

Performance Evaluation of IEEE 802.16 for Broadband Wireless Access

Shyamal Ramachandran, Charles W. Bostian, Scott F. Midkiff
Center for Wireless Telecommunications, Virginia Tech
Blacksburg, Virginia 24061
{shyamal, bostian, midkiff}@vt.edu

Abstract

Broadband Wireless Access (BWA) is drawing a great deal of interest from the engineering and research communities. The IEEE 802.16 task group is in the process of standardizing the physical (PHY) and media-access control (MAC) layers of BWA systems [1]. The proposed MAC supports a range of physical layer technologies and provides manufacturers with the option of switching modulation and coding schemes based on the channel conditions. We model IEEE 802.16 and evaluate the performance of its MAC layer over various physical layer options using OPNET. Our results show the delay and throughput performance of IEEE 802.16 when using quadrature phase shift keying (QPSK) and quadrature amplitude modulation (QAM) modulation schemes in an additive white Gaussian noise (AWGN) channel. Based on these results we demonstrate the need for an algorithm to dynamically switch between different PHY burst profiles in order to improve the protocol's performance under various channel conditions.

I. Introduction

Amidst the rapidly growing demand for high-speed Internet access in the residential and small office sectors, and an equally fast paced growth in last-mile access technologies, broadband wireless access has emerged as a promising solution. Despite being capable of high-speed multimedia services and rapid flexible deployment, BWA has fallen short of becoming a cost-effective option in the market. A lack of standardization and the relatively harsh operating conditions provided by the wireless channel are probably the prime reasons for the slow growth in the BWA industry. Until recently most BWA systems and vendors have used proprietary schemes. Most have adapted popular wireline standards for their wireless applications. Previous research has also looked in to developing efficient and adaptive MAC layers for BWA [2].

The IEEE 802.16 task group recently published its WirelessMANTM air interface standard [1] that standardizes the MAC and PHY layers for fixed BWA systems. A significant aspect of this standard is that it tries to retain much of what is already in use in the market and only tries to streamline the development process by recommending practices based on already widespread technology. Many proprietary BWA systems adapt the DOCSIS v. 1.1 [3] MAC to operate in a wireless environment using QPSK and QAM based PHY schemes. The WirelessMANTM air interface is also based on similar lines. It has been designed to address systems operating from 10 to 66 GHz. It standardizes a common DOCSIS-based MAC over a wide range of PHY options. The PHY supports various modulation and coding schemes in order to operate over this wide spectrum. In addition, IEEE 802.16 also defines a radio link control (RLC) layer enabling dynamic switching between

PHY configurations. This allows the system to adapt to various physical channel impairments such as rain fades and multipath.

Our research focuses on evaluating the performance of IEEE 802.16's adaptive PHY capabilities. Our interest is to study the performance of the various PHY options provided by IEEE 802.16 under various channel conditions. Based on this work we intend to propose an algorithm to implement adaptive modulation and coding in such systems. We have not yet achieved all of our research goals; in this paper we present our progress thus far. We have adapted OPNET's DOCSIS models to simulate an IEEE 802.16 system. We have also evaluated the downstream performance of the various supported modulation schemes.

The rest of this paper is organized as follows. In Section II we briefly introduce the IEEE 802.16 standard itself. In Section III we describe our simulation model and discuss how we implement PHY effects in OPNET. Section IV summarizes our findings. We conclude with a brief discussion of our future goals in Section V.

II. IEEE 802.16 Standard Overview

IEEE 802.16 defines the WirelessMANTM air interface standard for fixed BWA metropolitan area networks (MANs). The standard covers a wide range of spectrum from 10 to 66 GHz. The propagation mechanism in this region of spectrum restricts the use of this system to line-of-sight environments. The standard was designed for point-to-multipoint broadband networks where a base station (BS) transmits to multiple subscriber stations (SS) in its coverage area. The BS maintains a time-division multiplexed (TDM) link on the downlink. The subscriber stations use time-division multiple access (TDMA) on the uplink and transmit back to the BS in their allotted time slots.

