

White Paper

Overview of the 3GPP Long Term Evolution Physical Layer

Jim Zyren

jzyren@freescale.com

Dr. Wes McCoy, Technical Editor

Document Number: 3GPPEVOLUTIONWP
Rev 0
07/2007



Overview

Long Term Evolution (LTE) is the next step forward in cellular 3G services. Expected in the 2008 time frame, LTE is a 3GPP standard that provides for an uplink speed of up to 50 megabits per second (Mbps) and a downlink speed of up to 100 Mbps. LTE will bring many technical benefits to cellular networks. Bandwidth will be scalable from 1.25 MHz to 20 MHz. This will suit the needs of different network operators that have different bandwidth allocations, and also allow operators to provide different services based on spectrum. LTE is also expected to improve spectral efficiency in 3G networks, allowing carriers to provide more data and voice services over a given bandwidth.

This technical white paper provides an overview of the LTE physical layer (PHY), including technologies that are new to cellular such as Orthogonal Frequency Division Multiplexing (OFDM) and Multiple Input Multiple Output (MIMO) data transmission.

Contents

1 Introduction	1	3.1.4 Physical Signals	17
1.1 LTE Design Goals	1	3.1.5 Transport Channels	18
2 LTE Basic Concepts	2	3.1.6 Mapping Downlink Physical Channels to Transport Channels	19
2.1 Single Carrier Modulation and Channel Equalization	2	3.1.7 Downlink Channel Coding	19
2.2 OFDM	3	3.2 Uplink	20
2.2.1 Disadvantages of OFDM	5	3.2.1 Modulation Parameters	20
2.3 OFDMA	7	3.2.2 Multiplexing	20
2.3.1 Comparison of OFDMA with Packet-Oriented Protocols	7	3.2.3 Uplink Physical Channels	20
2.3.2 OFDMA and the LTE Generic Frame Structure	8	3.2.4 Uplink Physical Signals	21
2.4 MIMO and MRC	10	3.2.5 Uplink Transport Channels	21
2.5 SC-FDMA	13	3.2.6 Mapping Uplink Physical Channels to Transport Channels	22
3 LTE Physical Layer	14	3.2.7 Coding	22
3.1 Downlink	15	3.3 MB-SFN	22
3.1.1 Modulation Parameters	15	4 Conclusions	22
3.1.2 Downlink Multiplexing	16	5 References	23
3.1.3 Physical Channels	16	6 Acronyms and Abbreviations	24

1 Introduction

The 3GPP Long Term Evolution (LTE) represents a major advance in cellular technology. LTE is designed to meet carrier needs for high-speed data and media transport as well as high-capacity voice support well into the next decade. It encompasses high-speed data, multimedia unicast and multimedia broadcast services. Although technical specifications are not yet finalized, significant details are emerging. This paper focuses on the LTE physical layer (PHY).

The LTE PHY is a highly efficient means of conveying both data and control information between an enhanced base station (eNodeB) and mobile user equipment (UE). The LTE PHY employs some advanced technologies that are new to cellular applications. These include Orthogonal Frequency Division Multiplexing (OFDM) and Multiple Input Multiple Output (MIMO) data transmission. In addition, the LTE PHY uses Orthogonal Frequency Division Multiple Access (OFDMA) on the downlink (DL) and Single Carrier – Frequency Division Multiple Access (SC-FDMA) on the uplink (UL). OFDMA allows data to be directed to or from multiple users on a subcarrier-by-subcarrier basis for a specified number of symbol periods. Due to the novelty of these technologies in cellular applications, they are described separately before delving into a description of the LTE PHY.

Although the LTE specs describe both Frequency Division Duplexing (FDD) and Time Division Duplexing (TDD) to separate UL and DL traffic, market preferences dictate that the majority of deployed systems will be FDD. This paper therefore describes LTE FDD systems only.

