Project	**IEEE 802.16 Broadband Wireless Access Working Group**
Title	Proposal for BWA Physical Layer Protocol based on DVB Downstream and DOCSIS Upstream
Date Submitted	1999-10-29
Source	Karl Stambaugh Motorola 8220 E. Roosevelt R1106 Scottsdale, AZ 85257 Voice: (480) 441-7842 Fax: (480) 675-2116 E-mail: karl_stambaugh@email.mot.com
Re:	This is in response to the PHY Layer Task Group's Call for Contributions for meeting #4. Document 80216p-99_01
Abstract	The proposed PHY protocol is based on a modified version of DVB downstream and DOCSIS upstream protocols. This proposal is a summary of the minimum set of requirements that are needed to define the physical layer. The details of the requirements are to be provided later.
Purpose	The authors desire that the 802.16 working group incorporate all or part of the proposal into the 802.16.1 standard.
Notice	This document has been prepared to assist the IEEE 802.16. It is offered as a basis for discussion and is not binding on the contributing individual(s) or organization(s). The material in this document is subject to change in form and content after further study. The contributor(s) reserve(s) the right to add, amend or withdraw material contained herein.
Release	The contributor acknowledges and accepts that this contribution may be made public by 802.16.
IEEE Patent Policy	The contributor is familiar with the IEEE Patent Policy, which is set forth in the IEEE-SA Standards Board Bylaws <http://standards.ieee.org/guides/bylaws> and includes the statement: "IEEE standards may include the known use of patent(s), including patent applications, if there is technical justification in the opinion of the standards-developing committee and provided the IEEE receives assurance from the patent holder that it will license applicants under reasonable terms and conditions for the purpose of implementing the standard."

Proposal for BWA Physical Layer Protocol based on DVB Downstream and DOCSIS Upstream

Karl Stambaugh, Glen Sater
Motorola Satellite Communication Group

Introduction

This physical layer proposal is based on a Digital Audio Video (DVB) project standard and the Data Over Cable Service Interface Specification (DOCSIS) standard. Some modifications to these standards are proposed in order to improve the latency, improve the interference immunity, and increase the capacity of the referenced standards.

This proposal defines the physical layer protocol for a flexible, high-bandwidth, point-to-multipoint system for the delivery of Internet Protocol (IP), ATM, SDH (or PDH) and video distribution services. This protocol will enable the bi-directional delivery of ATM and IP traffic and SDH services between a central base station (BS) and a number of subscriber stations (SS).

The proposed physical layer also enables the distribution of MPEG-2 video from the BS to all SS. The MPEG-2 video frames and the MAC layer frames, which carry the other services, can both be multiplexed within the same downstream channel.

Downstream Reference Model

The proposed downstream physical protocol is a modified version of on the ITU-T J.83 Annex A specification. The interleaver and modulation formats are modified to reduce latency, improve the interference immunity, and increase the channel capacity.

The downstream is continuously transmitted by the BS on a single carrier. The physical layer multiplexes both video and MAC frames onto the single downstream channel. Both the video and MAC frames utilize the MPEG-2 frame format. The multiplexing of these two sources is controlled by the MAC layer. The downstream data is then processed according to the BS downstream physical layer reference model that is shown below.

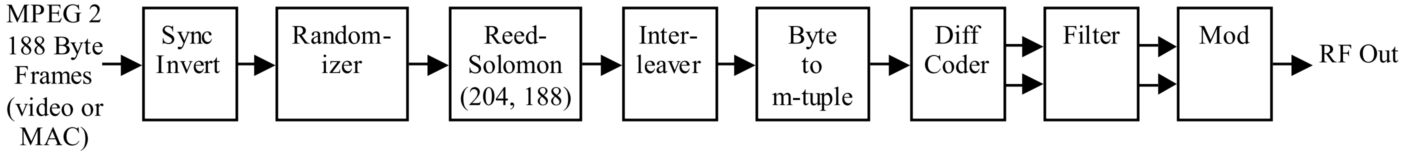

The PHY layer reference diagram for the downstream receiver is the reverse of the above one, except that an optional adaptive equalizer can be inserted at the appropriate point in the chain if desired. The specifics of the adaptive equalizer design is not specified within this document, and is left up to each manufacturer.