MAC Layer Overview

The IEEE 802.16 MAC regulates uplink channel access using a Demand Assigned Multiple Access (DAMA) TDMA system. The BS controls uplink bandwidth allocation. Subscriber stations request for transmission opportunities on the uplink channel. The BS collects these requests and grants permissions to the subscriber stations based on their quality and service agreements. These transmission opportunities may be pre-allocated for particular subscribers or may be made available for contention by all SSs. The uplink channel is divided in to a stream of mini-slots. The system divides time in to physical slots (PSs). Each physical slot is defined to be four modulation symbols. Each mini-slot is an integral number of PSs. Whenever a SS desires to use the uplink, it requests for transmission opportunities in the units of mini-slots. The BS accepts these requests over a period of time and compiles an allocation map describing the channel allocation for a certain period in to the

future called the map time. This message is broadcast on the downlink to the SSs. A MAP message may allocate a certain number of mini-slots for particular SSs to transmit and a certain number of slots open for contention-based transmission. Collisions that occur while attempting to transmit in the contention slots are resolved by using the truncated binary exponential backoff algorithm.

The MAP message itself, carries Information Elements (IEs) that define how a particular uplink interval may be used. Request IEs specify an uplink interval during which SSs can make requests for uplink bandwidth. Initial Maintenance IEs specify time intervals when new SSs may enter the network. Station Maintenance IEs specify time intervals when SSs may perform ranging and power control. Data Grant Burst Type IEs specify time intervals when SSs may transmit protocol data units (PDUs) on the uplink channel.

The IEEE 802.16 MAC defines various scheduling service classes. Each SS negotiates its service policies with the BS at the connection setup time. These scheduling policies are used by the BS while allocating uplink bandwidth to each SS. The scheduling services defined in IEEE 802.16 are based on those defined in the DOCSIS v. 1.1 standard [4]. The various scheduling services defined in IEEE 802.16 are, Unsolicited Grant Service (UGS), Real-Time Polling Service (rtPS), Non-Real-Time Polling Service (nrtPS) and Best Effort Service (BE). UGS is defined to support services that generate fixed sized packets on a periodic basis [2]. To reduce overhead involved in the request-grant process, this service pre-allocates periodic grants to the SSs. This grant size is a system parameter and is negotiated at connection setup and is a part of the service agreements. UGS targets applications such as E1/T1 links. Real-Time Polling Service targets applications that need real-time service flows but generate variable sized data packets, such as voice over Internet Protocol (VoIP), streaming audio, and streaming video. This service scheme ensures periodic request opportunities, which can be used by the SS to request for a specific grant in real-time. Non-Real-Time Polling Service is designed to cater to the needs of non-real-time applications such as bulk file transfers that may need high bandwidth allocations but can tolerate longer delays. This service requires the SSs to randomly access transmission opportunities to request for uplink bandwidth. The Best Effort service provides service to best effort traffic where no throughput or delay guarantees are provided. The SSs are required to contend for bandwidth during contention request opportunities. The availability of such opportunities also depends on network load and is not guaranteed. In addition to these scheduling services, the MAC also provides support for concatenation and fragmentation.

PHY Layer Overview

The IEEE 802.16 PHY supports data rates of up to 134.4 Mbps on radio frequency (RF) channels as wide as 28 MHz [2]. Both frequency division duplex (FDD) and time division duplex (TDD) variants are defined. The standard supports a wide variety of physical layer options including various modulation schemes and forward error correction (FEC) schemes. On the uplink, support for QPSK is mandatory while 16-QAM and 64-QAM are optional. The downlink supports QPSK and 16-QAM, while 64-QAM is optional. Support for such a variety of

modulation schemes permits vendors to tradeoff efficiency for robustness or vice versa. In addition to these modulation schemes, the PHY also defines various FEC coding schemes on the uplink as well as the downlink. These range from single stage Reed-Solomon (RS) Galois Field (GF) (256) codes to RS concatenated with inner Block Convolution Codes (BCC). The standard has even defined Block Turbo Codes to ensure robustness in extreme fading channels.

RLC and Adaptive PHY

The IEEE 802.16 standard provides adaptive PHY profile switching capability. The Radio Link Control (RLC) defined in the standard is capable of switching between PHY burst profiles on a per-frame and per-SS basis. Combinations of PHY modulation and FEC schemes used between the BS and SSs are called uplink and downlink burst profiles depending on the direction of their flow. Burst profiles are identified as Downlink Interval Usage Code (DIUC) and Uplink Interval Usage Code (UIUC). These burst profiles are negotiated between the BS and every SS based on decisions made at either the BS or the SS. The standard recommends that the SS use $C/(N+I)$ as a metric to decide if a burst profile change is essential. This metric defines the ratio of the received signal power to the sum of the interference power and noise floor. Every SS is required to determine the signal quality during initial maintenance and request the BS for a particular DIUC. After initial maintenance the SS is required to continuously monitor the transmission quality and control the downlink burst profile. The BS may also monitor the received transmission quality and instruct the SS to switch to a different uplink burst profile. Such adaptation capability permits the system to switch to a more robust PHY technology during harsh channel conditions and back to more efficient schemes when the channel is reasonably good. The RLC's adaptation process is designed to be a continuous process, attempting to strike an optimal balance between efficiency and robustness.

