FLO Physical Layer: An Overview

Murali R. Chari, Fuyun Ling, Ashok Mantravadi, Raghuraman Krishnamoorthi, Rajiv Vijayan, G. Kent Walker, and Rob Chandhok

***Abstract*—This paper provides an overview of the physical layer of the Forward Link Only (FLO) Air Interface. The FLO Air Interface is a key component of the MediaFLO system developed by QUALCOMM as an alternative mobile multicast technology for the efficient transmission of multiple multi-media streams to mobile devices using TV and multi-media channel bandwidths in VHF, UHF, or L-band. The main concepts and features of the FLO Air Interface including the modulation and coding techniques used, the frame structure, and the different sub-channels within the physical layer are described. The available data rates as well as other characteristics of FLO are also described. Finally, the performance of the FLO physical layer in representative channel environments is presented.**

***Index Terms*—Forward error correction, mobile communication, modulation, multicast channels, orthogonal frequency division multiplexing (OFDM).**

I. Introduction

In the last few years, multiple technologies have been proposed to address delivery of streaming multimedia to mobile devices, such as cell phones, handhelds, gaming consoles, etc. These technologies have typically leveraged upon either 3rd generation cellular/PCS [1], [2] or digital terrestrial TV broadcast technologies [3]–[5]. The former suffers from the problem that unicasting (one-to-one communication) is a spectrally inefficient and expensive way to solve this problem, while the latter approach is often found to be confined by the need to satisfy legacy constraints in the interest of backward compatibility.

In order to effectively address this problem, a new Air Interface has been developed based on Forward Link Only (FLO) technology for QUALCOMM's MediaFLO mobile multimedia multicast system. Since FLO technology was designed from ground up to be a multicasting network, which is overlaid upon a cellular network, it doesn't need to support any backward compatibility constraints. More specifically, the FLO physical layer targets transmission in the VHF/UHF/L-band frequency bands, over channel bandwidths of 5, 6, 7, and 8 MHz. Hence, both the network infrastructure and the receiver chain (in the device) are separate from those for the cellular/PCS network. Moreover, as the name suggests, the technology relies on the use of a forward link (network to device) only. FLO enables the efficient multicasting of multiple, multimedia services, including real-time (video/audio/tele-text), non real-time (i.e., clip casts, which are downloaded for later viewing), and IP datacast, to mobile (FLO) devices, e.g., cell phones, handhelds, gaming consoles, etc. An important characteristic of the FLO device is that it is operated using a limited battery source. Minimizing the power dissipation is a very important factor in order that the device can still be used for other tasks, e.g., receive or make cellular phone calls in addition to receiving FLO transmission.

FLO is targeted to achieve a capacity of 1 bit per second (i.e., 6 Mbps in a RF channel bandwidth of 6 MHz). Since the FLO device typically uses a small display, it is possible to achieve an average bit rate of 200–250 kbps for a real time video/audio service with the use of advanced compression techniques, such as H.264/AVC and its variants. Hence, FLO can support the transmission of about 20 real time services over a 6 MHz bandwidth.

The FLO system consists of two parts: (a) The FLO network, which includes the collection of transmitters and the backhaul network, and (b) The FLO device, which may be a cell phone, handheld, or gaming console. This paper is concerned with the air interface between the FLO network and the FLO device, and in particular provides an overview of the physical layer part of this interface.

The paper is organized as follows: In Section II, we describe the main concepts in the FLO physical layer. This serves to highlight the important points and features in the design without getting into details of the signal structure. These details are the subjects of Sections III and IV. The frame structure of the FLO waveform is described in Section III, while the operations required for the generation of each physical layer sub-channel is described in Section IV. Performance results for the FLO physical layer in different channel environments are provided in Section V, followed by some concluding remarks in Section VI.

II. Main Concepts

To achieve good receiver performance and high spectral efficiency of multimedia multicasting in a mobile communications environment, the FLO physical layer uses Orthogonal Frequency Division Multiplexing (OFDM) as the modulation technique. In addition, it incorporates advanced forward error correction techniques involving the concatenation of a parallel concatenated convolutional code (PCCC), also called Turbo code, and a Reed-Solomon erasure correcting code. Moreover, various parts of the physical layer have been carefully designed to further improve receiver performance and to ensure a most satisfactory user experience. These main concepts incorporated in the FLO air interface design are the subject of this section.

A. Single Frequency Network (SFN) Operation

To ensure that the user experience is as uniform as possible over the entire coverage area and optimize spectral efficiency

Manuscript received July 13, 2006; revised November 30, 2006.

M. R. Chari, F. Ling, A. Mantravadi, R. Krishnamoorthi, R. Vijayan, and G. K. Walker are with QUALCOMM Inc., San Diego, CA 92121 USA (e-mail: mchari@qualcomm.com; fling@qualcomm.com; ashokm@qualcomm.com; raghuram@qualcomm.com; rvijayan@qualcomm.com; gwalker@qualcomm.com).

R. Chandhok resides in Poway, CA 92064 USA (e-mail: rob@chandhok.net).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TBC.2007.891696

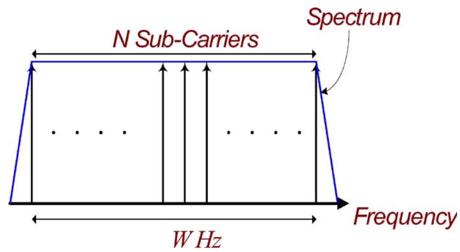

Fig. 1. Spectrum of an OFDM signal.

and network economics, FLO employs the concept of SFN operation [6], [7]. In such a network, the signals from multiple transmitters carry the same content and transmit identical waveforms. As a result, they can be viewed by the receiver as if they are signals from the same source with different propagation delays. In such a case, a FLO device receives the combined signal from the near-by transmitters. The FLO receiver demodulates the combined signal from multiple surrounding transmitters, ensuring a smooth transition from the coverage area of one transmitter to the next in an SFN. As a result, no explicit hand-off operation is necessary. On the other hand, under such SFN operations, the delay spread of the equivalent communication channel will be quite large, and equalization will be an essential component of the receiver. In order to perform effective equalization of the signal transmitted over such channels, FLO uses OFDM as its modulation technique.

B. OFDM

OFDM is a form of multi-carrier modulation [8]. Roughly speaking, the available bandwidth is divided into N bins, referred to as sub-carriers, with each sub-carrier modulated by a quadrature amplitude modulated (QAM) symbol, as is illustrated in Fig. 1. The bit stream to be transmitted is divided into N sub-streams, with each sub-stream having a fraction of the overall rate and being assigned to 1 of the N sub-carriers. Therefore, the OFDM symbol duration is nominally increased by a factor of N compared to the symbol rate of the composite bit stream, thereby making OFDM very resilient to operation even in channels with large multi-path delay spreads. Consequently, OFDM enables the achievement of higher capacity, compared to single-carrier modulation methods. Moreover, for high data rates, the implementation of equalization in single-carrier systems may not be practically feasible. For these reasons, the majority of standards for digital, terrestrial TV transmission also employ OFDM as the modulation technique [5], [9].

In FLO, transmission and reception are based on using 4096 (4K) subcarriers and the QAM modulation symbols are chosen from a QPSK or 16-QAM alphabet. For an RF channel bandwidth allocation of 6 MHz, the 4K subcarriers span a bandwidth of 5.55 MHz, which is also referred as the *chip* rate. The chip rates for 5, 7, and 8 MHz allocations are 4.625, 6.475, and 7.4 MHz, respectively. We note that the chip rate is smaller than the allocated RF bandwidth. For the 6 MHz case, this choice of chip rate enables meeting the regulatory (e.g., from the FCC in the United States) transmit spectral mask, which typically has very stringent requirements on side-bands outside the allocated bandwidth. Henceforth, for the purpose of this paper, we assume 6 MHz RF channel bandwidth unless stated otherwise.

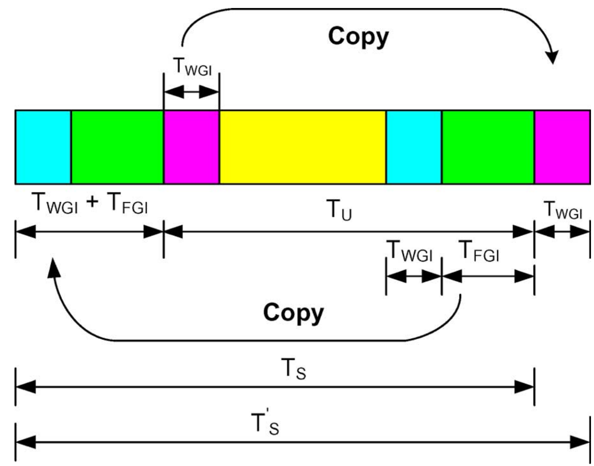

Fig. 2. OFDM symbol structure-time domain.

The *sub-carrier spacing*, $(\Delta f)_c$, is given by

$$\Delta f_c = 5.55\ \text{MHz}/4096 = 1.35498\ \text{KHz}$$

The sub-carrier spacing places limits on the receiver performance at high vehicular speeds. High Doppler results in inter-carrier interference (ICI) which can lead to packet error rate (PER) floors at high SNR. In the case of FLO, we desire any PER floors to be below 1% (see Section V). With the above sub-carrier spacing, this is ensured for all choices of constellation and code rate up to a mobile speed of 120 kph. For certain constellations and code rates, it is possible to operate at much higher vehicular speeds.

The frequency location of the 4K subcarriers at base-band is given by

$$f_i = (-2048 + i) \times \Delta f_c, \quad i = 0, 1, \ldots, 4095,$$

where $(\Delta f)_c$ is the sub-carrier spacing. Of the 4K subcarriers, 96 are unused and are referred to as guard subcarriers. The indices corresponding to the guard subcarriers are 0,. . ., 47, 2048, and 4049,. . ., 4095, where the sub-carrier 2048 corresponds to DC and is not transmitted. The remaining 95 guard subcarriers on the two sides serve the following two purposes in the device: (1) The base-band receive filter can have a larger transition band for meeting the adjacent channel interference (ACI) rejection specifications. (2) The power levels of noise and rejected ACI that is aliased within the signal bandwidth is reduced. The remaining 4000 subcarriers are referred to as active subcarriers, which are modulated by data or pilot symbols.

In the time-domain, each OFDM symbol consists of a number of (base-band) samples, called OFDM chips. These chips are transmitted at a rate of 5.55 MHz. The *total* OFDM symbol interval T'_S is comprised of four parts: a useful part with duration T_U, a flat guard interval with duration T_{FGI} and two windowed intervals of duration T_{WGI} on the two sides, as illustrated in Fig. 2. In the literature, the useful interval and the flat guard interval are also referred to as the *FFT interval* and the *cyclic prefix,* respectively. The total OFDM symbol interval

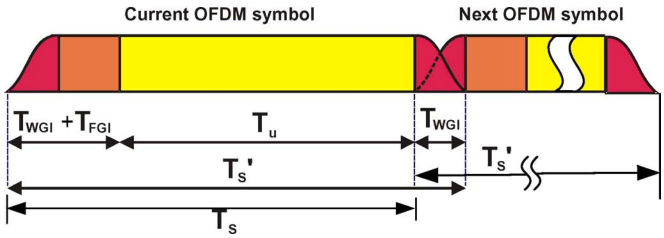

Fig. 3. Overlap of windowed OFDM symbols.

then has duration $T'_S = T_U + T_{FGI} + (2 \times T_{WGI})$, where $T_U = 4096$ chips, $T_{FGI} = 512$ chips and $T_{WGI} = 17$ chips. Hence, $T'_S = 4642$ chips.

