INTERNATIONAL JOURNAL OF COMMUNICATION SYSTEMS
Int. J. Commun. Syst. 2000; **00**:1–6 *Prepared using* dacauth.cls *[Version: 2002/09/18 v1.01]*

HomePlug 1.0 Powerline Communication LANs -*Protocol Description and Performance Results version 5.4* †

M. K. Lee[1], R. E. Newman[2], H. A. Latchman[1], S. Katar[3] and L. Yonge[3]

[1] *ECE Department, University of Florida, Gainesville FL 32611*
[2] *CISE Department, PO Box 116120 University of Florida, Gainesville, FL 32611*
[3] *Intellon Corporation, 5100 W. Silver Spring Blvd., Ocala, FL, 34482*

SUMMARY

Products implementing the HomePlug 1.0 standard allowing high-speed communication on low-voltage powerlines have recently started arriving on the U.S. market for home and office networking without the requirement for installing new wires. Effective use of the powerline bandwidth requires robust physical (PHY) and medium access control (MAC) protocols to mitigate the harsh conditions of the power line channel as well as the capability to support prioritized multimedia traffic. This paper describes powerline communications and the HomePlug 1.0 protocol, based on Orthogonal Frequency Division Multiplexing (OFDM) and Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), along with its changes to allow prioritized channel access. It then presents performance results for the HomePlug 1.0 protocol using a simulation model, ideal laboratory measurements with actual HomePlug 1.0 devices, and field tests in a residential building. Simulation and laboratory data rates were around 6 Mbps, and field tests gave rates from 1.6 to 5.3 Mbps at the application level. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS: Powerline Communications, OFDM, CSMA/CA, Virtual Carrier Sense, PHY, MAC, QoS, differentiated service

1. Introduction

Affordable broadband Internet communication to residential customers is now available via cable modems and various flavors of Digital Subscriber Lines (DSL). In turn there is a growing need for in-home networks to share this single full-time Internet access link, while supporting a wide range of digital data and multimedia communication services [5, 13]. While it is a

†Please ensure that you use the most up to date class file, available from the DAC Home Page at
`http://www.interscience.wiley.com/jpages/1074-5351/`

Contract/grant sponsor: Publishing Arts Research Council; contract/grant number: 98–1846389

Copyright © 2000 John Wiley & Sons, Ltd.
Received 17 January 2000
Revised 18 September 2002

Table I. Typical Costs for Various SOHO LAN Options

LAN	Characteristic	Limitations	Manufacturers	Costs(USA)
10/100 Base-T Ethernet	Fastest system and low cost NICs	Requires extensive wiring and retrofit	Numerous vendors of NIC cards and Hubs and Switchs	$75 to $200 per port for retrofit wiring. $60 -$100 new installation
Phone Line	Use existing phone wires	Computers must be near phone jack	Limited number of manufactures	$50 to $100
Wireless	Mobility	Wired infrastructure usually required	Growing number of suppliers	$60 to $250
Powerline	Data port is electrical outlet	New technology not yet widespread	Multiple OEMs using HomePlug standard specification	$74 to $150

Table II. Data Rates of Various SOHO LAN options

Network Type	Data Rate
Ethernet/IEEE 802.3	10/100 Mbps
IEEE 802.11a	55 Mbps
HomePlug 1.0	14 Mbps
IEEE 802.11b	11 Mbps
HomePNA 2.0	10 Mbps

simple matter to use a 10/100 Base-T network hub to link several computers in a single room or in a small office environment, it is much more challenging to provide network connections in several rooms in a typical home [8]. One option of course is to re-wire the home with network cabling (typically 10/100 Base-T CAT-5 UTP cabling), which is quite an expensive proposition, especially if existing homes are to be retrofitted with data communication cables (See Tables I and II for costs and rates as of fall 2002).

Another approach is to deploy a wireless LAN (Local Area Network) with wireless modems in each device connecting to one or more wireless hubs (infrastructure-based) or to each other (in an ad hoc network). The wireless option is certain viable, especially for small numbers of nodes in the home/office network, and several companies now offer wireless networking hardware and software typically based on the IEEE 802.11b standard. Unfortunately, experience suggests that a wired infrastructure connecting multiple access points is required to cover the entire home. A third solution is to establish a communication channel using the existing low-voltage (110/220 V) power lines that deliver electrical power to outlets and lights in every room (and, depending on the local building codes, to every wall) in the building [10]. A coalition of manufacturers, the HomePlug Powerline Alliance, has established a new protocol (the HomePlug 1.0 Standard) that will enable the establishment of an Ethernet-class network over powerline channels [9]. It is anticipated that manufacturers of computers, networking and communication devices, and peripherals will use low-cost integrated circuits based on the HomePlug 1.0 Standard to enable such equipment to act as a node on the home/office network by simply plugging it into the wall outlet. Thus the outlet, which would normally be a required connection for electrical power, now becomes simultaneously the point of connection for high speed data communication.

It is of course possible also to use the telephone wiring (with no new wires), in a similar

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

manner and perhaps with more attractive channel characteristics. However, phone lines, in common with Cable TV cabling, appear in only a few rooms in a home at best. Further, there is usually only one phone or cable jack in each room with service. On the other hand, power lines, though ubiquitous throughout the home, are notoriously bad as a communication channel due to electrical noise and interference as well as channel variability depending on the appliances that are in use from time to time. Despite these impediments, tests of the present version of the HomePlug 1.0 powerline devices in some 500 homes show that 80% of outlet pairs will be able to communicate with each other at about 5 Mbps or higher, and 98% will be able to support data rates greater than 1 Mbps. Field tests suggest that the powerline network will provide connectivity in situations where some wireless networks will fail due to large attenuations due to distance or obstructions such as intervening walls or furniture. The results of our comparative study seems to support this conclusion also.

The rest of the paper is structured as follows. Section 2 gives an overview of Powerline Communications (PLC), reviewing the nature of the environment and the modulation method. Sections 3 and 4 provide detailed description of both the PHY (Physical Layer) and the MAC (Medium Access Control) protocols of the HomePlug 1.0 specification. The next three sections present HomePlug1.0 throughput results, with simulations in Section 5, ideal lab experiments in Section 6, and field test results in Section 7. The paper concludes in Section 8 with some observations and comments on further work in this area.

2. Powerline Communications

2.1. Early power line communication technologies

Efforts to use to the powerline as a transmission medium were 160 years ago [5], but high speed transmission over 10 Mbps were archived in the mid 1990s. Before this time, the various types of power lines were assumed to be inherently low data rate transfer media. There are several Powerline Communication (PLC) protocols. The primary purpose of many of these protocols is not for home networking, but rather for powerline control protocols for home automation, home security, and lighting control. On the other hand, the HomePlug 1.0 protocol is a high speed in-home network standard. Although using the same powerline as home automation protocols, the HomePlug 1.0 device can coexist with device using these other protocols because by using a different frequency band than powerline control technologies such as X-10, CEBus, and Lonworks[8]. There are also some preliminary high-speed PLC network devices with limited deployment in Europe, where the interest is primarily on access ("last 100 feet") rather than in-home networks[7]. However, these are expensive and do not enjoy wide use yet; they also must contend with both different regulatory environments and power distribution topologies.

2.2. Power line medium

Powerlines were originally devised for transmission of power at 50-60 Hz and at most 400 Hz. At high frequencies the power line is very hostile for signal propagation [4, 17]. Powerline networks operate on standard in-building electrical wiring and as such consist of a variety of conductor types and cross sections joined almost at random. Therefore a wide variety of characteristic impedances will be encountered in the network. Further, the network terminal impedance will tend to vary both with communication signal frequencies and with time as

Copyright © 2000 John Wiley & Sons, Ltd.
Prepared using dacauth.cls

the consumer premises load pattern varies. This impedance mismatch causes a multi-path effect resulting in deep notches at certain frequencies [1]. In a typical home environment the attenuation on the power line is between 20 dB and 60 dB, and is a strong function of load.

The major sources of noise on power line are from electrical appliances, which utilize the 50 Hz electric supply and which generate noise components that extend well into the high frequency spectrum. Some common sources of electrical noise are certain types of halogen and fluorescent lamps, switching power supplies as well motors and variable resistance dimmer switches. Apart from these, induced radio frequency signals from broadcast, commercial, military, citizen band and amateur stations severely impair certain frequency bands on the powerline channel. Reliable data communication over this hostile medium requires powerful forward error correction (FEC) coding, interleaving, error detection and Automatic Repeat Request (ARQ) techniques, along with appropriate modulation schemes as well as a robust medium access protocol (MAC) to overcome these impairments.

