



HomePlug AV White Paper

Copyright © 2005, HomePlug® Powerline Alliance, Inc., All Rights Reserved

THE DOCUMENT IS PROVIDED "AS IS," AND THE HOMEPLUG POWERLINE ALLIANCE (INCLUDING ANY THIRD PARTIES THAT HAVE CONTRIBUTED TO THE DOCUMENT) MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE DOCUMENT ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

NEITHER THE HOMEPLUG POWERLINE ALLIANCE NOR ANY THIRD PARTY WILL BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR RELATING TO ANY USE OR DOCUMENT OF THE DOCUMENT.

HomePlug AV White Paper

Introduction

HomePlug AV (HPAV) represents the next generation of technology from the HomePlug Powerline Alliance. Its purpose is to provide high-quality, multi-stream, entertainment oriented networking over existing AC wiring within the home, while addressing interoperability with HomePlug 1.0. HPAV employs advanced PHY and MAC technologies that provide a 200 Mbps (million bits per second) class powerline network for video, audio and data. The Physical (PHY) Layer utilizes this 200 Mbps channel rate to provide a 150 Mbps information rate with robust, near-capacity communications over noisy power line channels. The Medium Access Control (MAC) Layer is designed to be highly efficient; supporting both TDMA and CSMA based access with AC line cycle synchronization. The TDMA access provides Quality of Service (QoS) guarantees including guaranteed bandwidth reservation, high reliability and tight control of latency and jitter. The CSMA access provides four priority levels. AC line cycle synchronization provides superior channel adaptation in the face of common line cycle-synchronized noise. The Central Coordinator (CCo) controls the activities of the network, allocating time for CSMA use and scheduling the TDMA use.

Homeplug AV also provides advanced capabilities consistent with new networking standards. Advanced Network Management functions and facilities are capable of supporting user plug-and-play configuration as well as service provider set-up and configuration. HPAV offers tight security based on 128-bit AES and makes provision for dynamic (automatic) change of the encryption keys and for several different user experiences in setting up security and admitting stations to the network. The design allows a station to participate in multiple AV networks. HPAV is backward compatible with HomePlug 1.0 and offers several mandatory and optional co-existence modes enabling multi-network operation, hidden node service and Broadband over Powerline (BPL) co-existence.

HPAV aims to be the network of choice for the distribution of data and multi-stream entertainment including HDTV, SDTV, and audiophile quality audio throughout the home. It is designed to provide the best connectivity at the highest QoS of the home networking technologies competing for these applications. HomePlug AV enables all devices with a power plug to have network access through HPAV. HPAV was designed to provide this capability at a cost that is competitive with other competing technologies.

A glossary at the end of the paper defines the acronyms used in the paper.

System Architecture

Figure 1 shows an architectural diagram of the HPAV system. The Higher Layer Entities (HLEs) above the H1 (Host) Interface may be bridges, applications or servers that provide off-chip services to clients below the H1 Interface. The Data Service Access Point (SAP) accepts Ethernet format packets, so all IP based protocols are easily handled.

The Architecture defines two planes as shown in Figure 1. The data plane provides the traditional layered approach with the M1 interface between the Convergence Layer (CL) and the MAC, and the PHY interface between the MAC and the PHY. In the control plane, the MAC is a monolith without conventional layering. In Figure 1 it is labeled as the Connection Manager (CM) since that is its primary function. The approach adopted for the control plane was chosen to provide more efficient processing and to provide implementers greater flexibility for innovation. Although part of the control plane in all stations, the Central Coordinator (CCo) entity will be active in one and only one station in a single HPAV network.

