-block pair.

In order to provide the possibility for distributed resource-block allocation in the case of resource-allocation type 2, as well as to allow for distributing the transmission of a single resource-block pair in the frequency domain, the notion of a *virtual resource block* (VRB) has been introduced for LTE.

What is being provided in the resource allocation is the resource allocation in terms of VRB pairs. The key to distributed transmission then lies in the mapping from VRB pairs to *physical resource block* (PRB) pairs—that is, to the actual physical resource used for transmission.

The LTE specification defines two types of VRBs: *localized* VRBs and *distributed* VRBs. In the case of localized VRBs, there is a direct mapping from VRB pairs to PRB pairs as illustrated in Figure 6.6.

However, in the case of distributed VRBs, the mapping from VRB pairs to PRB pairs is more elaborate in the sense that:

- consecutive VRBs are not mapped to PRBs that are consecutive in the frequency domain;
- even a single VRB pair is distributed in the frequency domain.

The basic principle of distributed transmission is outlined in Figure 6.7 and consists of two steps:

- A mapping from VRB pairs to PRB pairs such that consecutive VRB pairs are not mapped to frequency-consecutive PRB pairs (first step of Figure 6.7). This provides frequency diversity between consecutive VRB pairs. The spreading in the frequency domain is done by means of a block-based "interleaver" operating on resource-block pairs.
- A split of each resource-block pair such that the two resource blocks of the resource-block pair are transmitted with a certain frequency gap in between (second step of Figure 6.7). This also provides frequency diversity for a single VRB pair. This step can be seen as the introduction of frequency hopping on a slot basis.

Whether the VRBs are localized (and thus mapped according to Figure 6.6) or distributed (mapped according to Figure 6.7) is indicated as part of the scheduling assignment in the case

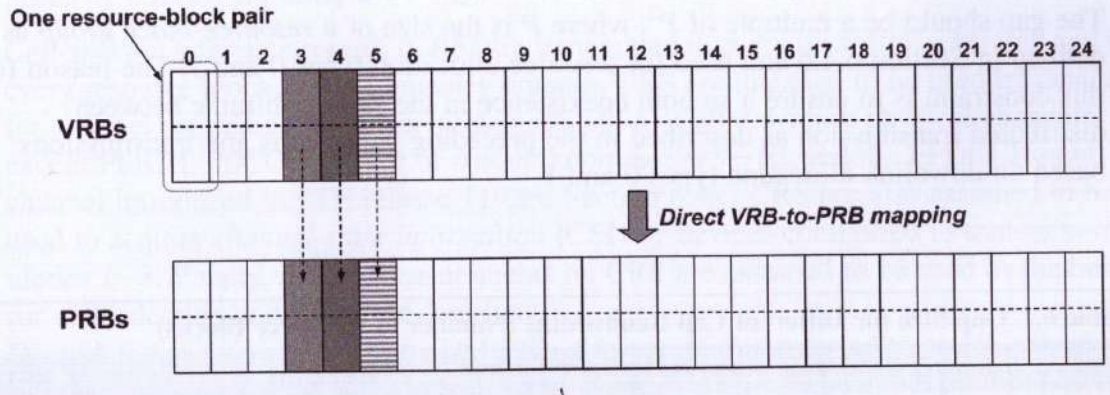


FIGURE 6.6

VRB-to-PRB mapping in case of localized VRBs. Figure assumes a cell bandwidth corresponding to 25 resource blocks.

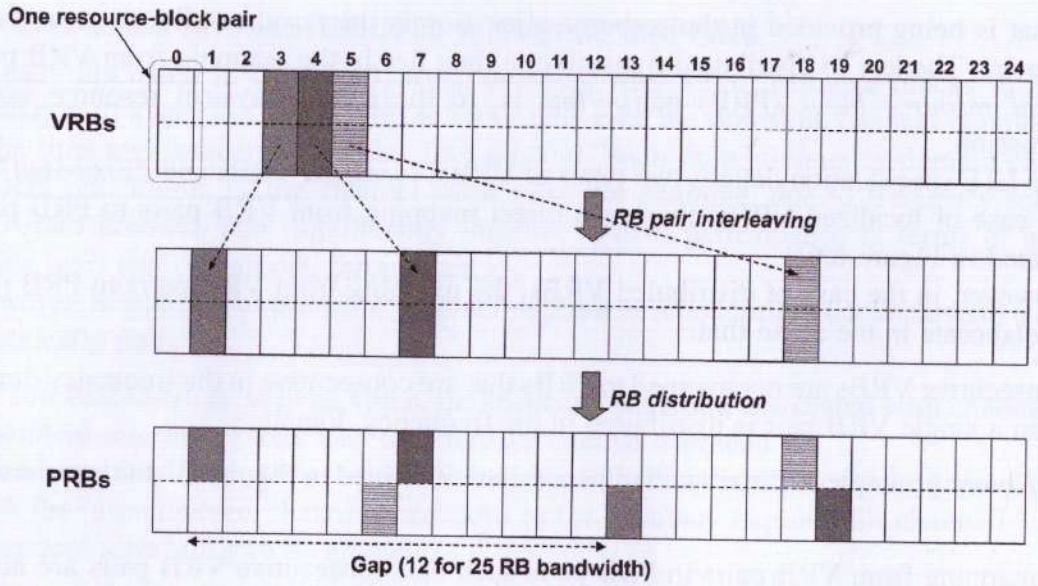


FIGURE 6.7

VRB-to-PRB mapping in case of distributed VRBs. Figure assumes a cell bandwidth corresponding to 25 resource blocks.

of type 2 resource allocation. Thus, it is possible to dynamically switch between distributed and localized transmission and also mix distributed and localized transmission for different devices within the same subframe.

The exact size of the frequency gap in Figure 6.7 depends on the overall downlink cell bandwidth according to Table 6.1. These gaps have been chosen based on two criteria:

1. The gap should be of the order of half the downlink cell bandwidth in order to provide good frequency diversity also in the case of a single VRB pair.
2. The gap should be a multiple of P^2, where P is the size of a *resource-block group* as defined in Section 6.4.6 and used for resource allocation types 0 and 1. The reason for this constraint is to ensure a smooth coexistence in the same subframe between distributed transmission as described in the preceding paragraphs and transmissions based on downlink allocation types 0 and 1.

Table 6.1 Gap Size for Different Cell Bandwidths (Number of Resource Blocks)

Bandwidth	6	7–8	9–10	11	12–19	20–26	27–44	45–63	64–79	80–110
P	1	1	1	2	2	2	3	3	4	4
Gap size	3	4	5	4	8	12	18	27	32	48

Table 6.2 Second Gap Size for Different Cell Bandwidth (Only Applicable for Cell Bandwidths of 50 RBs and Beyond)

Bandwidth	50–63	64–110
Gap size	9	16

Due to the constraint that the gap size should be a multiple of P^2, the gap size will in most cases deviate from exactly half the cell bandwidth. In these cases, not all resource blocks within the cell bandwidth can be used for distributed transmission. As an example, for a cell bandwidth corresponding to 25 resource blocks (the example in Figure 6.7) and a corresponding gap size equal to 12 according to Table 6.1, the 25th resource-block pair cannot be used for distributed transmission. As another example, for a cell bandwidth corresponding to 50 resource blocks (gap size equal to 27 according to Table 6.1) only 46 resource blocks would be available for distributed transmission.

In addition to the gap size outlined in Table 6.1, for wider cell bandwidths (50 RBs and beyond), there is a possibility to use a second, smaller frequency gap with a size of the order of one-fourth of the cell bandwidth (see Table 6.2). The use of the smaller gap enables restriction of the distributed transmission to only a part of the overall cell bandwidth. Selection between the larger gap according to Table 6.1 and the smaller gap according to Table 6.2 is indicated by an additional bit in the resource allocation.

6.2 DOWNLINK REFERENCE SIGNALS

Downlink reference signals are predefined signals occupying specific resource elements within the downlink time–frequency grid. The LTE specification includes several types of downlink reference signals transmitted in different ways and intended to be used for different purposes by a receiving device:

- *Cell-specific reference signals* (CRS) are transmitted in every downlink subframe and in every resource block in the frequency domain. CRS are intended to be used by devices for channel estimation for coherent demodulation of all downlink physical channels except PMCH, PDSCH in case of *transmission modes* 7–10, and the EPDCCH control channel introduced in LTE release 11 (see Section 6.4).[6] CRS are also assumed to be used to acquire *channel-state information* (CSI) by devices configured in transmission modes 1–8. Finally, device measurements on CRS are assumed to be used as the basis for cell-selection and handover decisions.
- *Demodulation reference signals* (DM-RS), also sometimes referred to as *UE-specific reference signals*, are intended to be used by devices for channel estimation for coherent

[6]See Section 6.3.1 for more details on LTE *transmission modes*.

of higher-order modulation, the device needs to know the power offset between the CRS and the PDSCH.

The device is informed about this power offset by means of higher-layer signaling. However, what is then provided is the offset between CRS power and the overall PDSCH power, including all layers. In the case of spatial multiplexing, the overall PDSCH power has to be divided between the different layers, and it is the relation between the CRS power and the per-layer PDSCH power that is relevant for demodulation.

In the case of pure spatial multiplexing (no MU-MIMO)—that is, transmission modes 3 and 4—the device knows about the number of layers and thus, indirectly, about the offset between the CRS power and the per-layer PDSCH power.

In the case of MU-MIMO, the total available power will typically also be divided between the transmissions to the different devices, with less PDSCH power being available for each transmission. However, devices are not aware of the presence of parallel transmissions to other devices and are thus not aware of any per-PDSCH power reduction. For this reason, transmission mode 5 includes the explicit signaling of an *additional power offset* of −3 dB to be used by the device in addition to the CRS/PDSCH power offset signaled by higher layers.

Transmission mode 5 is limited to single-rank transmission and, in practice, limited to two users being scheduled in parallel, as there is only a single −3 dB offset defined.

Note that the power-offset signaling is not needed for MU-MIMO based on DM-RS as, in this case, each transmission has its own set of reference signals. The power per reference signal will thus scale with the number of layers and transmissions, similar to the PDSCH power, and the reference signal to per-layer PDSCH power ratio will remain constant.

6.4 DOWNLINK L1/L2 CONTROL SIGNALING

To support the transmission of downlink and uplink transport channels, there is a need for certain *associated downlink control signaling*. This control signaling is often referred to as *downlink L1/L2 control signaling*, indicating that the corresponding information partly originates from the physical layer (Layer 1) and partly from Layer 2 MAC. Downlink L1/L2 control signaling consists of downlink scheduling assignments, including information required for the device to be able to properly receive, demodulate, and decode the DL-SCH[29] on a component carrier, uplink scheduling grants informing the device about the resources and transport format to use for uplink (UL-SCH) transmission, and hybrid-ARQ acknowledgments in response to UL-SCH transmissions. In addition, the downlink control signaling can also be used for the transmission of power-control commands for power control of uplink physical channels, as well as for certain special purposes such as MBSFN notifications.

The basic time–frequency structure for transmission of L1/L2 control signaling is illustrated in Figure 6.20 with control signaling being located at the beginning of each subframe

[29]L1/L2 control signaling is also needed for the reception, demodulation, and decoding of the PCH transport channel.

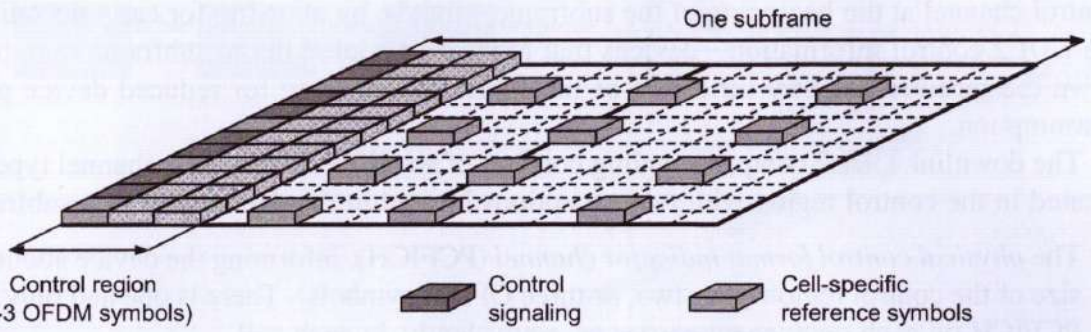


FIGURE 6.20

LTE time–frequency grid illustrating the split of the subframe into (variable-sized) control and data regions.

and spanning the full downlink carrier bandwidth. Each subframe can therefore be said to be divided into a *control region* followed by a *data region*, where the control region corresponds to the part of the subframe in which the L1/L2 control signaling is transmitted. Starting from release 11, there is also a possibility to locate parts of the L1/L2 control signaling in the data region as described later. However, the split of a subframe into a control region and a data region still applies.

To simplify the overall design, the control region always occupies an integer number of OFDM symbols, more specifically one, two, or three OFDM symbols (for narrow cell bandwidths, 10 resource blocks or less, the control region consists of two, three, or four OFDM symbols to allow for a sufficient amount of control signaling).

The size of the control region expressed in number of OFDM symbols, or, equivalently, the start of the data region, can be dynamically varied on a per-subframe basis. Thus, the amount of radio resources used for control signaling can be dynamically adjusted to match the instantaneous traffic situation. For a small number of users being scheduled in a subframe, the required amount of control signaling is small and a larger part of the subframe can be used for data transmission (larger data region).

The maximum size of the control region is normally three OFDM symbols (four in the case of narrow cell bandwidths), as mentioned in the preceding paragraphs. However, there are a few exceptions to this rule. When operating in TDD mode, the control region in subframes one and six is restricted to at most two OFDM symbols since, for TDD, the primary synchronization signal (see Chapter 11) occupies the third OFDM symbol in those subframes. Similarly, for MBSFN subframes (see Chapter 5), the control region is restricted to a maximum of two OFDM symbols.