2.3. Orthogonal Frequency Division Multiplexing

Orthogonal Frequency Division Multiplexing (OFDM) [2, 18, 19, 20] is one of the most promising techniques for data transmission over power lines [18]. OFDM is well known in the literature and in industry [2, 11]. It is currently used in DSL technology [19], terrestrial wireless distribution of television signals[23], and has also been adapted for IEEE's high rate wireless LAN Standards (802.11a and 802.11g[12]). The basic idea of OFDM is to divide the available spectrum into several narrowband, low data rate subcarriers. In this respect, it is a type of Discrete MultiTone modulation (DMT)[6]. To obtain high spectral efficiency the frequency response of the subcarriers are overlapping and orthogonal, hence the name OFDM. Each narrowband subcarrier can be modulated using various modulation formats. By choosing the subcarrier spacing to be small the channel transfer function reduces to a simple constant within the bandwidth of each subcarrier. In this way, a frequency selective channel is divided into many flat fading subchannels, which eliminates the need for sophisticated equalizers. OFDM has the following advantages. OFDM

(1) has excellent mitigation of the effects of time-dispersion;
(2) is very good at minimizing the effect of in-band narrowband interference;
(3) has high bandwidth efficiency;
(4) is scalable to high data rates;
(5) is flexible and can be made adaptive; different modulation schemes for subcarriers, bit loading, adaptable bandwidth/data rates are possible;
(6) has excellent ICI performance, so complex channel equalization is not required.

For these reasons it is an excellent candidate for powerline communication.

3. HomePlug 1.0 Physical Specifications

This section and the next provide a detailed overview of the HomePlug 1.0 specifications at a level appropriate for a protocol engineer. This section describes the Physical (PHY) layer, and the following section describes the Medium Access Control (MAC) layer. Briefly, the PHY uses adaptive Orthogonal Frequency Division Multiplexing (OFDM) with Cyclic Prefix

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

(CP). Both turbo product codes and Reed-Solomon concatenated with convolutional codes are used at various times for forward error correction. The PHY detects channel conditions using channel estimation, then adapts by avoiding poor subcarriers and selecting an appropriate modulation method and coding rate for the remaining subcarriers. Three variants of Phase Shift Keying (PSK) modulation [19, 21] are used: Coherent Binary PSK (BPSK), Differential BPSK (DBPSK), and Differential Quadrature PSK (DQPSK). A preamble and frame control combination are used as delimiters that start and end long frames, with only the payload portion adapted to the channel conditions. Physical carrier sense (PCS) is performed by the PHY layer, and helps the MAC determine when the medium is busy. Fine details necessary to implement a compliant system are available in the official specifications [11].

Connectors are assumed to contact only one line phase (L1 or L2) of the local power line network, and neutral. Connectors may or may not have ground contacts, and the user should be able to connect or disconnect at any time.

3.1. Signal Processing

The Physical Layer (PHY) of HomePlug 1.0 uses OFDM in a band from approximately 4.49 to 20.7 MHz. The band from 0 to 25 Mhz is divided into 128 evenly spaced carriers, of which 84 fall within the band used (subcarriers 23-106, inclusive). Additionally, eight of the subcarriers (tones) within the usable band are permanently masked to avoid 40 meter, 30 meter, 20 meter, and 17 meter amateur bands, which leaves 76 subbands for use in the U.S.A. The Tone Mask is intended be alterable to support regulatory requirements in different countries. Spectral compatibility is regulated through the FCC in the US (Part 15 rules), and compliance with radiated power requirements has resulted in -50 dBm/Hz of transmission spectral power density (PSD), much lower than that used by wireless technology.

In order to avoid ISI (Inter Symbol Interference), a cyclic prefix of the last 172 samples from the IFFT interval of 256 samples are prepended to the IFFT interval to form a 428-sample OFDM symbol. Using a 50 MHz clock, and 8 samples for roll-off interval, this results in 8.4 microseconds per symbol, with 5.12 μs. for the raw OFDM symbol and 3.28 μs. for the CP.

Before forming the OFDM symbol in the Analog Front End (AFE), data are scrambled, RS-encoded, convolution encoded, punctured, then interleaved on the transmitter. These processes will be discussed on more detail below. The AFE consists of a constellation mapping block, an IFFT block, a preamble block, a cyclic prefix block, and a raised cosine (RC) block. The mapping block groups data bits and maps them onto a constellation point of the modulation method; it selects the type of modulation and the carriers to be used in the IFFT block, as specified by the Tone Map and Tone Mask. The IFFT block modulates the constellation points onto the carrier waveforms (in discrete time), while the preamble block inserts the preamble. The CP is added by the Cyclic Prefix block, and RC shaping is employed to reduce out-of-channel energy.

The physical layer transmits four distinct entities (refer to Figure 1):

- the preamble,
- frame control (FC),
- the payload, and
- priority resolution signals (PRSs).

The first two of these are always sent together, and form a delimiter. A delimiter by itself, or a

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

payload surrounded by delimiters form a PHY protocol data unit (PPDU), which is discussed in more detail in section 1.

The preamble consists of seven and one-half special OFDM symbols without CP added, and lasts 38.4 μs. It is used for automatic gain control (AGC) and synchronization, as well as forming the phase reference for frame control encoding. The preamble is also used for early detection of the delimiter in Physical Carrier Sense (PCS), and the time needed to determine their presence determines the slot size used by the MAC in the contention period described in Section 4.6.

Priority resolution signals consist of six special OFDM symbols without CP, which are in fact the same as the preamble symbols with the polarity reversed. These take 30.72 μs. to send, and with processing delay determine the length of the priority resolution slots (PRS0 and PRS1) used by the MAC to establish prioritized access to the medium (see Section 4.5).

The data in the four OFDM symbols of the frame control (FC) are encoded using a specially designed Turbo Product Code (TPC)[**?**], and are interleaved with an interleaver distinct from the one used for the payload data. Coherent BPSK is always used for modulation of frame control symbols, and the field takes 33.6 μs. to send.

Together with the preamble, the FC forms a 72 μs. delimiter, of which there are three basic types: start, end, and response. Formats and functionality are discussed in more detail in Section 3.2. The FC of a start delimiter includes two fields needed by the PHY: a length field and the Tone Map Index (TMI). The receiver needs the TMI to know how to demodulate and to decode the payload, and the length tells the receiver how long that demodulation method must be used before the end delimiter arrives.

Delimiters and priority resolution signals must be detected and decoded correctly by all receivers, so they must use all subcarriers with the same modulation and encoding, no matter who is sending or receiving them. The payload can be and is adapted to the channel conditions by negotiation between the sender and receiver during channel estimation (see Section 4.11). Channel Estimation determines which subcarriers to use, and for these, which type of modulation and forward error control (FEC) rate to apply.

Depending on channel conditions, a number of combinations of modulation type and FEC rates are available, allowing the sender to adapt to the channel to improve the data and error rates. Modulation may be either Differential BPSK (DBPSK) or Differential Quadrature PSK (DQPSK), conveying 1 or 2 raw bits per carrier per symbol, respectively. All subcarriers use the same modulation method.

The PHY payload consists of some number of 20- and 40-OFDM symbol transmission blocks, encoded on a link-by-link basis using a Reed-Solomon/Convolutional concatenated code[15]. These block sizes are needed to combat impulse noise, which can easily damage several OFDM symbols (since differential modulation is used, at least two symbols at a time are lost). The convolutional encoder has constraint length 7, and rates of $\frac{1}{2}$ or $\frac{3}{4}$ (via puncturing) are selectable during channel adaptation. The Reed-Solomon (RS) code has coding rates ranging from $\frac{23}{39}$ to $\frac{238}{254}$. Each combination of modulation and convolutional code rate requires a minimum number of viable carriers to be selected. For DBSK and rate $\frac{1}{2}$, 32 carriers are needed; 16 are needed for DQPSK and rate $\frac{1}{2}$; while 11 suffice for DQPSK with rate $\frac{3}{4}$ convolutional coding.

Before channel adaptation has occurred, the receiver must be able to demodulate and to decode the payload with only *a priori* knowledge. Likewise, in order to multicast or broadcast a PHY frame, all the receivers must use a common demodulation and decoding method pair. For

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

this reason, and to handle those cases in which transmission using the existing TM have failed, a special form of modulation and FEC were developed. ROBO mode (for Robust OFDM) is based on DBPSK with extensive time and frequency diversity for robust operation under noisy conditions. ROBO mode always uses all carriers, and also uses a different interleaver than the other modes. Redundancy reduces the rate to $\frac{1}{4}$ bit per carrier per symbol for ROBO modulation. It also uses a different RS code with rates ranging from $\frac{31}{39}$ to $\frac{43}{51}$, and only supports 40-symbol physical transmission blocks.

Tone Maps (TMs) are used by sender-receiver pairs to adapt to varying channel conditions. A TM lists which carriers for a sender to use for unicast to that particular receiver, in order to avoid bad subbands where attenuation is severe or there is narrowband noise. Each TM also specifies the modulation method and convolutional coding rate to use for the symbols sent. (Note that in HomePlug 1.0, bit-loading is not used; the same modulation method and coding rate are used for all of the indicated carriers.) TMs are not used for frame delimiters, the preamble, or priority resolution symbols, nor is the TM obeyed when ROBO mode is employed (either for unicast or for multicast/broadcast). The FC frame length field indicates how long the receiver should use the demodulation and decoding methods specified by the tone map before searching for the next preamble. Altogether, eliminating the duplicate ways of achieving the same data rate, 139 distinct physical data rates are available from 1 Mbps. to 14.1 Mbps. With a nearly continuous range of data rates available, channel adaptation allows very good utilization of the bandwidth available, although there is room for improvement using full bit-loading.