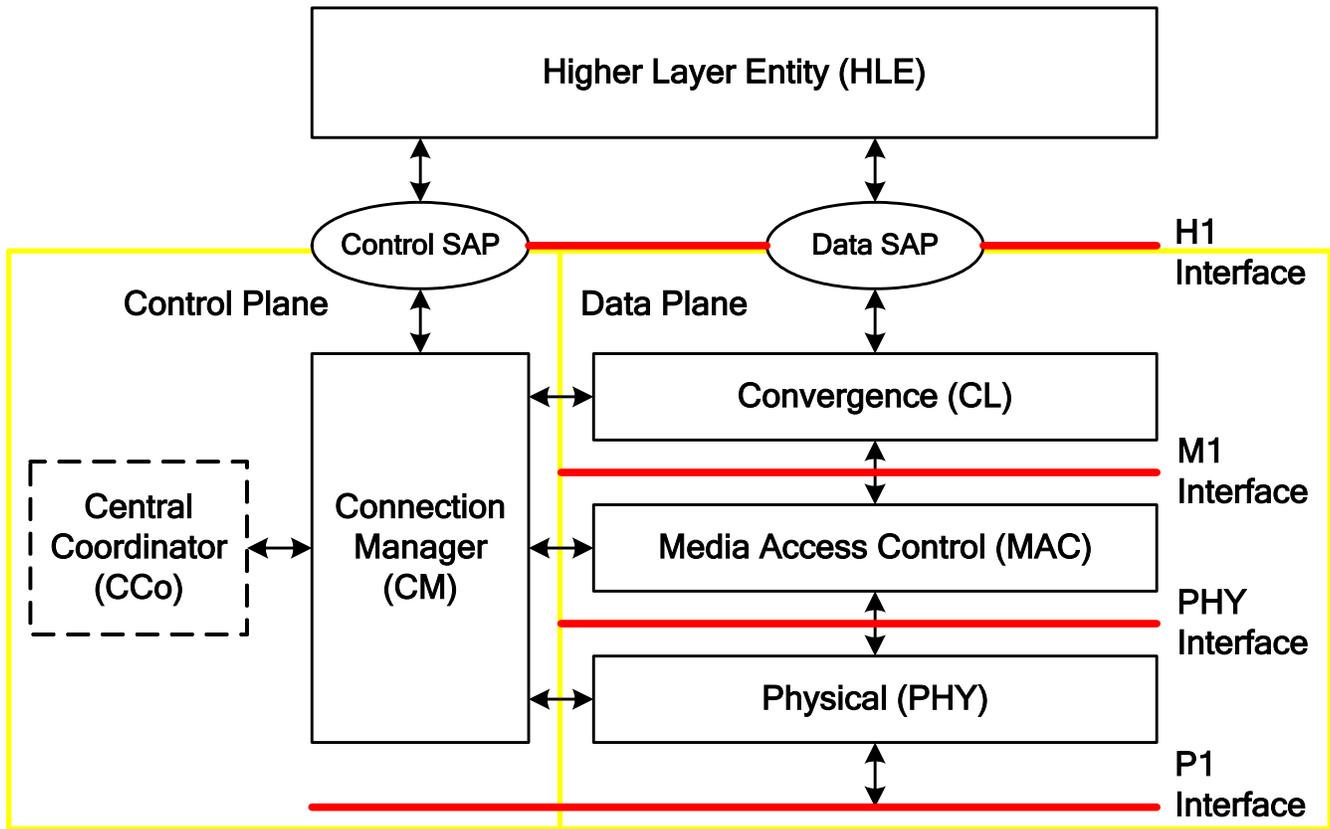


Figure 1 HPAV Architecture

Physical (PHY) Layer

The Physical Layer (PHY) operates in the frequency range of 2 - 28 MHz and provides a 200 Mbps PHY channel rate and a 150 Mbps information rate. It uses windowed OFDM and a powerful Turbo Convolutional Code (TCC), which provides robust performance within 0.5 dB of Shannon Capacity. Windowed OFDM provides flexible spectrum notching capability where the notches can exceed 30 dB in depth without losing significant useful spectrum outside of the notch. Long OFDM symbols with 917 usable carriers (tones) are used in conjunction with a flexible guard interval. Modulation densities from BPSK (which carries 1 bit of information per carrier per symbol) to 1024 QAM (which carries 10 bits of information per carrier per symbol) are independently applied to each carrier based on the channel characteristics between the transmitter and the receiver

Figure 2 shows a block diagram representation for the physical layer of a HPAV transmitter and receiver.