The reason for transmitting the control signaling at the beginning of the subframe is to allow for devices to decode downlink scheduling assignments as early as possible. Processing of the data region—that is, demodulation and decoding of the DL-SCH transmission—can then begin before the end of the subframe. This reduces the delay in the DL-SCH decoding and thus the overall downlink transmission delay. Furthermore, by transmitting the L1/L2

control channel at the beginning of the subframe—that is, by allowing for early decoding of the L1/L2 control information—devices that are not scheduled in the subframe may power down the receiver circuitry for a part of the subframe, allowing for reduced device power consumption.

The downlink L1/L2 control signaling consists of six different physical-channel types, all located in the control region with two exceptions located in the data region of a subframe:

- The *physical control format indicator channel* (PCFICH), informing the device about the size of the control region (one, two, or three OFDM symbols). There is one and only one PCFICH on each component carrier or, equivalently, in each cell.
- The *physical hybrid-ARQ indicator channel* (PHICH), used to signal hybrid-ARQ acknowledgments in response to uplink UL-SCH transmissions. Multiple PHICHs can exist in each cell.
- The *physical downlink control channel* (PDCCH), used to signal downlink scheduling assignments, uplink scheduling grants, or power-control commands. Each PDCCH typically carries signaling for a single device, but can also be used to address a group of devices. Multiple PDCCHs can exist in each cell.
- The *enhanced physical downlink control channel* (EPDCCH), introduced in release 11 to support DM-RS-based control signaling reception and carrying similar types of information as the PDCCH. However, in contrast to the PDCCH, the EPDCCH is located in the data region. Also, the EPDCCH can be subject to non-codebook-based precoding.
- The *MTC physical downlink control channel* (MPDCCH), introduced in release 13 as part of the improved MTC support, see Chapter 20. In essence it is a variant of the EPDCCH.
- The *relay physical downlink control channel* (R-PDCCH), introduced in release 10 to support relaying. A detailed discussion can be found in Chapter 18 in conjunction with the overall description of relays; at this stage it suffices to note that the R-PDCCH is transmitted in the data region.

In the following sections, the PCFICH, PHICH, PDCCH, and EPDCCH are described in detail while the MPDCCH and the R-PDCCH are described in Chapters 20 and 18, respectively.

6.4.1 PHYSICAL CONTROL FORMAT INDICATOR CHANNEL

The PCFICH indicates the instantaneous size of the control region in terms of the number of OFDM symbols—that is, indirectly where in the subframe the data region starts. Correct decoding of the PCFICH information is thus essential. If the PCFICH is incorrectly decoded, the device will neither know how to process the control channels nor where the data region starts for the corresponding subframe.[30] The PCFICH consists of two bits of information,

[30]Theoretically, the device could blindly try to decode all possible control channel formats and, from which format that was correctly decoded, deduce the starting position of the data region, but this can be a very complex procedure.

As mentioned, the device is monitoring *either* the device-specific PDCCH search spaces *or* the device-specific EPDCCH search spaces. The basic principle is that the device is monitoring EPDCCHs in all subframes whenever EPDCCH support is enabled, except for the special subframe in configurations where there are no DM-RS supported and hence no possibility to receive the EPDCCH. However, to provide additional control of when the EPDCCH is monitored, it is possible to provide a bitmap to the device, indicating in which subframes it should monitor EPDCCHs and in which subframes it should monitor PDCCHs. One possible use case for this is PDSCH transmission in MBSFN subframes. In MBSFN subframes where the PMCH is transmitted, it is not possible to transmit the EPDCCH and the device needs to be addressed using the PDCCH, while in MBSFN subframes not used for PMCH transmission, the EPDCCH can be used.

In the preceding paragraphs, the search spaces for PDCCH and EPDCCH in terms of control channel candidates are described. However, to determine whether a control channel candidate contains relevant downlink control information or not, the contents have to be decoded. If the CRC, which includes the device identity, checks, then the content of the control channel is declared as valid for this device and the device processes the information (scheduling assignment, scheduling grants, and so on). Thus, for each control-channel candidate, the device needs to attempt to decode the contents once for each supported DCI format. The downlink DCI formats to decode in the device-specific search spaces depend on the *transmission mode* configured for the device. Transmission modes are described in Section 6.3 and, in principle, correspond to different multi-antenna configurations. As an example, there is no need to attempt to decode DCI format 2 when the device has not been configured for spatial multiplexing, which helps to reduce the number of blind decoding attempts. The DCI formats a device addressed using the C-RNTI should monitor as a function of the transmission mode are listed in Table 6.5. Note that DCI format 1C is monitored in the common search spaces as well, although not with the C-RNTI identity. As seen from the table, there are two DCI formats to monitor in the device-specific search spaces and one in the common search spaces. In addition, a device also needs to monitor DCI format 1C in the common search space. Hence, with 16 PDCCH/EPDCCH candidates in the device-specific search spaces and six in the common search spaces, a device needs to execute $2 \cdot 16 + 2 \cdot 6 = 44$ blind decoding attempts in each subframe. With uplink spatial multiplexing, introduced in release 10, an additional uplink DCI format needs to be monitored in the device-specific search spaces, increasing the number of blind decoding attempts to $3 \cdot 16 + 2 \cdot 6 = 60$. These numbers are for a single component carrier; in case of carrier aggregation the number of blind decodes is increased further as discussed in Chapter 12. Finally, note that a certain device in some circumstances may be addressed with different RNTIs. For example, DCI format 1A in the common search spaces may be used either with the C-RNTI for normal scheduling purposes or with the SI-RNTI for scheduling of system information. This does not affect the number of blind decoding attempts, as they are related to

Table 6.5 Downlink DCI Formats Monitored in Different Search Spaces for C-RNTI. Note that DCI Format 1C is Monitored in the Common Search Space as Well, Although Not with the C-RNTI Identity

Mode	Search Space			Description	Release
	Common (PDCCH)	Device-Specific (PDCCH or EPDCCH)			
1	1A	1A	1	Single antenna transmission	8
2			1	Transmit diversity	
3			2A	Open-loop spatial multiplexing	
4			2	Closed-loop spatial multiplexing	
5			1D	MU-MIMO	
6			1B	Single-layer codebook-based precoding	
7			1	Single-layer transmission using DM-RS	
8			2B	Dual-layer transmission using DM-RS	9
9			2C	Multi-layer transmission using DM-RS	10
10			2D	Multi-layer transmission using DM-RS	11

the DCI format; checking two different RNTIs—that is, checking two different CRCs, after decoding is a very low-complexity operation.

Configuration of the transmission mode is done via RRC signaling. As the exact subframe number when this configuration takes effect in the device is not specified and may vary depending on, for example, RLC retransmissions, there is a (short) period when the network and the device may have different understandings of which transmission mode is configured. Therefore, in order not to lose the possibility of communicating with the device, it is necessary to have at least one DCI format that is decoded irrespective of the transmission mode. For downlink transmissions, DCI format 1A serves this purpose and the network can therefore always transmit data to the device using this DCI format. Another function of format 1A is to reduce overhead for transmissions when full flexibility in resource block assignment is not needed.

6.4.6 DOWNLINK SCHEDULING ASSIGNMENTS

Having described the transmission of downlink control information on PDCCH and EPDCCH, the detailed contents of the control information can be discussed, starting with the downlink scheduling assignments. Downlink scheduling assignments are valid for the same subframe in which they are transmitted. The scheduling assignments use one of the DCI formats 1, 1A, 1B, 1C, 1D, 2, 2A, 2B, 2C, or 2D and the DCI formats used depend on the transmission mode configured (see Table 6.5 for the relation between DCI formats and transmission modes). The reason for supporting multiple formats with different message sizes

for the same purpose is to allow for a trade-off in control-signaling overhead and scheduling flexibility. Parts of the contents are the same for the different DCI formats, as seen in Table 6.6, but there are also differences due to the different capabilities.

DCI format 1 is the basic downlink assignment format in the absence of spatial multiplexing (transmission modes 1, 2, and 7). It supports noncontiguous allocations of resource blocks and the full range of modulation-and-coding schemes.

Table 6.6 DCI Formats used for Downlink Scheduling

Field		DCI Format									
		1	1A	1B	1C	1D	2	2A	2B	2C	2D
Resource information	Carrier indicator	•	•	•		•	•	•	•	•	•
	Resource block assignment type	0/1	2	2	2'	2	0/1	0/1	0/1	0/1	0/1
HARQ process number		•	•	•		•	•	•	•	•	•
1st transport block	MCS	•	•	•	•	•	•	•	•	•	•
	RV	•	•	•		•	•	•	•	•	•
	NDI	•	•	•		•	•	•	•	•	•
2nd transport block	MCS						•	•	•	•	•
	RV						•	•	•	•	•
	NDI						•	•	•	•	•
Multi-antenna information	PMI confirmation			•							
	Precoding information			•		•	•	•			
	Transport block swap flag						•	•			
	Power offset					•					
	DM-RS scrambling								•		
	#Layers/DM-RS scrambling/ antenna ports									•	•
PDSCH mapping and quasi-colocation indicator											•
Downlink assignment index		•	•	•		•	•	•	•	•	•
PUCCH power control		•	•	•		•	•	•	•	•	•
SRS request[a]			F						T	T	T
ACK/NAK offset (EPDCCH only)			•	•		•	•	•	•	•	•
Flag for 0/1A differentiation			•								
Padding (only if needed)		(•)	(•)	(•)		(•)	(•)	(•)	(•)	(•)	(•)
Identity		•	•	•	•	•	•	•	•	•	

[a]*Format 1A for FDD and formats 2B, 2C, and 2D for TDD.*

DCI format 1A, also known as the "compact" downlink assignment, supports allocation of frequency-contiguous resource blocks only and can be used in all transmission modes. Contiguous allocations reduce the payload size of the control information with a somewhat reduced flexibility in resource allocations. The full range of modulation-and-coding schemes is supported. Format 1A can also be used to trigger a contention-free random access (see Chapter 11), in which case some bit fields are used to convey the necessary random-access preamble information and the remaining bit fields are set to a specific combination.

DCI format 1B is used to support codebook-based beam-forming described in Section 6.3.3, with a low control-signaling overhead (transmission mode 6). The content is similar to DCI format 1A with the addition of bits for signaling of the precoding matrix. As codebook-based beam-forming can be used to improve the data rates for cell-edge devices, it is important to keep the related DCI message size small so as not to unnecessarily limit the coverage.

DCI format 1C is used for various special purposes such as random-access response, paging, transmission of system information, MBMS-related signaling (see Chapter 19), and eIMTA support (Chapter 15). Common for these applications is simultaneous reception of a relatively small amount of information by *multiple* users. Hence, DCI format 1C supports QPSK only, has no support for hybrid-ARQ retransmissions, and does not support closed-loop spatial multiplexing. Consequently, the message size for DCI format 1C is very small, which is beneficial for coverage and efficient transmission of the type of system messages for which it is intended. Furthermore, as only a small number of resource blocks can be indicated, the size of the corresponding indication field in DCI format 1C is independent of the cell bandwidth.

DCI format 1D is used to support MU-MIMO (transmission mode 5) scheduling of one codeword with precoder information. To support dynamic sharing of the transmission power between the devices sharing the same resource block in MU-MIMO, one bit of power offset information is included in DCI format 1D, as described in Section 6.3.5.2.

DCI format 2 is an extension for DCI format 1 to support closed-loop spatial multiplexing (transmission mode 4). Thus, information about the number of transmission layers and the index of the precoder matrix used are jointly encoded in the precoding information field. Some of the fields in DCI format 1 have been duplicated to handle the two transport blocks transmitted in parallel in the case of spatial multiplexing.

DCI format 2A is similar to DCI format 2 except that it supports open-loop spatial multiplexing (transmission mode 3) instead of closed-loop spatial multiplexing. The precoder information field is used to signal the number of transmission layers only, hence the field has a smaller size than in DCI format 2. Furthermore, since DCI format 2A is used for scheduling of multi-layer transmissions only, the precoder information field is only necessary in the case of four transmit antenna ports; for two transmit antennas the number of layers is implicitly given by the number of transport blocks.

DCI format 2B was introduced in release 9 in order to support dual-layer spatial multiplexing in combination with beam-forming using DM-RS (transmission mode 8). Since scheduling with DCI format 2B relies on DM-RS, precoding/beam-forming is transparent to the device and there is no need to signal a precoder index. The number of layers can be

controlled by disabling one of the transport blocks. Two different scrambling sequences for the DM-RS can be used, as described in Section 6.2.2.

DCI format 2C was introduced in release 10 and is used to support spatial multiplexing using DM-RS (transmission mode 9). To some extent, it can be seen as a generalization of format 2B to support spatial multiplexing of up to eight layers. DM-RS scrambling and the number of layers are jointly signaled by a single three-bit field.

DCI format 2D was introduced in release 11 and is used to support spatial multiplexing using DM-RS (transmission mode 10). In essence it is an extension of format 2C to support signaling of quasi-colocation of antenna ports.

Many information fields in the different DCI formats are, as already mentioned, common among several of the formats, while some types of information exist only in certain formats. Furthermore, in later releases, some of the DCI formats are extended with additional bits. One example hereof is the addition of the carrier indicator in many DCI formats in release 10, as well as the inclusion of the SRS request in DCI formats 0 and 1A in release 10. Such extensions, for which different releases have different payload sizes for the same DCI format, are possible as long the extensions are used in the device-specific search spaces only where a single device is addressed by the DCI format in question. For DCI formats used to address multiple devices at the same time, for example, to broadcast system information, payload extensions are clearly not possible as devices from previous releases are not capable of decoding these extended formats. Consequently, the extensions are allowed in the device-specific search spaces only.

The information in the DCI formats used for downlink scheduling can be organized into different groups, as shown in Table 6.6, with the fields present varying between the DCI formats. A more detailed explanation of the contents of the different DCI formats is as follows:

- Resource information, consisting of:
  - Carrier indicator (0 or 3 bit). This field is present in releases 10 and beyond if cross-carrier scheduling is enabled via RRC signaling and is used to indicate the component carrier the downlink control information relates to (see Chapter 12). The carrier indicator is not present in the common search space as this would either impact compatibility with devices not capable of carrier aggregation or require additional blind decoding attempts.
  - Resource-block allocation. This field indicates the resource blocks on one component carrier upon which the device should receive the PDSCH. The size of the field depends on the cell bandwidth and on the DCI format, more specifically on the resource indication type, as discussed in Section 6.4.6.1. Resource allocation types 0 and 1, which are the same size, support noncontiguous resource-block allocations, while resource allocation type 2 has a smaller size but supports contiguous allocations only. DCI format 1C uses a restricted version of type 2 in order to further reduce control signaling overhead.