MAC Interface

The MAC layer generates a data stream that has the same framing format as MPEG-2. Each 188 byte frame contains a single synchronization byte (i.e. 47_{HEX}), a three-byte header, which contains service, scrambling, and control information, and then 184 bytes of data.

The PHY layer also can receive MPEG-2 video frames. The video frames and MAC layer frames can be multiplexed into the physical layer of a single downstream channel.

Sync Invert and Randomization

The MPEG-2 stream is scrambled by multiplying a pseudo-random bit sequence (PRBS) to the data stream. The PRBS code is reset after every 8th sync byte, which is bit-wise inverted as a means of indication to the decoder. During the sync bytes in the following seven frames, the PRBS code continues to run but is blanked.

Forward Error Correction

Following the energy dispersal randomization process, shortened Reed-Solomon encoding is performed on each randomized MPEG-2 transport packet, with $t = 8$ (8 erroneous bytes per transport packet can be corrected). This process adds 16 parity bytes to the MPEG-2 transport packet to give a 204 byte code word (204, 188). The MPEG-2 sync byte is included within the 188 bytes, but remains unaffected by the RS encoding.

Interleaver

The system uses a convolutional interleaver that is similar to that specified in Annex A, but which has 17 branches, each with a delay increment that is either 2,3,4,6, or 12 bytes depending on the delay and burst error immunity constraints. The interleaver design is modified here to allow its contribution to the latency to be tailored to modulation rate and the channel burst characteristics. The interleaver is synchronized with the MPEG-2 frame so that the sync byte is always the first byte through the transmitter interleaver's non-delayed branch.

Byte Slice and Differential Encoding

The byte stream is split into symbols of either 2, 4 or 6 bits depending on whether QPSK, 16 QAM or 64 QPSK is selected. The resulting symbols are then mapped into the transmitted constellation using a differential encoding scheme, which uses the two MSB bits of each symbol to differentially select the quadrant, and the remaining bits to select the actual transmit symbol within that quadrant. When the format is QPSK there are only two bits per symbol, and so there are no bits that bypass the differential encoder.

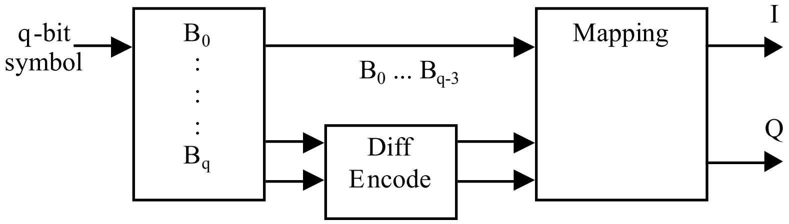

Spectral-Shaping Filter

After the differential encoding, the I and Q channels are each filtered using a square-root raised-cosine filter with $\alpha = 0.15$.

Modulation

The downstream PHY layer supports QPSK and 16-QAM formats, and optionally supports 64-QAM. The downstream PHY layer supports the modulation symbol rates in Table 1 , which also lists the associated channel bandwidths.

These symbol rates have been selected to obtain the listed channel bandwidths. If supporting a desired channel capacity takes precedence over operating within a given bandwidth, the symbol rates can be easily changed. All of the below modulation rates are sub-multiples of the highest symbol rate for each direction, and are also multiples of the basic 8000 Hz SDH frequency, which is done to enable SDH frequency synchronization.

Symbol Rate (Msym/sec)	Chan BW (MHz)
86.4	100
43.2	50
21.6	25
8.64	10
4.32	5
2.16	2.5
0.864	1

Table 1 Downstream Modulation Rates

Upstream Reference Model

The upstream PHY layer uses a TDMA burst modulation format, which includes pulse shaping for spectral efficiency, and has selectable output power level. The PHY layer format includes a variable-length modulated burst with precise timing that is controlled by the BS. The timing of each upstream channel is based on mini-slots. Each channel's mini-slot length is set to a multiple of 64 transmitted symbols. Each SS transmission is assigned as required by the BS (through the MAC layer) to a variable number of mini-slots.