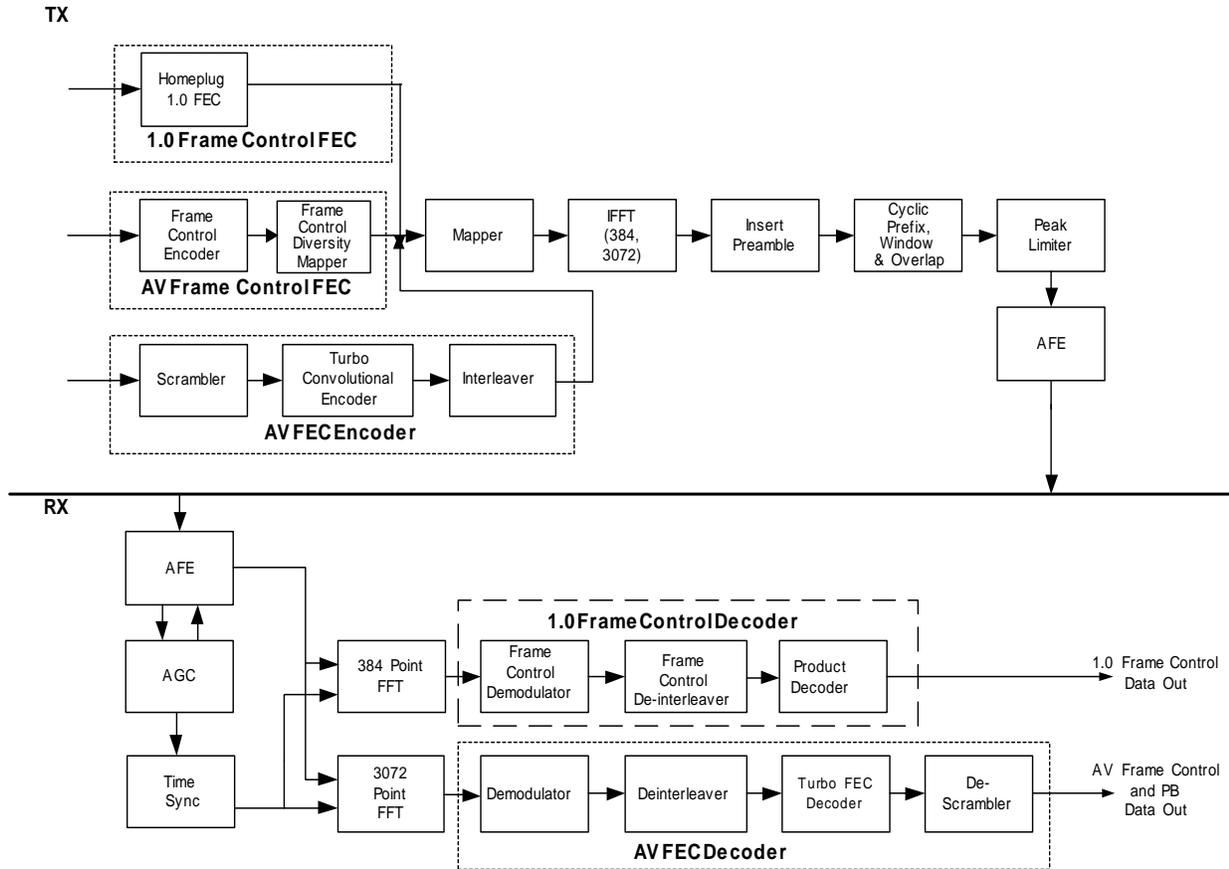


Figure 2 HPAV OFDM Transceiver

On the transmitter side, the PHY layer receives its inputs from the Medium Access Control (MAC) layer. There are separate inputs for HPAV data, HPAV control information, and HomePlug 1.0 control information (the latter in order to support HomePlug 1.0 compatibility). HPAV control information is processed by the Frame Control Encoder block, which has an embedded Frame Control FEC block and Diversity Interleaver. The HPAV data stream passes through a Scrambler, a Turbo FEC Encoder and an Interleaver. The outputs of the three streams lead into a common OFDM Modulation structure, consisting of a Mapper, an IFFT processor, Preamble and Cyclic prefix insertion and a Peak Limiter. This output eventually feeds the Analog Front End (AFE) module which couples the signal to the Powerline medium.

At the receiver, an AFE operates in conjunction with an Automatic Gain Controller (AGC) and a time synchronization module to feed separate data information and data recovery circuits. The HPAV Frame Control is recovered by processing the received stream through a 3072-point FFT, a Frame Control Demodulator and a Frame Control Decoder. The HomePlug 1.0 Frame Control, if present, is recovered by a 384-point FFT. In parallel, the data stream is retrieved after processing through a 3072-point FFT for HPAV, a demodulator with SNR estimation, a De-mapper, De-interleaver, Turbo FEC decoder, and a De-scrambler for HPAV data.

The HPAV PHY provides for the implementation of flexible spectrum policy mechanisms to allow for adaptation in varying geographic, network and regulatory environments. Frequency notches can be applied easily and dynamically, even in deployed devices. Region-specific keep-out regions can be set under software control. The ability to make soft changes to alter the device's tone mask (enabled tones) allows for implementations that can dynamically adapt their keep-out regions.

MAC Protocols/Services

HPAV provides connection-oriented Contention Free (CF) service to support the QoS requirements (guaranteed bandwidth, latency and jitter requirements) of demanding AV and IP applications. This Contention Free service is based on periodic Time Division Multiple Access (TDMA) allocations of adequate duration to support the QoS requirements of a connection.

