Project	**IEEE 802.16 Broadband Wireless Access Working Group**
Title	**A Proposed Approach to Defining an Interoperable MAC/PHY Layer Scheduler for 802.16**
Date Submitted	**1999-10-29**
Source	Ray Sanders CircuitPath Network Systems P O Box 24950 Los Angeles, CA 90024 Voice: (310) 476-5063 Fax: (310) 471-7854 E-mail: rws@CircuitPath.com
Re:	This contribution is in response for the MAC Task Group Call for Contributions – Session #4 dated 22 September 1999.
Abstract	A Broadband Wireless Access system standard can only meet multivendor interoperability requirements if it includes a precise description of ways of apportioning limited available system bandwidth. This document describes a MAC/PHY scheduling protocol that overcomes shortcomings of other approaches. The scheduler is based on a simple, universal principle that allows Time Division Multiplex frames to be defined from two points of view. One allows a frame to be viewed conceptually as headerless variable-length groups of time slots where each group is assigned to a specific communication channel of any type, agnostic to user and network protocols. The second spreads each group's time slots nearly uniformly at precisely defined positions throughout a Time Division Multiplex frame. Among the scheduler's advantages are that 1) scheduling requires only a small amount of information being sent over a Base Station common signaling broadcast channel, 2) channels can be grouped by class of service that assures each class will meet quality objectives, 3) channel delay is minimal and inversely proportional to a channel's data rate, 4) efficient multilevel Forward Error Correction does not require complex interleaving schemes to overcome burst errors, and 5) both unicast and multicast services are supported.
Purpose	The purpose of this contribution is to outline a MAC/PHY bandwidth management scheduler protocol, a key element of any Broadband Wireless Access system.
Notice	
Release	
IEEE Patent Policy	

A Proposed Approach to Defining an Interoperable MAC/PHY Layer Scheduler for 802.16

Ray W. Sanders
CircuitPath Network Systems Corporation

1 Introduction

This proposal deals with an important aspect of devising an 802.16 MAC/PHY protocol intended to meet or exceed the 802.16 System Requirements as enunciated in Document 802.16s0-99/4 and subsequent revisions discussed in Session #3 interim meeting of the 802.16 working group. The focus of this document is an approach to scheduling transmissions from a Base Station to Subscriber Stations and return.

The discussion is largely tutorial in nature without in-depth details of specifics. The intent of this contribution is to stimulate discussion leading to merging the ideas of the proposed approach with ideas of others in the 802.16 Working Group.

The author believes that one of the most important aspects of the 802.16 standard must be bandwidth scheduling algorithms that can support small incremental bandwidth changes and are simple to implement. This must be easily certifiable for interoperability among multiple vendor embodiments. The intent of this contribution is to further the achievement of these goals.

The protocol approach outlined herein is applicable to any PHY layer transmission method that can incorporate a time-based scheduler. In other words, these methods are applicable to any PHY layer modulation schema including Frequency, Code and Time Division Multiplexing (TDM) where the transmissions can be divided into discrete cell slots[1].

The proposed methods can address all of the evaluation issues listed in the Session #4 Call for Contributions although this submission does not explicitly address a few of them at this stage.

2 Reference Model

In this document it is assumed that a single Base Station communicates with at least one Subscriber Station and that transmissions are made by a modulation method defined in the PHY layer of the standard. For purposes of this discussion, it will be assumed that the modulation methods make use of sinusoidal RF carriers whose center frequencies and bandwidths are well-specified as a part of the PHY layer. It is further assumed that information transmitted in both directions between a Base Station and a Subscriber Station is in a Time Division format. Even though the emphasis is on Time Division Multiplexing, the protocols described herein may be extended to other possible implementations such as Frequency Hop Multiplexing and Orthogonal Frequency Division Multiplexing (OFDM).

Duplex methods can be either based on frequency or on time — that is, either Frequency Division Duplex (FDD) where different RF carrier frequencies (and often bandwidths) are used in the two directions of transmission, or Time Division Duplex (TDD) where a single RF carrier of fixed bandwidth is time shared in the two directions.

This document proposes that the 802.16 standard be based on a flow model for all transmissions. This means that the basic operation of an 802.16 system is to assign bandwidth dynamically to support any defined bearer service. This approach results in simple and efficient agnostic support for both legacy and emerging services, particularly those associated with the Internet.

[1] *Cell slots* are defined to be discrete periods of time within a Time Division Multiplex stream each of which contains the same number of bits throughout a Time Division Multiplex frame. A cell slot may be either a defined-length group of bits within a Time Division Multiplex frame or may be a mini-slot in an upstream return path transmission.

3 Background

There is a growing realization that the future of the Internet depends on supporting message flows as opposed to depending only upon the routed-datagram model that has been the Internet's foundation. Work is proceeding within the IETF and other forums to develop the ability to support both Integrated Services and Differentiated Services. These have resulted in prioritization methods embodied within, for example, RSVP and DiffServ. In addition, MPLS has been accepted as a method of enhancing the support of flows.

The 802.16 standard has the unique opportunity to support this trend in a way that can enhance the competitiveness of Broadband Wireless Access systems compared to cable modem, DSL and other approaches. The fundamental problem of supporting message flows within an access network is to apportion available bandwidth in conformance with assignable Quality of Service (QoS) objectives and, at the same time, be protocol agnostic to any bearer service.

In a shared medium network such as a Broadband Wireless Access system, data are divided into timed frames, often with different framing strategies in the downstream and upstream directions. For the Base Station to Subscriber Station downstream case, transmissions can be a continuous stream of frames. For the Subscriber Station to Base Station upstream case, transmissions must be broken into controllable time segments.

A simple way to envision shared medium multiplexing is to realize that, regardless of the number of Subscriber Stations, its fundamental operation is that of two time division multiplexers connected by a single two-way transmission link. Transmissions from a Base Station to its Subscriber Stations appear as a single Time Division Multiplex data stream equivalent to one end of a Time Division Multiplex link sending information to the distant end. Each Subscriber Station chooses only those portions of the transmissions that are intended for it. Return path transmissions from each Subscriber Station, however, must be precisely timed so that they appear as a single Time Division Multiplex stream.

This leads to the conclusion that it is necessary to define two aspects for scheduling transmissions within a shared medium system. One is to provide a synchronized timing mechanism between a Base Station and all Subscriber Stations. The other is to define a scheduling protocol for supporting many information channels. The synchronization aspect itself is usually divided into two parts. The first is initial acquisition; the second is maintaining synchronization after initial acquisition. These functions have been solved by reserving small time slots within upstream Subscriber Station transmissions to allow for initial acquisition and contention-mode signaling for sending requests for bandwidth grants to a Base Station to which the Subscriber Station is connected. This approach is assumed to be a part of the PHY layer protocol that will evolve within overall 802.16 discussions.