All of the upstream transmission parameters associated with the SS burst transmissions are configurable by the BS via MAC messaging. Many of the parameters are programmable on a burst-by-burst basis. These include the modulation format, symbol rate, preamble parameters and programmable FEC encoding.

The reference model for the PHY layer for the upstream path SS transmitter is shown below. The adaptive equalizer at the transmitter is optional, and is considered to be used to equalize multipath signals that change slowly. This is used mostly to perform a pre-equalization to minimize noise expansion at the receiver. If an equalizer is needed for a quickly changing multipath environment, the BS receiver must incorporate one that can adapt to each burst. The details of the equalizer is left to each manufacturer to decide.

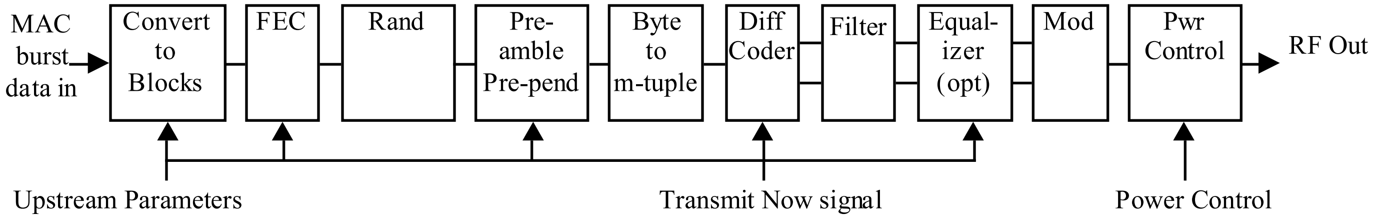

The reference model for the BS upstream receiver has the same sequence but in reverse and without the power control function.

The below diagram shows the various parts of the upstream burst. The burst always begins at a mini-slot boundary. The upstream data is split into blocks for FEC parity generation and the reassembled. After the main transmission comes a guard interval where no information is transmitted and the transmitter is turned off. This guard interval allows for some degree of timing error in the system, and its length is implementation dependent.

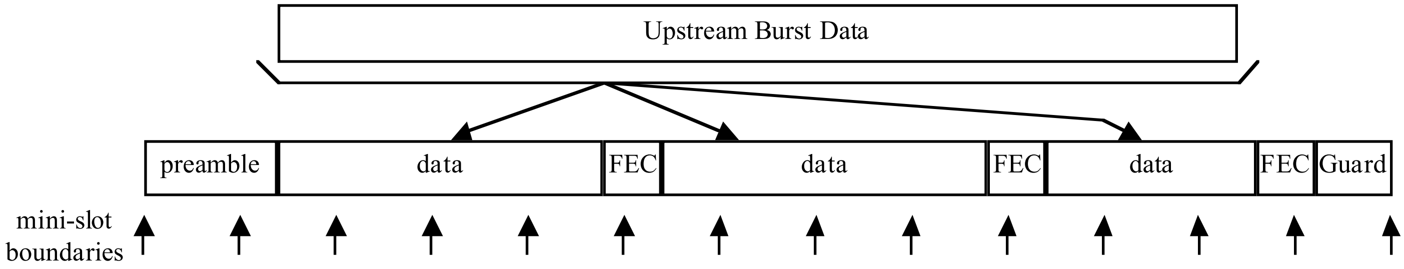

MAC/PHY Interface

The MAC/PHY interface consists of the MAC message that is to be transmitted during the burst, a Transmit Now signal, and a power control signal, which sets the modulator output power. The Transmit Now signal tells the PHY layer when to start transmitting. The exact timing of the Transmit Now signal and the value of the SS transmit power is controlled by the BS via MAC messages.