As shown in Fig. 3, there is an overlap of T_{WGI} between consecutive OFDM symbols. Hence, the *effective* OFDM symbol interval is defined to have a duration $T_S = T_U + T_{FGI} + T_{WGI} = 4625$ chips. The effective OFDM symbol interval is also referred to as the OFDM symbol interval. For FLO, with 6 MHz RF bandwidth, it is equal to $833.33\ldots\mu$s.

For FLO, the cyclic prefix is a fixed fraction (1/8th) of the useful interval of the OFDM symbol. The significance of the cyclic prefix in an SFN environment is that as long as the received paths from multiple transmitters have a differential distance of less than 27.7 km, the resulting multipath channel will not lead to either inter-symbol interference (ISI) or inter-carrier interference (ICI) at the receiver [10]. Furthermore, in order that the performance degrade gracefully for channels with multipath delay spreads greater than the cyclic prefix, the pilot sub-carrier indices are *staggered* (see Section IV-C-8 in consecutive OFDM symbols, with the number of pilot sub-carriers in each OFDM symbol closely matched to the cyclic prefix duration in chips. This enables a FLO receiver to use the pilot (sub-carrier) observations from multiple OFDM symbols in order to estimate channels with delay spreads up to two times the duration of the cyclic prefix. The selection of cyclic prefix duration and pilot subcarrier structure is based in part on coverage studies in SFNs using high power[1] transmitters, at UHF, that also accounted for receiver degradations (e.g., due to ISI/ICI).

The purpose of the windowed portions of the guard intervals is to improve the spectral mask of the base-band OFDM signal by attenuating the side-bands, over that obtained by using a rectangular window function. This is beneficial in reducing the attenuation requirements of the filtering after the (high) power amplifier at the transmitter and thus simplifies spectral planning in shared use RF allocations.

The main OFDM parameters in FLO for 6 MHz[2] RF channel bandwidths are summarized in Table 1.

Finally, it is noteworthy that for mobile multicasting applications, the choice of 4096 subcarriers in FLO is more appropriate than two alternative choices, namely 2048 and 8192 for FLO's intended application environment. For transmission in the UHF/VHF bands, using 8K sub-carriers impairs receiver performance at high speeds due to increased sensitivity to Doppler impairments. Also, the receiver and its memory size is almost doubled compared to using 4K subcarriers, since the FFT block accounts for most of the receiver size. On the other hand, using 2K subcarriers would imply having to double the overhead due to the cyclic prefix in order to support the same transmitter spacing as a system using 4K subcarriers.

TABLE 1
FLO OFDM PARAMETERS

Parameter	Value
Number of sub-carriers	4096
FFT bandwidth (Chip rate)	5.55 MHz
Sub-carrier spacing	1.355 KHz
FFT interval	738.02 us
Cyclic Prefix	92.25 us
Window interval	3.06 us
OFDM symbol interval	833.33 us
Chip duration	0.18 us
Number of Guard sub-carriers	96
Number of Pilot sub-carriers	500

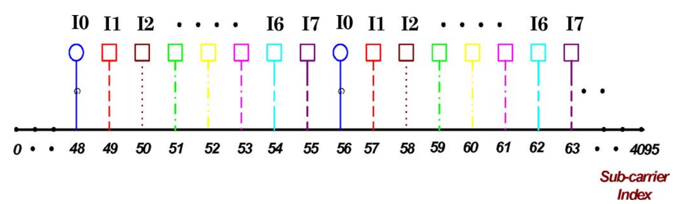

Fig. 4. FLO interlace structure.

C. Interlace Structure

As stated above, in each FLO OFDM symbol, there are 4000 active subcarriers. These active subcarriers are further equally divided into eight disjoint groups called interlaces. An interlace consists of 500 subcarriers that are evenly spaced across the FLO signal bandwidth, as illustrated in Fig. 4. The interlace index ranges from 0 through 7. Note that, between the adjacent subcarriers within each interlace, there are 7 subcarriers, each of which belongs to a different interlace. In each OFDM symbol, either interlace 2 or 6 is assigned to the FDM Pilot and is used for channel estimation [11]. Hence, 500 of the active subcarriers are modulated with known (pilot) modulation symbols. The remaining 7 interlaces, or 3500 subcarriers, are available for modulation with data symbols. In FLO, both the data and pilot subcarriers are modulated with symbols that have the same energy (see Section IV-D).

The main advantages of the interlace structure are:

1) It enables the frequency-division multiplexing of FLO logical channels, referred to as Multicast Logical Channels (MLCs), which are described below, within each OFDM symbol without the loss of frequency diversity. The minimum frequency allocation to an MLC, within a single OFDM symbol, is an interlace. Hence, at most 7 MLCs can be multiplexed within a single OFDM symbol. Since, the subcarriers within an interlace span the total FLO signal bandwidth there is no loss of frequency diversity, compared to the case where all the active subcarriers are used.

[1]The maximum Effective Radiated Power (ERP) used was 50 kW.

[2]For alternate RF channel bandwidths, the values for sub-carrier spacing and time duration(s) in Table 1. are easily derived from the corresponding (scaled) chip rate.

2) It enables the transmission of MLCs with finer granularity. For transmission at high spectral efficiency, tens of kbits can potentially be transmitted within a single OFDM symbol. Hence, having the ability to allocate a fraction of the subcarriers to MLCs enables supporting low data rate MLCs without incurring a large overhead expense.
3) The interlace structure is also beneficial from a receiver power consumption point of view. The FFT block in the receiver can be designed such that only the required subset of interlaces, corresponding to the desired MLCs, are demodulated. Hence, when combined with the frequency multiplexing of MLCs, the receiver need not always be performing a 4096-point FFT, thereby saving on power consumption.

D. Multicast Logical Channels (MLCs) and Statistical Multiplexing

A FLO-based system is capable of multicasting various services, such as live video and audio streams (News, Music, or Sports channels), subtitles or stock-quotes for broadcast data. A service can be viewed as an aggregation of one or more related data components, such as the video, audio, text or signaling associated with a service. Furthermore, services are classified into two types based on their coverage: Wide-area services and Local-area services. A local-area service is multicast for reception within a metropolitan area. By contrast, Wide-area services are multicast in one or more metropolitan areas.

Each FLO service is carried over one or more logical channels, referred to as Multicast Logical Channels or MLCs. An MLC has the attribute that it contains one or more decodable subcomponents of a service that is of independent reception interest. Furthermore, an important aspect is that MLCs are distinguishable at the physical layer. For example, the video and audio components of a given service can be sent on two different MLCs. A device that is interested in the audio component only can receive the corresponding MLC without receiving the MLC for the video component, thereby saving on battery resources.

The data rates required by these services are expected to vary over a wide range, depending on their multimedia content. While low to moderate data rates, i.e., tens of kbps, are sufficient for data and audio streams, video streams may require instantaneous rates ranging from a few kbps to a few Mbps even though the average rate is in the range of 200–300 kbps. Thus, effective use of statistical multiplexing [12] can significantly increase the number of services supported by a multicast system using a specified channel bandwidth.

In FLO, statistical multiplexing of different services, or MLCs, is achieved by varying *only* the MLC time and frequency allocations over prescribed time intervals to match the variability in the MLC's source rates. We *exclude* the possibility of varying the constellation and code rate assigned to an MLC in order to *maintain a constant coverage*[3] *area* for each MLC. Specifically, MLCs are transmitted over a certain number of OFDM symbols to achieve Time-Division Multiplexing (TDM) and a subset of the interlaces in these OFDM symbols to achieve Frequency-Division Multiplexing (FDM). Moreover, the time and frequency allocations for each MLC vary over the duration of transmission of the MLC. It is important to note that mobile devices are *constantly updated* on the time and frequency allocations of the MLCs so they know when to "wake up" to receive a specific MLC.

Furthermore, a video or audio stream can be sent in two layers, i.e., a *base* (B) layer that enjoys reception over a wide area and an *enhancement* (E) layer that improves the audio-visual experience provided by the base layer over a more limited coverage area, as discussed below.

However, we *restrict* the base and enhancement layers of a given service to be sent within a *single* MLC.

In contrast with other digital TV transmission standards, for example DVB-T and ISDB-T, the implementation of statistical multiplexing in FLO enables the receiver to demodulate and decode *only* the MLC(s) of interest. In the case of these other standards, the receiver must demodulate and decode the entire transmission, and only then de-multiplex the signal portion of interest. Hence, using FLO provides significant savings in power consumption due to the reduced receiver "on" time, which is very important for reception when using mobile units with limited battery power.

The choice of constellation and code rate for each MLC is based on various factors, including the service (wide-area/local-area) area, the content (video/audio/data), coverage requirements and whether layered modulation is used. FLO provides several choices for constellation and code rate that allow a service provider to tradeoff spectral efficiency against coverage. Furthermore, for any choice of constellation and code rate, the variable time and frequency allocations can be used to scale the achievable data rates.

E. Wide and Local Coverage

The transmitters in a FLO network are assumed to be distributed over a wide geographical area, e.g., the west coast of the continental United States. Since, the typical transmit power for the TV bands in UHF/VHF are assumed to be high, e.g., 50 kW ERP, it is anticipated that transmitter spacing can be fairly large, e.g., 30–40 km. In this setting, a few transmitters (2–3) may be sufficient to cover a metropolitan market area.

All transmitters within the FLO network are assumed to be synchronized in transmit time and carrier frequency. The term Local-area is used to denote the transmitters within a metropolitan area. Transmitters in one or more metropolitan areas that multicast the same Wide-area services constitute a Wide-area. Hence, a Wide-area consists of one or more Local-areas, with the transmitters in the different Local-areas multicasting different Local-area services and possibly using different RF center frequencies. As a result, the transmit waveform is not identical for all Wide-area transmitters in their local portion of transmission. Finally, transmitters in different Wide-areas also transmit different waveforms, since they multicast different Wide-area and Local-area services . An example of a FLO network consisting of two Wide-areas and four Local-areas is illustrated in Fig. 5.

[3]By coverage here, we mean the fraction of locations in a geographical region where a FLO device can receive service, assuming a specific deployment of FLO transmitters in that region. Reception of a service, in turn, can be based on the packet error rate (PER) exceeding a specific threshold, e.g., 1% for real-time video and audio services. The transmit mode determines the lowest Signal-to-Noise Ratio (SNR) at which the PER threshold can be achieved, thereby impacting coverage.

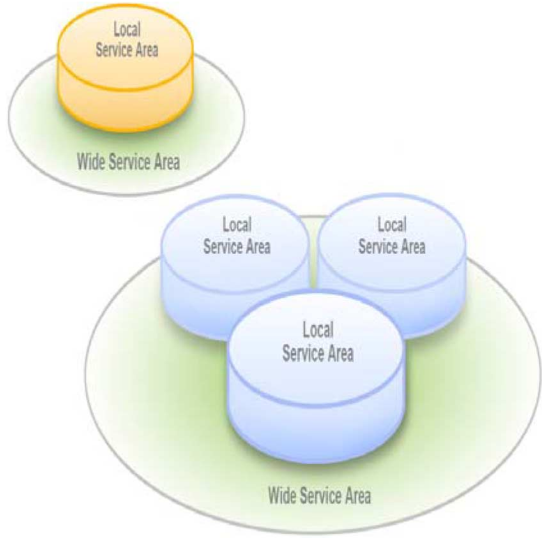

Fig. 5. Wide and local coverage areas.

Note that, in the case of FLO, within a local or wide area, the notion of an SFN may apply, while at the boundaries between local (or wide) areas there might be interference between neighboring transmissions. The case of different local-area transmissions occupying different RF center frequencies is commonly referred to as multi-frequency network (MFN). Even in a MFN scenario, it is possible for a FLO device to receive a wide-area service without interruption when moving between neighboring local-areas, implying "soft" handoff capability.