3.2. PHY Frames

While it is commonly the case to refer to physical Protocol Data Units (PHY PDUs) and Medium Access Control (MAC) PDUs when discussing protocols, this is not done in the HomePlug 1.0 Specification. The authors suspect that this is due to the heavy interaction between the PHY and the MAC in the frame control fields of the various delimiters, perhaps due as well as the lack of a physical address visible to the receiver at the PHY level. Since the physical addresses are so long (48 bits) and the net bandwidth efficiency of the modulation and encoding used for the delimiters is necessarily low, the overhead would be prohibitive if these were included at the PHY layer. Channel adaptation requires that each sender-receiver pair optimize the tone map to the channel conditions on that link, but there is information at both the PHY and the MAC levels that all nodes must see. Both MAC and PHY need length information, and the MAC also needs contention control and priority information. The compromise used in the HomePlug 1.0 standard is to violate strict layering for the sake of efficiency, and allow the universally readable FC field to hold information needed by both layers. For this reason, the delimiter FC fields appear in the PDUs of both layers. We will proceed to call the PHY frames PPDUs, and the MAC frames MPDUs, and not worry that the FC appears in both.

As shown in Figure 1, a long PHY Protocol Data Unit (PDU) starts with a Start-of-Frame delimiter (SOF) followed by the payload, a 1.5 μsec. End-of-Frame Gap (EFG) and an End-of-Frame frame delimiter (EOF). The EFG is a delay inserted to allow for processing time. The EOF helps in collision detection and recovery, in addition to its MAC functions. Payload size in data bits is determined not only by the length (in increments of 20 symbols as required by the PHY transmission blocks, or 40 symbol blocks for ROBO mode), but also by the

Copyright © 2000 John Wiley & Sons, Ltd.

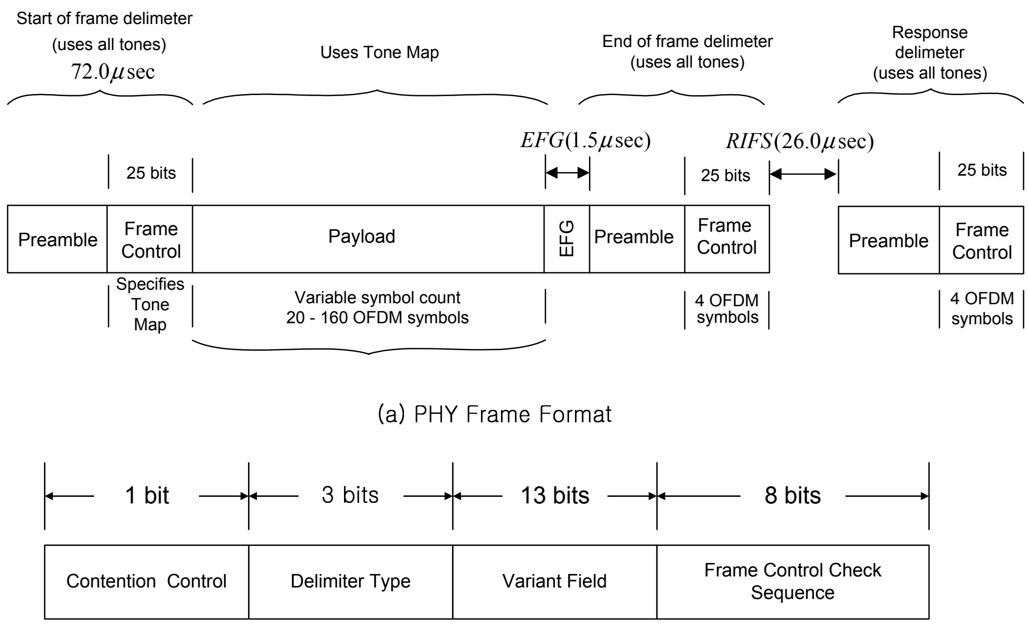

(a) PHY Frame Format

(b) Frame Control

Figure 1. HomePlug 1.0 PHY frame format

modulation method and the coding rate used for that transmission. Short PPDUs consist solely of a response delimiter.

Each 72 microsecond frame delimiter consists of a preamble (see Section 3.1) followed by a Frame Control (FC) field. The four OFDM symbols of frame control provide 25 bits of control data per frame control field. A type field distinguishes SOF, EOF and response delimiters, and in SOF and EOF, includes whether or not a response is expected. All delimiters have an 8-bit Cyclic Redundancy Check (CRC) as a Frame Check Sequence (FCS) to detect errors in the frame control field.

The SOF FC includes the length of frame (encoding the number and size of PHY transmission blocks) and the tone map index (TMI) that the intended receiver should use to find the Tone Map needed to demodulate and decode the payload. Others attempting to decode the payload using a different TM (with the same TMI) are very unlikely to do so correctly, while the intended receiver will be able to confirm its status by checking the destination address in the decoded payload contents.

It is necessary for the PHY to know what type of modulation and decoding is used to decode the PPDU payload properly. This information is contained in the SOF in the TMI, which indexes into each node's Tone Map Table where the modulation type and coding rates negotiated during channel adaptation are stored. The PHY also needs to know how long it should use this tone map before it looks for the EOF; this information is carried in the length

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

field of the SOF FC.

4. HomePlug 1.0 MAC Layer

The HomePlug 1.0 MAC Layer uses channel access based on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) to transport data from 46 to 1500 bytes long from encapsulated IEEE 802.3 frames as MAC Service Data Units (MSDUs).

A four-level priority scheme enforces strictly prioritized access (higher priority traffic will be able to gain access to the medium as soon as a lower priority segment has been sent). Segmentation limits delay for high priority traffic, and contention-free access modes support low jitter requirements. Contention for the channel is limited to those nodes that survive the Priority Resolution Period.

Stations inferring a collision must invoke the backoff procedure, by which they successively increase the amount of random delay they wait in the contention period up to some maximum, depending on the priority level of their data. Different than other CSMA/CA methods, HomePlug 1.0 also uses the number of times that a station has deferred to other stations at the same priority level to infer the amount of traffic present at that level and to adjust the backoff range accordingly.

The MAC segments longer MAC service data units (MSDUs) to limit transmit time as the PHY rate is lowered adaptively. Each unicast segment is acknowledged, and "Partial ARQ" is an option for multicast/broadcast segments to let the sender know that at least one station received the segment correctly. Segment bursting avoids contention for the medium on every segment, yet allows for preemption by higher priority traffic.

MAC management functions support channel estimation and rate adaptation, as well as key management for cryptographic isolation of logical networks. All stations in logical network share the same DES key, called a Network Encryption Key (NEK). This is needed since multiple apartments may share the same transformer, which allows nodes in one apartment to hear PPDUs sent by nodes in a neighboring apartment over this broadcast medium. An IEEE-registered Ethertype allows MAC management information to be passed transparently.

In the interest of brevity, MAC functions of bridging, link status, and parameter and statistics reporting have been omitted.

4.1. Carrier Sense and Collision Detection over Power Lines

Due to attenuation, noise, and channel adaptation, it is difficult to use only physical carrier sense (PCS) as is used in many other CSMA system such as Ethernet. The HomePlug 1.0 PHY layer reports physical carrier sense (PCS) by detecting preambles or priority slot assertions, while the MAC layer maintains virtual carrier sense (VCS) using the length field of the SOF frame control, along with information on whether a response is expected or not (present in both the SOF and the EOF frame control).

Likewise, direct collision detection as used in Ethernet is unreliable due to attenuation and other factors, so collisions can only be inferred from a lack of response after a frame is sent. This makes collisions very costly compared to CSMA/CD systems, so they must be avoided by being less aggressive when the medium is busy. Rather than transmit as soon as the medium becomes idle as in standard Ethernet, HomePlug 1.0 uses Carrier Sense Multiple

Copyright © 2000 John Wiley & Sons, Ltd.

Access with Collision Avoidance (CSMA/CA). Similar to IEEE 802.11, following each frame there is typically a contention period consisting of a succession of short slots (35.84 μsec. for HomePlug 1.0) during which a station may initiate transmission, provided that it has detected no other station that has started sending before that slot (i.e., deferred).

4.2. IEEE 802.3 Frame Encapsulation

MAC Service Data Units (MSDUs) are received from the host interface, and include the DA, SA, optional VLAN tag, type/length, and MSDU data. All of these except the source and destination addresses are placed into a service block (SB), which is the basic unit of information transfer by the MAC. Each SB is sent to the destination as the frame body of an MPDU (see Section 4.3), whose header holds the source and destination addresses.

Each SB consists of 9 bytes of Encryption Control (ECtl), followed by an optional 4-byte VLAN tag, an optional, variable-length MAC Management Information field, 2 bytes for Type if it contains an MSDU, the variable length frame data (if any), ending in 0 to 7 bytes of Encryption Pad (E-PAD) and 4 bytes of (ICV), as shown in Figure 2. Only the Encryption Control is sent in the clear; the rest are encrypted with the Network Encryption Key (NEK) for the Logical Network (LN) to which the sender and receiver belong (see Section 4.12).