The MAC layer also passes to the PHY layer a number of general upstream configuration parameters such as the preamble and FEC characteristics, the modulation rate, the mini-slot size and guard time. The MAC specifies for each burst its modulation type, the number of mini-slots that are assigned to that burst and whether to use a fixed or shortened code word.

Mini-slot Definition

Mini-slots are defined as an integral number of clock ticks. One clock tick is always equal to 32 upstream transmission symbols. Thus for QPSK, one clock tick is equal to 32 symbols or 8 bytes of data. Each upstream channel is configured so that one mini-slot contains 2^m clock ticks, where m is from 1 to 8. Thus, each mini-slot can transport between 8 and 1024 bytes of data with QPSK and 16 to 2048 bytes with 16-QAM. The choice of mini-slot size is based on the expected type of traffic and its required maximum latency.

Conversion to Blocks and Forward Error Correction

The forward error correction is realized using an encoder with a shortened Reed Solomon code that is over GF(256). Each upstream channel configuration defines the maximum code word size (n, 16 to 255 bytes) and the correction capability (t , 1 to 10 bytes). The values of n and t consequently define the number of information bytes k ($n - 2*t$).

Fixed Code Word Mode

In the fixed code word mode, the upstream burst data is divided into a number of blocks that are each k bytes long. If the final block contains less than k information bytes, zero-filled bytes are added to obtain a total of k bytes. Each block of k bytes is then encoded using the (k+2t, k) shortened Reed Solomon encoder to produce the FEC parity bytes.

Shortened Code Word Mode

In the shortened code word mode, the upstream burst data is first divided into a string of blocks that are k bytes long with the final block containing any remaining bytes ($k' < k$). Each block of k bytes is encoded using the above Reed Solomon encoder to produce the FEC parity bytes. The number of information bytes in the final code word (k'') is then calculated so that after encoding there exists the specified guard time between the end of the transmission and the next mini-slot boundary. The final information bytes are then created by adding zero-fill bytes to the final k' bytes to get k'' bytes. This block is then encoded using a (k'' +2t, k'') shortened code.

Scrambling

The scrambler randomizes the upstream FEC result by combining it with a pseudo-random bit sequence (PRBS) using an exclusive-or function. The PRBS sequence generator is initialized at the beginning of each burst. The PRBS sequence is generated using the polynomial $x^{15} + x^{14} + 1$.

Preamble

After scrambling, a preamble sequence, whose length and pattern is programmable for each upstream channel, is added to the front of the burst. The length of this preamble is within the range of 0 to 1024 bits. Thus, there are a maximum of 512 symbols with QPSK, 256 with 16QAM and 128 with 64QAM. This preamble is primarily used by the BS as a training sequence for an adaptive equalizer if it is present and possibly as an acquisition aid in some demodulator implementations. The characteristics of this preamble is implementation dependent, so is not specified here.

Spectral-Shaping Filter

After the differential encoding, the I and Q channels are each filtered using a square-root raised-cosine filter with $\alpha = 0.2$.

Modulation

The upstream PHY layer supports QPSK and 16-QAM formats, and optionally supports 64-QAM. The upstream PHY layer supports the modulation symbol rates listed in Table 2 along with the associated channel bandwidths.

These symbol rates have been selected to obtain the specified channel bandwidth goals. If supporting a desired channel capacity takes precedence over operating within a given bandwidth, the symbol rates can be easily changed. All of the below modulation rates are sub-multiples of the highest symbol rate for each direction, and are also multiples of the basic 8000 Hz SDH frequency, which is done to enable SDH frequency synchronization.

Symbol Rate (Msym/sec)	Chan BW (MHz)
83.2	100
41.6	50
20.8	25
8.32	10
4.16	5
2.08	2.5
0.832	1

Table 2 Upstream Modulation Rates

Frequency Control

The upstream and downstream channel frequencies are controlled by the BS via MAC messages. MAC messages are capable of setting the channel center frequencies to any multiple of 1 kHz up to 4,000 GHz. The upstream transmit offset frequency is separately controlled via the MAC with a control resolution of 1 Hz and a range capability of $\pm 2^{31}$ Hz (~2 GHz).