HPAV also provides a connectionless, prioritized Contention based service to support both best-effort applications and applications that rely on prioritized QoS. This service is based on Collision Sense Multiple Access/Collision Avoidance (CSMA/CA) technology which is applied to only traffic at the highest pending priority level after the pending traffic with lower priority levels has been eliminated during a brief Priority Resolution phase at the beginning of the contention window.

To efficiently provide both kinds of communication service, HPAV implements a flexible, centrally-managed architecture. The central manager is called a Central Coordinator (CCo). The CCo establishes a Beacon Period and a schedule which accommodates both the Contention Free allocations and the time allotted for Contention-based traffic. As shown in Figure 3, the Beacon Period is divided into 3 regions:

- Beacon Region
- CSMA Region
- Contention-Free Region

The CCo broadcasts a beacon at the beginning of each Beacon Period; it uses the beacon to communicate the scheduling within the beacon period. The beacons are extremely robust and reliable. The schedules advertised in the Beacon are persistent—i.e., the CCo promises not to change the schedule for a number of Beacon Periods—and the persistence is also advertised in the beacon so that the transmitting station for a connection can confidently transmit during its persistent allocation(s) even if it has missed several beacons within the advertised persistence of the schedule. This provides additional continuity even if a few beacons are missed. The CSMA periods are also persistent so that stations wishing to send CSMA traffic can do so even if they miss a few beacons.

The MAC layer provides both Contention (CSMA) and Contention Free (CF) services through the respective regions in the Beacon Period. The CCo-managed Persistent Contention Free (PCF) Region enables HPAV to provide a strict guarantee on Higher Layer Entity (HLE) QoS requirements. An HLE uses the Connection Specification (CSPEC) to specify its QoS requirements. The Connection Manager (CM) in the station evaluates the CSPEC and, if appropriate, communicates the pertinent requirements to the CCo and asks the CCo for a suitable Contention Free allocation. QoS features specified in the CSPEC include:

- Guaranteed bandwidth
- Quasi-Error free service
- Fixed Latency
- Jitter control

If the CCo is able to accommodate the connection request, it will ask the stations to “sound” the channel. This allows the stations to perform the initial channel estimation (i.e., establish a Tone Map specifying the optimal modulation on each OFDM tone). The Tone Map is communicated from the receiver to the transmitter; the channel estimation is also communicated in abbreviated form to the CCo to help it determine how much time should be allocated to the connection. Based on the CSPEC and the channel sounding results, the CCo provides one or more persistent time allocations—Transmit Opportunities (TXOPs)—for the connection within the PCF Region.