The scheduling methods proposed herein have a long history. The first embodiment within Time Division Multiplexers was built in 1970 and thousand of units were subsequently produced. The technology is described in a patent that issued in 1972. Although the original could be used, newer methods developed since then can much more simply and effectively meet the needs of Broadband Wireless Access systems. These methods form much of the basis of this proposal.

4 The Problems of Multiplexing

Time Division Multiplexing within world standards organizations has been almost exclusively focused on support of digitized voice switching within the public telephone network. The approaches chosen have been serviceable for the Public Switched Telephone Network (PSTN). Unfortunately, they have been restrictive in terms of available channel bandwidths for non-voice applications. Just as important, current standards exclude a channel's bandwidth from being changed quickly once a channel has been set up. On the positive side for all future networks, transmission within the digital voice network is based on isochronous or plesiochronous clocking, a critically important part of meeting stringent delay and delay variation requirements required for a voice (or video or interactive multimedia) network.

At the same time that the telephone industry was focusing on digital implementations of the voice network, there has emerged a large body of work within academic institutions and computer companies that developed packet-based approaches to multiplexing that do not depend on TDM methods. This approach has been great for the data market where delay variation of traffic within a network was not terribly important. However, today's work-in-process complex protocols result from attempting to provide support for voice, video and other delay-sensitive services using asynchronously-timed multiplexing as occurs in IP- and ATM-based multiplexing.

Modern shared medium networks must support traffic that may or may not be delay sensitive. This requires incorporating the best of packet- and cell-based multiplexing with the best that time division multiplexing can achieve.

A scheduling protocol for a Broadband Wireless Access system must meet four specific objectives to satisfy optimally the overall 802.16 System Requirements:

- Enable rapid setup and clearing of bandwidth-on-demand channels
- Enable rapid changing of channel bandwidth after it has been established
- Require very little signaling and framing overhead
- Provide for bandwidth increments that support ***any*** mix of bearer services

The MAC/PHY layer Time Division Multiplex scheduler proposed herein meets these objectives with simple protocols.

5 Basic Ideas

Cell slots embedded within Time Division Multiplex frames

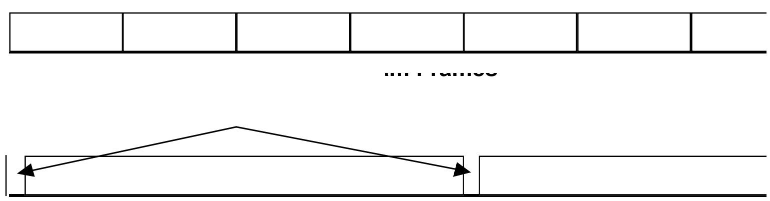

FIGURE 1 Typical Downstream and Upstream Time Division Multiplex Frames

For the purposes of this submission, a Time Division Multiplex system is assumed to include frames of data of defined length as illustrated in Figure 1. Each frame[2] is divided into an integer number of cell slots[3] (not shown in the figure) each of which is the same length (measured in bits). All downstream frames broadcast from a Base Station to all Subscriber Stations are assumed to be of the same length without any gaps. Data within any broadcast frame may consist of many channels of data meant for any number of Subscriber Stations. Any one channel is devoted to either a single Base Station to Subscriber Station connection or a multicast one-to-many connection from a Base Station to many Subscriber Stations.

---

[2] A frame of data in the context of this document relates to Time Division Multiplexing as distinguished from a frame of packet data or a frame of MPEG data. [Note that the word "frame" as used in DOCSIS and certain other protocols has a very different meaning and significance compared to the term as used in this document.] There is no requirement for a header to be present within a Time Division Multiplex frame to define how much information is allocated to each channel. Channel bandwidth assignments are made over a signaling channel that defines the number of cell slots per Time Division Multiplex frame assigned to the channel.

[3] In the discussion herein, the cell slots are taken to be for payload-bearing portions of a frame. Overhead symbols required for synchronization, contention/acquisition time slots, burst modem guard times, interframe time slots required for TDD operation, and error correction/control (e.g. For FEC) are not included. Including these parameters for more detailed system analysis is not difficult using the principles discussed in this contribution.

Subscriber Station to Base Station channels are all assumed to be one-to-one connections. The size of all upstream frames are of equal length, but they are separated by small interspersed time gaps that are dedicated to time slots used for contention signaling among the Subscriber Stations and for initial Subscriber Station timing acquisition. One frame may contain channels from many Subscriber Stations. There is no requirement that the frame sizes for downstream and upstream transmissions be the same. However, synchronizing the upstream frames with the downstream frames simplifies system design.

In both downstream and upstream directions, successive frames may be based on the same memory map. That is, channel bandwidth assignments may be persistent from frame to frame. Memory map changes are all made as required to meet the need to reallocate available bandwidth among the channels by transmitting the changes over a low error rate common signaling channel of efficient design.

Frames are made up of cell slots, the size of which depends on factors discussed below. Employing cell slots that are all of the same size in one direction of transmission simplifies the protocol design without loss of generality in supporting any bearer service. Channel bandwidth is managed by assigning an integer number of cell slots within a frame. Frame length (number of cell slots per frame) required depends on an acceptable frame rate. This, in turn, is determined by two parameters: 1) an increment of bandwidth that will support all desired channel data rates, and 2) the frame's cell slot size. The desired (maximum) frame rate equals the incremental bandwidth divided by the cell slot size.

Example of the relationship between cell slot size, incremental bandwidth and frame rate

In order to be able to assign channels without resorting to pad bits (or bytes) within cell slots (an approach that complicates a multiplexing protocol and wastes bandwidth), it is necessary that all channel rates to be supported are integer multiples of an incremental bandwidth. Starting with a set of desired channel rates, the largest possible incremental bandwidth for the set is the greatest common divisor (GCD) of the rates.

For example, suppose it is desired to support isochronous channel rates of 56,000, 64,000, 384,000, and 1,536,000 bps. The GCD of these rates is 8,000 bps which can be taken as the desired incremental bandwidth.

If a cell slot is 8 bits long, then the largest permissible frame rate would be 1,000 fps (frames per second). (In addition, any integer submultiple of 1,000 fps would also be an acceptable frame rate.)