- Hybrid-ARQ process number (3 bit for FDD, 4 bit for TDD), informing the device about the hybrid-ARQ process to use for soft combining. Not present in DCI format 1C as this DCI format is intended for scheduling of system information which does not use hybrid ARQ retransmissions.
- For the first (or only) transport block:[38]
  - Modulation-and-coding scheme (5 bit), used to provide the device with information about the modulation scheme, the code rate, and the transport-block size, as described later. DCI format 1C has a restricted size of this field as only QPSK is supported.
  - New-data indicator (1 bit), used to clear the soft buffer for initial transmissions. Not present in DCI format 1C as this format does not support hybrid ARQ.
  - Redundancy version (2 bit).
- For the second transport block (only present in DCI format supporting spatial multiplexing):
  - Modulation-and-coding scheme (5 bit).
  - New-data indicator (1 bit).
  - Redundancy version (2 bit).
- Multi-antenna information. The different DCI formats are intended for different multi-antenna schemes and which of the fields below that are included depends on the DCI format shown in Table 6.5.
  - PMI confirmation (1 bit), present in format 1B only. Indicates whether the eNodeB uses the (frequency-selective) precoding matrix recommendation from the device or if the recommendation is overridden by the information in the PMI field.
  - Precoding information, providing information about the index of the precoding matrix used for the downlink transmission and, indirectly, about the number of transmission layers. This information is present in the DCI formats used for CRS-based transmission only; for DM-RS-based transmissions the precoder used is transparent for the device and consequently there is no need to signal this information in this case.
  - Transport block swap flag (1 bit), indicating whether the two codewords should be swapped prior to being fed to the hybrid-ARQ processes. Used for averaging the channel quality between the codewords.
  - Power offset between the PDSCH and CRS used to support dynamic power sharing between multiple devices for MU-MIMO.
  - Reference-signal scrambling sequence, used to control the generation of quasi-orthogonal DM-RS sequences, as discussed in Section 6.2.2.
  - Number of layers, reference-signal scrambling sequence and the set of antenna ports used for the transmission (jointly encoded information in releases 10 and later, 3 bits, possibility for extension to 4 bits in release 13).

[38] A transport block can be disabled by setting the modulation-and-coding scheme to zero and the RV to 1 in the DCI.

- PDSCH resource-element mapping and quasi-colocation indicator (2 bit), informing the device which set of parameters to assume when demodulating the PDSCH. Up to four different sets of parameters can be configured by RRC to support different CoMP schemes as discussed in Chapter 13.
- Downlink assignment index (2 bit), informing the device about the number of downlink transmissions for which a single hybrid-ARQ acknowledgment should be generated according to Section 8.1.3. Present for TDD only or for aggregation of more than 5 carriers in which case 4 bits are used.
- Transmit-power control for PUCCH (2 bit). For scheduling of a secondary component carrier in the case of carrier aggregation, these bits are reused as *acknowledgment resource indicator* (ARI)—see Chapter 12.
- SRS request (1 bit). This field is used to trigger a one-shot transmission of a sounding reference signal in the uplink, a feature introduced in release 10 and an example of extending an existing DCI format in the device-specific search space with additional information fields in a later release. For FDD it is present in format 1A only, while for TDD it is present in formats 2B, 2C, and 2D only. For TDD, where short-term channel reciprocity can be used, it is motivated to include this field in the DCI formats used to support DM-RS-based multi-layer transmission in order to estimate the downlink channel conditions based on uplink channel sounding. For FDD, on the other hand, this would not be useful and the SRS request field is consequently not included. However, for DCI format 1A, including an SRS request bit comes "for free" as padding otherwise would have to be used to ensure the same payload size as for DCI format 0.
- ACK/NAK offset (2 bit). This field is present on EPDCCH only and thus supported in release 11 and later. It is used to dynamically control the PUCCH resource used for the hybrid-ARQ acknowledgment as discussed in Section 7.4.2.1.
- DCI format 0/1A indication (1 bit), used to differentiate between DCI formats 1A and 0 as the two formats have the same message size. This field is present in DCI formats 0 and 1A only. DCI formats 3 and 3A, which have the same size, are separated from DCI formats 0 and 1A through the use of a different RNTI.
- Padding. The smaller of DCI formats 0 and 1A is padded to ensure the same payload size irrespective of the uplink and downlink cell bandwidths. Padding is also used to ensure that the DCI size is different for different DCI formats that may occur simultaneously in the same search space (this is rarely required in practice as the payload sizes are different due to the different amounts of information). Finally, for PDCCH, padding is used to avoid certain DCI sizes that may cause ambiguous decoding.[39]

[39]For a small set of specific payload sizes, the control signaling on PDCCH may be correctly decoded at an aggregation level other than the one used by the transmitter. Since the PHICH resource is derived from the first CCE used for the PDCCH, this may result in incorrect PHICH being monitored by the device. To overcome this, padding is used if necessary to avoid the problematic payload sizes. Note that this padding applied for PDCCH only.

- Identity (RNTI) of the device for which the PDSCH transmission is intended (16 bit). As described in Sections 6.4.3 and 6.4.4, the identity is not explicitly transmitted but implicitly included in the CRC calculation. There are different RNTIs defined depending on the type of transmission (unicast data transmission, paging, power-control commands, etc.).

6.4.6.1 Signaling of Downlink Resource-Block Allocations

Focusing on the signaling of resource-block allocations, there are three different possibilities, types 0, 1, and 2, as indicated in Table 6.6. Resource-block allocation types 0 and 1 both support noncontiguous allocations of resource blocks in the frequency domain, whereas type 2 supports contiguous allocations only. A natural question is why multiple ways of signaling the resource-block allocations are supported, a question whose answer lies in the number of bits required for the signaling. The most flexible way of indicating the set of resource blocks the device is supposed to receive the downlink transmission upon is to include a bitmap with size equal to the number of resource blocks in the cell bandwidth. This would allow for an arbitrary combination of resource blocks to be scheduled for transmission to the device but would, unfortunately, also result in a very large bitmap for the larger cell bandwidths. For example, in the case of a downlink cell bandwidth corresponding to 100 resource blocks, the downlink PDCCH would require 100 bits for the bitmap alone, to which the other pieces of information need to be added. Not only would this result in a large control-signaling overhead, but it could also result in downlink coverage problems as more than 100 bits in one OFDM symbol correspond to a data rate exceeding 1.4 Mbit/s. Consequently, there is a need for a resource allocation scheme requiring a smaller number of bits while keeping sufficient allocation flexibility.

In resource allocation type 0, the size of the bitmap has been reduced by pointing not to individual resource blocks in the frequency domain, but to groups of contiguous resource blocks, as shown at the top of Figure 6.33. The size of such a group is determined by the downlink cell bandwidth; for the smallest bandwidths there is only a single resource block in a group, implying that an arbitrary set of resource blocks can be scheduled, whereas for the largest cell bandwidths, groups of four resource blocks are used (in the example in Figure 6.33, the cell bandwidth is 25 resource blocks, implying a group size of two resource blocks). Thus, the bitmap for the system with a downlink cell bandwidth of 100 resource blocks is reduced from 100 to 25 bits. A drawback is that the scheduling granularity is reduced; single resource blocks cannot be scheduled for the largest cell bandwidths using allocation type 0.

However, also in large cell bandwidths, frequency resolution of a single resource block is sometimes useful, for example, to support small payloads. Resource allocation type 1 address this by dividing the total number of resource blocks in the frequency domain into dispersed subsets, as shown in the middle of Figure 6.33. The number of subsets is given from the cell bandwidth with the number of subsets in type 1 being equal to the group size in type 0. Thus, in Figure 6.33, there are two subsets, whereas for a cell bandwidth of 100 resource blocks

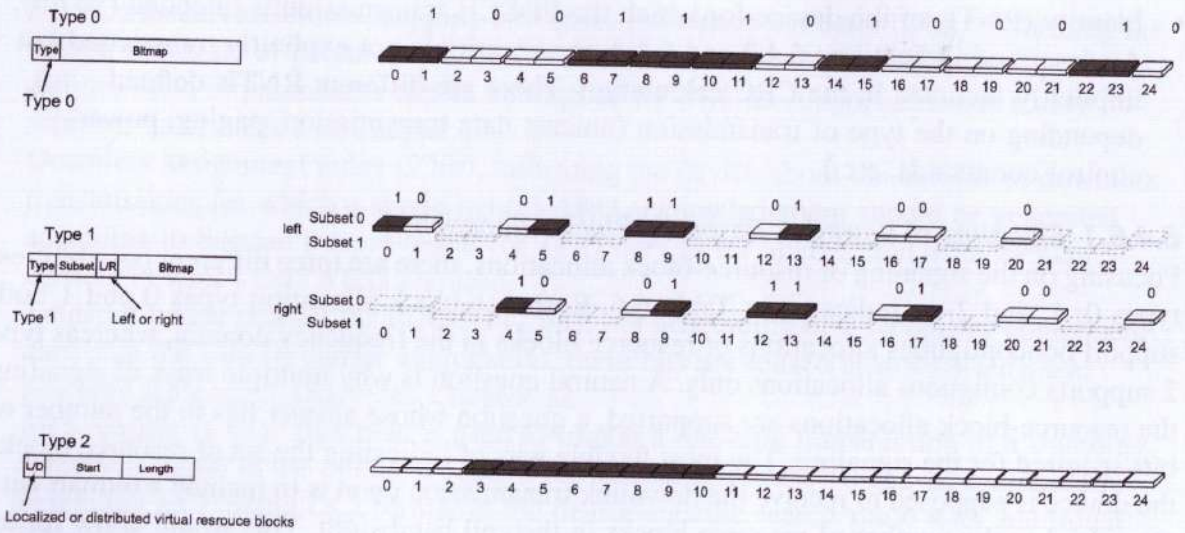


FIGURE 6.33

Illustration of resource-block allocation types (cell bandwidth corresponding to 25 resource blocks used in this example).

there would have been four different subsets. Within a subset, a bitmap indicates the resource blocks in the frequency domain upon which the downlink transmission occurs.

To inform the device whether resource allocation type 0 or 1 is used, the resource allocation field includes a flag for this purpose, denoted "type" in the leftmost part of Figure 6.33. For type 0, the only additional information is the bitmap discussed previously. For type 1, on the other hand, in addition to the bitmap itself, information about the subset for which the bitmap relates is also required. As one of the requirements in the design of resource allocation type 1 was to maintain the same number of bits in the allocation as for type 0 without adding unnecessary overhead,[40] the bitmap in resource allocation type 1 is smaller than in type 0 to allow for the signaling of the subset number. However, a consequence of a smaller bitmap is that not all resource blocks in the subset can be addressed simultaneously. To be able to address all resources with the bitmap, there is a flag indicating whether the bitmap relates to the "left" or "right" part of the resource blocks, as depicted in the middle part of Figure 6.33.

Unlike the other two types of resource-block allocation signaling, type 2 does not rely on a bitmap. Instead, it encodes the resource allocation as a start position and length of the resource-block allocation. Thus, it does not support arbitrary allocations of resource blocks but only frequency-contiguous allocations, thereby reducing the number of bits required for signaling the resource-block allocation. The number of bits required for resource-signaling type 2 compared to type 0 or 1 is shown in Figure 6.34 and, as shown, the difference is fairly large for the larger cell bandwidths.

[40]Allowing different sizes would result in an increase in the number of blind decoding attempts required in the device.

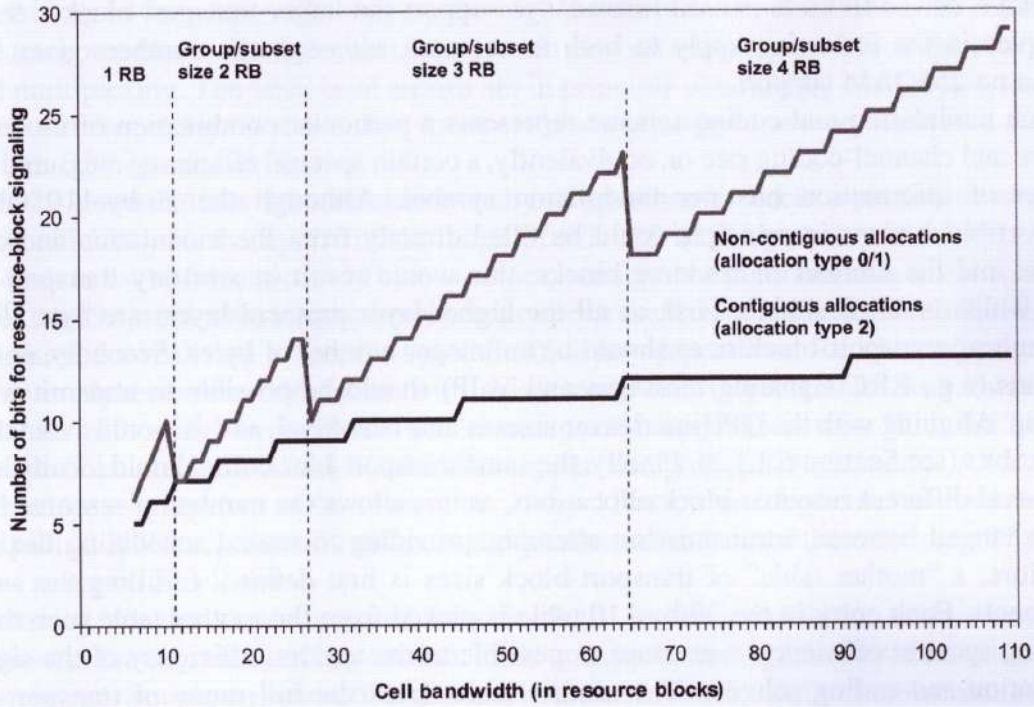


FIGURE 6.34

Number of bits used for downlink resource allocation signaling for downlink allocation types 0/1 and 2.