F. Forward Error Correction Coding

It is well known [8] that in order to gain the benefits of OFDM over time/frequency-selective channels the use of forward error correction is mandatory. The FLO design is based on the use of a concatenated coding scheme, consisting of an outer Reed-Solomon (RS) code and an inner parallel concatenated convolutional code (PCCC), also called turbo code.

More specifically, the outer code consists of an (N, K) Reed-Solomon code over the Galois Field with 256 elements, $GF(256)$, and is intended for erasure-correction. The value of N is fixed at 16, while the value of K can be chosen from the set $\{8, 12, 14, 16\}$. The case of $K = 16$ corresponds to the case when no RS encoding is actually performed. For the case of $K = 12$, 14, the generator polynomial and parity check matrix are identical to those in the EV-DO Gold Multicast (BCMCS, TIA-1006) standard [13].

For MLCs containing a base and enhancement layer, the encoding is done independently for each layer.

As shall be discussed below, when an MLC is transmitted, a minimum number of packets must be sent within a specified time unit called a superframe. Briefly speaking, a superframe has the duration of exactly 1 second, and consists of 4 frames of equal duration, each roughly 1/4th of a second. These packets are first RS-encoded and then Turbo-encoded. They are referred to as MAC layer packets. During the Reed-Solomon encoding process, $N - K$ parity packets are generated for every K information packets. CRC bits are generated for each of the N packets. The packets with data and CRC bits are Turbo encoded and transmitted. Thus, the minimum number of information packets of an MLC that can be transmitted in a superframe is K. The collection of K information packets and $N - K$ parity packets is referred to as an RS, or outer, code block. Finally, MLC transmissions in each superframe are always in integer multiples of outer code blocks.

During transmission, *each* RS code block is *split into 4 equal sub-blocks*, with *each sub-block sent in a unique frame* within a super-frame. The main purpose of utilizing RS-coding is to exploit the time diversity of the packets within a superframe. The time span of the packets of an RS code block is at least 0.75 seconds. Such a time span ensures decorrelation of these packets even at low vehicle, or pedestrian, speeds and improves the robustness of reception in impulsive noise environments.

The rationale for choosing $N = 16$ and the allowed range for K should now be apparent. Since, an MLC is always transmitted in integer units of code blocks, it is desirable to choose N small in order to enable transmission of MLCs with low granularity. Moreover, the value of N directly dictates the complexity of RS decoding in the device. A smaller value translates into lower complexity, which is beneficial from the viewpoint of device implementation. Finally, link-level simulations showed that the smallest value of N that guarantees acceptable performance in a variety of reception scenarios is 16. With regard to choosing K, the values 8, 12 enable recovery of the code block even when entire[4] sub-block(s) sent in a frame are received in error after turbo decoding. These choices for K are especially advantageous for reception at pedestrian speeds. Choosing $K = 14$ can provide improved performance at high mobile speeds, compared to the case when no RS encoding is performed, i.e., $K = 16$.

The MAC information or parity packets are Turbo coded. The structure of the encoder is the same as that used in the CDMA 2000 1X and 1X EV-DO standards [1], [2]. In FLO, the code rates used are 1/5, for transmitting critical overhead information, and {1/3, 1/2, 2/3} for transmitting MLCs. The higher code rates are obtained from the base code rate using puncturing.

The inner code exploits the frequency-diversity inherent in the channel. Compared to convolutional coding, it is well known that a system employing Turbo coding requires lower signal to noise ratio (SNR) and, thus, has a higher system capacity (more bits per Hertz). This advantage is especially significant for an OFDM system when the channel has spectral nulls, which are likely to occur in an SFN environment.

G. Regular and Layered Subcarrier Modulation

Each active subcarrier is modulated with a symbol chosen from a QAM constellation. In FLO, the choices are: (1) Regular (non-layered): the constellation could be QPSK and 16-QAM. (2) Layered: the constellation is one of the generalized 16-QAM constellations. Higher order constellations, e.g., 64-QAM, are not used in FLO because of the higher SNRs typically required for reception as well as the need for larger receiver implementation margins.

In FLO the notion of layering is supported across different layers of the associated protocol stack. Hence, at the application layer, the source encoders produce 2 components, base and enhancement layer bit streams, associated with the same service.

[4]More specifically, this implies that reception of the FLO signal is possible even in the presence of a (signal) fade or an impulse (burst) noise of duration 500 ms or 250 ms for K equal to 8 or 12, respectively.

The devices that can only reliably receive the base layer bits can provide the user this service, albeit at lower quality. A device that can reliably receive both base and enhancement layer bits can provide the same service at better, i.e., enhanced, quality.

These two components are independently processed in the Upper layers of the FLO Air Interface. Both the outer and the inner encoding operations are performed independently on the bit streams associated with the two components, with the encoded, interleaved bits streams finally getting combined when determining the constellation point. The modulation symbol is either chosen from a uniform or non-uniform 16-QAM constellation. In either case, the base layer bits are mapped to the more reliable bits of the constellation and the enhancement layer bits are mapped to the less reliable bits of the constellation. As a result, the base layer bits can be demodulated more reliably than the enhancement layer bits.

Such utilization of layered modulation is particularly suitable and advantageous for digital multicasting applications, where the signal strength is not evenly distributed in the entire coverage area. It provides more flexibility in coverage and link budget planning. Specifically, we don't have to design the coverage based on the worst case link budget. For example, we can let users in 90% of the coverage area receive the enhancement layer signal, while those in the remaining area can receive only the base layer signal. In such a manner, it is possible to improve the coverage or spectrum efficiency without increasing transmitter power while still providing an acceptable user experience.

H. Spectral Efficiency

The spectral efficiency of a transmission scheme is defined as the number of *information* bits transmitted per second per unit of bandwidth (Hz). For FLO, the calculation of this value is complicated by the fact that different MLCs are transmitted using only a fraction of the available time-frequency space. They may also use different constellations and code rates. The computation is therefore simplified by assuming that all MLCs use the same constellation and code rate and by disregarding the overhead incurred because of the outer code. This provides a valuable benchmark for each choice of constellation and code rate, from which the *actual* MLC data rate is obtained by scaling the benchmark with the average "on" time and bandwidth occupied by the MLC.

Based on the above assumptions, the *spectral efficiency*, η, is easily computed for the different constellations and code rates. For FLO, η is given by:[5]

$$\eta = \left(\frac{\text{Number of Information Bits per OFDM Symbol}}{\text{OFDM Symbol Duration} \times \text{Channel Bandwidth}}\right) = (N_{data} \times b_{\text{mod}} \times r_{inner}) \times \left(\frac{1}{T_{sym}}\right) \times \left(\frac{1}{6.0}\right)$$

In the above equation, b_{mod} is the number of bits needed to represent the modulation symbol on each subcarrier. For QPSK, $b_{\text{mod}} = 2$, while for 16-QAM, $b_{\text{mod}} = 4$. For a layered constellation the above equation computes the spectral efficiency for *each* layer (with $b_{\text{mod}} = 2$). The values for N_{data}, T_{sym}, and the channel bandwidth are known and fixed. Hence η is a function of two parameters, namely b_{mod} and r_{inner}, and is proportional to the product of those parameters.

TABLE 2
FLO DATA RATES, SPECTRAL EFFICIENCY, AND AWGN PERFORMANCE

Mode	Modulation	Inner Code Rate	Data Rate (Mbps)	Spectral Efficiency (bps/Hz)	Required C/N (dB) for AWGN
0	QPSK	1/3	2.8	0.47	-0.4
1	QPSK	1/2	4.2	0.7	1.8
2	16-QAM	1/3	5.6	0.93	4.5
3	16-QAM	1/2	8.4	1.4	7.3
4	16-QAM	2/3	11.2	1.86	10.0
5	QPSK	1/5	(for OIS* only)	(for OIS* only)	-3.0
6	Layered (ER = 4.0)	1/3	5.6	0.93	1.5(B), 6.6 (E)
7	Layered (ER = 4.0)	1/2	8.4	1.4	4.8(B), 9.0(E)
8	Layered (ER=4.0)	2/3	11.2	1.86	8.3(B), 11.5(E)
9	Layered (ER = 6.25)	1/3	5.6	0.93	0.8(B), 7.8(E)
10	Layered (ER = 6.25)	1/2	8.4	1.4	3.6(B), 10.5(E)
11	Layered (ER=6.25)	2/3	11.2	1.86	6.6(B).,12.6(E)

* ***OIS: Overhead Information Symbols (see Section III)***

Table 2 lists the different possible combinations of b_{mod} and r_{inner}. For each combination,[6] the spectral efficiency, the data rate (in a 6 MHz channel bandwidth), and the required Carrier-to-Noise ratio (C/N) for operation[7] in the Additive White Gaussian Noise (AWGN) channel is listed. The required C/N is based on the packet error rate without R-S coding, and is reflective of the inner turbo code performance. In the following sections, each row in Table 2 is referred to as a FLO *mode* and is referenced by its corresponding row number, which ranges from 0–11. Modes 0–5 are referred to as the *regular* modes, while modes 6–11 constitute the *layered* modes.

For the layered modes, the required C/N values for the base and enhancement layers are indicated using "(B)" and "(E)", respectively. The constellation is chosen from one of two options, distinguished by their energy ratio. The *energy ratio (ER)* is defined as the ratio of the energy of the base component to that of the enhancement component. For modes 6–8, the energy ratio is 4.0 and the constellation is identical to that for 16-QAM. For modes 9–11, the energy ratio is 6.25 and the constellation is a non-uniform 16-QAM constellation. The data rate and spectral efficiency of the layered modes is cumulative over base and enhancement layers.

From Table 2, note that a *minimum* spectral efficiency of 0.47 bits/sec/Hz is achieved with regular QPSK of code rate

[5]This equation computes the spectral efficiency of the physical layer only, and hence ignores the overhead due to padding of MAC layer packets, TDM pilots and overhead OFDM symbols and RS encoding.

[6]Since the row Mode 5 depicts a combination used only for the Overhead Information Symbols (OIS) OFDM symbols and not data MLCs, we do not list the spectral efficiency and corresponding data rate.

[7]A detailed discussion of the performance requirements for FLO operation and results in other channel scenarios are presented in Section V.

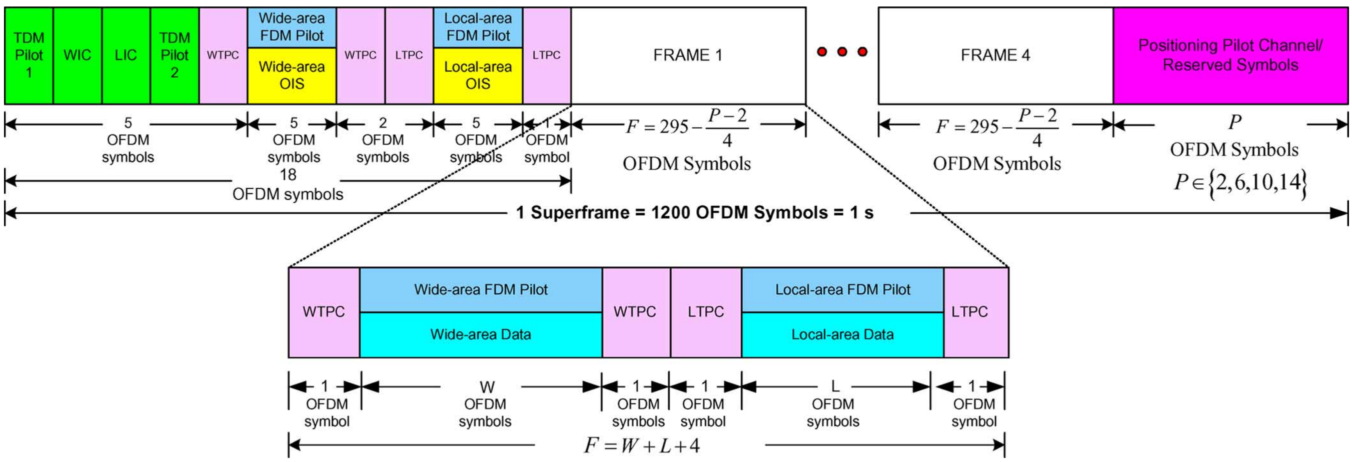

Fig. 6. FLO superframe structure.