Burst Timing Control

The upstream timing is set relative to MAC timing synchronization messages that are contained within the downstream data. The upstream burst transmission time is controlled by the BS via MAC messages. The transmission time is set relative to the downstream timing with a resolution of 1/4th the symbol period. The maximum one-way path delay is 800 µS.

Power Control

The output power of the SS is controlled by the BS via MAC messages. The output power has a minimum control range of 50 dB to allow for power variations due to range and rain fade. Each power increment is 1 ± 0.5 dB. The absolute power levels are implementation dependent, and so are not specified here.

Band Plan

The proposed frequency plan is shown in the figure below. This frequency plan is desirable because it maximizes the separation between the upstream and downstream frequencies while utilizing the entire Block A frequency allocation. The large separation between upstream and downstream frequency bands allows for a simpler and more cost-effective design for the customer premise transceiver.

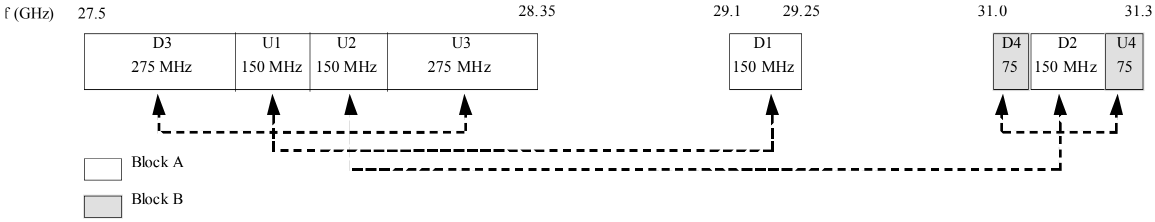

In this plan, the two smaller 150 MHz-wide Block A bands are paired with a corresponding bandwidth in the center of the large 850 MHz block. This leaves two 275 MHz-wide bands on the outer edges of the 850 MHz block which provide an additional upstream-downstream band pair for Block A having a minimum separation of 300 MHz. These 275 MHz-wide bands can be split into two sub-bands to achieve a greater minimum separation between upstream and downstream frequencies and provide more uniform bandwidths.

Performance

Spectral Efficiency

The spectrum efficiencies for the QPSK format are listed below for the upstream downstream channels. These are calculated from the modulation bit rate and the proposed channel bandwidth. These values double for 16-QAM and are three times as large for the 64-QAM format because of the greater channel capacity.

	Upstream	Downstream
	Mbps/MHz	Mbps/MHz
QPSK	1.66	1.73
16-QAM	3.32	3.46
64-QAM	4.98	5.19

Reference System Gain

The reference system gain (RSG) is defined as the output power of the transmitter minus the receiver threshold at a post coding BER of 10^{-6}, including back-off required for proper transmission. It is assumed that the 1 dB compression point of the transmitter (prior to back-off) is 0 dBW, and the receiver has an ideal noise figure of 0 dB.

The RSG curves for the QPSK format at various modulation symbol rates and for different levels of Reed Solomon encoding are graphed in the below chart. The downstream (DS) RSG is listed using the single proposed Reed Solomon encoder. The upstream (US) encoding however, is widely configurable depending on the application, so only two potential encoder configurations are shown.