All three resource-allocation types refer to *VRBs* (see Section 6.1.1.8 for a discussion of resource-block types). For resource-allocation types 0 and 1, the VRBs are of localized type and the VRBs are directly mapped to PRBs. For resource-allocation type 2, on the other hand, both localized and distributed VRBs are supported. One bit in the resource allocation field indicates whether the allocation signaling refers to localized or distributed resource blocks.

6.4.6.2 Signaling of Transport-Block Sizes

Proper reception of a downlink transmission requires, in addition to the set of resource blocks, knowledge about the modulation scheme and the transport-block size, information (indirectly) provided by a 5-bit field in the different DCI formats. Of the 32 combinations, 29 are used to signal the modulation-and-coding scheme whereas 3 are reserved, the purpose of which is described later. Together, the modulation-and-coding scheme and the number of resource blocks assigned provide the transport-block size on the DL-SCH. Thus, the possible transport-block sizes can be described as a table with 29 rows and 110 columns, one column for each number of resource blocks possible to transmit upon (the number of columns follows from the maximum downlink component carrier bandwidth of 110 resource blocks). For devices configured with support for 256QAM, 4 of the 32 combinations are reserved and an

alternative 28×110 table is used instead[41] to support the larger transport block sizes. The principles in the following apply to both these cases, although the numbers given below assume no 256QAM support.

Each modulation-and-coding scheme represents a particular combination of modulation scheme and channel-coding rate or, equivalently, a certain spectral efficiency measured in the number of information bits per modulation symbol. Although the 29-by-110 table of transport-block sizes in principle could be filled directly from the modulation-and-coding scheme and the number of resource blocks, this would result in arbitrary transport-block sizes, which is not desirable. First, as all the higher-layer protocol layers are byte aligned, the resulting transport-block sizes should be an integer number of bytes. Secondly, common payloads (e.g., RRC signaling messages and VoIP) should be possible to transmit without padding. Aligning with the QPP interleaver sizes is also beneficial, as this would avoid the use of filler bits (see Section 6.1.1.3). Finally, the same transport-block size should ideally appear for several different resource-block allocations, as this allows the number of resource blocks to be changed between retransmission attempts, providing increased scheduling flexibility. Therefore, a "mother table" of transport-block sizes is first defined, fulfilling the said requirements. Each entry in the 29-by-110 table is picked from the mother table such that the resulting spectral efficiency is as close as possible to the spectral efficiency of the signaled modulation-and-coding scheme. The mother table spans the full range of transport-block sizes possible, with an approximately constant worst-case padding.

From a simplicity perspective, it is desirable if the transport-block sizes do not vary with the configuration of the system. The set of transport-block sizes is therefore independent of the actual number of antenna ports and the size of the control region.[42] The design of the table assumes a control region of three OFDM symbols and two antenna ports, the "reference configuration." If the actual configuration is different, the resulting code rate for the DL-SCH will be slightly different as a result of the rate-matching procedure. However, the difference is small and of no practical concern. Also, if the actual size of the control region is smaller than the three-symbol assumption in the reference configuration, the spectral efficiencies will be somewhat smaller than the range indicated by the modulation-and-coding scheme signaled as part of the DCI. Thus, information about the modulation scheme used is obtained directly from the modulation-and-coding scheme, whereas the exact code rate and rate matching is obtained from the implicitly signaled transport-block size together with the number of resource elements used for DL-SCH transmission.

For bandwidths smaller than the maximum of 110 resource blocks, a subset of the table is used. More specifically, in case of a cell bandwidth of N resource blocks, the first N columns of the table are used. Also, in the case of spatial multiplexing, a single transport block can be

[41] The alternative table is only used for transmissions scheduled with the C-RNTI. Random-access response and system information cannot use 256QAM for backward-compatibility reasons.

[42] For DwPTS, the transport-block size is scaled by a factor of 0.75 compared to the values found in the table, motivated by the DwPTS having a shorter duration than a normal subframe.

mapped to up to four layers. To support the higher data rates this facilitates, the set of supported transport-block sizes needs to be extended beyond what is possible in the absence of spatial multiplexing. The additional entries are in principle obtained by multiplying the sizes with the number of layers to which a transport block is mapped and adjusting the result to match the QPP interleaver size.

The 29 combinations of modulation-and-coding schemes each represent a reference spectral efficiency in the approximate range of 0.1−5.9 bits/s/symbol (with 256QAM the upper limit is 7.4 bit/s/symbol).[43] There is some overlap in the combinations in the sense that some of the 29 combinations represent the same spectral efficiency. The reason is that the best combination for realizing a specific spectral efficiency depends on the channel properties; sometimes higher-order modulation with a low code rate is preferable over lower-order modulation with a higher code rate, and sometimes the opposite is true. With the overlap, the eNodeB can select the best combination, given the propagation scenario. As a consequence of the overlap, two of the rows in the 29-by-110 table are duplicates and result in the same spectral efficiency but with different modulation schemes, and there are only 27 unique rows of transport-block sizes.

Returning to the three reserved combinations in the modulation-and-coding field mentioned at the beginning of this section, those entries can be used for retransmissions only. In the case of a retransmission, the transport-block size is, by definition, unchanged and fundamentally there is no need to signal this piece of information. Instead, the three reserved values represent the modulation scheme, QPSK, or 16QAM or 64QAM,[44] which allows the scheduler to use an (almost) arbitrary combination of resource blocks for the retransmission. Obviously, using any of the three reserved combinations assumes that the device properly received the control signaling for the initial transmission; if this is not the case, the retransmission should explicitly indicate the transport-block size.

The derivation of the transport-block size from the modulation-and-coding scheme and the number of scheduled resource blocks is illustrated in Figure 6.35.

6.4.7 UPLINK SCHEDULING GRANTS

Uplink scheduling grants use one of DCI formats 0 or 4; DCI format 4 was added in release 10 to support uplink spatial multiplexing. The basic resource-allocation scheme for the uplink is single-cluster allocations where the resource blocks are contiguous in the frequency domain, although release 10 added support for multi-cluster transmissions of up to two clusters on a single component carrier.

DCI format 0 is used for scheduling uplink transmissions not using spatial multiplexing on one component carrier. It has the same size control-signaling message as the "compact" downlink assignment (DCI format 1A). A flag in the message is used to inform the device

[43]The exact values vary slightly with the number of resource blocks allocated due to rounding.
[44]Four reserved values representing QPSK, 16QAM, 64QAM, and 256QAM for the alternative table.

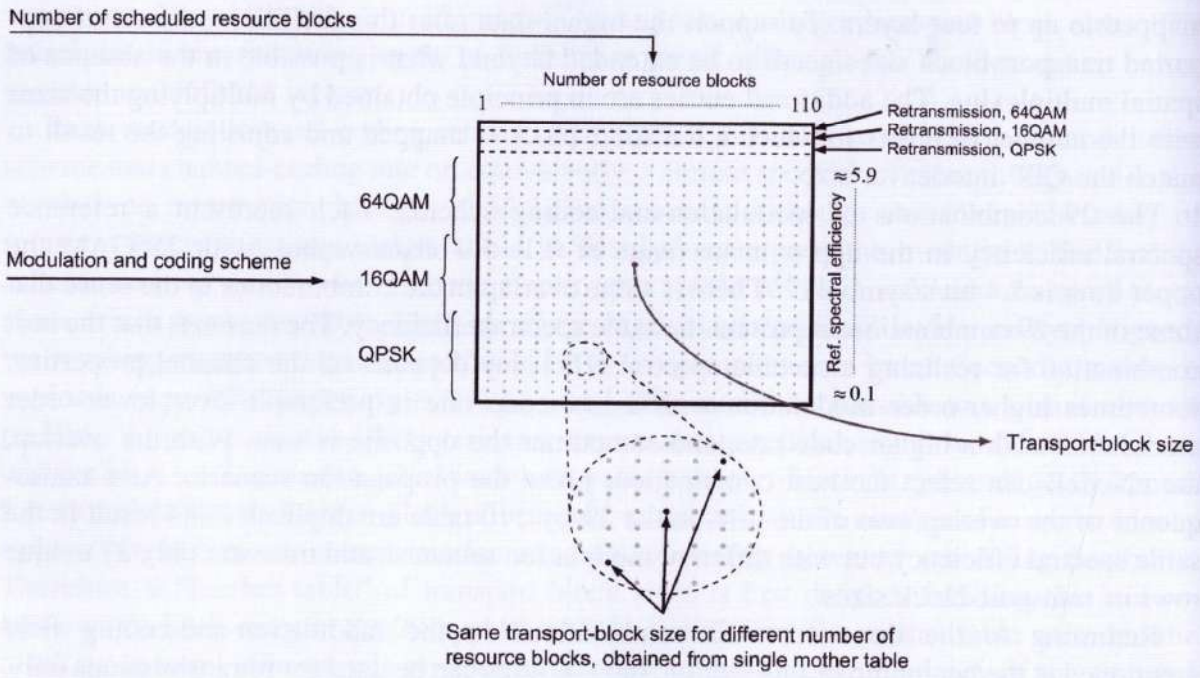


FIGURE 6.35

Computing the transport-block size (no 256QAM configured).

whether the message is an uplink scheduling grant (DCI format 0) or a downlink scheduling assignment (DCI format 1A).

DCI format 4 is used for uplink transmissions using spatial multiplexing on one component carrier. Consequently, the size of DCI format 4 is larger than that of DCI format 0, as additional information fields are required.

Many information fields are common to the two DCI formats, but there are also differences, as shown in Table 6.7. Similarly to the downlink scheduling assignments, some DCI formats have been extended with additional bits in later releases. The contents of the different DCI formats are explained in more detail in the following:

- Resource information, consisting of:
  - Carrier indicator (0 or 3 bit). This field is present in releases 10 and beyond only if cross-carrier scheduling is enabled via RRC signaling and is used to indicate the uplink component carrier the grant relates to (see Chapter 12). The carrier indicator is not present in the common search space as this would either impact compatibility with devices not capable of carrier aggregation or require additional blind decoding attempts.
  - Resource allocation type or multi-cluster flag (1 bit), indicating whether resource allocation type 0 (one cluster of resource blocks is used for uplink transmissions) or type 1 (two clusters of resource blocks are used for the uplink transmission) is used.

Table 6.7 DCI Formats for Uplink Scheduling Grants

Field		DCI Format 0	DCI Format 4
Resource information	Carrier indicator	•	•
	Resource allocation type	•	•
	Resource block assignment	0/(1)	0/1
1st transport block	MCS/RV	•	•
	NDI	•	•
2nd transport block	MCS/RV		•
	NDI		•
DM-RS phase rotation and OCC index		•	•
Precoding information			•
CSI request		•	•
SRS request		•	•
Uplink index/DAI (TDD only)		•	•
PUSCH power control		•	•
Flag for 0/1A differentiation		•	
Padding (only if needed)		(•)	(•)
Identity		•	•

This flag is not present in releases prior to release 10. In previous releases, the downlink bandwidth was, in practice, always at least as large as the uplink bandwidth, implying that one padding bit was used for DCI format 0 in those releases to align with the size of format 1A. The padding bit could therefore be replaced by the multi-cluster flag in release 10 without sacrificing backward compatibility. In DCI format 4, supported in releases 10 and later, the multi-cluster flag is always present.

- Resource-block allocation, including hopping indication. This field indicates the resource blocks upon which the device should transmit the PUSCH using uplink resource-allocation type 0 (DCI format 0) or type 1 (DCI format 4), as described in Section 6.4.7.1. The size of the field depends on the cell bandwidth. For single-cluster allocations, uplink frequency hopping, as described in Chapter 7, can be applied to the uplink PUSCH transmission.

- For the first (or only) transport block:
  - Modulation-and-coding scheme including redundancy version (5 bit), used to provide the device with information about the modulation scheme, the code rate, and the transport-block size. The signaling of the transport-block size uses the same transport-block table as for the downlink—that is, the modulation-and-coding scheme together with the number of scheduled resource blocks provides the transport-block size. However, as the support of 64QAM in the uplink is not mandatory for all devices,

devices not capable of 64QAM use 16QAM when 64QAM is indicated in the modulation-and-coding field. The use of the three reserved combinations is slightly different than for the downlink; the three reserved values are used for implicit signaling of the RV, as described later. A transport block can be disabled by signaling a specific combination of modulation-and-coding scheme and number of resource blocks. Disabling one transport block is used when retransmitting a single transport block only.
    - New-data indicator (1 bit), used to indicate to the device whether transmission of a new transport block or retransmission of the previous transport block is granted.
- For the second transport block (DCI format 4 only):
    - Modulation-and-coding scheme including redundancy version (5 bit).
    - New-data indicator (1 bit).
- Phase rotation of the uplink demodulation reference signal (3 bit), used to support MU-MIMO, as described in Chapter 7. By assigning different reference-signal phase rotations to devices scheduled on the same time–frequency resources, the eNodeB can estimate the uplink channel response from each device and suppress the interdevice interference by the appropriate processing. In releases 10 and later, it also controls the orthogonal cover sequence, see Section 7.2.
- Precoding information, used to signal the precoder to use for the uplink transmission in releases 10 and later.
- Channel-state request flag (1, 2, or 3 bit). The network can explicitly request an aperiodic channel-state report to be transmitted on the UL-SCH by setting this bit(s) in the uplink grant. In the case of carrier aggregation of up to five carriers, 2 bits are used to indicate which downlink component carrier the CSI should be reported for (see Chapter 10), a number increased to 3 bits if more than five carriers are configured in the device.
- SRS request (2 bit), used to trigger aperiodic sounding using one of up to three preconfigured settings, as discussed in Chapter 7. The SRS request, introduced in release 10, is supported in the device-specific search space only for reasons already described.
- Uplink index/DAI (2, 4 bit for carrier aggregation of more than five carriers). This field is present only when operating in TDD or when aggregating more than five carriers. For uplink–downlink configuration 0 (uplink-heavy configuration), it is used as an uplink index to signal for which uplink subframe(s) the grant is valid, as described in Chapter 9. For other uplink–downlink configurations, it is used as downlink assignment index to indicate the number of downlink transmissions the eNodeB expects hybrid-ARQ acknowledgment for.
- Transmit-power control for PUSCH (2 bit).
- DCI format 0/1A indication (1 bit), used to differentiate between DCI formats 1A and 0 as the two formats have the same message size. This field is present in DCI formats 0 and 1A only.
- Padding; the smaller of DCI formats 0 and 1A is padded to ensure the same payload size irrespective of the uplink and downlink cell bandwidths. Padding is also used to ensure

that the DCI size is different for DCI formats 0 and 4 (this is rarely required in practice as the payload sizes are different due to the different amounts of information). Finally, for PDCCH, padding is used to avoid certain DCI sizes that may cause ambiguous decoding.