Encryption Control (ECtl), E-PAD, and ICV are mandatory in all SBs. The first byte of ECtl is the one-byte Encryption Key Select (EKS) field, followed by eight bytes of Initialization Vector (IV) used by the encryption algorithm. A value of 0x00 for the EKS indicates the default key, which is different for each station and is set by the manufacturer. The other 255 values are for dynamically determined NEKs. The E-PAD is between 0 and 7 zeros added to make the encrypted portion of the frame body a multiple of the encryption algorithm's 8-byte block size. A 32-bit CRC over all the fields following the ECtl up to the E-PAD (in cleartext form) is used for the ICV, so that the receiver can confirm correct reception and decryption of the SB. E-PAD and ICV are required even if encryption is bypassed.

The IEEE 802.1Q VLAN tag is copied from the MSDU. MAC Management Information is optional, but if present, the first two octets must be 0x887B. Payload type and length in the Type field are upper layer type/length information copied from the MSDU. MSDU data must be between 46 and 1500 bytes long, if the SB carries an MSDU. With the maximum MSDU data, plus the required and optional SB fields, a node must be able to accept an SB of at least 1616 Bytes.

4.3. MPDU Frame Structure

Two MAC Protocol Data Unit (MPDU) formats are supported: a long format (taking 313.5 - 1489.5 μsec) used for data and control and a short format (taking 72 μsec) for responses. The short format MPDU is for responses, and consists only of a response delimiter (preamble with frame control). Figure 4(a) depicts the long MPDU format.

The 25 bits of the frame control consist of 1 bit for Contention Control (CC), followed by 3 bits of Delimiter Type (DT), a 13-bit Variant Field (VF), with the 8-bit Frame Control Check Sequence last. During segment bursting or contention free access, CC=1 until the last segment. In this case, only higher priority traffic may preempt the sender for access to the channel following transmission of the MPDU and its response. The DT field encodes whether the delimiter is an SOF, and EOF, or a response delimiter, and whether or not a response is expected if it is an SOF or EOF. The FCS is an 8-bit CRC checking the other 17 bits.

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

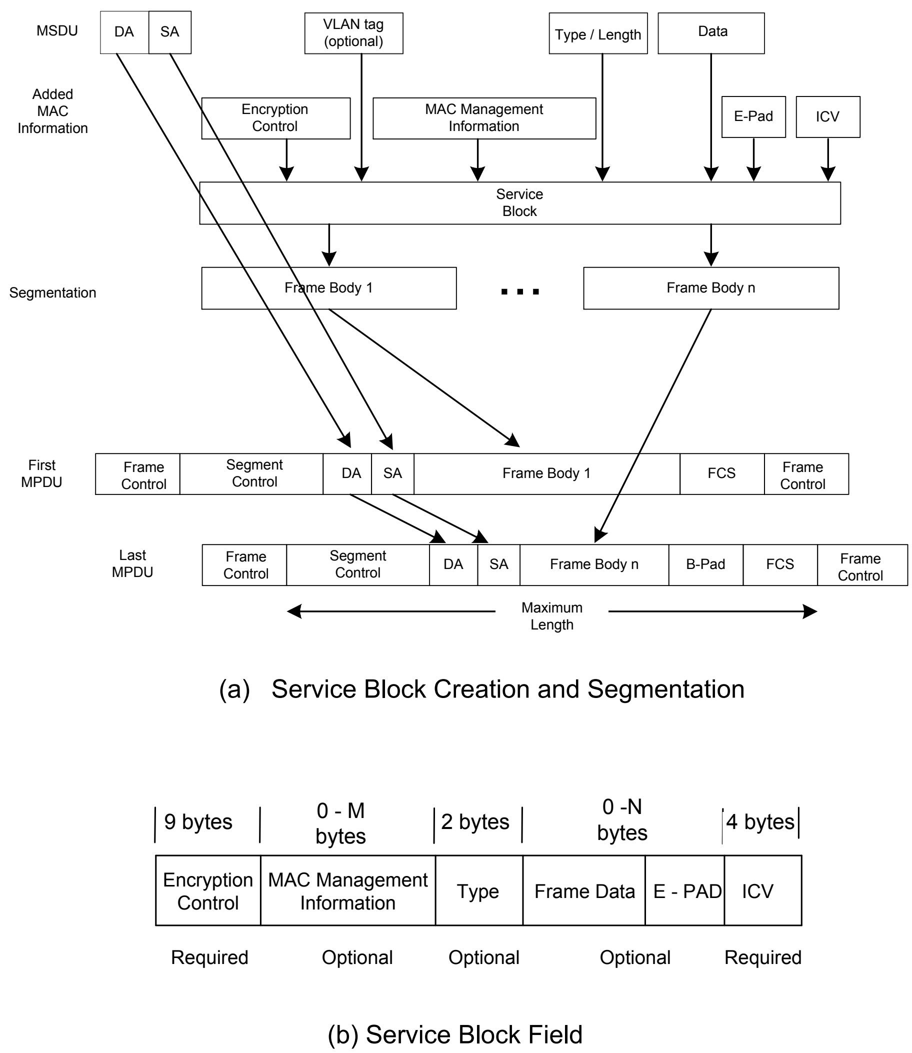

Figure 2. Service Block

In the SOF, the variant field holds an 8-bit frame length (FL) field and a 5-bit Tone Map Index (TMI) field. FL indicates the number of 40-symbol PHY transmission blocks in the payload of the long MPDU, and whether or not these are followed by a 20-symbol unit. The length encoding supports 8 lengths in 20 symbol increments from 20 symbols to 160 symbols, inclusive. TMI is the index that the receiver uses to find the tone map it needs to demodulate and to decode the payload, as set after channel estimation. Though there are five bits, indicating 32 possible TMs, half are reserved and one (0b00000) is used to indicate ROBO mode, leaving 15 available for negotiated tone maps per receiver. Note that use of ROBO mode restricts the FL field, as ROBO mode only supports 40-symbol physical transmission blocks.

Copyright © 2000 John Wiley & Sons, Ltd.

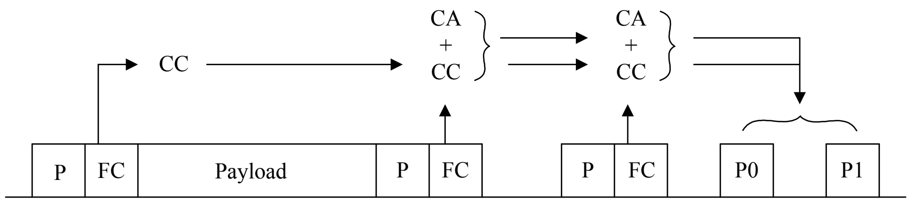

(a) CC + CA Fields used for Priority

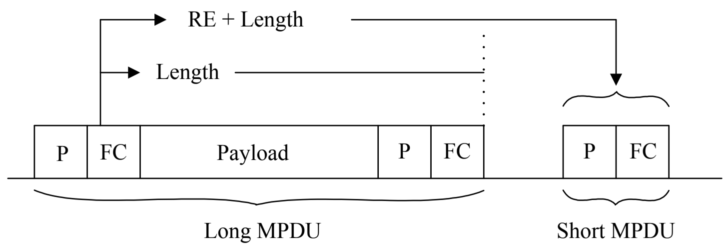

(b) Length & Response Expected used for VCS

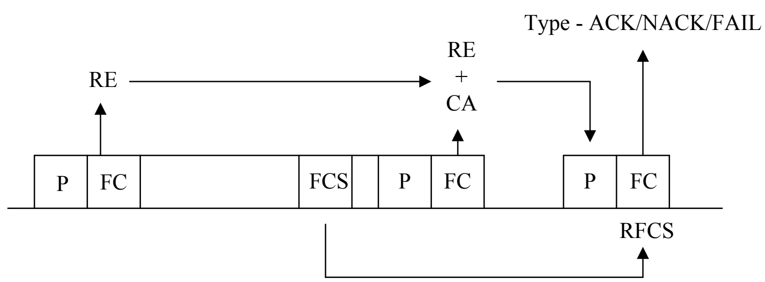

(c) FC & FCS fields used for ARQ

Figure 3. MAC Use of Frame Control Field

In the EOF, the variant field holds 2 bits of Channel Access Priority (CAP) of the frame. CAP is used during bursting to determine if a node can preempt the burst.

The variant field for ACKs also includes a 2-bit CAP that echoes the priority of the frame for which the response was generated. The remaining 11 bits are the Received FCS (RFCS), which are the 11 LSBs of the ACKed MPDU's FCS field. These allow the sender to confirm that the ACK was intended for the frame it sent, and not some other sender's frame that the receiver may have heard. If these bits do not match, then the sender treats it as a collision.

NACK and FAIL responses share a common VF. The first two bits are the CAP as with ACKs, followed by 1 bit of response type (0=NACK, 1=FAIL), and 10 bits of RFCS. NACK does not use the RFCS field, and the sender treats the FAIL as a collision if the RFCS bits do not match the sender's FCS bits.