1/3, while a *maximum* spectral efficiency of 1.86 bits/sec/Hz is achieved with either regular 16-QAM and code rate 2/3, or a layered constellation with a code rate of 2/3 on both the base and enhancement layers. These numbers imply that the physical layer can support data rates ranging from 2.8–11.2 Mbps in a 6 MHz channel bandwidth. Similarly, data rates ranging from 3.7–14.9 Mbps in a 8 MHz channel bandwidth can be supported.

III. Superframe Structure

In this section, we describe the FLO superframe structure in detail and discuss how its design contributes to a number of FLO's desired properties and advantages over competing technologies.

As mentioned in Section II-A, transmitters in the FLO network are synchronized with respect to transmit time and carrier frequency. This network synchronization can be achieved, for example, by using GPS. Transmit times should be synchronized to within a small fraction (2%) of the cyclic prefix duration of 92.25 μs. On the other hand, the maximum allowed carrier frequency offset is less than $\pm 1 \times 10^{-9}$ of the RF carrier frequency.

At each transmitter, the transmitted data are organized as *super-frames*, whose time duration is *exactly* 1 second. Since GPS is the most likely candidate for synchronizing multiple transmitters and GPS receivers produce a 1 pulse-per-second (pps) signal, timing synchronization of FLO signals from multiple transmitters can be achieved by simply aligning the super-frame boundary with the GPS 1 pps signal at each transmitter. This reduces the overhead information that must be sent over the backhaul[8] link in order to synchronize these transmitters.

Since the OFDM symbol duration is exactly 0.833 ..., ms, each superframe consists of 1200 OFDM symbols for 6 MHz RF bandwidth.[9] The structure of the super-frame is shown in Fig. 6. A superframe is divided into four portions: TDM pilots, OIS, data frames and Positioning Pilot Channel (PPC).

There are four TDM pilot OFDM symbols at the beginning of each super frame: TDM pilot 1 (TDM1), Wide-area Identification Channel (WIC), Local-area Identification Channel (LIC), and TDM pilot 2 (TDM2). TDM1 is the first symbol of a superframe and marks its beginning. In addition to being used for frame synchronization, TDM1 can also be used for initial time and frequency synchronization. Following TDM1 are the WIC and LIC symbols, which identify Wide and Local areas that correspond to the receiver location. The last symbol in this portion is TDM2. The main purpose for TDM2 is to provide more precise timing information so the receiver can immediately start decoding the information carried by the OIS.

Following the TDM pilot symbols are two sections of OIS, each consisting of five OFDM symbols that carry overhead information regarding the wide and local data channels. Specifically, they contain the time-frequency allocation for each MLC in the *current* super-frame.

Following the OIS portions are four data sections, called *frames*, of equal duration. When an MLC is transmitted in a super-frame the payload is divided into four equal bursts, with each burst transmitted in a *unique* frame. Each burst then consumes part of the time-frequency resource available within each frame. The time allocation is in units of OFDM symbols, while the frequency allocation is in units of interlaces. In FLO, an imposed simplification is that the time-frequency allocation, i.e., the collection of OFDM symbols along with the number of interlaces corresponding to each of these OFDM symbols, for an active[10] MLC is *identical* in the frames *within* a super-frame. However, the time-frequency allocation for an MLC can change from one superframe to the next.

As mentioned in Section II-E, FLO supports the transmission of both wide-area and local-area services. Because, a wide-area may consist of multiple local-areas, and there is the possibility of interference between transmissions received at the boundary between neighboring local-areas, the waveforms corresponding to the two types of services are time-division multiplexed. This enables the independent optimization of the transmit waveforms intended for the different coverage areas. Hence each frame is subdivided into two parts, as shown in Fig. 6. The first part is referred to as the Wide-area Data Channel

[8]The backhaul link is necessary in order to move content from a central NOC (Network Operations Center) to each transmitter site.

[9]For {5, 7, 8} MHz RF bandwidths, there are {1000, 1400, 1600} OFDM symbols, respectively, in each superframe.

[10]An MLC is said to be active in a particular superframe if data corresponding to that MLC is transmitted in the superframe.

and is dedicated to the transmission of wide-area services, and the second part is referred to as the Local-area Data Channel and is used solely for the transmission of local-area services. Correspondingly, the OIS is subdivided into two equal parts, as discussed above. The percentage of capacity allocated to the wide-area (or local-area) data channel can vary from 0 to 100%. Although the percentage can be set in every superframe, it is expected to vary infrequently.

In FLO, the available time-frequency (channel) resources are allocated once for both the wide-area and the local-area MLCs in each superframe. The OIS channels preceding the data frames contain the time-frequency allocation for each MLC in these frames. As mentioned earlier, the constellation and code rate are *typically not* altered[11] while transmitting an MLC; however, the time-frequency allocations for an MLC *typically change* from one superframe to the next.

With regard to the super-frame duration, there are several consequences to making the super-frame size large. The consequences are

(a) Though the *instantaneous* source rates of individual MLCs vary with time, the aggregate payload from all MLCs transmitted in a super-frame is *approximately* constant; the approximation becoming more exact with increasing super-frame duration. This yields greater statistical multiplexing gain;
(b) Larger time-diversity across frames is gained from larger frame sizes;
(c) Larger buffer sizes and longer latencies for source (video/audio) decoding;
(d) Longer acquisition time for the mobile when switching between MLCs.

A superframe duration on the order of 1 second is a good compromise. It enables deriving the benefits of (a), (b) without incurring significant penalty with (c), (d). More specifically, FLO can transmit tens of MLCs in each super-frame, thereby deriving the benefits of statistical multiplexing (a). Moreover, since FLO targets mobile users whose speeds can be as low as 3 kph, the coherence time[12] of the channel can be as large as 212 ms. Hence, a frame duration of 250 ms allows de-correlating the transmissions in multiple frames, providing time-diversity gains (b). It should be noted that such time diversity gain is achieved independent of the MLC's data rate. Since the instantaneous rates for video streams can be as large as 1 Mbps, a buffer size of 1 Mbits, which is feasible in mobile devices, is needed. Finally, with a superframe duration of 1 second, the average switching time between MLCs is 1.5 seconds, which is reasonable in terms of matching user expectations.

IV. Physical Layer Sub-Channels

A. TDM Pilot Channels

The first four symbols in a superframe are all TDM pilot channels that are used primarily for initial acquisition, which involves the estimation of local oscillator frequency offsets and receiver timing, i.e., the position of the FFT window, as well as determining the local and wide areas for the receiver location. These channels, i.e. TDM1, TDM2, WIC and LIC channels, share a common property: they are all periodic within an OFDM symbol. Equivalently, only part of the active subcarriers in an OFDM symbol have non-zero magnitude and there are an equal number of zeros between two non-zero subcarriers. Below, we provide a detailed description of these channels.

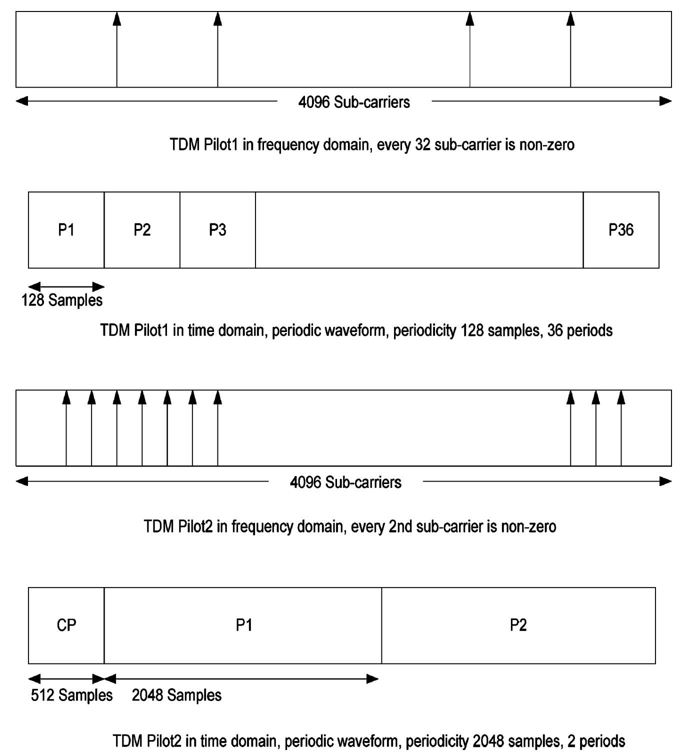

Fig. 7. Time and frequency domain characteristics of TDM1 and TDM2.

1) TDM Pilot 1 (TDM1) Channel: TDM1 is the first OFDM symbol in each superframe. It consists of 36 periods, each of which is 128 chips long. Note that 32 of the 36 periods correspond to the FFT duration of 4096 chips, and the remaining four correspond to the cyclic prefix of 512 chips. In the frequency domain, 124 out of 4000 active subcarriers are non-zero and there are 31 zeros between any two adjacent non-zero subcarriers (see Fig. 7). The non-zero subcarriers are modulated by QPSK symbols derived from a PN-sequence.

TDM1 can be used for: (a) superframe synchronization; (b) initial LO frequency offset estimation; and (c) coarse timing determination.

2) Wide-Area Identification Channel (WIC) and Local-Area Identification Channel (LIC): Following TDM1 are one WIC symbol and one LIC symbol. They are used by the receiver to determine the wide and local areas that correspond to the FLO device location. Both WIC and LIC symbols are similar in structure. They both have one non-zero interlace, i.e., the interlace with index 0, and the other 7 interlaces are inactive. In other

[11]Parameters that do not vary on a super-frame basis are not sent in the OIS. Instead, they are sent in designated MLCs, referred to as the Control Channel.

[12]At an RF frequency of 719 MHz, a mobile speed of 3 kmph results in a Doppler frequency of $f_d = 2$ Hz. The coherence time is given by $0.423/f_d$ [14].

words, there are 7 zeros between any two adjacent non-zero subcarriers.[13] Consequently, the WIC and LIC symbols consist of 9 periods, each of which is 512 chips long, and one of the periods corresponds to the cyclic prefix.

The non-zero interlace of the WIC and LIC symbols are modulated by a QPSK symbol sequence that depends on the wide- and local-area index assigned to the transmitter. Specifically, a PN sequence of length 1000 bits is generated using a Linear Feedback Shift Register (LFSR) whose seed depends on the wide-area and local-area indices (See Section IV-C). The PN sequence is then scrambled with a 1000 bit zero sequence, and the resulting 1000 bits are mapped to 500 QPSK symbols. The wide-area index can take on values between 0 and 15, and the same holds for the local area index as well. The receiver can determine the wide and local area indices by performing hypothesis testing on the received sequence in the WIC and LIC OFDM symbols.

3) TDM Pilot 2 (TDM2) Channel: TDM2 is the fourth OFDM symbol in a superframe. It has 2000 non-zero subcarriers, or 4 non-zero interlaces. Each interlace is modulated by QPSK symbols, which are zero data symbols scrambled by a PN sequence. There is one zero subcarrier between any two adjacent non zero subcarriers. In time domain, TDM2 is periodic with two periods, each of which is 2048 chips long, plus the cyclic prefix (see Fig. 7).