Our ultimate goal is to develop an algorithm to implement adaptive burst profiles involving the PHY burst profiles standardized in IEEE 802.16. To determine the optimal $C/(N+I)$ switching thresholds, we need to first study the performance of the system in each of the available burst profiles and understand the efficiency versus robustness tradeoff involved in operating at those burst profiles.

III. Simulation Model

In this section we describe our simulation model and methodology. We used OPNET's DOCSIS models to simulate the IEEE 802.16 MAC. The IEEE 802.16 standard has significant similarities with the DOCSIS v. 1.1 MAC [5][6]. Both standards target the same kind of applications. Therefore, they operate with similar architectures and protocols. Like DOCSIS, the IEEE 802.16 MAC uses a DAMA-TDMA scheme. The services classes supported by IEEE 802.16 are also exactly the same as those in DOCSIS [4]. Like DOCSIS, IEEE 802.16 uses mini-slots and MAP messages to manage medium access request-grant mechanism. Both IEEE 802.16 and DOCSIS support the same IEs in the MAP. The contention resolution mechanism employed is also exactly the same. In fact, quality of service (QoS) features, such as concatenation,

fragmentation and payload header suppression (PHS), are also common to both standards.

There are significant dissimilarities as well. IEEE 802.16 supports TDD mode operation in addition to the FDD mode. Also, the number of bytes transmitted in a mini-slot is kept constant as the symbol rates change, thereby varying the amount of time covered by a mini-slot. The most significant difference is that IEEE 802.16 operates on a wireless PHY as opposed to the hybrid fiber-coax medium supported by DOCSIS. This means that the PHY implementation of IEEE 802.16 has to be totally different from that of DOCSIS and that the model has to incorporate, in detail, the effects of various types of physical layers.

Since the OPNET DOCSIS models simulate the MAC functionality in great detail, we chose to build upon them and incorporate the necessary modifications to support our study of the WirelessMAN™ air interface. For this, we had to replace the bus-based PHY model in OPNET with the corresponding radio models. We also modified all the necessary pipeline stages to suit our project requirements. Although OPNET provides a detailed PHY by means of separating the various PHY effects in to different pipeline stages, we chose not to use most of them. This decision was based on the fact that we could ignore a few features, such as the antenna gains and interference, and could simulate the effect of various other issues, such as path loss and background noise, using tools like MATLAB. Also, we intend to incorporate fading and multipath in our channel models. Using MATLAB to simulate these aspects of the PHY provides greater flexibility. As a result of our MATLAB simulations we generate bit error rate (BER) tables, which we use in OPNET’s error calculation pipeline stages. Since we use MATLAB for the BER calculations, we can also incorporate the effect of FEC coding in our simulations. The resultant BER table is merely substituted for the modulation tables in OPNET. The subsequent error correction stages drop packets even if a single bit is in error.

The OPNET DOCSIS models have been designed to perform checks at various points to see if network connectivity exists via the bus transmitters and receivers. This has caused us some difficulty in modifying the node architectures to support radio links. Also, replacing the bus link with the radio link seems to cause some interrupt related errors during a simulation run. As a result, at this time, we have retained the Bus Receiver (OPC_OBJTYPE_BURX) and the Bus Transmitter (OPC_OBJTYPE_BUTX) objects in the node models and have simply modified the pipeline stages to resemble the stages of the radio pipeline. The most significant modification that we had to make to implement such a model was to support two BERs in the pipeline, one for the uplink and one for the downlink. Although this is the default paradigm in the radio models, the bus-based models typically support only a single BER, the BER of the bus. We have assigned two BERs to the bus link model used in our simulations. The error allocation pipeline stages first determine the direction of flow for each packet and then subject it to the corresponding BER. Therefore, although our model uses a bus link for convenience, the internal implementation of the pipeline stages make it behave like a radio link. As indicated above, the BER values in our OPNET models are computed in MATLAB.