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

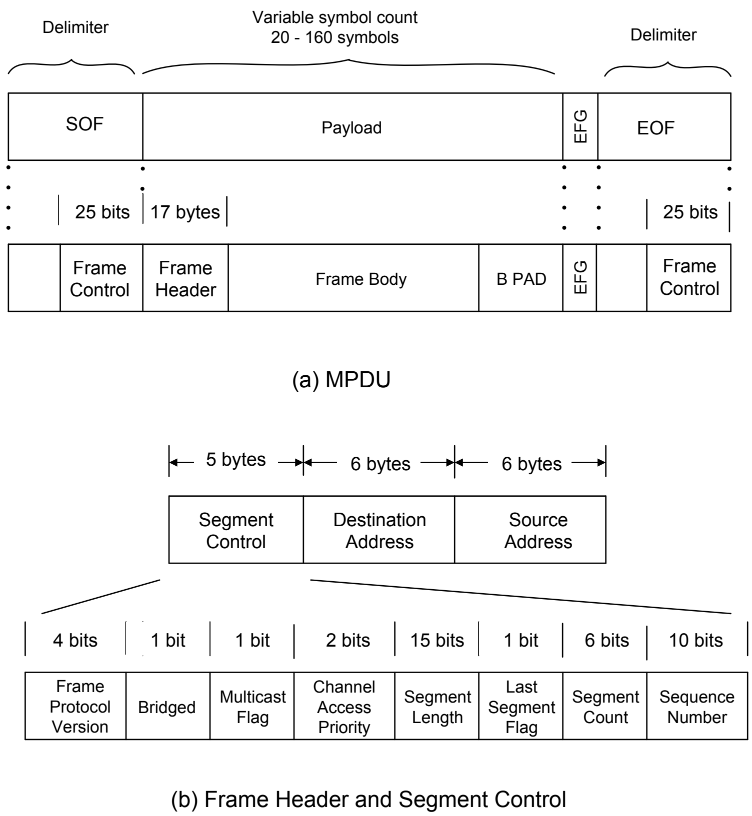

Figure 4. Long MPDU Format

Figure 3 shows the relevance of the Frame Control field to the MAC. There are three ways in which FC is used by the MAC. First, (a) shows its use for Virtual Carrier Sense. The length field of the SOF, and the type field of the SOF and the EOF (including whether or not a response is expected) are used to determine when the priority resolution period should end and contention begin. Second, (b) shows how the contention control bit of all delimiters and the channel access priority bits of the SOF and the response delimiters, are used for contention control purpose (see Section 4.9). Finally, (c) depicts how part of the FCS of the MPDU is copied into the response delimiter to confirm ACKs and FAILs to the sender in the ARQ mechanism (see Section 4.7).

The long MPDU consists of the SOF delimiter (preamble and FC), the 20 to 160 symbol payload, the EFG, and the EOF delimiter (refer to Figure 4). The MPDU coincides with the PPDU since the delimiter FC fields contain information used by both the PHY and the MAC, but in the MPDU, the contents of the PHY payload are exposed since the receiver has the correct tone map to demodulate and decode that PPDU field. In the long MPDU, the PHY payload is opened to reveal five bytes of segment control (SC), 48-bit IEEE 802.3 destination address (DA) and source address (SA) (6 bytes each), followed by the variable length frame body, zero block padding needed to fill out the physical transmission block (B-PAD), and a

Copyright © 2000 John Wiley & Sons, Ltd.
Prepared using dacauth.cls

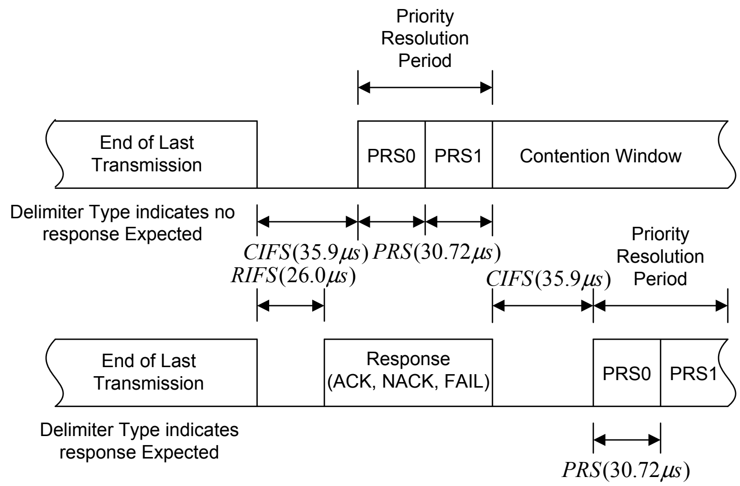

Figure 5. Interframe spacing in HomePlug 1.0

2-byte frame check sequence (FCS). B-PAD is only present in the last segment of a service block (refer to Section 4.8). The first three fields are also called the frame header (FH), and are sent unencrypted. FCS is a 16-bit CRC computed over bytes from the SC through the last byte of the B-PAD. Modulation and coding rate determine how many bytes can be carried in the 20 to 160 symbols allowed for the MPDU payload.

Segment Control's 40 bits are shown in Figure 4(b) and consist of 3 bits of Frame Protocol Version 2 reserved bits, a Multicast Flag (MCF, 0b1 for multicast or broadcast frames), 2 bits of CAP, 15 bits of Segment Length (SL) in bytes, a Last Segment Flag (LSF, 0b1 if this is the last segment of an SB), 6 bits of Segment Count (SC), and 10 bits of Sequence Number (SN). If the MCF is set, then the service block must have a MAC Management Multicast with Response entry; if the SB is segmented, then all segments have MCF=1. The SL field indicates the length of the frame body in bytes, not including the E-PAD or the ICV (see Sections 4.8 and 4.12). The SC and SN fields are used by segmentation and reassembly to reconstruct service blocks correctly. Each source maintains an SN per destination per priority class.

4.4. Interframe Spacing and Timing

There are several points at which the protocol requires a delay for processing or to change from receive to transmit mode. When a response is expected (as indicated by both the SOF and the EOF FC fields), the responding station waits for a Response InterFrame Space (RIFS) of 26.0 μsec before transmitting its response. After the response, if one is expected, or after the EOF, if no response is expected, all stations consider the medium to be busy for a Contention InterFrame Space (CIFS) of 35.84 μsec before the Priority Resolution Period begins. Each 35.84 μsec Priority Resolution Slot may contain a Priority Resolution Signal lasting 30.72 μs. sent by one or more stations. Due to tight synchronization and minimal propagation delays, these signals are additive, and the slots have their own processing delay built in.

As shown in Figure 5, the RIFS and CIFS are needed for propagation and processing times. The Extended InterFrame Space (EIFS) is used when a station is not sure of the state of

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

the medium while it listens for another station to start transmission, e.g. It accounts for the maximum time that a station could be transmitting by including the nonROBO maximum frame duration along with the various delimiters, priority slots, CIFS, RIFS, and EFG. All of this amounts to 1695.0 μsec. EIFS is also used to determine how long the channel is considered busy after a collision or a Frame Control error (since the length of the frame cannot be determined reliably for more accurate VCS).

4.5. Priority Resolution

Before the contention period, there is a priority resolution period (PRP) consisting of two 35.84 μsec. priority resolution slots (PRSs). Using the PRP, only the stations with the highest priority traffic to send may contend for the medium in the contention period.

Stations contend using PRS0 and PRS1 to determine maximum priority traffic on the network. Four priority levels are supported: CA3 and CA2 for time-sensitive, high priority traffic, and CA1 and CA0 for lower priority traffic. If the EOF delimiter or Response delimiter of the frame immediately preceding the PR period has the contention control bit set, then any nodes with the same or lower priority defer. Otherwise, CA3 and CA2 nodes assert PRS0, which will cause CA1 and CA0 nodes to defer. CA3 nodes also assert PRS1, which will cause CA2 nodes to defer; if PRS0 is not active, then CA1 nodes will assert PRS1, which will cause CA0 nodes to defer. Only nodes from the highest (winning) priority class contend for access to the medium using the contention window. Priorities should be assigned according to the 802.1D guidelines. This allows HP 1.0 networks to operate with RFC 2205 RSVP (Resource Reservation Protocol) and the internet draft standard Subnet Bandwidth Manager (SBM) to provide differentiated quality of service (diffserve) for multimedia traffic [3, 22]. CA3 is used for VLAN tag priority 6 or 7 (network control and extremely delay-sensitive traffic such as voice), and CA2 is used for VLAN tag priority 4 or 5 (delay-sensitive traffic such as video and audio, and important business applications subject to admission control). CA1 is used for "excellent effort" (highest quality best effort service) and is the default level. CA0 is for standard LAN traffic (when this is labeled as such) and background traffic. MAC management entries may also use the priority necessary (e.g., for channel estimation requests and replies).

4.6. Channel Access

The contention period is used by the Backoff Procedure, which is disabled during segment bursting or contention-free access (see below). First, PCS/VCS is used to detect the medium state and wait either until it is idle, or until the priority resolution period (in which case the station asserts its priority unless it must defer), or until it has won the PRP and can contend for the medium in the contention window.