Based on the coarse timing determined from TDM1, the receiver utilizes TDM2 to accurately refine the timing, or FFT window position, to within a few chips. TDM2 can also be used to generate an initial estimate of the channel.

4) Transition Pilot Channel: Each continuous portion of local/wide-area OIS and local/wide-area data channels is flanked by a pair of known symbols, called Transition Pilot Channel (TPC) symbols. The TPC symbols are used to assist channel estimation for demodulation of the data OFDM symbols adjacent to them. All eight interlaces are non-zero in each TPC symbol; hence, they do not consist of any periodic components, other than the cyclic prefix.

5) Reserved OFDM Symbols/Positioning Pilot Channel: The last portion of a superframe consists of 2, 6, 10, or 14 OFDM symbols that can be either reserved symbols to fill the superframe or Positioning Pilot Channel (PPC) Symbols. These PPC symbols are unique for each transmitter and can be used for transmitter identification and/or position location of the receiving device using triangulation methods. They have not been completely defined at the time of writing this paper.

B. OIS Channels

Different types of overhead or control information are transmitted over the FLO network.

One of the more critical pieces of overhead information conveyed to the mobile receivers is the time-frequency allocations for each MLC. This overhead information is transmitted through the overhead information symbols (OIS) and also by *embedding* it in the traffic payload.

The OIS consists of 10 OFDM symbols separated into two parts: (a) wide-area, and (b) local-area. The Wide-area OIS contains information about those MLCs that are common to the wide area, while the local-area OIS contains information about those MLCs that are common to specific local coverage areas. By separating local and wide-area information, coverage areas for the OIS are ensured to be commensurate with coverage areas for the corresponding MLCs.

The OIS channels are special cases of data channels. The details of the generation of data channel symbols described in the next section generally apply to the OIS as well. Here, we only describe properties specific to the OIS.

The modulation used for the OIS is QPSK. OIS is only turbo encoded at the rate of 1/5 with no RS encoding. Such modulation and the low code rate ensure robust recovery of the OIS at SNRs below the threshold SNR for data traffic transmitted in a superframe, even with the loss of time-diversity due to absence of RS encoding.

In both the wide-area and local-area OIS, exactly seven packets are transmitted, with each occupying exactly five slots in either one or two OFDM symbols. Each information packet contains 976 bits, or 6832 overhead information bits in both the wide-area and local-area OIS. This is sufficient to support the transmission of a maximum of 170 MLCs, assuming that the OIS uses about 40 bits per MLC, in the wide-area and local-area segments of a super-frame. Since 10 OFDM symbols are used in each super-frame to convey the OIS, the result is a small spectral efficiency penalty of $(10/1200) \times 100 = 0.83\%$.

In addition to the OIS, the time-frequency allocations of MLCs are also carried by *embedded overhead information* in the traffic payload. Two different mechanisms are used for sending the allocation information because the mobile uses each mechanism in a different set of scenarios. The OIS is used (a) during initial acquisition, (b) if the mobile is not able to decode the contents of a super-frame, (c) if the user desires to switch to a different service and thereby needs to decode a different set of MLCs, and (d) if the mobile is in "sleep" mode and needs to occasionally wake up to receive any updates. On the other hand, while successfully[14] decoding an MLC, the mobile uses the embedded overhead information to determine the "on" time for this MLC in the *next* super-frame. Hence, in this case, the mobile does not have to wake up to decode the OIS in the next superframe. The use of embedded overhead thus leads to reduced power consumption.

Note that the embedded overhead in a particular MLC contains information pertaining only to that MLC, while the OIS carries allocation information for all MLCs.

C. Data Channels

The data channels occupy most of the superframe and carry the multicast multimedia information to the FLO devices.

Fig. 8 shows the procedures for the generation of OFDM symbols in the data channels.

1) Form MAC Layer Packets and RS Encoding: As discussed above, the multimedia information bit streams are organized as MLCs. Each MLC consists of a number of MAC layer packets, each of which are of size 976 bits. During the encoding process,

[13] Strictly speaking this is true at all carriers except for those around the DC carrier.

[14] By successful decoding, *all* erasures within *each* RS code block are corrected by the RS decoding algorithm.

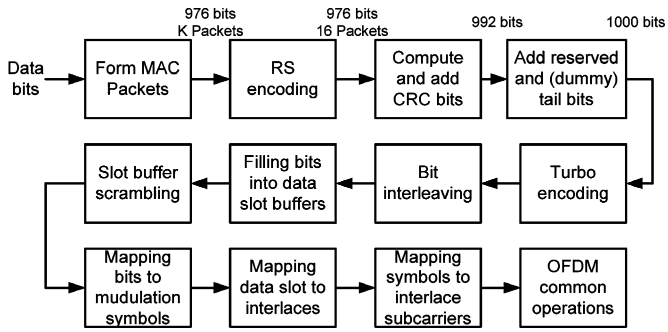

Fig. 8. Generation of data channel.

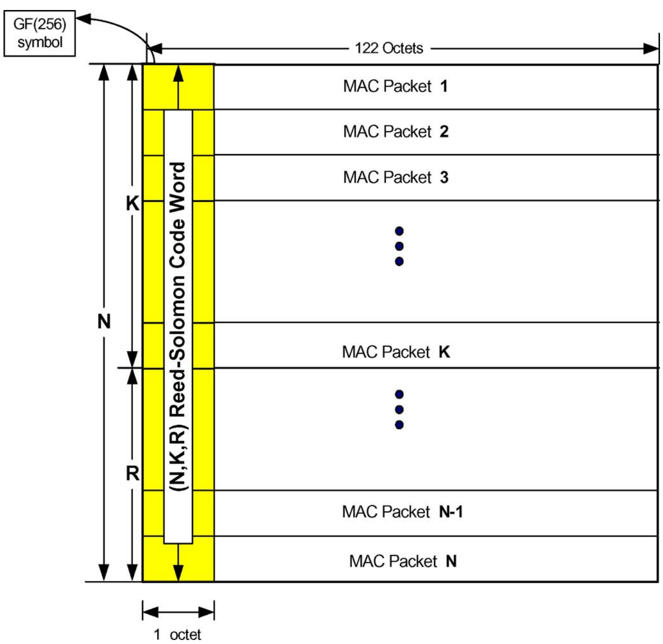

Fig. 9. RS-code block.

K MAC layer information packets are collected together. The bits in each packet are grouped into 8-bit octets. Reed-Solomon encoding is then performed on each column of K octets to produce codewords of length N octets (see Fig. 9). The parity octets of each codeword form $(N-K)$ new packets. 16 bit CRC values are generated for each of the K MAC layer information packets as well as the parity packets. These packets are subsequently turbo encoded. The collection of K MAC layer packets and $N-K$ parity packets is referred to as an RS, or outer, code block as shown by Fig. 9. Finally, MLC transmissions in each superframe are always in integer multiples of outer code blocks.

During transmission, each RS code block is split into 4 equal sub-blocks, with each sub-block sent in a unique frame within a super-frame. This is illustrated in Fig. 10, for the case of $\mathrm{R} = 12/16$. Each sub-block consists of 1/4th of the rows in the RS code block, i.e., either 4 MAC layer or 4 parity packets (Fig. 10). The time-spacing of an RS code block over a super-frame increases the time-diversity gained across each code block. In particular, if $K = 12$ the entire contents in a frame can be lost, yet the information transmitted in the super-frame can still be recovered, providing significant gains in slow fading channels and impulsive noise environments.

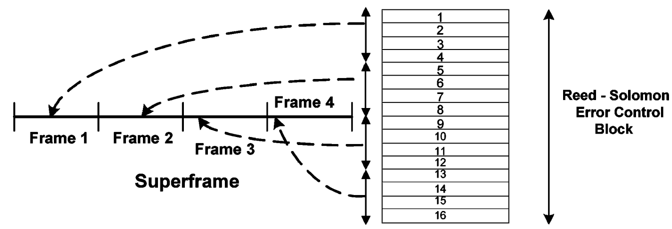

Fig. 10. Division of an outer code block into sub-blocks across the frames.

When an MLC containing multiple outer code blocks needs to be transmitted in a superframe, then the process is repeated for each code block. A sub-block from each code block is sent in every frame with all of the sub-blocks assigned to a frame transmitted contiguously. The following algorithm is utilized while transmitting the MAC layer and parity packets: Let the MAC packets in a code block be denoted by m-n, where m refers to the code block number and n refers to MAC packet number within a code block. In the i-th frame, $i \in [0, 3]$, in a superframe, n varies from $4\mathrm{i}+1$ to $4\mathrm{i}+4$, where i is the frame number while m ranges from 1 to G (G is the number of code blocks of the MLC sent in a superframe).

As an example, for $\mathrm{G} = 2$ the MAC layer and parity packets are transmitted in the order 1-1, 2-1, 1-2, 2-2, 1-3, 2-3, 1-4, 2-4, 1-5, 2-5, . . ., 1-16 and 2-16. This interleaving of MAC layer and parity packets increases the time diversity within a frame for each RS code block, which can be significant for large values of G.

2) CRC and Physical Layer Packets (PLPs): After the RS code blocks are formed, 16 CRC bits are generated for each row of an RS-code block using the standard CRC-CCITT generator polynomial:

$$g(x) = x^{16} + x^{12} + x^5 + 1$$

These CRC bits are used to determine, at the receiver, whether the packet was correctly decoded. This information is then used by the decoder for the outer RS code to perform erasure correction.

Following the addition of the CRC, each row of an RS-code block is padded with 2 reserved bits and 6 dummy tail bits, so each padded row becomes exactly 1000 bits. This is also referred to as a turbo information packet or a physical layer packet (PLP).

3) Turbo Encoding: The PLPs are then coded by a Turbo code, also referred as the inner code. The main purpose of the inner code is to exploit the frequency-diversity inherent in the channel. It is the same turbo code defined in the cdma2000 and 1x-EV-DO standards described in [1], [2]. Multiple code rates are supported in FLO. The code rates for transmitting MLCs are {1/3, 1/2, 2/3}. For transmitting the OIS, the code rate is 1/5.

The encoding operation is essentially a block operation, with an encoded output packet generated for each input packet. The encoder discards the dummy tail bits of the input PLP and encodes the remaining bits with a parallel turbo encoder. If R is the code rate, then the turbo encoder adds an internally generated tail of 6/R output code bits. The total number of turbo encoded bits at the output is 1000/R. Thus, for MLC transmissions, an

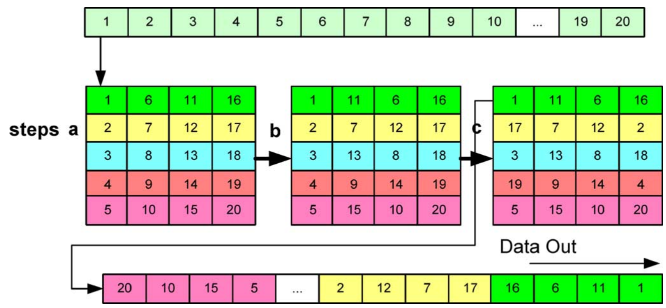

Fig. 11. The block bit interleaving process.

encoded packet contains 1500–3000 bits, while for the OIS, an encoded packet contains 5000 bits.

4) Block Bit Interleaving: The output bits of the inner encoder are bit-interleaved, modulated into symbols, with a further level of (symbol) interleaving prior to assignment to sub-carriers. The overall purpose of interleaving is to try and ensure that consecutive bits in the decoder input experience frequency fading that is uncorrelated, enabling the recovery of the encoded packet as whole.