If the incremental bandwidth were equal to the bandwidth of a single DS0 channel (64,000 bps), the required frame rate would be 8,000 fps (or integer submultiple thereof) for an 8-bit cell slot size.

A simple method of defining channels within a frame is to allocate an integer number of contiguous cell slots to a single channel. One cell slot per frame equates to the incremental bandwidth. By assigning $\boldsymbol{k}$ cell slots per frame to a channel, the channel's data rate is $\boldsymbol{k}$ times the incremental bandwidth.

IEEE 802.16 System Requirements set limits on system delay. A key part of achieving small delays is the Time Division Multiplex strategy. The minimum amount of delay possible may be constrained by the multiplexing strategy. For example, if one were to define a system where the frame rate is 1,000 fps and the cell slot size is 8 bits, the delay for all channels would be 1 millisecond (msec) using a *time-contiguous* method of assigning cell slots to channels. In other words, the multiplexing delay is the same for all values of $\boldsymbol{k}$. The incremental bandwidth in this case is 8,000 bps. [Note that, in this case, by sending a single cell slot per frame ($\boldsymbol{k} = \mathbf{1}$) the system supports directly compressed voice channels corresponding to the ITU G.729 recommendation without resorting to encapsulating data within an IP packet][4].

Delays that obtain from time-contiguous cell slot assignment are much larger than necessary. One of the main attributes of the proposed protocol is its ability to reduce channel delay, often dramatically.

---

[4] Several factors go into determining total delay. For the simplicity, delays given in this dodument are approximate, but in all cases are not so far from reality that they should not be taken as representative. They stated delays may be shorter or longer than reality, but in most cases are conservative (longer).

Limitations imposed by minimum cell slot size of Subscriber Station to Base Station upstream transmissions

An 8-bit cell slot size is satisfactory for Base Station to Subscriber Station transmissions. Unfortunately, it is not feasible for return path Subscriber Station to Base Station transmissions.

The minimum length of an upstream burst is limited by modem technology. This minimum length is further limited by the requirement of assuring that the transmissions from one Subscriber Station do not overlap (and thus interfere with) the transmissions from another Subscriber Station. Since there must be some guard time allowed between the end of one Subscriber Station's transmission and the beginning of that from another Subscriber Station, large burst sizes result in a smaller proportion of time being allocated to interburst periods than do short bursts. Also, long bursts make modem design simpler insofar as achieving carrier and symbol synchronization and determining required dynamic signal amplitude compensation for each Subscriber Station to Base Station path. (Amplitude compensation for fading and multipath effects is necessary to maintain high quality service.)

Short modem bursts decrease delay while longer bursts tend to improve bandwidth efficiency. The desirable cell slot size for different systems may vary widely implying that cell slot size should not be restricted as a part of the protocol, but should be a parameter that can be set by the design requirements for specific sets of applications.

Smallest supportable upstream data rates

The smallest supportable upstream channel data rate is equal to the frame rate multiplied by the (minimum) burst size of a modem. The minimum practical burst size for a bandwidth-efficient modem is currently of the order of 16 symbols and it is unlikely that shorter burst modems will become bandwidth efficient in the near future. For a QPSK modem, there are two bits per symbol resulting in a burst length of 32 bits. A 16-QAM modem with the same number of symbols has a burst size of 64 bits; a 64-QAM modem has a burst size of 96 bits. A modem designed to transmit a single ATM cell has a burst size of 384 bits. (The header bits do not count in this discussion since we are only interested here in payload bearing bits.).

If one were to specify that the frame rate for the system is to be 1,000 fps (a rate that assures reasonable delay in many applications), the minimum bandwidth (and hence, incremental bandwidth as defined above) supportable by these modems would be 32,000 bps, 64,000 bps, 96,000 bps and 384,000 bps respectively. Furthermore, if one were to specify that these modem burst sizes are to be taken as cell slot sizes, it would be necessary to limit supportable channels to only integer multiples of these rates.

An approach to solving the dilemma

There is a simple approach to solving the contradictory-sounding upstream transmission requirements. If one allows longer frame sizes with concomitant frame rates of less than 1,000 frames per second, a good compromise between delay and efficiency can occur.

The place to begin is to set the frame rate to be equal to the desired incremental bandwidth divided by the upstream cell slot size, that is, the minimum burst size. For 16-symbol QPSK, 16-QAM and 64-QAM cases, the frame rates required to support an 8,000 bps incremental bandwidth (and minimum channel data rate) are 250 fps, 125 fps and 41.67 fps, corresponding to frame periods of 4, 8 and 12 msec., respectively. For ATM cell size bursts (384 bits), the required frame rate would be 0.0208 fps corresponding to a frame period of 48 msec. To send an 8,000 bps channel one needs to completely fill a buffer whose length is equal to the cell slot size. If the data comes from an isochronously clocked 8,000 bps source, it takes 48 msec just to accumulate the bits to fill a buffer for an ATM-cell-length cell slot. If the accumulation of the bits in the buffer is not timed to be completed just before the assigned cell slot occurs within a frame, the multiplexing delay for sending a single-cell-slot-per-frame channel is could be any value between zero and to the frame period (48 msec). If the filling of a cell slot length buffer occurs at some random time with respect to the occurrence of the cell slot within a frame, the mean multiplexing delay is one-half the frame rate, which, in the case of filling an ATM-length cell with 8,000 bps data is 24 msec. It is not being suggested that using cell slot sizes equal to an ATM cell length should be incorporated in a Broadband Wireless Access system. However, the concepts just described pertain to cell slot sizes of any length. The concepts just work better in supporting low speed channels if cell slot sizes are kept small.

If such long frame periods were implemented and if all upstream transmissions were based on time-contiguous assignment of cell slots per channel, long latencies would result for all channels regardless of data rate. All data would have delay that can be equal to or greater than the frame period. Clearly, to minimize delay, it is necessary to distribute cell slots as uniformly as possible throughout long frames. This is the case even with shorter modem burst sizes than a complete ATM cell[5]. The proposed scheduler protocol accomplishes this objective in a simple and flexible manner.

6 The Proposed Scheduler Protocol

An ideal scheduling protocol would accomplish two objectives that, on the surface, appear incompatible. One is to allow for assignment of bandwidth as a set of contiguous cell slots. The other is to distribute cell slots assigned to a single channel nearly uniformly throughout a Time Division Multiplex frame. A major attribute of the proposed protocol is the simultaneous accomplishment of both of these objectives.