1.1 LTE Design Goals

The LTE PHY is designed to meet the following goals [1]:

1. Support scalable bandwidths of 1.25, 2.5, 5.0, 10.0 and 20.0 MHz
2. Peak data rate that scales with system bandwidth
 - a. Downlink (2 Ch MIMO) peak rate of 100 Mbps in 20 MHz channel
 - b. Uplink (single Ch Tx) peak rate of 50 Mbps in 20 MHz channel
3. Supported antenna configurations
 - a. Downlink: 4x2, 2x2, 1x2, 1x1
 - b. Uplink: 1x2, 1x1
4. Spectrum efficiency
 - a. Downlink: 3 to 4 x HSDPA Rel. 6
 - b. Uplink: 2 to 3 x HSUPA Rel. 6
5. Latency
 - a. C-plane: <50 – 100 msec to establish U-plane
 - b. U-plane: <10 msec from UE to server
6. Mobility
 - A. Optimized for low speeds (<15 km/hr)
 - B. High performance at speeds up to 120 km/hr
 - C. Maintain link at speeds up to 350 km/hr
7. Coverage
 - a. Full performance up to 5 km
 - b. Slight degradation 5 km – 30 km
 - c. Operation up to 100 km should not be precluded by standard

2 LTE Basic Concepts

Before jumping into a detailed description of the LTE PHY, it's worth taking a look at some of the basic technologies involved. Many methods employed in LTE are relatively new in cellular applications. These include OFDM, OFDMA, MIMO and Single Carrier Frequency Division Multiple Access (SC-FDMA). Readers familiar with these technologies can skip this material and proceed directly to Section 3.

LTE employs OFDM for downlink data transmission and SC-FDMA for uplink transmission. OFDM is a well-known modulation technique, but is rather novel in cellular applications. A brief discussion of the basic properties and advantages of this method is therefore warranted.

When information is transmitted over a wireless channel, the signal can be distorted due to multipath. Typically (but not always) there is a line-of-sight path between the transmitter and receiver. In addition, there are many other paths created by signal reflection off buildings, vehicles and other obstructions as shown in Figure 2.0-1. Signals traveling along these paths all reach the receiver, but are shifted in time by an amount corresponding to the differences in the distance traveled along each path.

2.1 Single Carrier Modulation and Channel Equalization

To date, cellular systems have used single carrier modulation schemes almost exclusively. Although LTE uses OFDM rather than single carrier modulation, it's instructive to briefly discuss how single carrier systems deal with multipath-induced channel distortion. This will form a point of reference from which OFDM systems can be compared and contrasted.

The term delay spread describes the amount of time delay at the receiver from a signal traveling from the transmitter along different paths. In cellular applications, delay spreads can be several microseconds. The delay induced by multipath can cause a symbol received along a delayed path to "bleed" into a subsequent symbol arriving at the receiver via a more direct path. This effect is depicted in Figure 2.1-1 and is referred to as inter-symbol interference (ISI). In a conventional single carrier system symbol times decrease as data rates increase. At very high data rates (with correspondingly shorter symbol periods), it is quite possible for ISI to exceed an entire symbol period and spill into a second or third subsequent symbol.

Figure 2.0-1 Multipath is Caused by Reflections Off Objects Such as Buildings and Vehicles



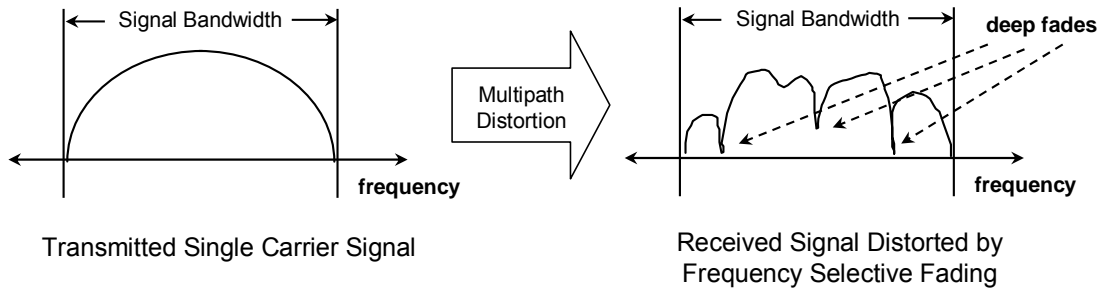


Figure 2.1-1 Multipath-Induced Time Delays Result in ISI



It's also helpful to consider the effects of multipath distortion in the frequency domain. Each different path length and reflection will result in a specific phase shift. As all of the signals are combined at the receiver, some frequencies within the signal passband undergo constructive interference (linear combination of signals in-phase), while others encounter destructive interference (linear combination of signals out-of-phase). The composite received signal is distorted by frequency selective fading (see Figure 2.1-2).