The specific purpose of bit interleaving is twofold. The code bits of each turbo encoded packet are interleaved such that:

1) Adjacent coded bits are mapped into different constellation symbols; and
2) For 16-QAM modulation, adjacent coded bits are alternately mapped into the more and less "reliable" bits in the constellation, which uses a Gray mapping.

To perform the interleaving operation, an interleaver buffer consisting of $N/4$ rows by 4 columns is used, where N denotes the number of bits in a turbo encoded packet, which ranges from 1500–3000, depending on the inner code rate chosen for an MLC.

The various steps in bit interleaving are illustrated in Fig. 11 for the simple case of a 20 bit packet. In this example, the interleaver buffer consists of five rows and four columns. Code bits are written into the buffer column by column (see Step a). The middle two columns are then swapped, for every alternate row, starting with the top row (see Step b). For rows not processed in Step (b), the entries in the first and last column are swapped (see Step c). Finally, the bits are read out from left to right from each row starting with the top row (see Step d).

Note that, Steps (b) and (c) increase the separation of adjacent code bits for QPSK. Step (c) also helps achieve criterion (2) above for 16-QAM.

5) Filling Slot Buffers: The smallest unit of bandwidth allocated to an MLC over an OFDM symbol corresponds with a group of 500 constellation symbols. This group is called a slot. Note that the term *slot* refers to a group of constellation symbols, while the term *interlace* refers to a group of subcarriers, and each slot is mapped to one interlace. For QPSK and 16-QAM constellations, the 500 constellation symbols correspond to 1000 and 2000 bits respectively. In the OFDM symbols of each frame, up to 7 data slots are allocated per OFDM symbol for the transmission of the multiple turbo encoded packets associated with one or more MLCs, and the remaining slot is dedicated to pilot symbols. For some modes, a turbo encoded packet occupies a fraction of the slot. However, slot allocations are done so that multiple MLCs do not share slots within OFDM symbols. This is possible because each MLC contains an integer number (n) of RS blocks, i.e., $16 \times n$ turbo packets, in each superframe, or $4 \times n$ turbo packets in a frame. It is easy to verify that, for any FLO mode, $4 \times n$ turbo packets always occupy an integer number of slots. Slot allocation decisions are made at the MAC layer and remain in effect for all four frames in a super-frame. It should be noted that each MLC can span multiple OFDM symbols.

6) Scrambling: The bits in each slot buffer are XORed sequentially with the output bits of a scrambler to randomize the bits prior to modulation. The scrambling operation is performed for two reasons: (a) Minimize the magnitude of the peaks in the FLO waveform, resulting in Peak-to-Average Power Ratio (PAPR) properties similar to that of a Gaussian distribution. (b) Enable the FLO device to distinguish the FDM Pilot Channel of the desired[15] signal at the boundary between neighboring wide/local-areas (see Section IV-D).

The scrambling sequence used for any slot buffer depends on the OFDM symbol index and slot index. The sequence is produced by taking the modulo-2 inner product of the output of a 20-tap LFSR with the generator polynomial $h(D) = D^{20} + D^{17} + 1$, and a 20-bit mask associated with the slot index.

The scrambler is reinitialized at the start of each slot. The initial state of the LFSR is changed after each OFDM symbol. The state is denoted as $[\mathrm{d_3d_2d_1d_0c_3c_2c_1c_0b_0a_{10}a_9a_8a_7a_6a_5a_4a_3a_2a_1a_0}]$. Bits $\mathrm{d_3 - d_0}$ are set to the 4-bit Wide-area Differentiator (WID). Bits $\mathrm{c_3 - c_0}$ is set to 0000 for TDM2 and wide area-related channels, and to the 4-bit Local-area Differentiator (LID) for local area related channels. Bit $\mathrm{b_0}$ is a reserved bit and shall be set to '1'. Bits $\mathrm{a_{10}}$ through $\mathrm{a_0}$ shall correspond to the OFDM symbol index number in a superframe, which ranges from 0 through 1199. The shift register is reloaded with a new state $[\mathrm{d_3 - a_0}]$ for each slot at the start of every OFDM symbol.

The mask values used for each slot are pre-calculated to divide the state space of the LFSR into 8 non-overlapping segments.

7) Subcarrier Modulation: Three modulation modes, i.e., QPSK, 16QAM and layered modulation, are employed in FLO. Two or four scrambled coded bits map to an QPSK or 16QAM constellation, respectively. For layered modulation, 2 base layer bits and 2 enhancement layer bits are combined and mapped to a uniform or non-uniform 16-QAM constellation. The bits corresponding to the base component determine the quadrant that the point lies in, while the enhancement component bits determine the point, within the quadrant selected by base component bits, as is shown in Fig. 12.

In Fig. 12, the average energy per constellation point is $2(\alpha^2 + \beta^2)$, of which $2\alpha^2$ is the contribution from the base component, while $2\beta^2$ is the contribution from the enhancement component. The ratio (α^2/β^2) is referred to as the energy ratio of the layered constellation, and, for FLO, is either equal to 4.0 or 6.25.

[15]The desired signal in this context is the stronger (in received power) of the interfering signals.

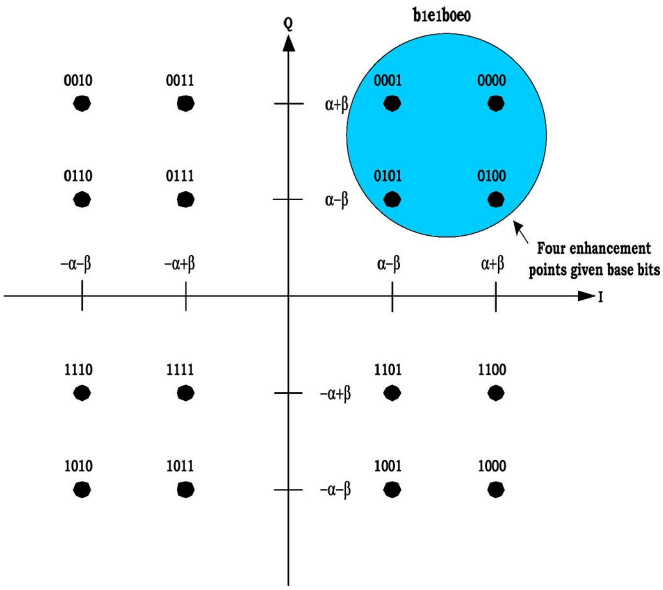

Fig. 12. Bit mapping in layered modulation.

8) Slot to Interlace Mapping: In FLO either interlace 2 or 6 in each OFDM symbol is used as the FDM pilot, alternating from one OFDM symbol to the next. The pilots are staggered in two consecutive OFDM symbols to allow for doubling the length of the channel estimate in the receiver from 512 to 1024 chip-spaced taps. This allows channel estimates with delay spreads up to twice as large as the length of the cyclic prefix. This is important in the context of improved performance in channels with delay spread exceeding the cyclic prefix.

As a result, only 7 interlaces in each OFDM symbol can carry constellation symbols of MLCs. From the viewpoint of the scheduler, there are eight slots per OFDM symbol, with Slot 0 always assigned to the pilot and the remaining slots available for allocation to MLCs.

The FDM Pilot Channel is assigned to interlace 2 and 6 for even and odd OFDM symbol indices, respectively. The remaining 7 interlaces in each OFDM symbol are assigned to slots 1 through 7.

The slot to interlace mapping for slots 1 though 7 can be described as follows:

Let I_j denote the j^{th} interlace. Permute the interlace sequence $\{I_0 I_1 I_2 I_3 I_4 I_5 I_6 I_7\}$ by replacing the index i in I_i with the bit reverse value of i, to generate the permuted sequence, $\mathrm{PS} = \{I_0 I_4 I_2 I_6 I_1 I_5 I_3 I_7\}$.

Club interlaces I_2 and I_6 in the PS to generate shortened interlace sequence, $\mathrm{SIS} = \{\mathrm{I}_0 \mathrm{I}_4 \mathrm{I}_2/\mathrm{I}_6 \mathrm{I}_1 \mathrm{I}_5 \mathrm{I}_3 \mathrm{I}_7\}$.

From one OFDM symbol to next in a superframe, perform a right hand cyclic shift by a value of $(2 \times j)\mathrm{mod}7$, where $j \in \{1, 2, \ldots, 1199\}$ is the OFDM symbol index in a superframe, to generate the permuted shortened interlace sequence.

Interlace I_2 or I_6, whichever is not occupied by the pilot slot shall be used in the permuted interlace sequence.

Fig. 13 illustrates the interlace assignment to all 8 slots over 15 consecutive OFDM symbol intervals. The mapping pattern from slots to interlaces repeats after 14 consecutive OFDM symbol intervals.

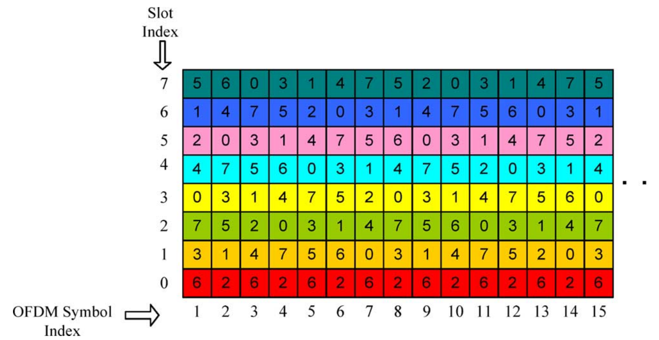

Fig. 13. Slot to interlace mapping at the start of the superframe.

The above mapping ensures that the MLC payload, even though allocated to a fixed set of slots in each frame, gets distributed over all interlaces in multiple OFDM symbols. This results in improved demodulation performance on channels with periodic nulls.

9) Mapping of Modulation Symbols Into Subcarriers: After a slot is mapped to one of the eight interlaces, the constellation symbols of that slot are mapped to the subcarriers of that interlace. To achieve maximum frequency diversity, the indices of the 500 subcarriers within a interlace are first interleaved using a bit-reversal mapping. Subsequently, the 500 constellation symbols in a slot are assigned sequentially to the subcarriers of the interlace, using the interleaved index. This ensures that the adjacent symbols of a slot are mapped to distant subcarriers within each interlace.

Since 500 is not a power of 2, the following *reduced-set* bit reversal operation is used:

$$\Re\{x\} = \mathfrak{R}\{x + f(x)\}$$

where $f\{x\} = \sum_{n=0}^{x} g(n)$, $g(n) = \begin{cases} 0, & if\ \mathfrak{R}(n) < 500 \\ 1, & if\ \mathfrak{R}(n) \geq 500 \end{cases}$, and $\mathfrak{R}\{n\}$ is the regular bit-reversal function.

D. FDM Pilot Channel

The FDM pilot used for channel estimation is always allocated to Slot 0 by the scheduler. As shown in Fig. 13, Slot 0 is either assigned to interlace 2 or 6. The subcarriers of the pilot interlace are modulated with QPSK symbols with the *same* energy as the MLC constellation symbols on the other interlaces. The justification for this selection is as follows: Let R denote the energy ratio of the pilot to data sub-carriers. Intuitively, the selection R of involves trading off the channel estimation accuracy with the SNR of the data sub-carriers. Using link-level simulations, we evaluated the PER performance of FLO in both Additive White Gaussian Noise (AWGN) and fading channels for different R. The results showed that the PER performance for $R = 1$ was as good as, if not better than, that for other choices of R.

The QPSK constellation points for the (FDM) pilot sub-carriers are generated using the same LFSR used by the scrambler. The starting value for the LFSR and the output mask are chosen according to the procedure outlined in Section IV-C. Before running the LFSR, the slot (zero) buffer is filled with a 1000-bit all-zero pattern.