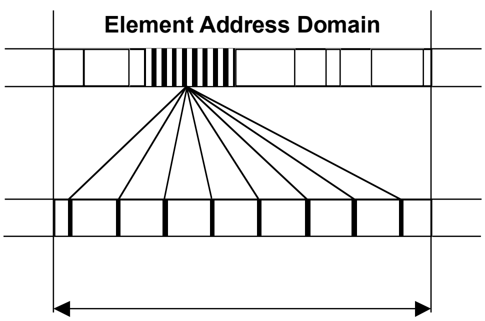

The Idea underlying the proposed protocol

The basic idea behind the proposed scheduler is its ability to address Time Division Multiplexing simultaneously in two domains that are connected by a simply defined and simple-to-implement mathematical transform. This transform is in some ways analogous to Fourier transforms that enable envisioning a single (continuous or discrete) time series in both a frequency and a time domain. A major benefit is the ability to define a channel as a set of contiguous cell slots per frame in one domain (channel bandwidth assignment) and, by applying the transform, produce cell slots in a second (time) domain that are nearly uniformly spaced as shown in the diagram on the right.[6]

Delays caused by assigning cell slots contiguously in the time domain

As described above, if cell slots assigned to a channel are placed contiguously in time in a Time Division Multiplex frame, it is necessary to buffer enough payload data to fill the contiguous cell slots. The delay caused by this approach can be large. The buffer delay consists of 1) time required to fill the buffer with the amount of payload data to be carried in the cell slots, and 2) wait for the contiguous cell slots to occur within the Time Division Multiplex frame. Spreading the cell slots nearly uniformly throughout a Time Division Multiplex frame can drastically reduce this delay. The multiplexing delay now becomes associated with the inter cell slot period as opposed to the delay imposed by defining channels as time-contiguous cell slots and the buffer size needed now is only that for a single cell slot, not the period of those contiguously assigned. The mean multiplexing delay[7] is reduced to the frame period divided by the number of cell slots per frame assigned to a channel.

[5] ATM cells themselves could be fragmented using the proposed protocols to overcome the disadvantages of using an ATM-cell-length cell slot size.

[6] It is not mathematically possible to distribute cell slots exactly uniformly throughout a frame in all cases. This causes a small amount of jitter that, for isochronous and plesiochronous channels, requires a dejitter buffer to be included within a system. However, the amount of the delay added by such a dejitter buffer is small and mathematically bounded. The result is simple-to-implement traffic shaping with zero delay variation for isochronous and plesiochronous services.

[7] The word "mean" in the term 'mean multiplexing delay' relates to the statistical average of a single channel's delay, *not* the delay averaged across an aggregate number of channels at a single data rate. The reason for including the word relates to the small amount of cell slot jitter for a channel at a receiving station as discussed in Footnote 6.

Figure 2 illustrates this smaller mean delay for various cell slot sizes within a Time Division Multiplex system. The delay is inversely proportional to the bearer channel data rate and directly proportional to the cell slot size.

For a cell slot size of 8-bits (assumed to be used for the downstream broadcast path), the mean delay is 1,000 µsec (1.0 msec) for an 8,000 bps channel. For a DS0 (64,000 bps) channel it is 125 µsec. The mean delay for (return path) cell slot sizes of 32, 64, 96 and 384 bits (the size of a single ATM cell's payload field) increases by factors of 4, 8, 12, and 48, respectively, compared to delay for the 8-bit case.

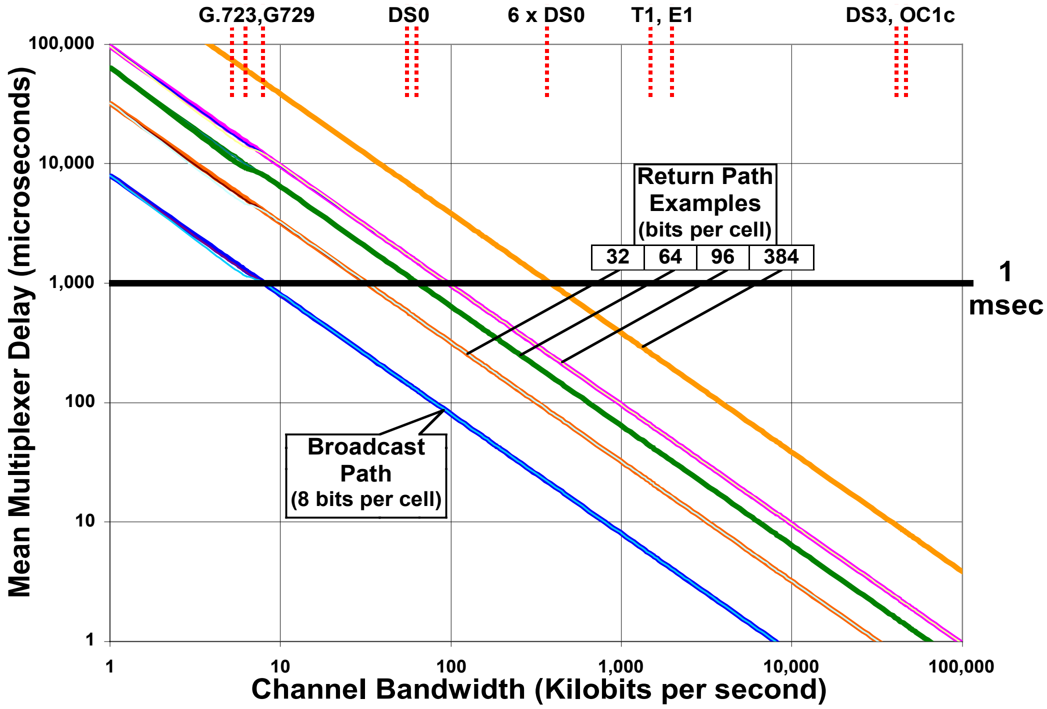

FIGURE 2 MEAN MULTIPLEXER DELAY VS BEARER CHANNEL RATE

Supporting low speed channels

Shown at the top of Figure 2 are various bearer channels that many service providers often need to support. Included is support for G.723 (data rates of 5.3 and 6.3 Kbps) and G.729 (with a data rate of 8.0 Kbps). These data rates are important for both cellular telephone service and Voice Over IP (VOIP). And Voice over Frame Relay (VOFR).

Frame sizes can be made very large in order to support very small incremental data rates and assure very high bandwidth utilization efficiency. For an extreme example, consider simultaneously supporting 5,300, 6,300 and 8,000 bps G.72x channels in a common frame format. The greatest common divisor of these rates is 100 bps so that a frame containing a single cell slot that provides a data rate of 100 bps will support the desired rates (as well as any other integer multiples of 100 bps). Channel assignments of 53, 63 and 80 cell slots per frame will exactly support the desired rates.