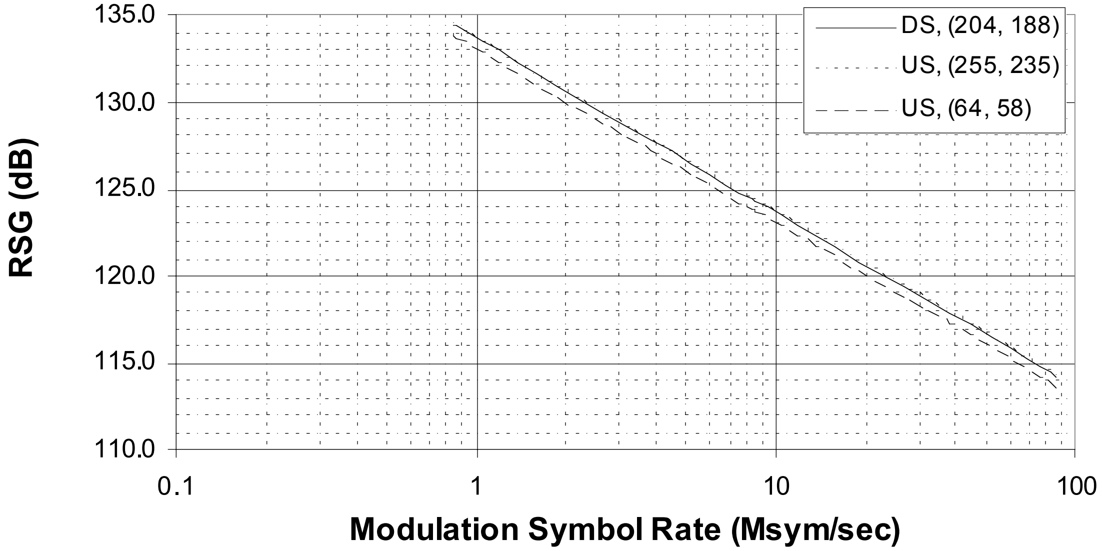

Benefits & Unique Features

Integral Video MPEG-2 in Downstream

The downstream PHY protocol's use of MPEG-2 framing to carry the MAC layer data enables the efficient distribution of video signals in their native format. The standard sync byte within the MPEG-2 frame is also used to synchronize the MAC signals.

Upstream Flexibility

The proposed PHY layer upstream protocol can efficiently transport packets of different sizes and different latency requirements on a single upstream channel from any particular SS. Because the length and time of each upstream burst is assigned in terms of mini-slots, the BS can assign to each SS whatever amount of bandwidth is required by the SS at any particular time. If one IP datagram is to be transported on a best effort basis, the BS can assign the required number of mini-slots to that SS to carry only that datagram. On the other hand, if an SS requires a dedicated ATM channel at a constant bit rate, the BS can regularly assign the needed number of mini-slots to that SS.

SDH Timing Synchronization

In SDH systems, the clock jitter and wander of equipment must meet stringent requirements. BWA systems must provide a means to control the frequency and timing of the SDH signal that is reconstructed at the SS from packets sent from the BS.

In the proposed PHY protocol, the BS transmitted modulation symbol clock can be synchronized with the SDH interface timing source at the BS network interface. To facilitate this, all of the proposed modulation symbol rates are at integer multiples of 8 kHz, which is the fundamental frequency for all SDH signals in both Europe and the United States. Circuitry within the SS then generates an SDH clock that is phase locked to the modulation symbol rate. This clock then is used to shift out the data from the jitter-elimination buffer.

Drawbacks

The proposed downstream channel is based on an MPEG-2 format, which is continuously transmitted by the BS. Therefore, the downstream is applicable to a full-duplex, FDD system, but not to a TDD system.

Heritage

Downstream

The downstream PHY protocol is derived from the ITU-T J.83 Annex A specification. This heritage protocol is used because it can be adapted to support QPSK, 16QAM and 64QAM formats, because it is a proven PHY protocol that is presently in use. A side benefit is that the downstream is based on the MPEG-2 format, which also neatly supports the distribution of video that is required. Annex B was not selected because the trellis coded modulation approach specified is specific to 64QAM and 256QAM. If a 7/8 punctured trellis coded modulation scheme was used with the QPSK format, it would increase the overhead by approximately 14%, which is deemed too severe of a penalty.

Instead, the QPSK format is added to the Annex A specification, and various data rates are also proposed. In addition, the interleaver design is altered to allow the upstream latency and the burst error protection to be selected in order to optimize the downstream channel for various types of services.