- Identity (RNTI) of the device for which the grant is intended (16 bit). As described in Sections 6.4.3 and 6.4.4, the identity is not explicitly transmitted but implicitly included in the CRC calculation.

There is no explicit signaling of the redundancy version in the uplink scheduling grants. This is motivated by the use of a synchronous hybrid-ARQ protocol in the uplink; retransmissions are normally triggered by a negative acknowledgment on the PHICH and not explicitly scheduled as for downlink data transmissions. However, as described in Chapter 8, there is a possibility to explicitly schedule retransmissions. This is useful in a situation where the network will explicitly move the retransmission in the frequency domain by using the PDCCH instead of the PHICH. Three values of the modulation-and-coding field are reserved to mean redundancy version 1, 2, and 3. If one of those values is signaled, the device should assume that the same modulation and coding as the original transmission is used. The remaining entries are used to signal the modulation-and-coding scheme to use and also imply that redundancy version zero should be used. The difference in usage of the reserved values compared to the downlink scheduling assignments implies that the modulation scheme, unlike the downlink case, cannot change between uplink (re)transmission attempts.

6.4.7.1 Signaling of Uplink Resource-Block Allocations

The basic uplink resource-allocation scheme is single-cluster allocations—that is, allocations contiguous in the frequency domain—but releases 10 and later also provide the possibility for multi-cluster uplink transmissions.

Single-cluster allocations use uplink resource-allocation type 0, which is identical to downlink resource allocation type 2 described in Section 6.4.6.1 except that the single-bit flag indicating localized/distributed transmission is replaced by a single-bit frequency hopping flag. The resource allocation field in the DCI provides the set of VRBs to use for uplink transmission. The set of PRBs to use in the two slots of a subframe is controlled by the hopping flag, as described in Chapter 7.

Multi-cluster allocations with up to two clusters were introduced in release 10, using uplink resource-allocation type 1. In resource-allocation type 1, the starting and ending positions of two clusters of frequency-contiguous resource blocks are encoded into an index. Uplink resource-allocation type 1 does not support frequency hopping (diversity is achieved through the use of two clusters instead). Indicating two clusters of resources naturally requires additional bits compared to the single-cluster case. At the same time, the total number of bits used for resource-allocation type 1 should be identical to that of type 0. This is similar to the situation for the downlink allocation types 0 and 1; without aligning the sizes a new DCI format with a corresponding negative impact on the number of blind decodings is necessary. Since frequency hopping is not supported for allocation type 1, the bit otherwise

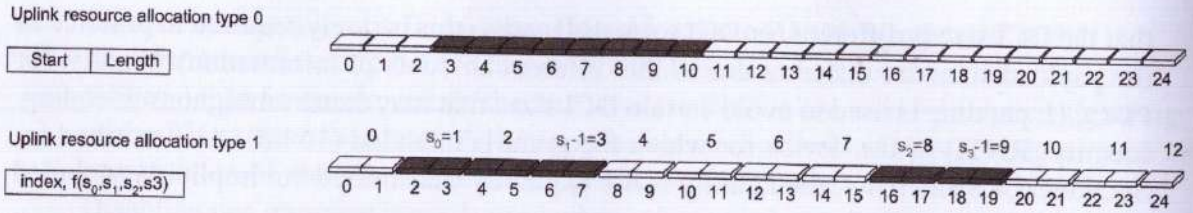


FIGURE 6.36

Illustration of uplink resource-block allocation types (uplink bandwidth corresponding to 25 resource blocks used in this example).

used for the hopping flag can be reused for extending the resource-allocation field. However, despite the extension of the resource-allocation field by one bit, the number of bits is not sufficient to provide a single-resource-block resolution in the two clusters for all bandwidths. Instead, similar to downlink resource-allocation type 0, groups of resource blocks are used and the starting and ending positions of the two clusters are given in terms of group numbers. The size of such a group is determined by the uplink carrier bandwidth in a similar way as for the downlink. For the smallest bandwidths there is only a single resource block in a group, implying that an arbitrary (as long as the limit of at most two clusters is observed) set of resource blocks can be scheduled, whereas for the largest cell bandwidths groups of four resource blocks are used. In the example in Figure 6.36, the cell bandwidth is 25 resource blocks, implying a group size of two resource blocks.

6.4.8 POWER-CONTROL COMMANDS

As a complement to the power-control commands provided as part of the downlink scheduling assignments and the uplink scheduling grants, there is the potential to transmit a power-control command using DCI formats 3 (2-bit command per device) or 3A (single-bit command per device). The main motivation for DCI format 3/3A is to support power control for semi-persistent scheduling. The power-control message is directed to a group of devices using an RNTI specific for that group. Each device can be allocated two power-control RNTIs, one for PUCCH power control and the other for PUSCH power control. Although the power-control RNTIs are common to a group of devices, each device is informed through RRC signaling which bit(s) in the DCI message it should follow. No carrier indicator is used for formats 3 and 3A.

CHAPTER

8 RETRANSMISSION PROTOCOLS

Transmissions over wireless channels are subject to errors, for example due to variations in the received signal quality. To some degree, such variations can be counteracted through link adaptation as is discussed in the next chapter. However, receiver noise and unpredictable interference variations cannot be counteracted. Therefore, virtually all wireless communications systems employ some form of *forward error correction* (FEC), adding redundancy to the transmitted signal allowing the receiver to correct errors, based on the concept whose roots can be traced to the pioneering work of Claude Shannon in 1948 [7]. In LTE, Turbo coding is used as discussed in Chapter 6.

Despite the error-correcting code, there will be data units received in error. *Hybrid automatic repeat request* (ARQ), first proposed in [18] and relying on a combination of error-correcting coding and retransmission of erroneous data units, is therefore commonly used in many modern communication systems. Data units in error despite the error-correcting coding are detected by the receiver, which requests a retransmission from the transmitter.

In LTE, there are *two* mechanisms responsible for retransmission handling, namely the MAC and RLC sublayers. Retransmissions of missing or erroneous data units are handled primarily by the hybrid-ARQ mechanism in the MAC layer, complemented by the retransmission functionality of the RLC protocol. The reasons for having a two-level retransmission structure can be found in the trade-off between fast and reliable feedback of the status reports. The hybrid-ARQ mechanism targets very fast retransmissions and, consequently, feedback on success or failure of the decoding attempt is provided to the transmitter after each received transport block. Although it is in principle possible to attain a very low error probability of the hybrid-ARQ feedback, it comes at a cost in transmission power. Keeping the cost reasonable typically results in a feedback error rate of around 1%, which results in a hybrid-ARQ residual error rate of a similar order. Such an error rate is in many cases far too high; high data rates with TCP may require virtually error-free delivery of packets to the TCP protocol layer. As an example, for sustainable data rates exceeding 100 Mbit/s, a packet-loss probability less than 10^{-5} is required [35]. The reason is that TCP assumes packet errors to be due to congestion in the network. Any packet error therefore triggers the TCP congestion-avoidance mechanism with a corresponding decrease in data rate.

Compared to the hybrid-ARQ acknowledgments, the RLC status reports are transmitted relatively infrequently and thus the cost of obtaining a reliability of 10^{-5} or better is relatively

4G, LTE-Advanced Pro and The Road to 5G. http://dx.doi.org/10.1016/B978-0-12-804575-6.00008-X
Copyright © 2016 Erik Dahlman, Stefan Parkvall and Johan Sköld. Published by Elsevier Ltd. All rights reserved.

small. Hence, the combination of hybrid-ARQ and RLC attains a good combination of small round-trip time and a modest feedback overhead where the two components complement each other—fast retransmissions due to the hybrid-ARQ mechanism and reliable packet delivery due to the RLC. As the MAC and RLC protocol layers are located in the same network node, a tight interaction between the two protocols is possible. Hence, to some extent, the combination of the two can be viewed as *one* retransmission mechanism with *two* feedback channels. However, note that, as discussed in Chapter 4 and illustrated in Figure 8.1, the RLC operates per logical channel, while the hybrid-ARQ operates per transport channel (that is, per component carrier). One hybrid-ARQ entity may therefore retransmit data belonging to multiple logical channels.

In the following section, the principles behind the hybrid-ARQ and RLC protocols are discussed in more detail.

8.1 HYBRID ARQ WITH SOFT COMBINING

The hybrid-ARQ operation described in the preceding section discards erroneously received packets and requests retransmission. However, despite it was not possible to decode the packet, the received signal still contains information, which is lost by discarding erroneously received packets. This shortcoming is addressed by *hybrid ARQ with soft combining*. In hybrid ARQ with soft combining, the erroneously received packet is stored in a buffer

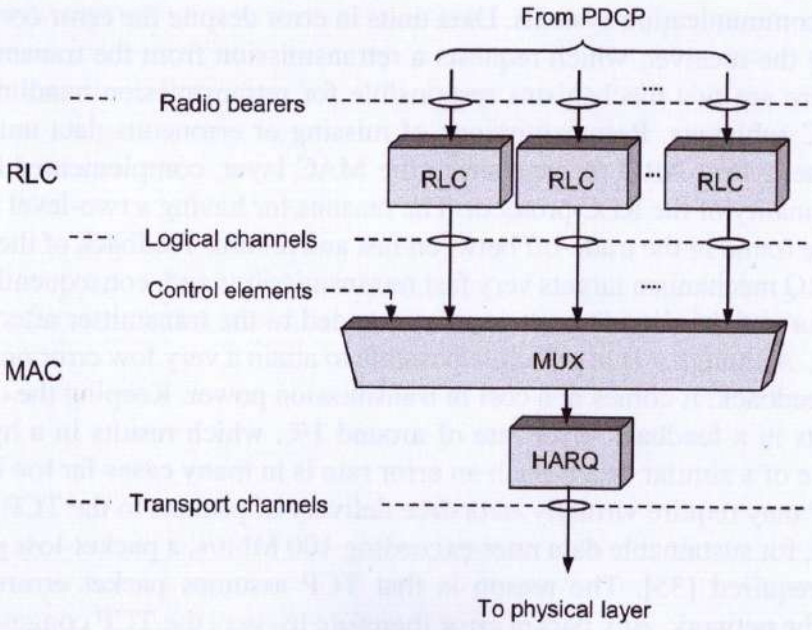


FIGURE 8.1

RLC and hybrid-ARQ retransmission mechanisms in LTE.

memory and later combined with the retransmission to obtain a single, combined packet that is more reliable than its constituents. Decoding of the error-correction code operates on the combined signal.

The hybrid-ARQ functionality spans both the physical layer and the MAC layer; generation of different redundancy versions at the transmitter as well as the soft combining at the receiver are handled by the physical layer, while the hybrid-ARQ protocol is part of the MAC layer.

The basis for the LTE hybrid-ARQ mechanism is a structure with multiple stop-and-wait protocols, each operating on a single transport block. In a stop-and-wait protocol, the transmitter stops and waits for an acknowledgment after each transmitted transport block. This is a simple scheme; the only feedback required is a single bit indicating positive or negative acknowledgment of the transport block. However, since the transmitter stops after each transmission, the throughput is also low. LTE therefore applies *multiple* stop-and-wait processes operating in parallel such that, while waiting for acknowledgment from one process, the transmitter can transmit data to another hybrid-ARQ process. This is illustrated in Figure 8.3; while processing the data received in the first hybrid-ARQ process the receiver can continue to receive using the second process and so on. This structure, multiple hybrid-ARQ processes operating in parallel to form one hybrid-ARQ entity, combines the simplicity of a stop-and-wait protocol while still allowing continuous transmission of data.

There is one hybrid-ARQ entity per device. Spatial multiplexing, where two transport blocks can be transmitted in parallel on the same transport channel as described in Chapter 6, is supported by one hybrid-ARQ entity having two sets of hybrid-ARQ processes with independent hybrid-ARQ acknowledgments. The details for the physical-layer transmission of the downlink and uplink hybrid-ARQ acknowledgments are described in Chapters 6 and 7.

Upon receiving a transport block for a certain hybrid-ARQ process, the receiver makes an attempt to decode the transport block and informs the transmitter about the outcome through a hybrid-ARQ acknowledgment, indicating whether the transport block was correctly decoded or not. The time from reception of data until transmission of the hybrid-ARQ acknowledgment is fixed, hence the transmitter knows from the timing relation which hybrid-ARQ process a received acknowledgment relates to. This is beneficial from an overhead perspective as there is no need to signal the process number along with the acknowledgment.

An important part of the hybrid-ARQ mechanism is the use of *soft combining*, which implies that the receiver combines the received signal from multiple transmission attempts.

By definition, retransmission in any hybrid-ARQ scheme must represent the same set of information bits as the original transmission. However, the set of coded bits transmitted in each retransmission may be selected differently as long as they represent the same set of information bits. Depending on whether the retransmitted bits are required to be identical to

the original transmission or not, the soft combining scheme is often referred to as *chase combining*, first proposed in [19], or *incremental redundancy* (IR), which is used in LTE. With IR, each retransmission does not have to be identical to the original transmission. Instead, *multiple sets* of coded bits are generated, each representing the same set of information bits [20,21]. The rate matching functionality of LTE, described in Chapter 6, is used to generate different sets of coded bits as a function of the redundancy version as illustrated in Figure 8.2. In addition to a gain in accumulated received E_b/N_0, IR also results in a coding gain for each retransmission. The gain with IR compared to pure energy accumulation (chase combining) is larger for high initial code rates [22]. Furthermore, as shown in [23], the performance gain of IR compared to chase combining can also depend on the relative power difference between the transmission attempts.

In the discussion so far, it has been assumed that the receiver has received all the previously transmitted redundancy versions. If all redundancy versions provide the same amount of information about the data packet, the order of the redundancy versions is not critical. However, for some code structures, not all redundancy versions are of equal importance. One example here is Turbo codes, where the systematic bits are of higher importance than the parity bits. Hence, the initial transmission should at least include the systematic bits and some parity bits. In the retransmission(s), parity bits not in the initial transmission can be included.