The frame sizes in this case are indeed large. For cell slots that contain 8 bits, the frame period is 80 msec; for 32-bit, 64-bit, 96-bit and 384-bit cell slots, the frame periods are 320, 640, 960 and 3,840 msec respectively. The point of this example is to illustrate the flexibility of the proposed protocols and not to suggest adopting specific parameters as standards, but instead, perhaps only as examples for clarity. There is no need to specify specific frame sizes or rates within the standard, but, instead, only specify structural rules for choosing parameters to meet any network need. (Obviously, the proposed approach will allow vendors to field systems that can be made field programmable without hardware modification to meet changing user needs, making the proposed approach truly non-obsolescent.)

As shown in Figure 2, the, the mean delay for upstream low speed channels is high. For a single 5,300 bps channel, it is 6.0, 12.1, 18.1 and 72.5 msec for 32-bit, 64-bit, 96-bit and 384-bit cell slot sizes respectively. However, for interactive traffic (such as compressed voice) these large numbers are mitigated somewhat when an 8-bit cell slot size is chosen for downstream broadcast transmissions. This is so since Broadband Wireless Access systems supporting interactive applications act as two-way connections. For the 5,300 bps case, the downstream delay is 1.51 msec using 8-bit cell slots. Thus, the two-way mean delay for the cell slot sizes cited are 7.6, 13.6 and 19.6 for 32-bit, 64-bit and 96-bit, respectively. Even using ATM cell sizes in the upstream return path may not be completely out of the question since the two-way delay in this case is 74.0 msec (equivalent to 37.0 msec one-way delay in each of two directions) — not totally unacceptable if there are truly compelling reasons for using cell slot sizes equal to an ATM cell (an instance that the author does not expect to occur in the context of this example, particularly since the delays stated are only multiplex delays and do not include buffer fill time delays).

Using channel aggregation to lower delay

The negative effects of large cell slot sizes in upstream return path transmissions can be mitigated for those cases where supporting single low speed channels is not required. Return traffic from a Subscriber Station can consist of an aggregation of channels. In this case, the multiplexing delay for the aggregation is determined by the cell slot size and the *aggregate* data rate. Suppose that channels within the aggregation are defined using 8-bit cell slot size multiplexing. The delay for a single channel is then equal to the aggregate data rate delay plus the delay associated with the 8-bit cell slot multiplexing.

For example, suppose that a service provider offers premium service that includes a minimum full duplex Committed Information Rate (CIR) of 384,000 bps. (6 DS0 bandwidth equivalent) and further, agrees to provide additional bandwidth on an “as available” basis. The mean delay for this CIR using 32-bit, 64-bit, 96-bit and 384-bit cell slot sizes is 83.3, 166.7, 250.0 and 1,000.0 µsec, respectively. The 384,000 bps can be used as transmission bandwidth for aggregated channels, each of which can be defined using 8-bit cell slots. The total mean delay for any single channel would be (approximately) the sum of the CIR delay plus the mean delay for an 8-bit cell-slot-defined channel. The chart shown in Figure 2 can be used to determine the approximate delay for any combination of aggregate channel delay and embedded channel multiplexing delay based on 8-bit cell slots. For example, the mean round-trip multiplexing delay for an 8,000 bps embedded channel embedded within this 384,000 bps CIR bandwidth would be 2.1, 2.2, 2.3 or 3.0 msec for the combined 8-bit cell slot downstream path plus either the 32-bit, 64-bit, 96-bit or 384-bit upstream path — much smaller than round-trip delay for an individually-supported 8,000 bps channel.

Applicability to ATM centric networks

A number of vendors are incorporating ATM technology into their Broadband Wireless Access systems. Although committing to use ATM cells as the basis for systems limits system flexibility, there may be current market reasons for using ATM. Also, there is an installed base of ATM systems that could make effective use of the proposed protocols without requiring “fork-lift” upgrades.

There are three major benefits to be gained by using the proposed approaches within ATM-based Broadband Wireless Access systems. One is to eliminate the ATM bandwidth “head tax”. ATM cells can be allocated as band-width-on-demand channels where there is no need to send header addresses over the Broadband Wireless Access links. Headers necessary for interconnection with ATM backbones can be inserted (or, in the reverse direction, deleted) at a Base Station.

A second benefit is to schedule ATM cells to be nearly uniformly spaced so that there is no need for additional traffic shaping within the Broadband Wireless Access system. This not only makes system implementation simpler, but it also reduces the need for traffic shaping buffers (for data originating within the Broadband Wireless Access system’s sources of ATM-encapsulated data) with their concomitant contribution to higher system delay.

The third benefit is the use of ATM channels to send traffic as aggregations of lower speed channels. The protocols described in this contribution can be used to make certain that the number of pad bits included within ATM cells is minimized and that channel delay is made much smaller for lower speed channels than current ATM practice.

A fourth benefit could be listed. As mentioned above, ATM cells can be fragmented within the proposed approach that overcomes the delay limitations of ATM cell size in some instances.

Although not explicitly discussed in this document, similar points exist for applying the proposed approaches in an MPEG environment, such as that being used within the DOCSIS environment.

7 The Persistence of Flows

In the beginning of the Internet's development, the metaphor for building large packet networks was that the Internet is much like the post office. Information is divided into small packets, each of which contains a destination address. This address allows a network node to act like a post office that can read the packet's address (header) and send it to another post office whose job, in turn, is to send the packet closer to its intended destination. The concept of a "connection" between a the post office where a user deposits his packet (a source network edge node) and a destination post office that delivered the packet (a destination network edge node) was not (and, in most respects, still is not) a part of the Internet's design. This connectionless network concept, important for the military-oriented initial applications of the Internet, has prevailed within networks that no longer place priority only on assuring survivability in the sense intended in the 1970s. The concept of survivability as an issue has now been replaced by specifications for how many "9's" of availability are required for information from a source node to be delivered to a destination node within some specified period of time. Delivering *flows* of packets and not merely individual packets has now become a dominant theme in designing the next generation Internet.

Our intuitive idea of a flow of information involves assuring timely delivery of information as being as important as assuring that delivery occurs at all. This intuitive idea can be converted to concrete parameters of network design specifications by paraphrasing the abstract idea into a question: "How many bits of information must be delivered in a defined amount of time?" In other words: "What is the bandwidth required of a connection between a source and a destination node to achieve delivery of the information in a defined amount of time?" Posing the question in this way gets us closer to finding ways to build an Internet that meets today's users demands. However, this question is still not complete. We are not dealing today only with static messages of known total size. We must be able to deliver information where the instantaneous required bandwidth of a connection can change during the information flow. (Video and multimedia services are examples.) Merely setting up a connection of fixed bandwidth is not enough.