The different transmitters generate the *same* pilot (constellation) symbol sequence in a particular wide-area OFDM symbol. This is needed to estimate the composite channel and avail the gains from the SFN nature of wide-area transmissions. However, the pilot symbol sequence changes from one wide-area OFDM symbol to the next. The same is true for different transmitters within a local coverage area for local-area OFDM symbols, i.e., transmitters in different local coverage areas transmit different pilot symbol sequences in the same local-area OFDM symbol. This allows the receiver to estimate the channel from a single local coverage area, thereby enabling the demodulation and decoding of the content from that particular local area.

V. Performance

In this section, we present representative performance results for FLO in a few different channel environments. In terrestrial mobile multicasting, using high power transmitters, the channels are characterized by the presence of both frequency-selective and time-selective fading. The former results from the large possible (channel) delay spreads, while the latter is due to device motion. In addition, the device must receive transmissions in the presence of adjacent/co-channel interference, RF receiver and fixed-point quantization floors, etc.

In evaluating the performance of the FLO physical layer, we use packet error rate (PER) as the primary criterion. Two reference points within the device are used at which PER is measured: (a) Pre-RS decoding, and (b) Post-RS decoding. Note that, at reference point (a), a packet is a PLP (1000 bits), while for (b), a packet denotes a MAC layer packet (976 bits). Furthermore, for coverage purposes, we target a PER of 1%. This error rate (and below) has been verified as adequate[16] for the operation of video and audio decoders.

In Section II-H, we listed the possible FLO modes. In the initial deployment of FLO systems, the preferred modes are 7 (Layered, energy ratio 4, code rate 1/2) for video streams and 2 (16-QAM, code rate 1/3) for audio streams, while mode 1 (QPSK, CR 1/2) is intended for use in sending overhead information sent in an MLC. In the sequel, we focus on the performance of these modes. Another consideration is the selection of RS (outer) code rate. The values of interest are 12/16 and 16/16, with no RS encoding performed in the latter case.

We first examine the performance of FLO on the Additive White Gaussian Noise (AWGN) channel. On this channel, the only impairment is noise, and the PER is plotted against the Carrier-to-Noise (C/N) power ratio, in dB, in Fig. 14. The PER curves in Fig. 14 are obtained without the use of RS encoding, and, hence, are reflective of the performance of the inner (Turbo) code. Note that performance of the base layer and the enhancement layer in Mode 7 is shown separately. In order to achieve a PER of 1%, the required C/N ranges from 1.8 dB to 9 dB over the FLO modes shown in Fig. 14. The required C/N for all the modes is listed in Section II-H (Table 2).

[16]The video and audio decoders used typically rely on advanced error concealment strategies in order to operate at such high PER.

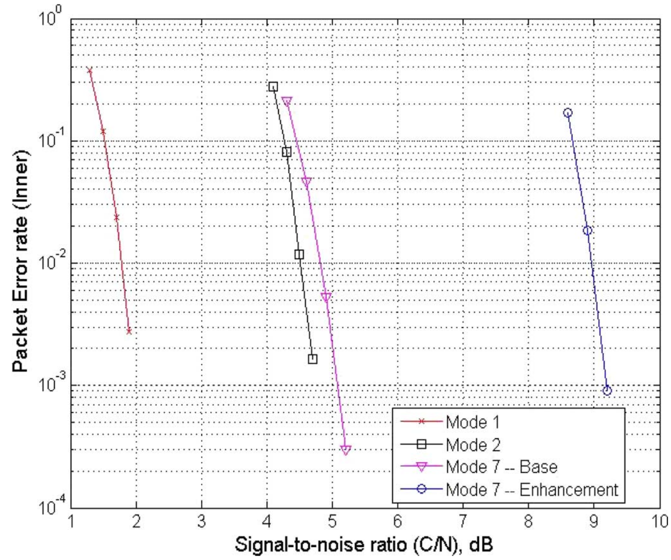

Fig. 14. Performance of typical FLO modes in AWGN channel.

TABLE 3
Modified PEDB Multipath Profile

Path	Relative delay (ns)	Average power (dB)
1	0	-5.1
2	200	-6.0
3	800	-10.0
4	1200	-13.1
5	2300	-12.9
6	3700	-29
7	40000	-10.1
8	40200	-11.0
9	40800	-15.0
10	41200	-18.1
11	42300	-17.9
12	43700	-34.0

Next, we examine the performance of FLO on a channel with both frequency-selective and time-selective fading. The model considered is a modified Pedestrian B (PEDB) profile [15] in order to represent a SFN deployment. The basic PEDB profile, consisting of 6 channel taps, is modified by adding an additional cluster of 6 taps that is both delayed in time (40 us) and attenuated by 5 dB relative to the main cluster. Hence, this channel profile models reception at a particular point in an SFN with two transmitters, such that the path differential between signals from the two transmitters is 12 km. The channel profile is given in Table 3, with the powers of the taps normalized to sum to 1. Each tap experiences independent Rayleigh fading, with the Doppler spectrum dependent on the speed of the mobile (device).

Fig. 15 depicts the performance of FLO for the modified PEDB profile. The RS code rate was set to 12/16 in generating these results.

In Fig. 15, we plot the C/N required to achieve a PER of 1% against the Doppler frequency, $f_d = (v/c)f_c$ where v is the vehicle speed, f_c is the carrier center frequency, and c is the

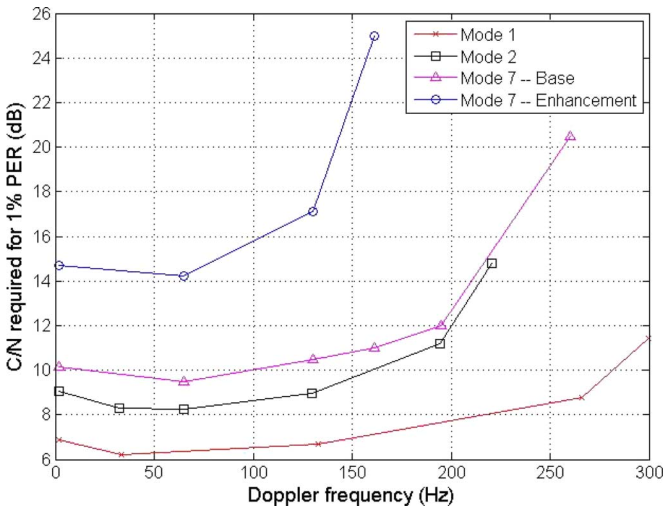

Fig. 15. Required C/N as a function of Doppler frequency for modified PEDB profile.

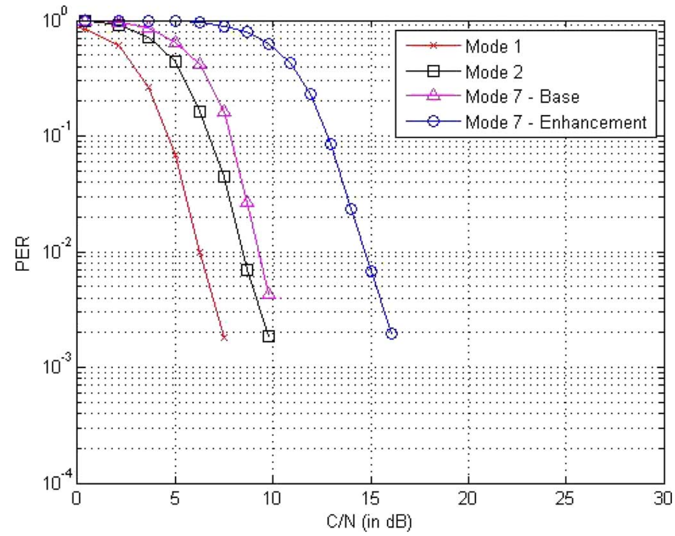

Fig. 16. Results from drive tests along highways.

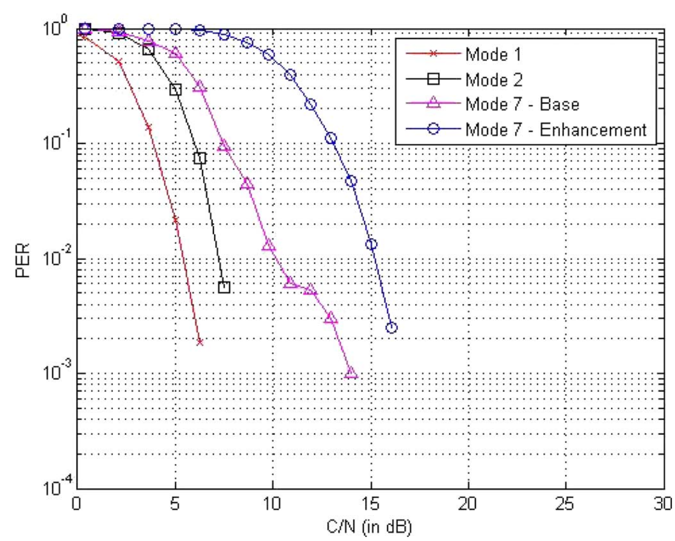

Fig. 17. Results from pedestrian indoor field tests.

speed of light. The carrier center frequency is assumed to be 700 MHz (UHF), and the RF channel bandwidth is 6 MHz.

We observe that all the depicted FLO modes perform very well over a wide range of Doppler frequencies, ranging from 2–130 Hz. At a carrier frequency of 700 MHz, this corresponds to vehicle speeds ranging from 3–200 km/hr. The excellent performance at low (pedestrian) speeds is due to the RS code exploiting the time-diversity inherent across MLC transmissions in a super-frame. With a RS code rate of 12/16, the outer code can correct packet erasures due to fades lasting the duration of a frame. The good performance at high speed is the result of the choice of sub-carrier spacing, FDM pilot channel design, and the use of a powerful (Turbo) inner code. It is also noteworthy that, for speeds ranging from 3–120 km/hr, the performance of FLO over the modified PEDB channel model agrees closely (within 1–1.5 dB) with the performance over the Typical Urban, 6-path, channel model (TU6) [16].

The results presented in Figs. 14 and 15 have been corroborated in the laboratory with actual handsets that contain dedicated RF and base-band chips, along with software implementation of the Upper layer protocol stack for FLO. In addition, QUALCOMM performed extensive field testing in San Diego using two high power transmitters and actual handsets. Each transmitter radiates the FLO signal at 50 kW effective radiated power (ERP) and the distance between the transmitters is 36 kilometers. The two transmitters were set up to create an SFN with an overlapping coverage area.

The methodology used in collecting field results, as well as the performance results obtained are summarized in a white paper [17]. In Figs. 16 and 17, we show some representative results. The curves in Fig. 16 were obtained from drive tests along highways, where the speed was about 105 km/hr, while those in Fig. 17 are from pedestrian indoor tests.

It can be seen from the figures that the required C/N for 1% PER is about 8–9 dB for Mode 2, 10 dB for the base layer of Mode 7 and about 15 dB for the enhancement layer. These numbers are in excellent agreement with those from the modified PEDB profile in Fig. 15, despite the fact the we have no control over the channel profiles from each transmitter in the field. Hence, the modified PEDB channel model provides a good benchmark for predicting field performance.

The required C/N does not vary for the two extremes of highway and pedestrian speeds. Drive tests at intermediate speeds in urban and suburban environments also yielded similar results [17]. Furthermore, these results from field testing in the overlapping (coverage) area of the transmitters confirm the practical feasibility of FLO SFNs, which result in improved coverage and spectral efficiency. The field results show that the performance of FLO is robust to varying reception conditions in the real world.