Upstream

The upstream PHY protocol is derived from the DOCSIS 1.1 RF Interface upstream specification. The DOCSIS upstream protocol is preferred because it carries variable length datagrams very well, which is optimal for IP traffic while still transporting ATM and SDH data well.

The upstream QPSK and 16QAM formats are the same as in DOCSIS, but 64QAM is added as an optional format for greater channel capacity. The most significant modification to the DOCSIS specification involves the length of the mini-slot, which in DOCSIS is always a multiple of 6.25 µS, but is now proposed as a multiple of 32 modulation symbols. Using DOCSIS at high modulation rates, the assignment of upstream bursts becomes very coarse since the mini-slots length is always based on a fixed time of 6.25 µS. Because the proposed protocol specifies the mini-slot length relative to the modulation rate, the burst size granularity stays relatively constant as the modulation increases.

Scalability

The scalability of the proposed PHY layer is improved over the DOCSIS upstream PHY. As described immediately above, the proposed mini-slot length is specified relative to the modulation rate instead of being a fixed length. This allows the burst length assignment granularity to remain constant as the modulation rate changes.

Both the downstream and upstream modulation rates are also independent of each other, while the proposed PHY layer protocol is independent of the modulation rate. This enables the upstream and downstream capacities to be scaled independently as needed. The only requirement on the modulation rate selection is that there be enough downstream capacity to carry both the information and the MAC messages that contain the upstream burst assignments. The only limiting factor on the modulation rate is the speed of the hardware itself.

The proposed PHY layer does not specify the performance or structure of the adaptive equalizers in either the BS or SS. This allows this protocol to be used at carrier frequencies other than the primary target, which is around 30 GHz. Equipment that is designed to operate at lower frequencies may require more stringent adaptive equalizers than equipment designed for 60 GHz.

This proposal also does not specify the absolute transmit power of either the BS or SS. Only the minimum power control range is specified. This allows this protocol to flexibly operate with future equipment that is designed with higher transmit powers than are envisioned today.

Intellectual Property Rights

The ITU patent list identifies no intellectual property with regards to the Annex A part of J.83, which is the basis for the proposed downstream protocol. However, there may be intellectual property that has not been disclosed to the ITU. The authors believe that this is unlikely given the straight forward nature of the proposed downstream physical layer coding scheme.

Four companies are identified by Cable Labs as holding intellectual property rights to aspects of the DOCSIS protocol. These companies are General Instruments, Broadcom, 3-Com, and Bay Networks. Since the proposed upstream protocol is derived from the DOCSIS protocol, there may be intellectual property that is owned by those companies. At this time, the nature of this intellectual property is not known to the authors. All proposed modifications to the existing DOCSIS upstream physical layer are free of intellectual property restrictions.

Evaluation Table

The evaluation table is included in Appendix A. The authors have addressed the various criteria within the table.

Conclusion

The proposed physical layer protocol meets all of requirements of the 802.16.1 System Requirements document. It flexibly supports the transport of ATM, SDH, and IP data, and also the distribution of MPEG-2 video in an efficient manner. It is based on proven standards, which are modified slightly to improve the latency and capacity. The protocols are insensitive to interference from other systems, but also offer optional modulation formats that improve the channel capacity in situations where the interference is minimal.