Figure 2.1-2 Longer Delay Spreads Result in Frequency Selective Fading




Single carrier systems compensate for channel distortion via time domain equalization. This is a substantial topic by itself, and beyond the scope of this paper. Generally, time domain equalizers compensate for multipath induced distortion by one of two methods:

1. Channel inversion: A known sequence is transmitted over the channel prior to sending information. Because the original signal is known at the receiver, a channel equalizer is able to determine the channel response and multiply the subsequent data-bearing signal by the inverse of the channel response to reverse the effects of multipath.
2. CDMA systems can employ rake equalizers to resolve the individual paths and then combine digital copies of the received signal shifted in time to enhance the receiver signal-to-noise ratio (SNR).

In either case, channel equalizer implementation becomes increasingly complex as data rates increase. Symbol times become shorter and receiver sample clocks must become correspondingly faster. ISI becomes much more severe—possibly spanning several symbol periods.

Figure 2.1-3 Transversal Filter Channel Equalizer



The finite impulse response transversal filter (see Figure 2.1-3) is a common equalizer topology. As the period of the receiver sample clock (τ) decreases, more samples are required to compensate for a given amount of delay spread. The number of delay taps increases along with the speed and complexity of the adaptive algorithm. For LTE data rates (up to 100 Mbps) and delay spreads (approaching 17 μ sec), this approach to channel equalization becomes impractical. As we will discuss below, OFDM eliminates ISI in the time domain, which dramatically simplifies the task of channel compensation.

2.2 OFDM


Unlike single carrier systems described above, OFDM communication systems do not rely on increased symbol rates in order to achieve higher data rates. This makes the task of managing ISI much simpler. OFDM systems break the available bandwidth into many narrower sub-carriers and transmit the data in parallel streams. Each subcarrier is modulated using varying levels of QAM modulation, e.g. QPSK, QAM, 64QAM or possibly higher orders depending on signal quality. Each OFDM symbol is therefore a linear combination of the instantaneous signals on each of the sub-

carriers in the channel. Because data is transmitted in parallel rather than serially, OFDM symbols are generally MUCH longer than symbols on single carrier systems of equivalent data rate.

There are two truly remarkable aspects of OFDM. First, each OFDM symbol is preceded by a cyclic prefix (CP), which is used to effectively eliminate ISI. Second, the sub-carriers are very tightly spaced to make efficient use of available bandwidth, yet there is virtually no interference among adjacent sub-carriers (Inter Carrier Interference, or ICI). These two unique features are actually closely related. In order to understand how OFDM deals with multipath distortion, it's useful to consider the signal in both the time and frequency domains.

To understand how OFDM deals with ISI induced by multipath, consider the time domain representation of an OFDM symbol shown in Figure 2.2-1. The OFDM symbol consists of two major components: the CP and an FFT period (T_{FFT}). The duration of the CP is determined by the highest anticipated degree of delay spread for the targeted application. When transmitted signals arrive at the receiver by two paths of differing length, they are staggered in time as shown in Fig. 2.2-2.