We use the scenario shown in Figure 1 to study the performance of the standard. We simulate a simple point-to-multipoint topology involving a single BS and two SSs. The system operates using FDD mode. When the simulation commences, the BS sends a downlink MAP (DL_MAP) message with the synchronization information. The SSs listen for the synchronization information and, upon its receipt, process the message to learn about the BS. The BS also broadcasts the uplink channel descriptor (UCD) to all the SSs. The SSs learn the uplink channel configuration upon receiving the UCD and subsequently send a range request (RNG_REQUEST) to the BS. The BS uses this to determine the propagation delay between each SS and the BS and resends this information in the range response (RNG_RESP) to the SSs as a frame timing-advance to compensate for the delay. The BS then sends a MAP with just contention opportunities. Depending on the kind of scheduling service these opportunities may or may not be used. The SSs request the BS for transmission opportunities and the BS transmits MAP messages containing information about grants. The SSs then transmit data on the uplink based on the MAP message.




Figure 1: Simulation model.

The SSs in this scenario operate using UGS scheduling. This scenario simulates File Transfer Protocol (FTP) sessions between the BS and the SSs. OPNET’s Application Config and Profile Config models are used to define the parameters for the FTP session. The BS is designated as the FTP server and the SSs perform FTP “get” operations to download files from the server. Our current interest is to study the performance of the IEEE 802.16 downlink. This application profile exactly suites our requirements. This scenario permits both fragmentation and concatenation. PHS is not used.

Since we are currently providing proof-of-concept results, our PHY burst profiles are simple. The PHY burst profiles used here are QPSK, 16-QAM and 64-QAM, each with no FEC. We assume AWGN channel with no fading or interference effects. The BER results for these profiles are derived from closed form solutions using MATLAB. BER results for different SNR values are used in the OPNET pipeline stages to drop packets. We assume that both SSs are equidistant from the BS and are thus

subjected to the same path loss and, therefore, the same SNR. Since we are interested primarily in the downlink channel performance in the system, the uplink is assumed to be error free. We use QPSK on the uplink and assume good SNR values. Table 1 summarizes the parameters used in the simulation.

Simulation Parameters	Values
Uplink Data Rate	40 Mbps
Downlink Data Rate	QPSK – 40 Mbps 16 QAM – 80 Mbps 64 QAM – 120 Mbps
Bytes per Mini-slot	8
Duplex Mode	FDD
Scheduling Scheme	UGS
Grant Size	1024 bytes
PHY Burst Profiles	QPSK, 16QAM, 64QAM, with no FEC
Channel Model Used	AWGN with no fading
Application Profile	FTP File Size - 100 Kbytes Request Inter-arrival – 30 s

Table 1: Simulation Parameters

IV. Results and Discussion

To study the performance of IEEE 802.16's downlink, we subject our simulation model to FTP traffic. We then vary the downlink burst profile and the corresponding channel data rates. We monitor three key metrics; end-to-end link layer delay, TCP delay and link layer throughput. The link layer delay is the time between when a packet is received by the link layer at the transmitter and the time when the same packet leaves the link layer and is passed on to higher layers at the receiver. It mainly comprises the queuing delay at the link layer and the transmission delay in the medium. The TCP delay includes any additional delays caused due to packet losses and re-transmissions. It is important to keep in mind that we assume constant symbol rates throughout and vary bit rates on the downlink when the modulation scheme changes. BER, not SNR, is the metric that directly affects TCP performance. BER depends on both the modulation scheme and the SNR. Therefore, in our simulation a given BER value corresponds to different SNR values for different modulation schemes.

Figure 2 shows the end-to-end Link Layer Delay experienced by a SS as a function of the SNR and the modulation scheme used. These results show the expected trend in delay performance at the MAC layer. As the downstream burst profile changes from QPSK to 64-QAM, the data rate on the air interface increases as a function of the number of bits per symbol. This results in subjecting MAC frames to smaller transmission delays. This is clearly observed with the delay curves for QPSK and 16-QAM. Since the data rate of the former is half that of the latter, its link level delay is almost double. Figure 3 shows the Transmission Control Protocol (TCP) delay experienced by a SS. Again, the expected trend is evident. At similar BER values the different burst profiles result in similar TCP delays. The changes in the data rates when burst profiles change do not reflect in these results simply because the link layer delays are far too insignificant and do not appreciably contribute to changes in the TCP delays.