An access network is not the whole Internet. But, how an access network deals with the bandwidth management question has a major influence on end-to-end network performance. Handling flows to and from legacy PSTN, IP and ATM networks with minimum delay is crucial to overall network performance.

In the context of a Broadband Wireless Access system, we can restate the critical basic question in the following terms: "How do we define the 802.16 protocol to enable variable bandwidth-on-demand connections to be made that maximizes efficient utilization of available system bandwidth and minimizes user channel delay?" Working with others in the 802.16 Working Group to provide a 'best of breed answer' to this question is the focus of the proposed protocol.

Flow persistence allows design of Broadband Wireless Access systems that make effective use of a bandwidth control mechanism that is the essence of the proposed scheduler protocol. The reason is that flow persistence enables assignment of bandwidth to a channel that need ***not*** be changed every few microseconds (as discussed later in this section). Simulation studies so far have not uncovered applications where bandwidths need be changed in periods less than tens of microseconds. It is likely that none will be found that relate to the delivery of information to human users — humans do not have sensory capabilities to detect shorter periods.

That flows are an integral part of today's Internet needs no formal proof. Measurements on the backbone network of a major U.S. carrier show that over 75% of the traffic (in bits per second) delivers World Wide Web pages that are flows. Obviously, FTP sessions involve flows as do VOIP, video, multimedia and some E-Mail traffic. The conclusion is that the most traffic across the Internet results from persistent flows.

Flows may be recognized in two ways. One is from the nature of traffic from a single source. The other is the nature of traffic from an aggregation of sources.

Flows for various traffic types

The simplest case of a flow for a single source is a Constant Bit Rate (CBR) channel. Here bandwidth is allocated for periods of time often measured in seconds, minutes, hours, days or years. (For example, CBR periods may be allocated for seconds in the case of Time Actuated Speech Interpolation (TASI) operation applied to voice and other highly interactive channels. Full period Committed Information Rate channels may persist for years.)

Variable Bit Rate (VBR) channels are similar in that they usually persist for extended periods. In these cases, the bandwidth assigned to any channel can change over relatively short periods of time (tens of milliseconds or seconds).

Available Bit Rate (ABR) traffic is a significant part of the Internet. Persistent ABR traffic flows occur for file transfers of many types including World Wide Web applications.

With the growing emphasis on Quality of Service categories, an important adjunct to the persistence of flows from a single data source is the persistence of flows for an aggregation of traffic from multiple channels belonging to a single QoS class of service. By implementing MPLS, DiffServ and similar protocols within traffic sources (particularly for ABR traffic), simple buffering mechanisms can be used to aggregate packets by service class so that system bandwidth can be apportioned appropriately among the classes and the intensity of the traffic within each class.

The advantage of dealing with flow persistence rather than attempting to build Broadband Wireless Access systems that deal with traffic only on a packet-by-packet basis (which is still the case with packet protocols that rely on priority bits for QoS scheduling) is that the amount of signaling bandwidth required to set up Time Division Multiplex frames to handle flows need be only a fraction of that required for assigning bandwidth in packet-by-packet cases.

In addition to the above cases, there are many options that can be implemented to handle packets that are not a part of persistent flows that require only a small amount of signaling bandwidth.

Signaling requirements for assigning bandwidths to flows

Figure 3 shows the duration of flows for various message sizes that are transmitted over a wide range of bandwidths. This simple chart is helpful in assessing the importance of flows as they affect service quality. It gives insight into how quickly bandwidth assigned to flows must be changed to satisfy quality objectives within the constraints of available bandwidth.

One example is to consider an individual Web page that contains 50,000 bytes of information. The amount of bandwidth required to send this information in two seconds is 200,000 bps. In other words, to handle a flow from a Web server to a subscriber would require setting up bandwidth from the server to the subscriber for a period of two seconds. For servers that are directly connected to a Broadband Wireless Access system, this would provide truly impressive service to subscribers using today's access networks!

Unfortunately, the Internet as presently constructed does not work this way. Packets from distant nodes do not have dedicated 200,000 bps channels over which to send a 50,000 byte Web page. Packets arrive in clumps where the average data rate is often not predictable. Setting up a 200,000 bps channel within a Broadband Wireless Access system to support the 50,000-byte flow where the average arrival rate of packets does not equate to that much bandwidth would be wasteful. To use bandwidth more efficiently, one could employ a smoothing buffer between the Internet and the Broadband Wireless Access system. The bandwidth allocated could be adjusted many times per second to match the allocated bandwidth to the varying average arrival rate of the 50,000-byte flow.

Even though it is not contemplated that setting up channels on a packet-by-packet basis will be the major application of the contemplated protocols described, it is instructive to consider the case of setting up bandwidth to send individual packets of either 500-byte or 1500-byte lengths that might be used to transfer the 50,000-byte flow of the example.

The amount of time consumed by sending a 500-byte packet at 200,000 bps is 20 msec while the time to send a 1500-byte packet is 60 msec. Using the proposed protocols, the amount of signaling information required to set up bandwidth for each packet need be only a few bytes — far less that assumed for a protocol such as SS7 or Q93B. (An achievable goal appears to be to use no more than five percent of the total system bandwidth for bandwidth management signaling. As the Internet becomes better behaved in its ability to smooth out the arrival rate of packets from a network within a flow, the amount of signaling bandwidth will be far less than this. Also, in those instances where an aggregation of flows is an option, the amount of signaling bandwidth can be very small.)