Figure 2.2-2 OFDM Eliminates ISI via Longer Symbol Periods and a Cyclic Prefix

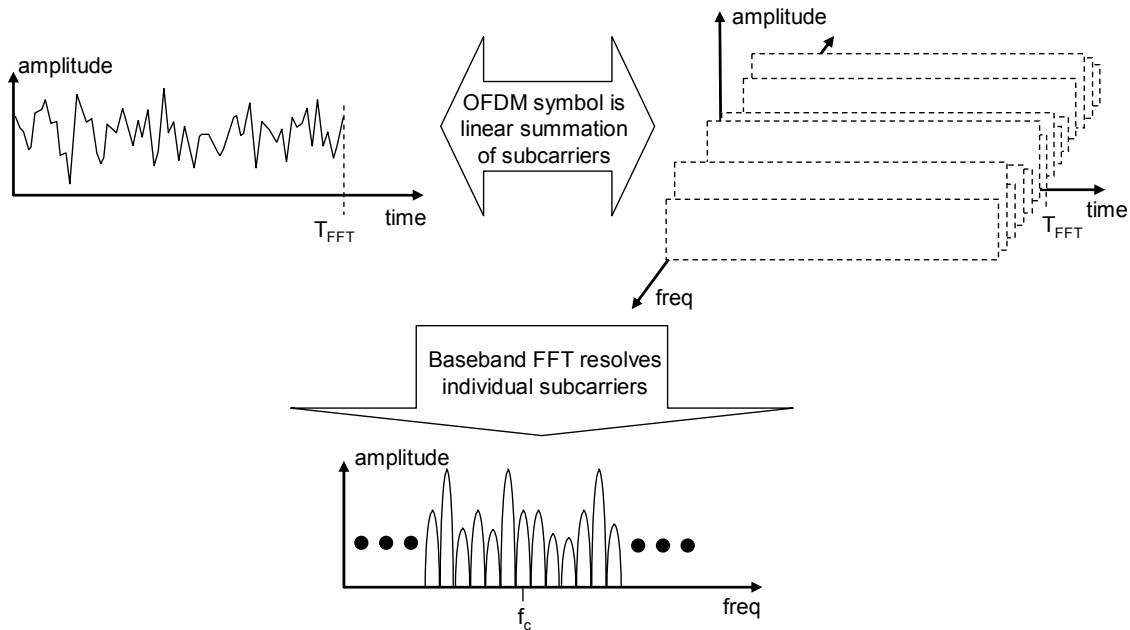


Within the CP, it is possible to have distortion from the preceding symbol. However, with a CP of sufficient duration, preceding symbols do not spill over into the FFT period; there is only interference caused by time-staggered “copies” of the current symbol. Once the channel impulse response is determined (by periodic transmission of known reference signals), distortion can be corrected by applying an amplitude and phase shift on a subcarrier-by-subcarrier basis.

Note that all of the information of relevance to the receiver is contained within the FFT period. Once the signal is received and digitized, the receiver simply throws away the CP. The result is a rectangular pulse that, within each subcarrier, is of constant amplitude over the FFT period.

The rectangular pulses resulting from decimation of the CP are central to the ability to space subcarriers very closely in frequency without creating ICI. Readers may recall that a uniform rectangular pulse (RECT function) in the time domain results in a SINC function ($\sin(x) / x$) in the frequency domain as shown in Fig. 2.2-3. The LTE FFT Period is 67.77 μsec. Note that this is simply the inversion of the carrier spacing ($1 / \Delta f$). This results in a SINC pattern in the frequency domain with uniformly spaced zero-crossings at 15 kHz intervals—precisely at the center of the adjacent subcarrier. It is therefore possible to sample at the center frequency of each subcarrier while encountering no interference from neighboring subcarriers (zero-ICI).

Figure 2.2-3 FFT of OFDM Symbol Reveals Distinct Subcarriers



2.2.1 Disadvantages of OFDM

As we have seen, OFDM has some remarkable attributes. However, like all modulation schemes, it suffers from some drawbacks. OFDM has two principle weaknesses relative to single carrier systems: susceptibility to carrier frequency errors (due either to local oscillator offset or Doppler shifts) and a large signal peak-to-average power ratio (PAPR).

As discussed above, OFDM systems can achieve zero-ICI if each subcarrier is sampled precisely at its center frequency. The time-sampled OFDM signal is converted into the frequency domain by means of a fast Fourier transform (FFT)—which is a highly efficient means of implementing a discrete Fourier transform (DFT). The DFT renders a discrete finite sequence of complex coefficients which are given by:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-j2\pi nk / N}, \quad k = 0, 1, \dots, N-1$$

The resulting Fourier spectrum has discrete frequencies at:

$$k/NT_s, \quad k = 0, 1, \dots, N-1$$

where T_s is the sample interval in the time domain and N is the number of samples. Thus, the frequencies in the Fourier representation are completely defined by the sample frequency ($1 / T_s$) and the number of samples taken within the FFT period.