When a station first starts contending for the channel, it randomly picks one of the first eight contention slots following the PRP to start its transmission, setting a Backoff Counter (BC) to the number of slots to leave empty for others to use. If no other SOF delimiters are detected before the selected slot arrives, then the station starts transmission in that contention slot. Otherwise it defers and sets its Deferral Counter (DC) according to a priority-dependent schedule. As unused slots go by in contention periods at the same priority level as the traffic waiting to be sent by the station, the backoff counter is decremented until one of two things happens. If the BC reaches zero, the station starts transmission, then awaits a response if one is expected. If the DC reaches zero, then so many stations at the same priority level have made

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

Table III. Backoff Schedule

CAP →	High	High	Low	Low
BPC	CW	DC	CW	DC
0	7	0	7	0
1	15	1	15	1
2	15	3	31	3
> 2	31	15	63	15

their presence known that the sender randomly picks a new value for the BC from the next larger range of values (again, depending on the priority level of its data), in much the same way that it would if it inferred a collision. This allows the extra information from deferrals to be used to avoid costly collisions.

When a station sends a frame and either does not expect a response or receives a matching ACK from its destination, then the frame has successfully been sent and the next frame is readied for transmission. If it was the last segment in a service block, then success is reported to the host interface.

If a valid FAIL response is received, and the maximum number of FAIL retries (between 0 and 6) has not been exceeded, then the station waits an extended time (10 ms.) before resuming its efforts.

Otherwise, if it either receives a NACK or it infers that a collision has occurred, the station invokes the backoff procedure by increasing the contention window size in accordance with the backoff schedule, up to a maximum of 32 for CA3 and CA2, or 64 for CA1 and CA0), and picking a new delay time for the BC. If the station has tried the maximum number of attempts that can be made using the destination-specific Tone Map, then the mode is switched to ROBO and further attempts are made to deliver the frame. Switching to ROBO mode may require resegmenting the service block (see Section 4.8). If the limit is exceeded on the number of transmission attempts that may be made (either in total or for ROBO mode, for NACKs or for collisions), then the frame is discarded and failure is reported.

In Table III, high priority is CA3 and CA2, low priority is CA1 and CA0. BPC is the Backoff Procedure Counter, which counts the number of times that the Backoff Procedure has been invoked due to collisions or deferrals. The BPC is reset to 0 after a FAIL, as the amount of traffic is expected to change after the long delay. DC is the Deferral Counter and CW is the Contention Window maximum (the minimum is zero, so the CW size is actually CW+1).

Collisions are inferred under several circumstances. First, if no response is detected when one is expected, then a collision is assumed, although this could be due to a bad channel. Even if a delimiter is detected, if the DC field of the delimiter is bad, or if the FC does not indicate that it is an ACK, NACK, or FAIL when a response is expected, then then a collision is inferred. Finally, FAIL and ACK responses contain a Response FCS (RFCS) field that echos the 10 or 11 (respectively) LSBs of the FCS from frame they acknowledge. If the RFCS field does not match the one sent, then the sender assumes a collision has occurred.

4.7. MAC Error Control

The MAC implements a stop-and-wait ARQ error control method. Acknowledgement and retransmission are performed on a per-segment basis. Expectation of a response is indicated in the FC fields of both the SOF and the EOF delimiters. The response is sent by the destination

Copyright © 2000 John Wiley & Sons, Ltd.

in a short format MPDU, consisting of a solitary response delimiter. Three types of responses are used: ACK, NACK, and FAIL. All three include a 10- or 11-bit Response FCS (RFCS) field, which is a copy of the LSBs of the 16-bit FCS of the MPDU for which the response was sent, as well as the channel access priority (CAP), which is used for contention control and preemption.

An ACK response has an 11-bit RFCS, and only if the RFCS matches the least significant 11 bits of the FCS of the transmitted frame is that frame considered acknowledged. Otherwise, the sender treats it as a collision. NACK and FAIL have 10-bit RFCS fields. NACK is used to indicate that a frame was received with errors, as indicated by the FCS, while FAIL is used to indicate that the receiver has no buffers available for reassembly of the service block or that the segment was received out of order. NACK and FAIL with valid RFCS surely indicate failure, and the absence of a valid ACK when one is expected is also taken to indicate a failure or collision. Frames receiving a NACK or no response may contend for retransmission right away, but a FAIL response requires the sender to wait a longer time (the 10 msec FAIL delay) before retrying, in the hope that the receiver has finished reassembling the service block on which it was working earlier, and will have the resources available to begin reassembly of the service block to which the frame belongs. In fact, only higher priority frames may be sent by that station to the destination address that responded with a FAIL. If the segment receiving the FAIL response was not the first segment, then the entire service block transmission attempt is aborted and the sender starts over. Stations attempt to transmit the frame until the retry limit is exceeded or the transmit lifetime is exceeded.

Optional, partial ACKs are available for multicast and broadcast frames; the sender can specify any station in its Logical Network (not just those in the multicast group) to respond to a multicast or broadcast frame. The Multicast with Response MAC management entry of the service block has the actual 6-byte multicast destination address, while the DA in the layer management MAC frame is a proxy for the multicast and will generate a response if the delimiter type indicates that one is expected. In the case of partial ACKs, FAIL may be treated as an ACK.

A tight response timeout interval is used to decide how long to wait for a response before inferring that a collision has occurred. A valid ACK inspires the node to send the next segment (if there is more of the SB to send) or to report success (if this was the last segment of the SB).

4.8. Segmentation and Reassembly

A service block is segmented into MPDU frame bodies as needed, depending on the number of bytes that an MPDU can carry due to the 160-symbol payload limitation and the data rate determined by the Tone Map. Each MPDU includes the MSDU's DA and SA in its frame header, and each has its own SC and FCS. Only the last segment has B-Pad, as all the others should completely fill their MPDU frame body. The Segment Control (SC) field is always present, even if the SB is not segmented (i.e., it is a one segment service block). SC includes the segment length in bytes exclusive of B-pad, SC, and LSF. Segment bursting is used to send all segments of an SB in as short a time as possible (modulo preemption by higher priority - see Section 4.9).

The Segment Control (SC) field of the frame header includes a sequence number (SN), a segment count, and a last segment flag (LSF). These are used to detect duplicates and to

Copyright © 2000 John Wiley & Sons, Ltd.

recover. The SN is the same for all the segments in same service block, and is incremented for each new service block. The Segment Count indicates which segment within an SB the MPDU holds, starting with a value of 0 for the initial segment. LSF indicates the last segment in a service block.

For each SA and priority, a receiver stores the most recent ⟨SA,SN,SC,LSF⟩ tuple. List entries are compared to incoming frame header fields to detect duplicates and omissions. New ⟨SA,priority⟩ pairs create new list entries, with an ACK if it is the first segment. Likewise, the first segment of a new received SN causes the list to be updated. Segments received in correct order naturally generate ACKs and update the list; and when the LSF is set, indicating that reassembly is complete, the reassembled service block is passed up for decryption. Most recent segment duplicates overwrite the reassembly buffer segment and generate an ACK, while older segments or segments that would leave a gap in the reassembly buffer cause the receiver to send a FAIL response and flush the reassembly buffer. This forces the sender to start transmission of the SB over from the first segment.

Responses are only sent for MPDUs requiring a response, and a reassembly timer is used to prevent a partially reassembled SB from occupying resources indefinitely. Every station is required to have at least one MPDU reassembly buffer with timer. More than one is allowed, but each buffer must have its own timer.

4.9. Segment Bursting and Contention-Free Access

Segment Bursting uses the contention control and channel access priority fields in the delimiter FC to allow a station to send segments belonging to the same service block without repeatedly contending for the channel. Stations with higher priority traffic can preempt a segment burst by asserting the priority bits between segments in the burst. Segment Bursting is limited to two consecutive MPDUs at CA3, in order to provide low jitter for CA3 traffic. This mechanism allows for efficient use of the medium while still preserving high quality differentiated service.

To use segment bursting, a station contends as usual. Once it gains access, it sets the CC bits in the SOF and EOF FCs to 1, and inserts its MPDU's priority in the CAP field of the EOF FC. After each segment is sent (and acknowledged, if required), the priority resolution period (PRP) is open to stations with traffic at a higher priority level. The sender also asserts its priority in the PRP, listening for a station with higher priority traffic. As long as it is not preempted, the station continues to transmit the rest of the segments in the service block right after the PRP. In the MPDU containing the last segment, the CC bit in FC is cleared to allow normal contention to occur. If the sender is preempted, then it goes back to priority resolution and contention as usual.

A more restricted way to limit contention is also supported. Contention-free access (CFA) is only available to stations with CA3 traffic, and allows a station to send all of its service blocks (even to different DAs) at level CA3 using CC=1 in FC. It is limited to seven consecutive MPDUs per CFA burst. CC is reset to 0 in the FC of the last segment of the last MPDU in a CFA burst.

4.10. MAC Management

The optional MAC Management Information Field of a service block must start with the two byte MTYPE field with value 0x887B, which is the IEEE Assigned Ethertype. This is followed by 1 byte of MAC Ctl Field (MCTRL), then the variable number of MAC Management

Copyright © 2000 John Wiley & Sons, Ltd.

Entries (MMENTRYs). Each MMENTRY consists of 1 byte MAC Management Entry Header (MEHDR), 1 byte MAC Management Entry Length (MELEN), then a variable length MAC Management Entry Data field (MEENTRY). The MSB of MCTRL is reserved, the remaining 7 bits indicate the number of MM entries to follow. The MELEN field allows extension of MAC Management functions, with older nodes able to skip over the fields they do not understand.