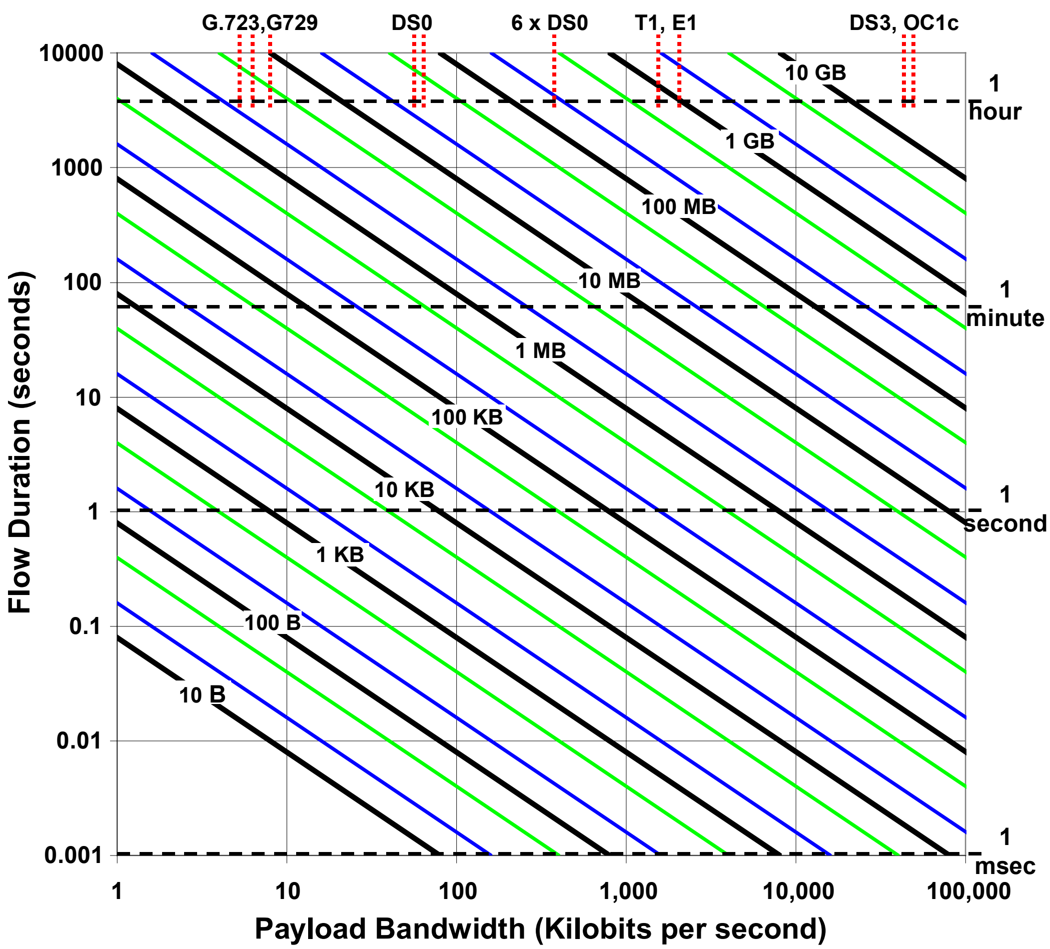

FIGURE 3 FLOW DURATION VS PAYLOAD BANDWIDTH FOR FLOW SIZES (MESSAGE SIZES) FROM 10 BYTES TO 10 GIGABYTES

Currently, on an Internet backbone network, about one-third of the packets are about 40 bytes long. These are UDP packets that are not a part of flows. In the downstream direction, these packets can be sent within a single broadcast channel that can be received by each Subscriber Station where the bandwidth of the channel is adjusted to the average arrival rate of these short packets. In the upstream direction, if only a single flow is being sent from a Subscriber Station to the Base Station, these short packets might best be embedded along with the longer packets associated with the flow. Alternatively, a separate limited bandwidth channel may be used to aggregate the small packets.

Distributing channel bandwidth assignment by flow persistence of various classes of service

Taking into account flow persistence of various classes of service in making channel assignments can make a substantial difference in overcoming the effects of transmission errors on overall service quality.

As discussed briefly in Section 6, cell slot positions within a Time Division Multiplex frame are referenced in two domains that are related by a simple transform between them. The channel assignment domain (called the Element Address Domain) enables an integer number of cell slots to be assigned to a specific channel (or aggregation of channels) using a simple metaphor. Within the context of a single Time Division Multiplex frame in the Element Address domain, channels assigned at the beginning and the end of the frame enjoy the uniqueness of having one of their two boundaries fixed. This suggests that channels assigned to either of these two positions be chosen as those with the greatest persistence. There will be fewer changes made to the multiplex structure for these channels than for those assigned to the middle of the frame in the Element Address domain.

The reason this is important is that in changing the multiplex map of channel assignments, there is always some small chance that an unrecoverable error will be committed in the transmission of channel assignment information from a Base Station to the Subscriber Stations. Even though the proposed protocol contains simple error recovery mechanisms for these rare events, every effort should be made to minimize their invocation.

The proposed approach to making channel assignments and bandwidth changes is to use the “less is more” principle. That is, the protocol should require the transmission of the smallest possible amount of information (number of bits) from a Base Station to Subscriber Stations to unambiguously define Time Division Multiplex map changes. This both minimizes signaling bandwidth required and minimizes the likelihood that a transmission error affects map synchronization among the stations.

Restricting the location of channels that have short (expected) flow persistence to the middle of the Element Address domain frames will enhance memory map reliability. For example, all default fixed rate common signaling and framing channels located at the ends of an Element Address frame (for a given system embodiment) are infinitely persistent. CBR subscriber channels of long persistence, followed by channels with progressively less persistence fill toward the center of the Element Address frame. UBR channels will usually be relegated to the middle of the Element Address frame.

A further discussion of error control topics independent of multiplexer map signaling occurs in Section 9.

8 Effects of Time Division Duplex (TDD) Operation

A Broadband Wireless Access system employing Time Division Duplex operation uses a single radio frequency for both downstream and upstream transmissions. Thus, while a Base Station is transmitting, it cannot receive signals from any Subscriber Station. In like manner, no Subscriber Station can be assured of receiving Base Station signals while any Subscriber Station is transmitting. Added to the transmit periods is a period of silence where neither a Base Station nor a Subscriber Station can transmit. This period is equal to the two-way propagation delay from the Base Station to the most distant Subscriber Station and return. This period of silence ensures that no station begins transmitting before all in-transit signals have been received by a distant end. (For a Broadband Wireless Access system designed for a maximum Base Station to Subscriber Station distance of 5 km, the period of silence must be about 34.0 μsec.)

9 Effects of Proposed Protocol on Error Control

A major concern within Broadband Wireless Access systems is the errors caused by the transmission environment. Forward Error Correction (FEC) methods are an appropriate approach to overcoming these errors. Well-known FEC methods exist (using Reed-Solomon codes for example) that can enhance transmission quality. However, there are two factors that need be considered in designing MAC/PHY protocols that can satisfy the requirements stated in the 802.16 System Requirements document. One is the fact that the more robust the FEC technique, the more system bandwidth is required. The other is that not all services require the same level of correction; some services such as voice transmission can tolerate error rates that are much higher than services such as video transmission using MPEG compression. The proposed protocol can enhance the effectiveness of standard FEC techniques by addressing both of these factors.