In Figure 2, it is interesting to observe the delay curves for individual burst profiles. All three burst profiles show exactly how the link layer delay varies with TCP's reaction to packet losses. In each curve, as the SNR deteriorates (along with the corresponding BER), the delay begins to decrease until a particular SNR value is reached. When the SNR drops to levels where the BER is around 10^{-5} or 10^{-4} , the delay starts to increase. This is because of TCP's fast retransmit phase [7]. At BERs of 10^{-5} or 10^{-4} , packets get dropped frequently enough to make an appreciable difference in TCP's fast retransmits. During this phase TCP does not drop its throughput and continues to retransmit the lost packets. Thus, the queue size and the corresponding queuing delay at the MAC layer increases, thereby increasing the end-to-end delay. When the SNR drops further, the errors are so frequent that TCP timeouts occur and TCP's congestion control mechanism brings the congestion window size to one and drastically drops the throughput. This causes the rapid drop in the end-to-end delay. As one would expect, these points in the curve coincide with the sudden rise in TCP delay.




Figure 2: Link layer end-to-end delay.




Figure 3: TCP delay.

From Figure 3, it is clear that as the BER values increase, TCP delay increases greatly below a particular SNR value. For example, from these curves it is clear that, it is absolutely essential that the downstream burst profile change from 64-QAM to 16-QAM if the SNR drops below 25 dB. Similarly, a second switch to QPSK must occur if the SNR drops below 19 dB. This is where IEEE 802.16's RLC should play its part to tradeoff efficiency for robustness.

Figure 4 shows the throughput curves measured at a SS in the downlink. These results show that the throughput saturates at around 34 Kbps. This is an extremely low value given the high data rates at the PHY. The reason for this observation is that the FTP load offered to the system is simply too low to drive the system beyond this throughput. Since we only concentrate on producing proof-of-concept results here and since simulation run times with larger FTP loads are extremely long, we chose to use a relatively lightly loaded system. However, we do expect to see the expected trend with throughput being higher for higher modulation schemes while operating at high SNRs. These throughput curves show the effects of TCP's behavior. An increase in throughput is seen during the fast retransmit phase and a drastic drop in throughput is observed during the slow start phase of TCP's congestion control algorithm.




Figure 4: Traffic received at a subscriber station.

V. Conclusion and Future Work

This paper presents our progress to date in evaluating the performance of IEEE 802.16. We have developed an OPNET simulation model to simulate the WirelessMAN™ Air Interface.

We used simple burst profiles to test our model and to demonstrate the need for an algorithm to facilitate adaptive burst profile change on the downlink. Our results show that link level end-to-end delay does not significantly contribute to the TCP delay. The TCP delay seems to be a plausible metric to mark the burst profile operation limits in our algorithm.

As mentioned before, we have currently adapted a bus-based model with radio like pipeline stages to serve our purpose. We intend to move to a permanent radio model. The results shown in this paper incorporated simple burst profiles with no FEC. We plan to look at the capabilities of the various FEC schemes standardized in IEEE 802.16. The switch in burst profiles used here is drastic or coarse grain. We need to explore a large number of burst profiles to be able to incorporate fine-tuning capability in our algorithm. The PHY profiles assumed a simple AWGN channel with no interference or fading. In future we need to work with reference multipath channel models recommended in the standard. As explained above, this simulation scenario was lightly loaded. We need to work on larger networks subjected to much higher traffic loads before we finalize our algorithm.

References

- [1] IEEE 802.16-2001, "IEEE Standard for Local and Metropolitan Area Networks – Part 16: Air Interface for Fixed Broadband Wireless Access Systems," April 2002.
- [2] T. J. Eshler, "Adaptive Protocols to improve TCP/IP Performance in an LMDS Network Using a Broadband Channel Sounder," M.S. Thesis, Virginia Polytechnic Institute and State University, April 2002.
- [3] SCTE DSS 00-05, "Data-Over-Cable Service Interface Specification (DOCSIS) SP-RF1v1.1-I05-000714, "Radio Frequency Interface 1.1 Specification," July 2000.
- [4] C. Eklund, R. B. Marks, K. L. Stanwood, and S. Wang, "IEEE Standard 802.16: A Technical Overview of the WirelessMAN™ Air Interface for Broadband Wireless Access," *IEEE Communications Mag.*, June 2002, vol. 40, no. 6, pp. 98-107.
- [5] IEEE 802.16.1mc-00/01, "Media Access Control Protocol based on DOCSIS 1.1," 1999.
- [6] G. Sater, Personal communication, May 2002.
- [7] L. L. Peterson and B. S. Davie, *Computer Networks: A Systems Approach*, 2nd Edition, Morgan Kaufman Publishers, 1999.