Appendix A, Annotated Evaluation Table

#	Criterion	Discussion
1	Meets system requirements	*How well does the proposed PHY protocol meet the requirements described in the current version of the 802.16 System Requirements (Document IEEE 802.16s0-99/n)?* The proposed PHY protocol meets the requirements very well. The protocol is capable of transporting ATM, IP, SDH traffic as well as MPEG-2 video without the need to encapsulate everything into ATM cells, while maintaining the required QoS levels.
2	Spectrum efficiency	*Defined in terms of single sector capacity assuming all available spectrum is being utilized (either in terms of Gbps/Available Spectrum or in terms of Mbps/MHz)* The spectrum efficiency of the QPSK format for the upstream channel is 1.66 Mpbs/MHz and 1.73 Mpbs/MHz for the downstream channel. This value is calculated using the proposed channel bandwidths (not the 3 dB bandwidth)
3	Simplicity of implementation	*How well does the proposed PHY allow for simple implementation or how does it leverage on existing technologies?* The downstream protocol is very similar to the existing DVB downstream protocol, and therefore leverages the existence of modems that already are built for that protocol. The upstream protocol is derived from the existing DOCSIS upstream protocol, which also has a broad base of existing hardware being deployed.
4	CPE cost optimization	*How does the proposed PHY affect CPE cost?* QPSK modulation requires a smaller amount of transmitter back-off than other modulation formats. Thus, the CPE transmitter for a given output power will have a lower cost than one with other formats.
5	Spectrum resource flexibility	*Flexibility in the use of the frequency band (i.e., minimum frequency band required to operate and migration capabilities)* The proposed PHY protocol's modulation rate and channel bandwidth are settable over a wide range of values. Operators, who deploy a system in regions where the available paired bandwidths are either small or large, can set the channel bandwidths accordingly. Also, the upstream and downstream channel bandwidths are independent of each other, and so can be configured to operate in unsymmetrical spectrums.
6	System diversity flexibility	*How flexible is the proposed PHY to any other system variations and future technology improvements or new services?* The proposed PHY is relatively independent of the modulation rate, carrier frequency, or MAC message size. The mini-slot size is specified as a multiple of the modulation symbol rate, so the burst length granularity scales well. The scheduling burst algorithms, which may change when different services are added, reside in the MAC layer above and so have little effect on the PHY.

7	Protocol Interfacing complexity	*Interaction with other layers of the protocol, specifically MAC and NMS* Most of the upstream burst parameters are controlled by the MAC layer, either burst-by-burst or for all burst on an upstream channel. The downstream modulation characteristics are likewise controlled by the MAC. Physical layer performance statistics such as transmit power and receiver error rates are also available to upper layers.
8	Implication on other network interfaces	*Intrinsic transport efficiency of telecomm and datacomm services* This physical protocol is well suited to the flexible transport of both ATM, SDH and IP data because the upstream burst length and rate can be flexibly adapted to the needs of each service. Constant bit rate services as well as a single FTP file transfer can be efficiently carried by either the downstream or the upstream.
9	Reference system gain*	*Sector coverage performance for a typical BWA deployment scenario (supply, reference system gain)* The reference system gain varies with the modulation frequency and is shown for QPSK format in this contribution. QPSK has the one of the highest system gains of any digital modulation format for a given channel modulation rate and channel bandwidth. The higher order QAM formats that are also proposed here will have smaller reference system gains, but will only be used if operators desire.
10	Robustness to interference	*Resistance to intra-system interference (i.e., frequency re-use) and external interference cause by other systems.* The QPSK modulation is one of the most robust modulation formats when operating with interference from any source. It operates at the lowest carrier to interference ratios of almost any other digital modulation format. This proposal higher order modulation formats are available only in situations where the system can tolerate their slightly higher sensitivity to interference. Because the proposed PHY protocol uses FDD, the upstream and downstream channels in any cell are never at the same frequency. Therefore, there is no need to synchronize neighboring cells that operate in order to control the interference between SS that are close to each other and at the same frequency.
11	Robustness to channel impairments	*Rain fading, multipath, atmospheric effects* The proposed protocol's power adjustment range is adequate to handle the effects of rain and atmospheric fading. Upstream and downstream receiver adaptive equalizers, which will remove the quickly varying multipath effects, are left up to each manufacturer to specify and design. The protocol has provisions for the BS to control an adaptive pre-equalizer at the upstream transmitter. However, its coefficients are expected to be updated at a fairly slow rate.

* In order to compare between PHY proposals, we define the reference system gain (RSG) as the output power of the transmitter minus the receiver threshold at a given working point, including back-off required for proper transmission. We will assume a 0 dBW transmitter (prior to back-off), and an ideal LNA (0 dB NF) and BER of 10^{-6} post coding.