Supporting multiple levels of service quality and error control

Since the proposed protocol relies on an Element Address domain where channel bandwidth assignments are defined as a set of contiguous cell slots per frame, it is possible to group cell slots by the error rate requirement for various service types.

This is particularly important for the downstream broadcast direction. This is the source of multicast (and the expected usual source of unicast) of error-sensitive MPEG transmissions and, more importantly, common channel bandwidth management signaling from a Base Station to Subscriber Stations. For example, suppose that it is required to achieve better than a 10^{-11} bit error rate for signaling and MPEG transmission channels. Each of the Element Address ranges assigned to these services can be subdivided into error correction blocks that can be encoded independent of specific bearer channel assignments. Such encoding is performed at a Base Station and decoded at each (relevant) Subscriber Station. The error correction blocks need not be restricted to individual channel assignment boundaries. (This does not preclude sending unicast MPEG streams upstream using the same methods.)

Automatic cell slot interleaving

FEC techniques such as Reed-Solomon encoding rely on appending a set of code bits to a block of data where the bit states depends on the content of the data block. The number of bit errors within a block that can be corrected by the appended code bits depends on the number of code bits added. If bit transmission errors are randomly distributed and the bit error rate is known with assurance, the probability that errored bits can be corrected can be calculated. However, if errors occur in bursts, it is much more difficult to estimate the effectiveness of the error correction coding.

The conventional way to deal with burst errors is to interleave signals in a manner that distributes bytes from a single channel or aggregate multiplex stream in time. The effect of the interleaving is to divide a long block of, say, ***nM*** bytes into ***n*** blocks of ***M*** bytes each. By dividing the long block into ***n*** interleaved shorter blocks, a burst that caused ***k*** bits of the long block to be in error within a short contiguous set of bits, result in roughly ***k/n*** error bits in each of the short blocks. Error correction decoders for a small number of errors in a block are simpler to mechanize and faster than are decoders for a large number of errors. Hence, overall system performance is improved by interleaving. (Once the errors are corrected, the interleaved bits must be reassembled into the block size of the original signal.)

Within the environment of the proposed protocol, individual cell slots are already spaced apart in time so that they can be encoded in small blocks. Further interleaving is not required. A thorough presentation of this capability is beyond the scope of this document, but it is an important attribute of the proposed protocol.

10 Conclusions

This documents proposes that the 802.16 standard be based on a Time Division Multiplex protocol that can efficiently schedule the limited available bandwidth of a Broadband Wireless Access system. There are two primary goals that the proposed approach meets. One is to define a simple protocol that can be implemented by any vendor with high assurance that each such embodiment will interoperate with any other. The second is to produce a protocol that meets or exceeds requirements specified in IEEE document 802.16s0-99, “Preliminary Draft Working Document for 802.16 Broadband Wireless Access System Requirements”. As of the date of this document, the current version is version 5, dated 21 October 1999.

The specifics of the proposed protocol include means of dealing with Time Division Multiplex frames from two different perspectives — one where groups of cell slots can be assigned to individual channels and another where cell slots of any one group are distributed nearly uniformly throughout a frame. The second perspective allows the scheduler to provide unique capabilities compared to other known methods of apportioning Time Division Multiplex bandwidth among a plurality of channels.

The protocol includes a downstream Base Station to Subscriber Stations broadcast transmissions based on a Time Division Multiplex structure of cell slots that can be of any size, but preferably those composed of only a few bits of information. Examples included in this document assume that these cell slots are all 8 bits long. Short cell slots enhance system performance since they enable very low multiplexing and access buffer delays.

Upstream Subscriber Station to Base Station transmissions also include minimal length cell slots, but in this case, assuming that each cell slot is equal in length to a transmission burst from a Subscriber Station, cell slot size is limited by modem technology to a length larger than can be used in the downstream broadcast direction. For purposes of illustration in this document, upstream bursts are assumed to be 16 symbols long, where each symbol can be QPSK, 16-QAM or 64-QAM. Examples using ATM cell length bursts (384 bits) are also included.

Results shown in this document include an outline of the protocol Time Division Multiplex structure and the supportable system modes and their performance metrics that result therefrom. The supportable system modes and performance metrics include the following:

- Multiplexer delay for all channels is inversely proportional to cell slot size.
- Multiplexer delay for an individual channel is inversely proportional the channel's bandwidth.
- Low speed channels (e.g., 5,300 bps) and high speed channels (Mbps) can be supported in a single Time Division Multiplex structure with no compromise in performance.
- Where cell slots sizes cannot be reduced to values that result in a desired channel delay for low speed channels, such channels can be aggregated as sub-channels within a higher speed channel with small delay so that the protocol may be applied to multiplex the low speed sub-channels using small cell slot sizes (such as 8 bits). Per channel delay is much reduced in many cases.
- Bandwidth allocation can be applied to information flows of any type (both packet-based and STM based) such as CBR, VBR, ABR and UBR channels.
- Both unicast and multicast services are supported.
- Both Frequency Division and Time Division Duplex (FDD and TDD) are supported.
- It is necessary only to send bandwidth scheduling information (over a broadcast common signaling channel from the Base Station) when the data rate assigned to any channel needs to be changed.
- The frequency with which bandwidth need be set up for a flow rarely will need to be on a packet by packet basis, although that mode is supported.
- Bandwidth scheduling supports current and emerging ATM and IP QoS protocols including RSVP, DiffServ, MPLS and others.
- The protocol can be used where it is mandated that broadcast and return transmissions be based on ATM cells or MPEG frames.
- Aggregation of channels by class of service is supported so that the included channels can be assured of meeting the class's quality objectives.

- Forward Error Correction (FEC) can be applied to channels according to their performance needs so that there is no need to either 1) protect all transmissions to the high performance specification of a system's most error-critical channel, or 2) implement tandem FEC algorithms.
- Interleaving traffic to achieve high performance FEC capability in the presence of burst noise is not required (unless cell slot sizes are large) since the protocol already distributes cell slots nearly uniformly throughout a frame.

Even though no detailed description yet exists for the proposed protocol as of the date of this document, there is a large legacy of efforts incorporated in other protocols that are directly applicable to proposed approach. A dedicated effort of a small group of 802.16 members should be able to propose a detailed protocol within the time constraints of the 802.16 efforts. The author believes that the universality, performance advantages and non-obsolescence nature of the protocol are sufficient to justify that an 802.16 working party be formed to complete a protocol design.