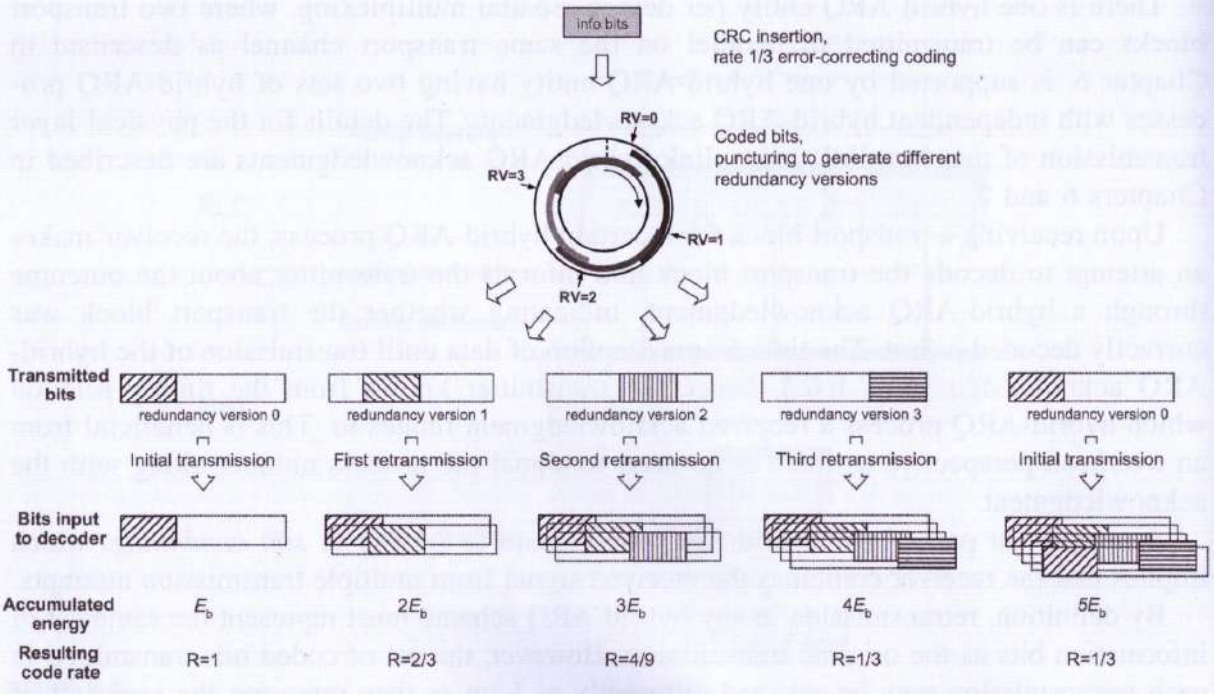


FIGURE 8.2

Example of incremental redundancy.

This is the background to why systematic bits are inserted first in the circular buffer discussed in Chapter 6.

Hybrid ARQ with soft combining, regardless of whether chase combining or IR is used, leads to an implicit reduction of the data rate by means of retransmissions and can thus be seen as implicit link adaptation. However, in contrast to link adaptation based on explicit estimates of the instantaneous channel conditions, hybrid ARQ with soft combining implicitly adjusts the coding rate based on the result of the decoding. In terms of overall throughput this kind of implicit link adaptation can be superior to explicit link adaptation, as additional redundancy is only added *when needed*—that is, when previous higher-rate transmissions were not possible to decode correctly. Furthermore, as it does not try to predict any channel variations, it works equally well, regardless of the speed at which the terminal is moving. Since implicit link adaptation can provide a gain in system throughput, a valid question is why explicit link adaptation is necessary at all. One major reason for having explicit link adaptation is the reduced delay. Although relying on implicit link adaptation alone is sufficient from a system throughput perspective, the end-user service quality may not be acceptable from a delay perspective.

For proper operation of soft combining, the receiver needs to know when to perform soft combining prior to decoding and when to clear the soft buffer—that is, the receiver needs to differentiate between the reception of an initial transmission (prior to which the soft buffer should be cleared) and the reception of a retransmission. Similarly, the transmitter must know whether to retransmit erroneously received data or to transmit new data. Therefore, an explicit *new-data indicator* is included for each of the one or two scheduled transport blocks as part of the scheduling information transmitted in the downlink. The new-data indicator is present in both downlink assignments and uplink grants, although the meaning is slightly different for the two.

For downlink data transmission, the new-data indicator is toggled for a new transport block—that is, it is essentially a single-bit sequence number. Upon reception of a downlink scheduling assignment, the device checks the new-data indicator to determine whether the current transmission should be soft combined with the received data currently in the soft buffer for the hybrid-ARQ process in question, or if the soft buffer should be cleared.

For uplink data transmission, there is also a new-data indicator transmitted on the downlink PDCCH. In this case, toggling the new-data indicator requests transmission of a new transport block, otherwise the previous transport block for this hybrid-ARQ process should be retransmitted (in which case the eNodeB should perform soft combining).

The use of multiple hybrid-ARQ processes operating in parallel can result in data being delivered from the hybrid-ARQ mechanism out of sequence. For example, transport block 5 in Figure 8.3 was successfully decoded before transport block 1, which required two retransmissions. Out-of-sequence delivery can also occur in the case of carrier aggregation, where transmission of a transport block on one component carrier could be successful while a retransmission is required on another component carrier. To handle out-of-sequence delivery

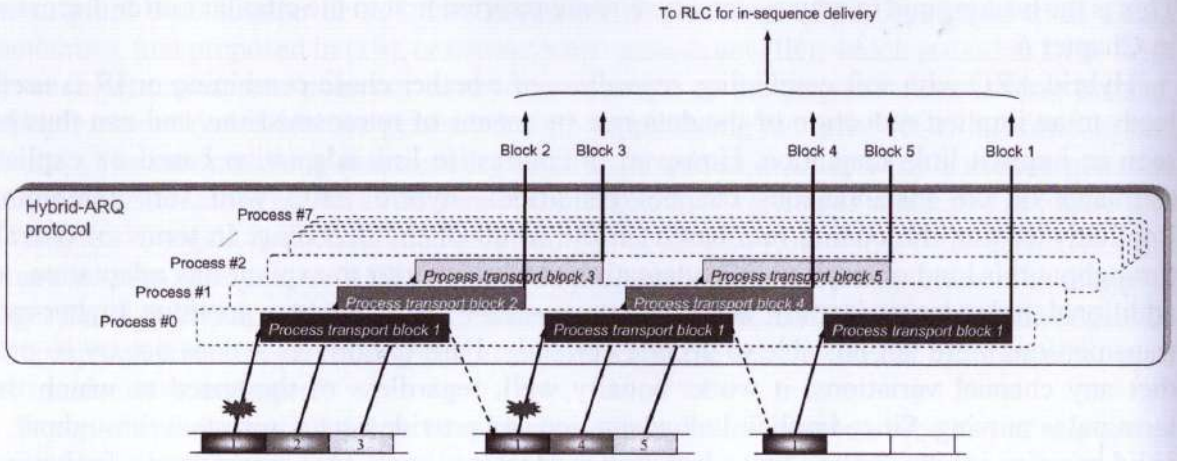


FIGURE 8.3

Multiple parallel hybrid-ARQ processes forming one hybrid-ARQ entity.

from the hybrid-ARQ protocol, the RLC protocol includes an in-sequence-delivery mechanism, as described in Section 8.2.

Hybrid-ARQ protocols can be characterized as synchronous versus asynchronous, related to the flexibility in the time domain, as well as adaptive versus nonadaptive, related to the flexibility in the frequency domain:

- An *asynchronous* hybrid-ARQ protocol implies that retransmissions can occur at any time, whereas a *synchronous* protocol implies that retransmissions occur at a fixed time after the previous transmission (see Figure 8.4). The benefit of a synchronous protocol is

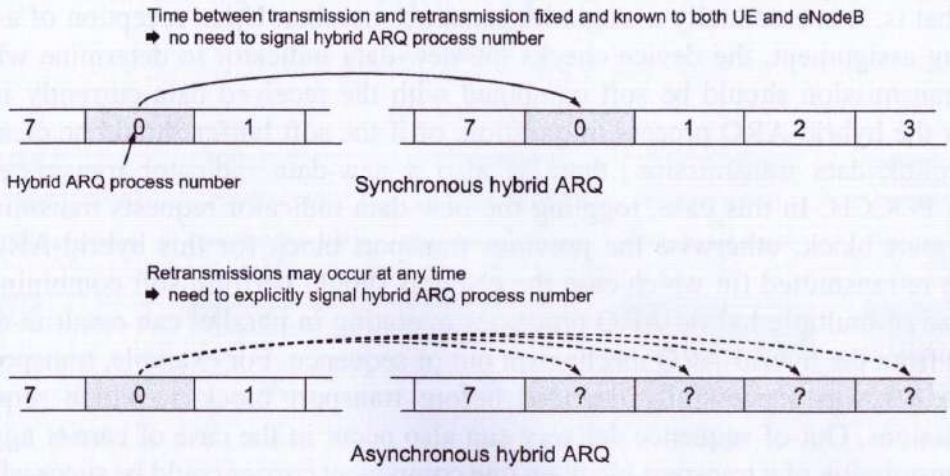


FIGURE 8.4

Synchronous and asynchronous hybrid ARQ.

that there is no need to explicitly signal the hybrid-ARQ process number as this information can be derived from the subframe number. On the other hand, an asynchronous protocol allows for more flexibility in the scheduling of retransmissions.

- An *adaptive* hybrid-ARQ protocol implies that the frequency location and possibly also the more detailed transmission format can be changed between retransmissions. A *nonadaptive protocol*, in contrast, implies that the retransmission must occur at the same frequency resources and with the same transmission format as the initial transmission.

In the case of LTE, asynchronous adaptive hybrid ARQ is used for the downlink. For the uplink, synchronous hybrid ARQ is used. Typically, the retransmissions are nonadaptive, but there is also the possibility to use adaptive retransmissions as a complement.

8.1.1 DOWNLINK HYBRID ARQ

In the downlink, retransmissions are scheduled in the same way as new data—that is, they may occur at any time and at an arbitrary frequency location within the downlink cell bandwidth. Hence, the downlink hybrid-ARQ protocol is *asynchronous* and *adaptive*. The support for asynchronous and adaptive hybrid ARQ for the LTE downlink is motivated by the need to avoid collisions with, for example, transmission of system information and MBSFN subframes. Instead of dropping a retransmission that otherwise would collide with MBSFN subframes or transmission of system information, the eNodeB can move the retransmission in time and/or frequency to avoid the overlap in resources.

Support for soft combining is, as described in the introduction, provided through an explicit new-data indicator, toggled for each new transport block. In addition to the new-data indicator, hybrid-ARQ-related downlink control signaling consists of the hybrid-ARQ process number (three bits for FDD, four bits for TDD) and the redundancy version (two bits), both explicitly signaled in the scheduling assignment for each downlink transmission.

Downlink spatial multiplexing implies, as already mentioned, transmission of two transport blocks in parallel on a component carrier. To provide the possibility to retransmit only one of the transport blocks, which is beneficial as error events for the two transport blocks can be fairly uncorrelated, each transport block has its own separate new-data indicator and redundancy-version indication. However, there is no need to signal the process number separately as each process consists of two sub-processes in the case of spatial multiplexing or, expressed differently, once the process number for the first transport block is known, the process number for the second transport block is given implicitly.

Transmissions on downlink component carriers are acknowledged independently. On each component carrier, transport blocks are acknowledged by transmitting one or two bits on the uplink, as described in Chapter 7. In the absence of spatial multiplexing, there is only a single transport block within a TTI and consequently only a single acknowledgment bit is required in response. However, if the downlink transmission used spatial multiplexing, there are two transport blocks per TTI, each requiring its own hybrid-ARQ acknowledgment bit. The total

number of bits required for hybrid-ARQ acknowledgments thus depends on the number of component carriers and the transmission mode for each of the component carriers. As each downlink component carrier is scheduled separately from its own PDCCH, the hybrid-ARQ process numbers are signaled independently for each component carrier.

The device should not transmit a hybrid-ARQ acknowledgment in response to reception of system information, paging messages, and other broadcast traffic. Hence, hybrid-ARQ acknowledgments are only sent in the uplink for "normal" unicast transmission.

8.1.2 UPLINK HYBRID ARQ

Shifting the focus to the uplink, a difference compared to the downlink case is the use of synchronous, nonadaptive operation as the basic principle of the hybrid-ARQ protocol, motivated by the lower overhead compared to an asynchronous, adaptive structure. Hence, uplink retransmissions always occur at an a priori known subframe; in the case of FDD operation uplink retransmissions occur eight subframes after the prior transmission attempt for the same hybrid-ARQ process. The set of resource blocks used for the retransmission on a component carrier is identical to the initial transmission. Thus, the only control signaling required in the downlink for a retransmission is a retransmission command, transmitted on the PHICH as described in Chapter 6. In the case of a negative acknowledgment on the PHICH, the data are retransmitted.

Spatial multiplexing is handled in the same way as in the downlink—two transport blocks transmitted in parallel on the uplink, each with its own modulation-and-coding scheme and new-data indicator but sharing the same hybrid-ARQ process number. The two transport blocks are individually acknowledged and hence two bits of information are needed in the downlink to acknowledge an uplink transmission using spatial multiplexing. TDD is another example, as is discussed later, where uplink transmissions in different subframes may need to be acknowledged in the same downlink subframe.[1] Thus, multiple PHICHs may be needed as each PHICH is capable of transmitting a single bit only.