The first 3 bits of each MEHDR is the MAC Entry Version (MEV), which must be 0b000 or the entry is discarded using the MELEN length field. The last 5 bits are the MAC Entry Type (METYPE). MAC entries may come from host via the M1 interface, or may be generated within the MAC itself. Order of entries matters for efficient processing; they must follow the order:

1. Multicast with Response METYPE,
2. Request Channel Estimation METYPE,
3. Channel Estimation Response METYPE,
4. Replace Bridge Address METYPE,
5. all other METYPEs in any order.

METYPEs include

- Multicast with Response,
- Channel Estimation Request and Response (from the MAC itself only),
- Replace Bridge Address,
- Set/Confirm Network Encryption Key,
- Parameters and Stats Request/Response,
- Vendor Specific, and
- Manufacturer-specific METYPE.

The manufacturer-specific METYPE must be received over the host interface, while the vendor specific METYPE may appear over the host interface. The first 3 bytes of the vendor-specific METYPE are the vendor's IEEE assigned Organizationally Unique Identifier (OUI). This entry may be ignored by the receiver if the receiver's OUI does not match the sender OUI. Unless otherwise specified, the SA of the MPDU is taken as the requester/responder for all requests/responses in an SB, and the DA is taken as the station for which the request/response is intended.

4.11. Channel Estimation Control Function

Stations must maintain fresh channel information on the channels over which they transmit unicast data. This is accomplished by the CE Control Function using the CE Request and CE Response MAC Management entries. A Tone Map is considered to be stale if the link is new, or if it has expired (more precisely, it has not been used successfully for 30 seconds), or if the sender had to revert to ROBO mode during backoff. A stale TM may not be used to send a unicast SB with MPDU in it. All broadcast or multicast MPDUs are sent using ROBO modulation with a TMI of zero, and so are not subject to this requirement.

Channel estimation is requested by sending a Request Channel Estimation MMENTRY. This one-byte MAC data entry contains the 4-bit channel estimation version capability of requester (0x0 for HP 1.0), and causes the receiver to send a CE response. CE requests should not be performed more frequently than once every 4.5 sec. (on average over any 5 min. interval),

Copyright © 2000 John Wiley & Sons, Ltd.

excluding CE requests made during recovery (after having to use ROBO mode or exceeding the retry limit). A station forced to drop to ROBO mode does not have to obtain a new TM before it finishes sending the current SB. A station having to drop to ROBO mode repeatedly does not have to continually request CE, but it must not use a TM that has not worked for 30 sec.

The Request CE MMENTRY must be sent using ROBO modulation, and it is recommended to keep the SB small due to the low bandwidth efficiency of ROBO mode. A station should avoid sending a Request CE MMENTRY unless it also has a SB to send. Once a new CE and TMI are received via the CE Response, a source must start using the new TM for all unicasts to that destination by the third SB following receipt of the CE Response.

The CE Response MMENTRY data field includes the 4-bit CE Response version number (must be 0x0 for HP 1.0 or else the requester discards the entry), the 5-bit TMI for the requester to use in for future unicasts to this destination, and 84 Valid Tone Flags (0b1 indicates that the tone is to be used, subject to the Tone Mask). In addition, there is a 1-bit FEC Rate selector and a 2-bit Modulation Method selector. A 1-bit Bridge Proxy bit indicates that the TM is being proxied for the DAs that follow. A 7-bit Number Bridged DAs field (NBDAS) indicates the number of DAs for which this TM is proxied, with a 6-byte MAC address for each of the Bridged DAs concluding the entry. CE Response MPDUs must be unicast with response expected. To avoid confusion, the TMI used in the CE Response should be different from the TMI last sent for that source to use. A receiver may send an unsolicited CE Response with a new TM if one is needed, as long as it has been at least 1 second since that station sent a CE Response to that source. Each station can hence support up to 15 different senders at a time using channel adaptation, while excess senders must use ROBO mode.

4.12. Privacy and Key Management

Power lines are shared from the transformer to all of the residences served by the transformer, so it is possible for a residence to hear the PLC transmissions of a nearby residence. It is therefore necessary to protect the privacy of users cryptographically, since installing low-pass filters would to some extent negate the cost advantages of the technology. To this end, nodes form Logical Networks (LNs) based on cryptographic isolation.

HP 1.0 uses DES (FIP Std. Publication 46-3) in cipher block chaining (CBC) mode. Keys are generated from passwords using the PBKDF1 function from PKCS#5 v2.0, Password-based Cryptography Std., with MD5 as the underlying hash algorithm. Stations store and retain both their default key (for re-key operations only) and any Network Encryption Keys (NEK) received (for any other transmissions) in nonvolatile memory.

All transmissions within a logical network are encrypted with the NEK that defines that logical network. To participate in a LN, a station must have the NEK for it. Stations may obtain a NEK by

1. Password Entry by the user and generation as described above, or
2. Network Entry by receipt of a Set NEK MMENTRY from another station encrypted using any key known to both stations.

A station may be a member of more than one LN, and may be required to store more than one (EKS,NEK) pair. Stations are not required to support participation in more than one LN at a time, however.

Copyright © 2000 John Wiley & Sons, Ltd.

Table IV. Simulation Results for HomePlug1.0

UDP Simulation Results for HomePlug 1.0			
	1 UDP Node	2 UDP Nodes	3 UDP Nodes
MAC Throughput	8.08 Mbps	7.46 Mbps	7.46 Mbps
TCP Simulation Results for HomePlug 1.0			
	1 TCP Node	2 TCP Nodes	3 TCP Nodes
MAC Throughput	6.16 Mbps	6.15 Mbps	6.12 Mbps
TCP Throughput	5.92 Mbps	5.91 Mbps	5.88 Mbps

A station without any NEK can use Network Entry to obtain an NEK and enter the LN by means of a default key generated from a manufacturer-determined password unique to the particular station. This default password must be entered at another station participating in the LN, and using Password Entry, that station must generate the default NEK for the new station and use that NEK with EKS=0x00 to encrypt and send the LN's NEK and EKS identifier to the new station in a Set NEK MMENTRY. The new station returns a Confirm NEK MMENTRY to the station that sent it the Set NEK MMENTRY.

Set NEK MMENTRY is never be sent in cleartext, and if a Set NEK MMENTRY is received in cleartext, it must be rejected. The 9 bytes of the Set NEK MMENTRY data field contain the 1-byte EKS and the corresponding 8-byte NEK. The Confirm NEK MMENTRY has an empty data field. The EKS values associated with an LN's NEK must be the same for all stations in the LN.

5. HomePlug1.0 Simulation Results

This section presents performance results[§] for the HomePlug 1.0 Powerline network using an event-based C program simulation. All scenarios assume QPSK with a $\frac{3}{4}$ coding rate on various links, and a maximum TCP segment size of 1460 bytes. In this simulation, we use UDP and TCP traffic sources. UDP traffic is generated with an exponential inter-arrival time with a 100 microsecond average. The UDP packet size is assumed to be a constant 1460 bytes with priority 0. TCP traffic is also generated with exponential inter-arrival time with 100 microsecond average at priority 0, and we assume that TCP traffic sources always have data to send. Every time a node has a chance to send, it is allowed to send the maximum segment size of 1460 bytes without headers. Table IV summarizes the simulation results [14]. The single node UDP traffic simulation scenario shows the best throughput in our simulations since there is no contention at all. Table IV also shows that channel contention with 2 and 3 UDP nodes causes a modest reduction in channel throughput.

In the TCP traffic simulation, though the single node scenario has only one traffic source, the bandwidth must be shared with data and response frames (e.g., ACK packets) thus resulting in a lower performance than the UDP traffic simulation. The MAC throughput represents the total number of transmitted bytes divided by the simulation time regardless of successful delivery. The TCP throughput includes only the successfully delivered data and ACKs.

[§] An early version of these results were presented in [14].

Copyright © 2000 John Wiley & Sons, Ltd.

Table V. Laboratory Measurement of FTP Results for HomePlug1.0

Measured FTP Results for HomePlug 1.0			
	1 FTP Node	2 FTP Nodes	3 FTP Nodes
MAC Throughput	6.21 Mbps	6.15 Mbps	6.27 Mbps

6. Performance Measurements in Ideal Laboratory Setting

In addition to simulating the performance§ of the HomePlug 1.0 protocol, we were also able to construct a real PLC network using actual commercial and reference design PLC devices. In this experiment, there were 4 desktop computers. A 450 MHz Pentium II desktop computer (PC-2 configured as a file server) was equipped with 128 MB RAM, a 3-COM fast Ethernet card, and a PLC PCI card. Two 700 MHz Pentium III desktop computers (PC-3 and PC-4) were each equipped with 256 MB RAM and a PLC PCI card. A 266 MHz Pentium MMX desktop computer (PC-1) was equipped with 64 MB RAM and a 3-COM fast Ethernet card. The PC-1 computer is connected to an Ethernet-to-power line bridge, which converts packets generated from the Ethernet card into PLC compatible packets, and vice versa. All computers were connected to power lines via wall outlets in the same room to give an almost ideal powerline channel.