Let's consider a specific LTE example. LTE defines transmission bandwidths from 1.25 MHz up to 20 MHz. In the case of 1.25 MHz transmission bandwidth, the FFT size is 128. In other words, 128 samples are taken within the FFT period of 66.67 μ sec. Therefore, $T_s = 0.52086 \mu$ sec, and the received signal is represented by frequencies at 15 kHz, 30 kHz, 45 kHz... These frequencies are the exact center frequencies of the signal subcarriers—unless frequency errors are encountered in the downconversion process.

The FFT is done at baseband frequency, after the received signal has been downconverted from the RF carrier frequency. Downconversion is typically performed by means of direct conversion. The received signal is mixed with a signal produced by the receiver's local oscillator (LO). Ideally, the carrier signal and the receiver LO are at the identical frequency. Unfortunately, this is not always the case.


The transmitter and receiver local oscillators will invariably drift, so active means must be taken to keep them synchronized. Each base station periodically sends synchronization signals which are used by the UE for this purpose, among other things (synchronization signals are also used for initial acquisition and handover). Even so, other sources

such as Doppler shifts and oscillator phase noise can still result in frequency errors. Uncorrected frequency errors will result in ICI as shown in Figure 2.2.1-1. For these reasons, the signal frequency must be tracked continuously. Any offsets must be corrected in the baseband processor to avoid excessive ICI that might result in dropped packets.

Figure 2.2.1.1 Uncorrected Frequency Errors Cause ICI



Demodulated Signal without Frequency Offset (Zero ICI)



Demodulated Signal with Frequency Offset Causing ICI

The other major drawback to OFDM is a high PAPR. The instantaneous transmitted RF power can vary dramatically within a single OFDM symbol. As stated above, the OFDM symbol is a combination of all of the subcarriers. Subcarrier

voltages can add in-phase at some points within the symbol, resulting in very high instantaneous peak power—much higher than the average power.

A high PAPR drives dynamic range requirements for A/D and D/A converters. Even more importantly, it also reduces efficiency of the transmitter RF power amplifier (RFPA). Single carrier systems sometimes use constant envelope modulation methods, such as Gaussian Minimum Shift Keying (GMSK) or Phase Shift Keying (PSK). The information in the signal of a single carrier system is conveyed by varying the instantaneous frequency or phase while the signal amplitude remains constant. The RFPA does not require a high degree of linearity. In fact, the PA can be driven so hard that the signal is “clipped” as the signal swings between the minimum and maximum voltages. Harmonic distortion due to clipping can be eliminated by output filtering. When RFPAs are operated in this manner, they can achieve efficiencies on the order of 70 percent.

In contrast, OFDM is not a constant envelope modulation scheme. Within each symbol, the amplitude and phase of each sub-carrier is constant. Over the duration of an OFDM symbol, there can be several large peaks. The RFPA must be capable of handling peak voltage swings without clipping, thus requiring a larger amplifier to handle a given average power. Efficiency is therefore lower. RFPA efficiencies for OFDM signals can be less than 20 percent. Although there are measures that can be taken to reduce voltage peaks, PAPR for OFDM results in RFPA efficiencies that are generally lower than for single-carrier constant envelope systems.


2.3 OFDMA

OFDMA is employed as the multiplexing scheme in the LTE downlink. Perhaps the best way to describe OFDMA is by contrasting it with a packet-oriented networking scheme such as 802.11a. In 802.11a, Carrier-Sense Multiple Access (CSMA) is the multiplexing method. Downlink and uplink traffic from the fixed access point (AP) to mobile user stations (STAs) is by means of PHY layer packets. As explained below, OFDMA makes much more efficient use of network resources.

2.3.1 Comparison of OFDMA with Packet-Oriented Protocols

Like 3GPP LTE, IEEE 802.11a uses OFDM as the underlying modulation method. However, 802.11a uses CSMA as the multiplexing method. CSMA is essentially a listen-before-talk scheme. For example, when the AP has queued traffic for a STA, it monitors the channel for activity. When the channel becomes idle, it begins to decrement an internal timer that is randomized within a specified window. The timer will continue to be decremented as long as the network remains idle. When the timer reaches zero, the AP will transmit a PHY layer packet of up to 2000 bytes addressed to a particular STA (or all STAs within the cell in the case of broadcast mode). The randomized back-off period is designed to minimize collisions, but it cannot eliminate them entirely.