Despite the fact that the basic mode of operation for the uplink is synchronous, nonadaptive hybrid ARQ, there is also the possibility to operate the uplink hybrid ARQ in a synchronous, *adaptive* manner, where the resource-block set and modulation-and-coding scheme for the retransmissions are changed. Although nonadaptive retransmissions are typically used due to the very low overhead in terms of downlink control signaling, adaptive retransmissions are sometimes useful to avoid fragmenting the uplink frequency resource or to avoid collisions with random-access resources. This is illustrated in Figure 8.5. A device is scheduled for an initial transmission in subframe n; a transmission that is not correctly received and consequently a retransmission is required in subframe $n + 8$ (assuming FDD; for TDD the timing depends on the uplink–downlink allocation, as discussed later). With nonadaptive hybrid ARQ, the retransmissions occupy the same part of the uplink spectrum as

[1] A third example is carrier aggregation, see Chapter 12.

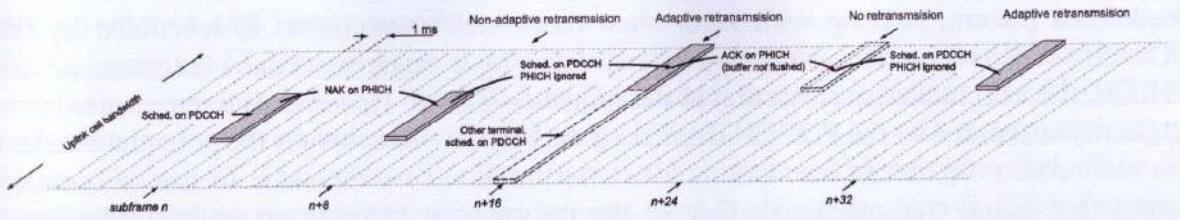


FIGURE 8.5

Nonadaptive and adaptive hybrid-ARQ operation.

the initial transmission. Hence, in this example the spectrum is fragmented, which limits the bandwidth available to another device (unless the other device is capable of multi-cluster transmission). In subframe $n + 16$, an example of an adaptive retransmission is found; to make room for another device to be granted a large part of the uplink spectrum, the retransmission is moved in the frequency domain. It should be noted that the uplink hybrid-ARQ protocol is still synchronous—that is, a retransmission should always occur eight subframes after the previous transmission.

The support for both adaptive and nonadaptive hybrid ARQ is realized by *not* flushing the transmission buffer when receiving a positive hybrid-ARQ acknowledgment on PHICH for a given hybrid-ARQ process. Instead, the actual control of whether data should be retransmitted or not is done by the new-data indicator included in the uplink scheduling grant sent on the PDCCH. The new-data indicator is toggled for each new transport block. If the new-data indicator is toggled, the device flushes the transmission buffer and transmits a new-data packet. However, if the new-data indicator does not request transmission of a new transport block, the previous transport block is retransmitted. Hence, clearing of the transmission buffer is not handled by the PHICH but by the PDCCH as part of the uplink grant. The negative hybrid-ARQ acknowledgment on the PHICH should preferably be seen as a single-bit scheduling grant for retransmissions where the set of bits to transmit and all the resource information are known from the previous transmission attempt. In Figure 8.5 an example of postponing a transmission is seen in subframe $n + 24$. The device has received a positive acknowledgment and therefore does *not* retransmit the data. However, the transmission buffer is not flushed, which later is exploited by an uplink grant requesting retransmission in subframe $n + 32$.

A consequence of the above-mentioned method of supporting both adaptive and nonadaptive hybrid ARQ is that the PHICH and PDCCH related to the same uplink subframe have the same timing. If this were not the case, the complexity would increase as the device would not know whether to obey the PHICH or wait for a PDCCH overriding the PHICH.

As explained earlier, the new-data indicator is explicitly transmitted in the uplink grant. However, unlike the downlink case, the redundancy version is *not* explicitly signaled for each retransmission. With a single-bit acknowledgment on the PHICH, this is not possible. Instead, as the uplink hybrid-ARQ protocol is synchronous, the redundancy version follows a

predefined pattern, starting with zero when the initial transmission is scheduled by the PDCCH. Whenever a retransmission is requested by a negative acknowledgment on the PHICH, the next redundancy version in the sequence is used. However, if a retransmission is explicitly scheduled by the PDCCH overriding the PHICH, then there is the potential to affect the redundancy version to use. Grants for retransmissions use the same format as ordinary grants (for initial transmissions). One of the information fields in an uplink grant is, as described in Chapter 6, the modulation-and-coding scheme. Of the 32 different combinations this five-bit field can take, three of them are reserved. Those three combinations represent different redundancy versions; hence, if one of these combinations is signaled as part of an uplink grant indicating a retransmission, the corresponding redundancy version is used for the transmission. The transport-block size is already known from the initial transmission as it, by definition, cannot change between retransmission attempts. In Figure 8.5, the initial transmission in subframe n uses the first redundancy version in sequence as the transport-block size must be indicated for the initial transmission. The retransmission in subframe $n + 8$ uses the next redundancy version in the sequence, while the explicitly scheduled retransmission in subframe $n + 16$ can use any redundancy scheme as indicated on the PDCCH.

Whether to exploit the possibility to signal an arbitrary redundancy version in a retransmission scheduled by the PDCCH is a trade-off between incremental-redundancy gain and robustness. From an incremental-redundancy perspective, changing the redundancy value between retransmissions is typically beneficial to fully exploit the gain from IR. However, as the modulation-and-coding scheme is normally not indicated for uplink retransmissions, as either the single-bit PHICH is used or the modulation-and-coding field is used to explicitly indicate a new redundancy version, it is implicitly assumed that the device did not miss the initial scheduling grant. If this is the case, it is necessary to explicitly indicate the modulation-and-coding scheme, which also implies that the first redundancy version in the sequence is used.

8.1.3 HYBRID-ARQ TIMING

The receiver must know to which hybrid-ARQ process a received acknowledgment is associated. This is handled by the timing of the hybrid-ARQ acknowledgment being used to associate the acknowledgment with a certain hybrid-ARQ process; the timing relation between the reception of data in the downlink and transmission of the hybrid-ARQ acknowledgment in the uplink (and vice versa) is fixed. From a latency perspective, the time between the reception of downlink data at the device and transmission of the hybrid-ARQ acknowledgment in the uplink should be as short as possible. At the same time, an unnecessarily short time would increase the demand on the device processing capacity and a trade-off between latency and implementation complexity is required. The situation is similar for uplink data transmissions. For LTE, this trade-off led to the decision to have eight hybrid-ARQ processes per component carrier in both uplink and downlink for FDD. For TDD, the number of processes depends on the uplink–downlink allocation, as discussed later.

8.1.3.1 Hybrid-ARQ Timing for FDD

Starting with the FDD case, transmission of acknowledgments in the uplink in response to downlink data transmission is illustrated in Figure 8.6. Downlink data on the DL-SCH is transmitted to the device in subframe n and received by the device, after the propagation delay T_p, in subframe n. The device attempts to decode the received signal, possibly after soft combining with a previous transmission attempt, and transmits the hybrid-ARQ acknowledgment in uplink subframe $n + 4$ (note that the start of an uplink subframe at the device is offset by T_{TA} relative to the start of the corresponding downlink subframe at the device as a result of the timing-advance procedure described in Section 7.6). Upon reception of the hybrid-ARQ acknowledgment, the eNodeB can, if needed, retransmit the downlink data in subframe $n + 8$. Thus, eight hybrid-ARQ processes are used—that is, the hybrid-ARQ round-trip time is 8 ms.

The description in the previous paragraph, as well as the illustration in Figure 8.6, describe the timing for downlink data transmission—that is, data on DL-SCH and acknowledgments on PUCCH (or PUSCH). However, the timing of uplink data transmission—that is, data on PUSCH and acknowledgments on PHICH—is identical, namely uplink data transmission in subframe n results in a PHICH transmission in subframe $n + 4$. This is also the same timing relation as for uplink scheduling grants in general (as described in Section 9.3) and allows a scheduling grant on the PDCCH to override the PHICH, as shown in Figure 8.5.

In Figure 8.6 it is seen that the processing time available to the device, T_{UE}, depends on the value of the timing advance or, equivalently, on the device-to-base-station distance. As a device must operate at any distance up to the maximum size supported by the specifications, the device must be designed such that it can handle the worst-case scenario. LTE is designed to handle at least 100 km, corresponding to a maximum timing advance of 0.67 ms. Hence, there is approximately 2.3 ms left for the device processing, which is considered a reasonable

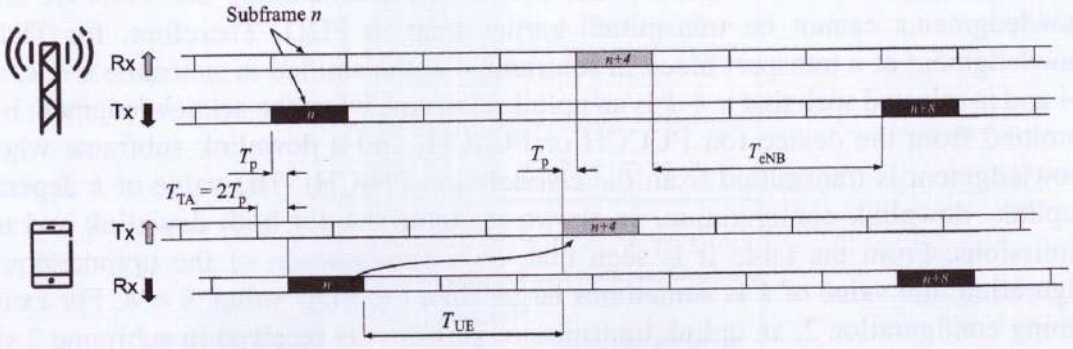


FIGURE 8.6

Timing relation between downlink data in subframe n and uplink hybrid-ARQ acknowledgment in subframe $n + 4$ for FDD.

trade-off between the processing requirements imposed on a device and the associated delays.

For the eNodeB, the processing time available, denoted as T_{eNB}, is 3 ms and thus of the same order as for the device. In the case of downlink data transmission, the eNodeB performs scheduling of any retransmissions during this time and, for uplink data transmission, the time is used for decoding of the received signal. The timing budget is thus similar for the device and the eNodeB, which is motivated by the fact that, although the eNodeB typically has more processing power than a device, it also has to serve multiple devices and to perform the scheduling operation.

Having the same number of hybrid-ARQ processes in uplink and downlink is beneficial for half-duplex FDD operation, discussed in Chapter 9. By proper scheduling, the uplink hybrid-ARQ transmission from the device will coincide with transmission of uplink data and the acknowledgments related to the reception of uplink data will be transmitted in the same subframe as the downlink data. Thus, using the same number of hybrid-ARQ processes in uplink and downlink results in a 50:50 split between transmission and reception for a half-duplex device.

8.1.3.2 Hybrid-ARQ Timing for TDD

For TDD operation, the time relation between the reception of data in a certain hybrid-ARQ process and the transmission of the hybrid-ARQ acknowledgment depends on the uplink–downlink allocation. An uplink hybrid-ARQ acknowledgment can only be transmitted in an uplink subframe and a downlink acknowledgment only in a downlink subframe. Furthermore, starting from release 11, different uplink–downlink allocations can be used on component carriers operating in different frequency bands. Therefore, the hybrid-ARQ timing relations for TDD are more intricate than their FDD counterparts.

For simplicity, consider the situation when the same uplink–downlink allocation is used across all component carriers (the case of different allocations is described later). Since the minimum processing time required in the device and the eNodeB remains the same in FDD and TDD as the same turbo decoders are used and the scheduling decisions are similar, acknowledgments cannot be transmitted earlier than in FDD. Therefore, for TDD the acknowledgment of a transport block in subframe n is transmitted in subframe $n + k$, where $k \geq 4$ and is selected such that $n + k$ is an uplink subframe when the acknowledgment is to be transmitted from the device (on PUCCH or PUSCH) and a downlink subframe when the acknowledgment is transmitted from the eNodeB (on PHICH). The value of k depends on the uplink–downlink configuration, as shown in Table 8.1 for both downlink and uplink transmissions. From the table it is seen that, as a consequence of the uplink–downlink configuration, the value of k is sometimes larger than the FDD value, $k = 4$. For example, assuming configuration 2, an uplink transmission on PUSCH received in subframe 2 should be acknowledged on PHICH in subframe $2 + 6 = 8$. Similarly, for the same configuration, downlink transmission on PDSCH in subframe 0 should be acknowledged on PUCCH (or PUSCH) in subframe $0 + 7 = 7$.

Table 8.1 Number of Hybrid-ARQ Processes and Acknowledgment Timing k for Different TDD Configurations

Configuration (DL:UL)	Downlink											Uplink									
	Proc	PDSCH Reception in Subframe n										Proc	PUSCH Reception in Subframe n								
		0	1	2	3	4	5	6	7	8	9		0	1	2	3	4	5	6	7	8	9
0 (2:3)	4	4	6	–	–	–	4	6	–	–	–	6	–	–	4	7	6	–	–	4	7	6
1 (3:2)	7	7	6	–	–	4	7	6	–	–	4	4	–	–	4	6	–	–	–	4	6	–
2 (4:1)	10	7	6	–	4	8	7	6	–	4	8	2	–	–	6	–	–	–	–	6	–	–
3 (7:3)	9	4	11	–	–	–	7	6	6	5	5	3	–	–	6	6	6	–	–	–	–	–
4 (8:2)	12	12	11	–	–	8	7	7	6	5	4	2	–	–	6	6	–	–	–	–	–	–
5 (9:1)	15	12	11	–	9	8	7	6	5	4	13	1	–	–	6	–	–	–	–	–	–	–
6 (5:5)	6	7	7	–	–	–	7	7	–	–	5	6	–	–	4	6	6	–	–	4	7	–

From the table it is also seen that the number of hybrid-ARQ processes used for TDD depends on the uplink–downlink configuration, implying that the hybrid-ARQ round-trip time is configuration dependent for TDD (actually, it may even vary between subframes, as seen in Figure 8.7). For the downlink-heavy configurations 2, 3, 4, and 5, the number of downlink hybrid-ARQ processes is larger than for FDD. The reason is the limited number of uplink subframes available, resulting in k-values well beyond 4 for some subframes (Table 8.1).

The PHICH timing in TDD is identical to the timing when receiving an uplink grant, as described in Section 9.3. The reason is the same as in FDD, namely to allow a PDCCH uplink grant to override the PHICH in order to implement adaptive retransmissions, as illustrated in Figure 8.5.