VI. Conclusion

In this paper, we have described the major features of the physical layer of the FLO (Forward Link Only) air interface. A detailed description of the various physical sub-channels and the resulting transmit waveform has also been provided. As a technology for delivering multimedia content to mobile handheld devices, FLO is a ground up effort, unlike other technologies that have evolved from existing or legacy systems. Consequently, FLO is not constrained to be backward compatible and can address the problem in the most efficient manner. For instance, the MediaFLO system is optimized across all layers of the protocol stack and is an integrated solution for multicast delivery. At the physical layer, the lack of compatibility constraints allows FLO to use turbo coding instead of convolutional coding and leads to a significant improvement in the required link budget.

More generally, the physical layer design has been driven by constraints specific to the mobile multicast scenario, such as minimizing the device power consumption and minimizing channel switch times while maximizing time & frequency diversity gains. In addition, it maximizes the spectral efficiency of the network, while providing sufficient flexibility in deployment. We have described the key features of FLO that achieve these goals, e.g. multiplexing the services in time and frequency to minimize receiver ON time; supporting variable rate codecs & statistical multiplexing for capacity gains; the use of outer R-S coding across the frames in a superframe to capture time diversity and minimize channel switch times; the use of layered codecs & modulation for trading off coverage and data rate of a service; and so on.

It is also important to note that FLO, while being a new technology, is already in a mature state of design, development and standardization. A consortium of companies, the FLO Forum, was established for promoting FLO. The member companies approved the FLO Air Interface specification, which was subsequently submitted to and recently standardized by TIA[17] (TR47.1[18]) as Technical Standard TIA-1099 [18].

On the development front, both the high power transmitter and the receiver for FLO have been implemented in hardware and integrated with the rest of the protocol stack in software. The performance of the system has been tested thoroughly in the lab and the field. As shown in this paper, the results from the field tests confirm that the performance of FLO matches the predicted results and achieves our design goals.

Acknowledgment

The authors would like to thank the engineers of the MediaFLO Technology and Corporate R & D Departments at QUALCOMM, who participated in, and contributed to, the specification of the FLO Air Interface and the development of the algorithms, software and hardware of the FLO transmitter and receiver. The authors also thank the editors for their support and the anonymous reviewer for helpful comments.

References

[1] *Physical Layer Standard for cdma2000 Spread Spectrum Systems-Release C*, 3GPP2 C.S2002-C v1.0, May 2002.

[2] *CDMA2000 High Rate Packet Data Air Interface Specification*, 3GPP2 C.S2024-A v1.0, March 2004.

[3] G. Faria, J. A. Henriksson, E. Stare, and P. Talmola, "DVB-H: Digital broadcast services to handheld devices," *Proceedings of the IEEE*, vol. 94, no. 1, pp. 194–209, Jan. 2006.

[4] S. Cho *et al.*, "System and services of terrestrial digital multimedia broadcasting (T-DMB)," *IEEE Trans. Broadcasting*, vol. 53, no. 1A, Mar. 2007.

[5] M. Takada and M. Saito, "Transmission systems for ISDB-T," *Proceedings of the IEEE*, vol. 94, no. 1, pp. 251–256, Jan. 2006.

[6] A. Mattson, "Single frequency networks in DTV," *IEEE Trans. Broadcasting*, vol. 51, no. 4, pp. 413–422, Dec. 2005.

[7] K. Salehian, Y. Wu, and B. Caron, "Design procedures and field test of a distributed-translator network, and a case study for an application of distributed-transmission," *IEEE Trans. Broadcasting*, vol. 52, no. 3, Sept. 2006.

[8] H. Sari, G. Karam, and I. Jeanclaude, "Transmission techniques for digital terrestrial TV broadcasting," *Communications Magazine, IEEE*, vol. 33, no. 2, pp. 100–109, Feb. 1995.

[9] U. Ladebusch and C. A. Liss, "Terrestrial DVB (DVB-T): A broadcast technology for stationary portable and mobile use," *Proceedings of the IEEE*, vol. 94, no. 1, pp. 183–193, Jan. 2006.

[10] M. Speth, S. A. Fechtel, G. Fock, and H. Meyr, "Optimum receiver design for wireless broad-band systems using OFDM. I," *IEEE Trans. Communications*, vol. 47, no. 11, Nov. 1999.

[11] Y. Li, "Pilot symbol aided channel estimation for OFDM in wireless systems," *IEEE Trans. Vehicle Technology*, vol. 49, pp. 1207–1215, July 2000.

[12] M. Balakrishnan, R. Cohen, E. Fert, and G. Keesman, "Benefits of statistical multiplexing in multi-program broadcasting," in *Broadcasting Convention, 1997. International*, Sept. 12–16, 1997, pp. 560–565.

[13] *CDMA2000 High Data Rate Broadcast-Multicast Packet Data Air Interface Specification*, 3GPP2 Contribution C30-20050926-012R4C_C.P0054-A, Dec. 12, 2005.

[14] W. C. Jakes, Ed., *Microwave Mobile Communications*. Piscataway, NJ: IEEE Press, 1993, (reissue).

[15] *Guidelines for Evaluation of Radio Transmission Technologies for IMT-2000*, Recommendation ITU-R M.1225, ITU.

[16] "Digital Land Mobile Radio Communications (Final Report)," Commission of the European Communities, Directorate General Telecommunications, Information Industries and Innovation, COST207, 1989, pp. 135–147.

[17] "MediaFLO Field Test Report," QUALCOMM White Paper, March 2006 [Online]. Available: http://www.qualcomm.com/mediaflo

[18] *Forward Link Only Air Interface Specification for Terrestrial Mobile Multimedia Multicast*, TIA-1099, Aug. 2006.

[17]TIA is the Telecommunications Industry Association of North America.

[18]TR47.1 is a subcommittee of TR47, within the TIA, that is dedicated to the development of technical specifications on FLO.

Murali R. Chari was born in Chennai, India, on May 28, 1968. He received the B. Tech. degree in electrical engineering from the Indian Institute of Technology, Chennai, in 1989, and the M.S.E. and Ph.D. degrees in electrical and computer engineering from The Johns Hopkins University, Baltimore, Maryland, in 1991 and 1996, respectively.

From 1995 to 2003, he worked for ComStream Corp. and later Conexant Systems on the design and development of digital broadcast satellite (DBS), cable modem and digital cable set tops, and digital TV receivers. In 2003, he joined QUALCOMM, Inc. where is currently a Principal Engineer/Manager in the Corporate Research and Development group. Since joining QUALCOMM, Murali has been working on the MediaFLO project contributing to the Physical layer specification and the standardization efforts for FLO technology.

Murali currently chairs the Technical Committee of the FLO Forum, a consortium of companies that have joined to advance the standardization and deployment of FLO-based systems worldwide. In this role, Murali leads the committee's work in producing technical specifications on various aspects of FLO-based systems, and is an active participant in organizations, such as TIA TR47.1, working on the development of FLO-based standards.

Fuyun Ling received his B. S. degree from Tsinghua University, Beijing, China, in 1968, and his MS and Ph.D. degrees from Northeastern University, Boston, Massachusetts, USA, in 1981 and 1984 respectively.

From 1984 to 1998, he was with Motorola Inc. working on research and development of various wireline and wireless communication technologies and products. In 1998, he joined QUALCOMM Inc., San Diego, California, as a Vice President of Technology and worked on research and development of new wireless communication technologies. Since 2003, he has been working on the MediaFLO project and is currently responsible for FLO PHY Layer specification and development.

Throughout his career, Dr. Ling has been working on digital communications and digital signal processing and their application to wireless and wireline communications. He holds over 60 US Patents with a number of other patent applications pending. He has also published over 80 papers and two books all in his technical areas. He is an Adjunct Professor of Tsinghua University and a Fellow of IEEE.

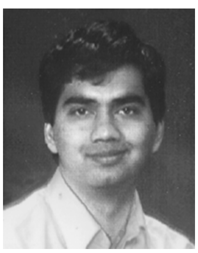

Ashok Mantravadi (S'97–M'02) has been with the Corporate R&D group at QUALCOMM since March 2002 where he has been involved in the development of Wireless LAN systems and the FLO multicast technology. He received the B. Tech degree in Electronics and Communications engineering from the Indian Institute of Technology, Madras in 1996. He joined Cornell University thereafter, and obtained his M.S and Ph.D. degrees in Electrical Engineering in January 1999 and 2002, respectively. While at Cornell, he was a recipient of the John McMullen and Schlumberger Fellowships. From August 2000 to December 2001, he was also a visiting scholar at the University of Illinois, Urbana-Champaign. His research interests include wireless multicast systems, multiuser communications and signal processing.

Raghuraman Krishnamoorthi received his B. Tech from the Indian Institute of Technology, Madras in 1996 and his masters in electrical engineering from the University of Illinois at Urbana-Champaign in 1998. From 1998 to 2000, he worked as a systems engineer at Conexant Systems on the development of voice band modems. He has been at QUALCOMM since December 2001, where he has been working on OFDM communication systems, including the FLO system for terrestrial mobile multicast.

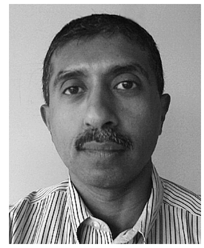

Rajiv Vijayan is a Vice President of Engineering in the Corporate R&D department of QUALCOMM Inc., San Diego, California, where he has been employed since 1993. He is co-lead of the system engineering department and also leads a team developing mobile broadband wireless technology. At QUALCOMM, Rajiv has worked on the design, implementation, and field deployment of different wireless technologies. His assignments at QUALCOMM have included being the lead engineer in charge of physical layer development for the Globalstar satellite communication system and leading the team that designed the air interface of the FLO mediacast technology for delivery of multimedia to cell phones. Prior to joining QUALCOMM, Rajiv was a postdoctoral associate with the Wireless Information network Laboratory (WINLAB) at Rutgers University. Rajiv has a Bachelor of Technology from the Indian Institute of Technology, Madras, and Masters and Ph.D. degrees from the University of Illinois at Urbana-Champaign, all in Electrical Engineering.

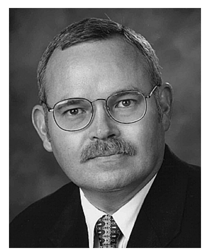

G. Kent Walker is a VP of Technology at QUALCOMM Inc. His current role is technical leadership in MediaFLO. Mr. Walker earned his BEE and MSEE degrees at Georgia Institute of Technology. Prior to QUALCOMM, he has worked at General Instrument, M/A-COM Linkabit, Scientific Atlanta and Harris Corporation(s). He has numerous granted and pending patents in the areas of digital communications, digital audio, and image transmission. Products developed under his technical leadership include analog and digital video satellite receivers, digital microwave radios, audio signal processors, digital satellite modems, video scrambling, access control, and digital cinema equipment. Mr. Walker is a member of SMPTE, AES, and IEEE.

Rob Chandhok is currently Chief Product Officer for SunRocket Inc.

Previously he was responsible for the technical architecture and development of the MediaFLO System as the Vice President of Engineering and Market Development for MediaFLO.

Chandhok joined QUALCOMM in 2000 as director of engineering in QUALCOMM's Eudora Products group. He then served as senior director of business development for the QUALCOMM Internet Services (QIS) division, followed by his role as vice president of technology for QChat, a push-to-talk VoIP technology for 3G CDMA networks.

Before joining QUALCOMM, Chandhok served as president of Within Technology for seven years after founding the company in 1993. He worked as a research computer scientist at the Carnegie Mellon School of Computer Science for more than 10 years, before leaving to found Within Technology.

Chandhok holds a Bachelor of Science degree in electrical engineering from Carnegie Mellon University.