To evaluate the throughput of this real PLC network, we conducted the following experiment. A 215,502,106 byte file was placed on the server running an FTP daemon (the large file size was chosen to minimize hardware uncertainty and human error). Client computers made FTP requests for the file. We tested different numbers of FTP connections, up to 3, using individual client machines in our PLC network.

The experimental results are also given in Table 6. The aggregated traffic in the table is calculated by adding all observed data rates of all connections. The experimental results show that the real PLC network performance is about 6 Mbps. When there is only one FTP connection, the observed throughput is 6.21 Mbps. The results suggest that one TCP connection will not fully utilize the PLC network, because the server has to stop if no ACK packets are received from the client. Aggregated traffic volume decreased as the number of connections increased from 1 to 2. This phenomenon is because of the ACK packets and the packet overhead increase as the number of connections increased. Although the network utilization improves, the improvement cannot compensate for the loss due to these overheads. When we increased the number of connections from 2 to 3, the PLC network had the highest throughput of 6.27 Mbps. This is because the network utilization increased as the number of connections increased which compensated for the packet overhead and ACK overhead.

7. Field Performance of HomePlug 1.0 in a Residential Setting

This section reports on the result of a field test conducted in a typical Florida home‡. While the maximum throughput of the powerline network is 14 Mbps, this rate is the net PHY data rate and is only possible under ideal conditions. Indeed for the Powerline network we have

‡A short version of some of these results was presented in [16].

already seen that simulation and ideal measured network throughput was of the order of 5-7 Mbps. This section shows real world performance in a residential setting.

7.1. Test setup

The performance test was conducted in a 10 year old, 2700 sq. ft. house in Gainesville, Florida. This residential location was selected for the field test since Powerline LAN technology is targeted at the standard home.

For this test, we used two laptop computers. One machine was a 700 MHz Pentium III equipped with a powerline USB bridge. The other machine was a 500 MHz Pentium II equipped with a powerline Ethernet bridge. The USB bridge is intended for connection to the power line through the USB port. The Ethernet bridge is intended for connection to the power line through an Ethernet card. These tests were designed to measure the data transfer capability of power line networks. To obtain results that would accurately reflect user experience, we used the following scenarios for testing.

(1) Scenario-1: FTP – We used the WSFTP program to do a file transfer. This utility is very widely used. We set up this utility with the following parameters.
Transmit Buffer Size: 4096 bytes
Receive Buffer Size: 4096 bytes
File transfer size: 40 Mbytes;
(2) Scenario-2: TCP – We used the WSTTCP program to test the TCP performance. WSTTCP is a popular program for TTCP. We set up this utility with the following parameters.
Buffer Length: 4096 bytes
Number of Buffers Sent: 5000
Total data exchange: 20 Mbytes
Protocol: TCP.

7.2. Test results

Table 6 shows the HomePlug 1.0 network throughput. The powerline network had full connectivity. Furthermore, powerline transfers were always at a near constant rate. For longer distance separations the performance dropped. Throughput was not the same for the forward link and the reverse link. Sometimes, the throughput on a given link gives a performance twice that achieved in the other direction. Varying channel characteristics on the power line cause the different performance results.

8. Conclusions

This paper provided a description of the powerline environment and the PHY and MAC protocols for HomePlug 1.0 Standard. Simulation and laboratory testing demonstrate a throughput of about 6-8 Mbps. HomePlug 1.0 powerline LAN operating in a typical residential environment had excellent coverage and reliable throughput of 1.6 to 5.3 Mbps at the application level. Throughput for the powerline network is in general not symmetric for a pair of nodes. Additional studies and more extensive testing are currently under way in both

Copyright © 2000 John Wiley & Sons, Ltd.

Table VI. Field Test Throughput Results.

Transmitter	Receiver	Distance between transmitter &receiver	Powerline	
			WSFTP Throughput (Mbps)	TTCP Throughput (Mbps)
Laptop 1 (Dining Room)	Laptop 2 (Dining Room)	∼2 feet	4.2	5.2
Laptop 1 (Den)	Laptop 2 (Dining Room)	∼23 feet	4.5	5.3
Laptop 2 (Office)	Laptop 1 (Kitchen)	∼35 feet	4.0	4.5
Laptop 1 (Kitchen)	Laptop 2 (Office)	∼35 feet	3.1	3.1
Laptop 1 (Children Room	Laptop 2 (Office)	∼70 feet	1.9	1.8
Laptop 2 (Office)	Laptop 1 (Children Room)	∼70 feet	4.1	3.9
Laptop 1 (Swimming Pool)	Laptop 2 (Office)	∼60 feet	2.0	1.6
Laptop 2 (Office)	Laptop 1 (Swimming Pool)	∼60 feet	2.4	2.8

laboratory and field conditions, and we expect to learn more about these networks at their conclusion.

In-home powerline communications have the potential for significant improvements. These enhancements include higher frequency bands, higher level modulations, better forward error correction code, etc. In October 2002, the HomePlug Powerline Alliance announced plans for the development of next generation specifications. Named HomePlug AV, the new specification will be designed to support distribution of data and multi-streaming entertainment including High Definition television (HDTV) and Standard Definition television (SDTV) throughout the home.

REFERENCES

1. Barnes JS. A physical multipath model for power distribution network propagation. *Proceedings of International Symposium on Power-line Communications and its Applications,* 1998; 76–89.
2. Bingham JAC. Multicarrier modulation for data transmission: an idea whose time has come. *IEEE Communications Magazine,* May 1990; 5–14.
3. Braden R (Ed.), Zhang L, Berson S, Herzog S, Jamin S. RFC 2205 Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification. September 1997; *IETF.* (Updated by RFC2750) (Status: PROPOSED STANDARD)
4. Brown PA. Some key factors influencing data transmission rates in the power line environment when utilizing carrier frequencies above 1 MHz. *Proceedings of International Symposium on Power-line Communications and its Applications,* 1998; 67–75.
5. Brown PA. Power line communications - past, present, and future. *Proceedings of International Symposium on Power-line Communications and its Applications,* Sept 1999; 1–8.
6. Chow JS, Tu JC, Cioffi JM. A discrete multitone transceiver system for HDSL applications. *IEEE Journal on Selected Areas in Communications,* **9**(6): Aug 1991; 895–908.
7. Design of Systems on Silicon. Home page. http://www.ds2.es/home/index_total.php [Oct 2002].
8. Dutta-Roy A. Networks for home. *IEEE Spectrum,* Dec 1999; **36**(12): 26–33.

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls

9. Gardner S, Markwalter B, Yonge L. HomePlug standard brings networking to the home. *Communication Systems Design Magazine,* http://www.commsdesign.com/main/2000/12/0012feat5.htm, Dec 2000.
10. Gershon R, Propp D, Propp M. A token passing network for powerline communications. *IEEE Transactions on Consumer Electronics,* May 1991; **37**(2): 129–134.
11. HomePlug Alliance. HomePlug 1.0 Specification. Jun 2001.
12. IEEE 802.11 WG. Part II: wireless LAN medium access control (MAC) and physical layer (PHY) specifications: high-speed physical layer in the 5 GHz band, *Supplement to IEEE 802.11 Standard,* Sep 1999.
13. Kaizawa Y, Marubayashi G. Needs for the power line communications. *Proceedings of International Symposium on Power-line Communications and its Applications,* 1998; 153–157.
14. Lin YJ, Latchman HA, Lee MK, Katar S. A power line communication network infrastructure for the smart home *IEEE Wireless Communications,* Dec 2002; **9**(6): 104–111.
15. Lin S, Costello DJ. Error Control Coding. 1983; *Prentice Hall,* Englewood Cliffs, NJ.
16. Newman R, Lee MK, Latchman H, Katar S, Yonge L. Field performance comparison of IEEE 802.11b and HomePlug 1.0. *Proccedings of the 27th Annual IEEE Conference on Local Computer Networks,* 2002.
17. Philipps H. Performance measurements of powerline channels at high frequencies. *Proceedings of International Symposium on Power-line Communications and its Applications,* 1998; 229–237.
18. Prasad R, van New R. OFDM Wireless Multimedia Communications. 2000; *Artech House,* City.
19. Proakis JG. Digital Communications. 1995; *McGraw Hill,* New York.
20. Starr T, Cioffi J, Silverman P. Understanding Digital Subscriber Line Technology. 1999; *Prentice Hall,* Englewood Cliffs, NJ.
21. Stallings W. Data and Computer Communications (6th Ed.). 2000; *Prentice Hall,* Englewood Cliffs, NJ.
22. Yavatkar R, Hoffman D, Bernet Y, Baker F, Speer M. RFC 2814 SBM (Subnet Bandwidth Manager): A protocol for RSVP-based admission control over IEEE 802-style networks. May 2000; *IETF.* (Status: PROPOSED STANDARD)
23. Zhou WY, Wu Y. COFDM: an overview. *IEEE Transactions on Broadcasting,* Mar 1995; **41**(1): 1–8.

Copyright © 2000 John Wiley & Sons, Ltd.

Prepared using dacauth.cls