The uplink–downlink allocation for TDD has implications on the number of transport blocks to acknowledge in a single subframe. For FDD there is always a one-to-one relation between uplink and downlink subframes. Hence, in absence of carrier aggregation, a subframe only needs to carry acknowledgments for one subframe in the other direction. For TDD, in contrast, there is not necessarily a one-to-one relation between uplink and downlink subframes. This can be seen in the possible uplink–downlink allocations described in Chapter 5.

For uplink transmissions using UL-SCH, each uplink transport block is acknowledged individually by using the PHICH. Hence, in uplink-heavy asymmetries (uplink–downlink configuration 0), the device may need to receive two hybrid-ARQ acknowledgments in downlink subframes 0 and 5 in the absence of uplink spatial multiplexing; for spatial multiplexing up to four acknowledgments per component carrier are needed. There are also some subframes where no PHICH will be transmitted. The amount of PHICH groups may therefore vary between subframes for TDD.

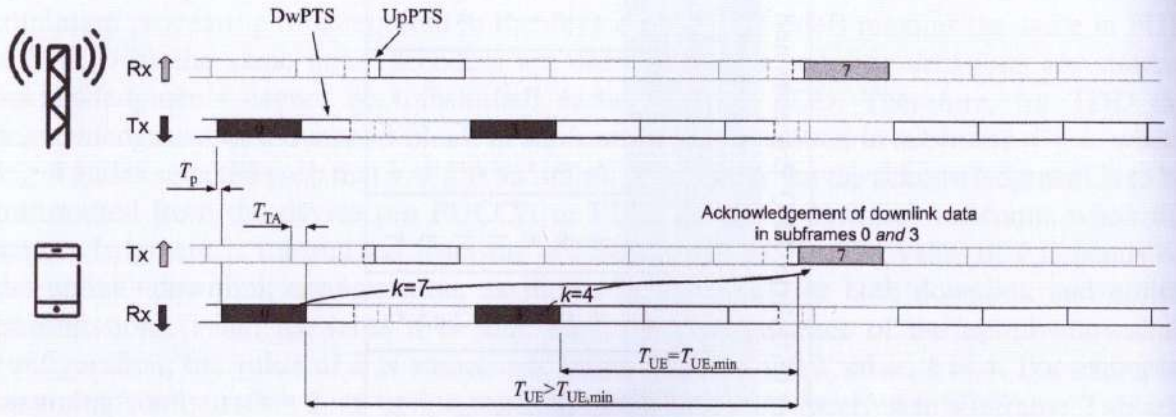


FIGURE 8.7

Example of timing relation between downlink data and uplink hybrid-ARQ acknowledgment for TDD (configuration 2).

For downlink transmissions, there are some configurations where DL-SCH receipt in multiple downlink subframes needs to be acknowledged in a single uplink subframe, as illustrated in Figure 8.7 (the illustration corresponds to the two entries shown in bold in Table 8.1). Two different mechanisms to handle this are provided in TDD: multiplexing and bundling.

Multiplexing implies that independent acknowledgments for each of the received transport blocks are fed back to the eNodeB. This allows independent retransmission of erroneous transport blocks. However, it also implies that multiple bits need to be transmitted from the device, which may limit the uplink coverage. This is the motivation for the bundling mechanism.

Bundling of acknowledgments implies that the outcome of the decoding of downlink transport blocks from multiple downlink subframes can be combined into a single hybrid-ARQ acknowledgment transmitted in the uplink. Only if both of the downlink transmissions in subframes 0 and 3 in the example in Figure 8.7 are correctly decoded will a positive acknowledgment be transmitted in uplink subframe 7.

Combining acknowledgments related to multiple downlink transmissions into a single uplink message assumes that the device has not missed any of the scheduling assignments upon which the acknowledgment is based. Assume, as an example, that the eNodeB scheduled the device in two (subsequent) subframes, but the device missed the PDCCH transmission in the first of the two subframes and successfully decoded the data transmitted in the second subframe. Without any additional mechanism, the device will transmit an acknowledgment based on the assumption that it was scheduled in the second subframe only, while the eNodeB will interpret acknowledgment as the device successfully received both transmissions. To avoid such errors, the *downlink assignment index* (see Section 6.4.6) in the scheduling assignment is used. The downlink assignment index in essence informs the device about the number of transmissions it should base the combined acknowledgment upon. If there is a mismatch between the assignment index and the number of transmissions the device received, the device concludes at least one assignment was missed and transmits no hybrid-ARQ acknowledgment, thereby avoiding acknowledging transmissions not received.

8.2 RADIO-LINK CONTROL

The *radio-link control* (RLC) protocol takes data in the form of RLC SDUs from PDCP and delivers them to the corresponding RLC entity in the receiver by using functionality in MAC and physical layers. The relation between RLC and MAC, including multiplexing of multiple logical channels into a single transport channel, is illustrated in Figure 8.8. Multiplexing of several logical channels into a single transport channel is mainly used for priority handling, as described in Section 9.2 in conjunction with downlink and uplink scheduling.

power headroom can very well be negative, indicating that the per-carrier transmit power was limited by $P_{CMAX,c}$ at the time of the power headroom reporting, that is, the network has scheduled a higher data rate than the device can support given the available transmission power. As the network knows what modulation-and-coding scheme and resource size the device used for transmission in the subframe to which the power-headroom report corresponds, it can determine the valid combinations of modulation-and-coding scheme and resource size M, assuming that the downlink path loss PL_{DL} and the term δ have not changed substantially.

Type 1 power headroom can also be reported for subframes where there is no actual PUSCH transmission. In such cases, $10 \cdot \log_{10}(M)$ and Δ_{MCS} in the expression above are set to zero. This can be seen as the power headroom assuming a default transmission configuration corresponding to the minimum possible resource assignment ($M = 1$) and the modulation-and-coding scheme associated with $\Delta_{MCS} = 0$ dB.

Similarly, Type 2 power headroom reporting is defined as the difference between the maximum per-carrier transmit power and the sum of the PUSCH and PUCCH transmit power (Eqs. (7.4) and (7.3), respectively), once again not taking into account any maximum per-carrier power when calculating the PUSCH and PUCCH transmit power.[3]

Along the lines of Type 1 power headroom reporting, the Type 2 power headroom can also be reported for subframes in which no PUSCH and/or PUCCH is transmitted. In that case a virtual PUSCH and PUCCH transmit power is calculated, assuming the smallest possible resource assignment ($M = 1$) and $\Delta_{MCS} = 0$ dB for PUSCH and $\Delta_{Format} = 0$ for PUCCH.

9.4 TIMING OF SCHEDULING ASSIGNMENTS/GRANTS

The scheduling decisions, downlink scheduling assignments and uplink scheduling grants, are communicated to each of the scheduled devices through the downlink L1/L2 control signaling as described in Chapter 6, using one (E)PDCCH per downlink assignment. Each device monitors a set of (E)PDCCHs for valid scheduling assignment or grants and, upon detection of a valid assignment or grant, receives PDSCH or transmits PUSCH, respectively. The device needs to know which subframe the scheduling command relates to.

9.4.1 DOWNLINK SCHEDULING TIMING

For downlink data transmission, the scheduling assignment is transmitted in the same subframe as the data. Having the scheduling assignment in the same subframe as the corresponding data minimizes the latency in the scheduling process. Also, note that there is no possibility for dynamically scheduling future subframes—data and the scheduling assignment are always in the same subframe. This holds for FDD as well as TDD.

[3]In case of carrier aggregation, type 2 reports are supported for the primary component carrier only, as PUCCH cannot be transmitted on a secondary component carrier (prior to release 13).

9.4.2 UPLINK SCHEDULING TIMING

The timing for uplink scheduling grants is more intricate than the corresponding downlink scheduling assignments, especially for TDD. The grant cannot relate to the same subframe it was received in as the uplink subframe has already started when the device has decoded the grant. The device also needs some time to prepare the data to transmit. Therefore, a grant received in subframe n affects the uplink transmission in a later subframe.

For FDD, the grant timing is straightforward. An uplink grant received in a subframe n triggers an uplink transmission in subframe $n + 4$, as illustrated in Figure 9.7. This is the same timing relation as used for uplink retransmission triggered by the PHICH, motivated by the possibility to override the PHICH by a dynamic scheduling grant, as described in Chapter 8.

For TDD, the situation is more complicated. A grant received in subframe n in TDD may not necessarily trigger an uplink transmission in subframe $n + 4$—the timing relation used in FDD—as subframe $n + 4$ may not be an uplink subframe. Hence, for TDD configurations 1–6 the timing relation is modified such that the uplink transmission occurs in subframe $n + k$, where k is the smallest value larger than or equal to 4 such that subframe $n + k$ is an uplink subframe. This provides at least the same processing time for the device as in the FDD

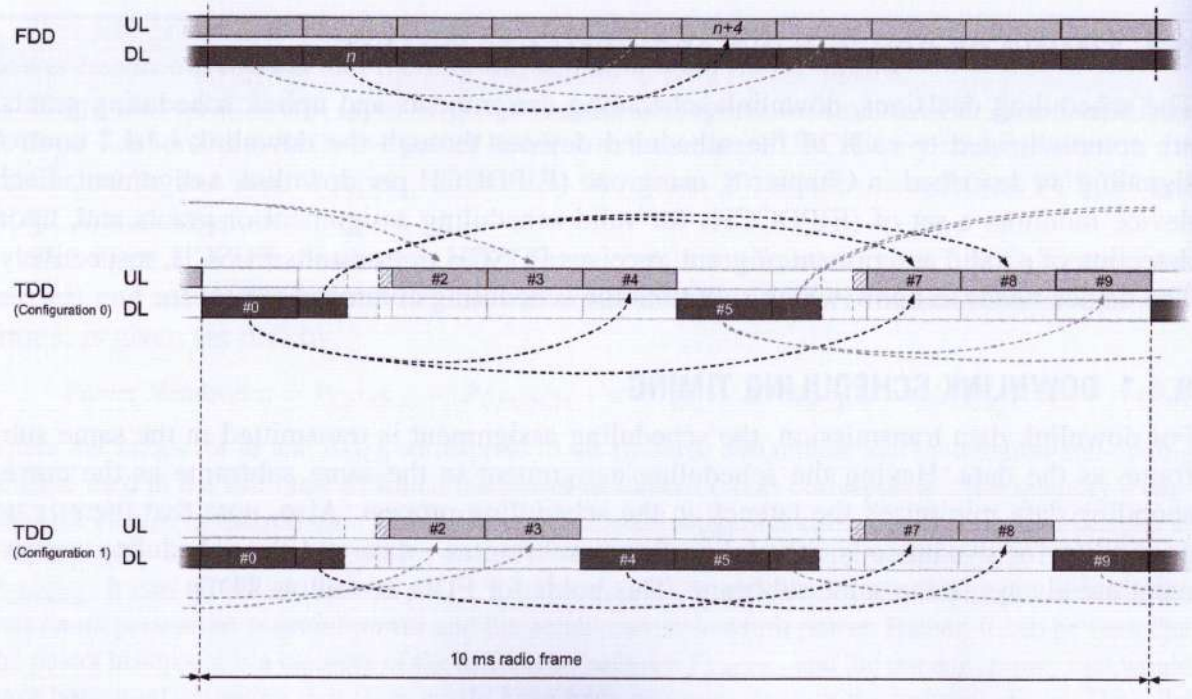


FIGURE 9.7

Timing relation for uplink grants in FDD and TDD configurations 0 and 1.

case while minimizing the delay from receipt of the uplink grant to the actual transmission. Note that this implies that the time between grant reception and uplink transmission may differ between different subframes. Furthermore, for the downlink-heavy configurations 1–5, another property is that uplink scheduling grants can only be received in some of the downlink subframes.

For TDD configuration 0 there are more uplink subframes than downlink subframes, which calls for the possibility to schedule transmissions in multiple uplink subframes from a single downlink subframe. The same timing relation as for the other TDD configurations is used but with slight modifications. Recall from Section 6.4.7 that the grant transmitted in the downlink contains an uplink index consisting of two bits. For uplink-downlink configuration 0, the index field specifies which uplink subframe(s) a grant received in a downlink subframe applies to. For example, as illustrated in Figure 9.7, an uplink scheduling grant received in downlink subframe 0 applies to one or both of the uplink subframes 4 and 7, depending on which of the bits in the uplink index are set.

9.5 SEMI-PERSISTENT SCHEDULING

The basis for uplink and downlink scheduling is dynamic scheduling, as described in Sections 9.2 and 9.3. Dynamic scheduling with a new scheduling decision taken in each subframe allows for full flexibility in terms of the resources used and can handle large variations in the amount of data to transmit at the cost of the scheduling decision being sent on an (E)PDCCH in each subframe. In many situations, the overhead in terms of control signaling on the (E)PDCCH is well motivated and relatively small compared to the payload on DL-SCH/UL-SCH. However, some services, most notably voice-over IP, are characterized by regularly occurring transmission of relatively small payloads. To reduce the control signaling overhead for those services, LTE provides semi-persistent scheduling in addition to dynamic scheduling.

With semi-persistent scheduling, the device is provided with the scheduling decision on the (E)PDCCH, together with an indication that this applies to every nth subframe until further notice. Hence, control signaling is only used once and the overhead is reduced, as illustrated in Figure 9.8. The periodicity of semi-persistently scheduled transmissions, that is, the value of n, is configured by RRC signaling in advance, while activation (and deactivation)

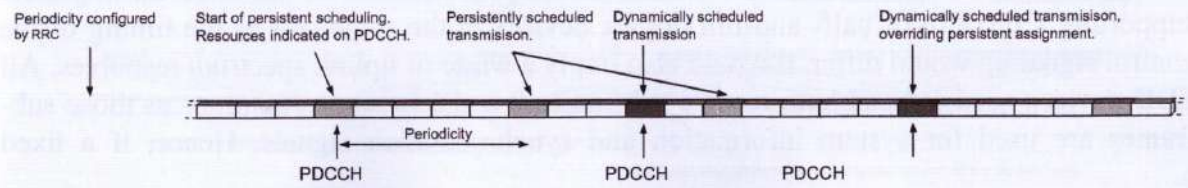


FIGURE 9.8

Example of semi-persistent scheduling.