Figure 2.3.1-1 Conventional Packet Oriented Networks Like IEEE 802.11a Precede Each Data Transmission with a PHY Layer Preamble and Header



Each 802.11a PHY packet utilizes all of the PHY layer bandwidth for the duration of the packet. Consider the 802.11a PHY packet format shown in Figure 2.3.1-1. Each 802.11a packet has a data payload of varying length from 64 to 2048 bytes. If the packet transmission is successful, the receiving station transmits an ACK. Unacknowledged packets are assumed to be dropped. Note that each packet is preceded by a PHY preamble which is 20 μsec in duration. The purposes of the PHY preamble are:

- Signal detection
- Antenna diversity selection

- Setting AGC
- Frequency offset estimation
- Timing synchronization
- Channel estimation

The address of the intended recipient is not in the PHY preamble. It is actually in the packet data and is interpreted at the MAC layer. From a networking perspective, the packet-oriented approach of 802.11a has the advantage of simplicity. Each packet is addressed to a single recipient (broadcast mode notwithstanding). However, the randomized backoff period of the CSMA multiplexing scheme is idle time and therefore represents an inefficiency. The PHY preamble is also network overhead and further reduces efficiency, particularly for shorter packets.


The typical real-world efficiency of an 802.11a system is approximately 50 percent. In other words, for a network with a nominal data rate of 54 Mbps, the typical throughput is about 25 – 30 Mbps. Some of the inefficiencies can be mitigated by abandoning the CSMA multiplexing scheme and adopting a scheduled approach to packet transmission. Indeed, subsequent versions of the 802.11 protocol include this feature. Inefficiencies due to dedicated ACK packets can also be reduced by acknowledging packets in groups rather than individually.

In spite of potential improvements, it remains difficult to drive packet-oriented network efficiency much beyond 65 to 70 percent. Further, because each packet completely consumes all network resources during transmission and acknowledgement, the AP can provide addressed (non-broadcast) traffic to user terminals only on a sequential basis. When many users are active within the cell, latency can become a significant problem. Clearly, the objective of cellular carriers is to create as much network demand as possible for a wide variety of traffic that includes voice, multimedia, and data. Efficiency and low latency are therefore paramount. As we will see in the following section, OFDMA is superior to packet-oriented schemes in both of these critical dimensions.

2.3.2 OFDMA and the LTE Generic Frame Structure

OFDMA is an excellent choice of multiplexing scheme for the 3GPP LTE downlink. Although it involves added complexity in terms of resource scheduling, it is vastly superior to packet-oriented approaches in terms of efficiency and latency. In OFDMA, users are allocated a specific number of subcarriers for a predetermined amount of time. These are referred to as physical resource blocks (PRBs) in the LTE specifications. PRBs thus have both a time and frequency dimension. Allocation of PRBs is handled by a scheduling function at the 3GPP base station (eNodeB).

Figure 2.3.2-1 LTE Generic Frame Structure



In order to adequately explain OFDMA within the context of the LTE, we must study the PHY layer generic frame structure. The generic frame structure is used with FDD. Alternative frame structures are defined for use with TDD. However, TDD is beyond the scope of this paper. Alternative frame structures are therefore not considered.

As shown in figure 2.3.2-1, LTE frames are 10 msec in duration. They are divided into 10 subframes, each subframe being 1.0 msec long. Each subframe is further divided into two slots, each of 0.5 msec duration. Slots consist of either 6 or 7 OFDM symbols, depending on whether the normal or extended cyclic prefix is employed.

Table 2.3.2-1 Available Downlink Bandwidth is Divided into Physical Resource Blocks

Bandwidth (MHz)	1.25	2.5	5.0	10.0	15.0	20.0
Subcarrier bandwidth (kHz)	15					
Physical resource block (PRB) bandwidth (kHz)	180					
Number of available PRBs	6	12	25	50	75	100

The total number of available subcarriers depends on the overall transmission bandwidth of the system. The LTE specifications define parameters for system bandwidths from 1.25 MHz to 20 MHz as shown in Table 2.3.2-1. A PRB is defined as consisting of 12 consecutive subcarriers for one slot (0.5 msec) in duration. A PRB is the smallest element of resource allocation assigned by the base station scheduler.

Figure 2.3.2-2 Downlink Resource Grid