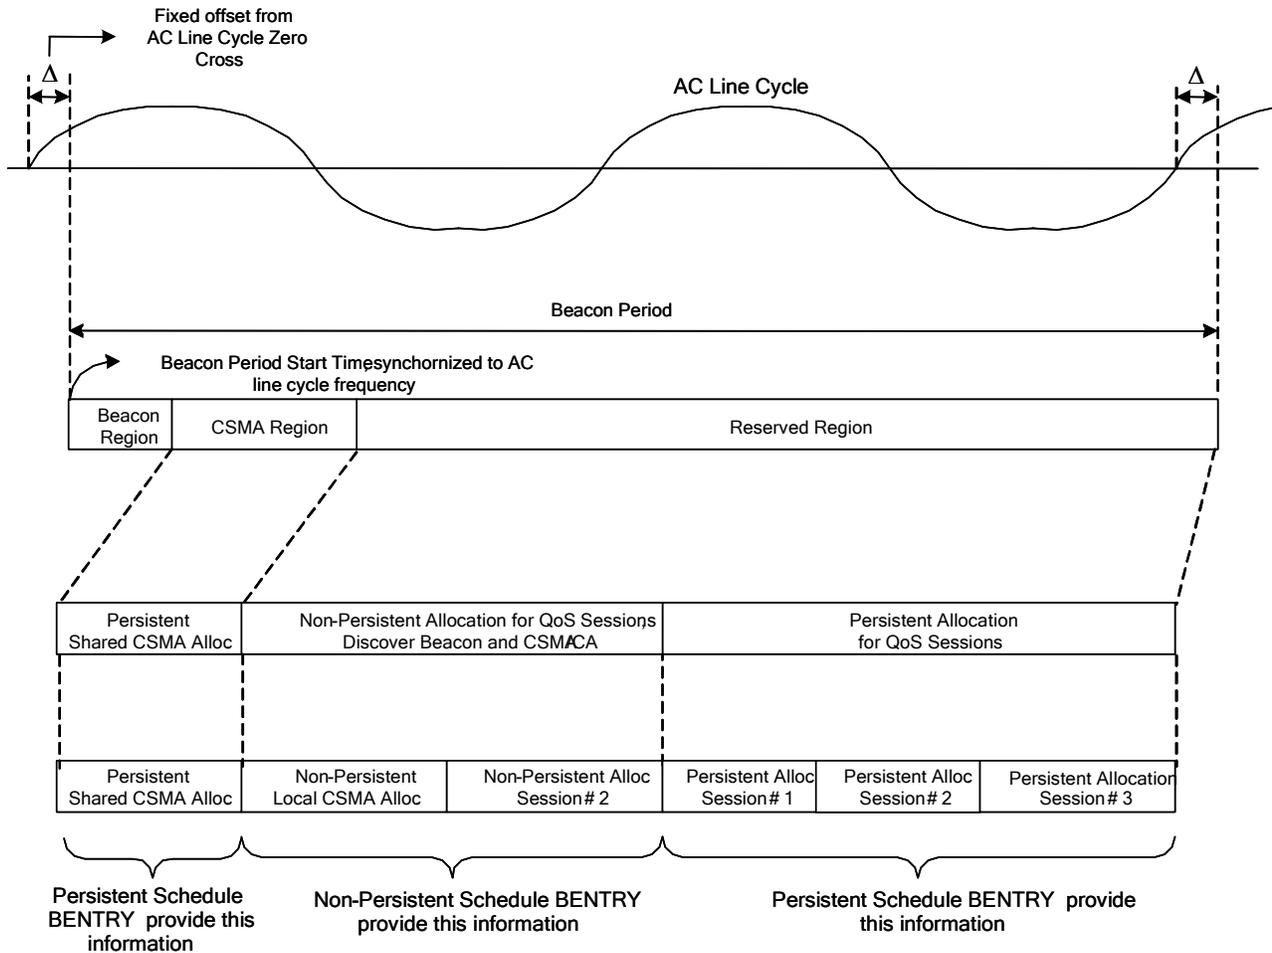
The PCF Region also contains time for non-persistent allocations good only in the current beacon period. These non-persistent allocations are used to provide additional short term bandwidth to connections that require it (e.g., because of transient errors or changing channel conditions) to meet their QoS requirements, providing that the transmitting station hears the beacon at the beginning of the Beacon Period. When this time is not used for non-persistent CF allocations, it may be used for CSMA traffic. Again, stations must hear the beacon in order to know whether the time is available for CSMA traffic.

Messaging in HPAV is direct from station to station; however, the CCo monitors the messages. The header of each message contains information about how much data is pending for transmission on the connection; if this amount becomes large on a given connection, the CCo may allocate additional non-persistent time to the connection in the PCF Region.

The Persistent CSMA Region provides prioritized contention-based communication. It is used where there is no CSPEC and/or the traffic is of short duration. When operating in 1.0 Coexistence mode, or “Hybrid Mode”, AV coordinates with HomePlug 1.0 devices and permits them to communicate during the CSMA period.

As shown in Figure 3, the Beacon Period is synchronized to the AC line cycle. By synchronizing to the line cycle, HPAV provides stability of the periodic allocations relative to the line cycle. This, in turn, provides better channel adaptation to the synchronous (to the line cycle) interference, resulting in improved throughput. The beacon provides announcements of where the beacon will occur over the next few beacon periods—i.e., beacon persistence—to enable continued communications by stations that miss an occasional beacon.

Figure 3 Example of Beacon Period Structure



MAC Control Plane

The Medium Access Control (MAC) Layer contains an integrated Connection Manager (CM). HLEs provide a Connection Specification (CSPEC) that details QoS requirements for application data. For bridged traffic, CSPECs may be generated dynamically by the Auto Connection Service (ACS) or by a higher layer QoS Manager that coordinates QoS over multiple network segments; otherwise the traffic is transmitted as prioritized CSMA traffic.

The Control Plane provides a seamless interface to the application layer. Application requirements are received at the H1 Control SAP in the CSPEC and are interpreted by the CM. The CM is responsible for evaluating the CSPEC and setting up the appropriate connection in conjunction with the CM in the station at the other end of the connection and with the CCo. It is the Connection Manager's responsibility to ensure that the appropriate AV mechanisms are engaged in order to provide the application with the bandwidth it requires. It must also monitor the level of service that the connection is receiving and take remedial action if the guaranteed QoS is not being provided.

The MAC also maintains a clock that is tightly synchronized to the CCo's clock (the CCo includes a timestamp in the beacon). This means that the entire HPAV network shares a common network clock for use by HLEs that have tight timing constraints (e.g., to synchronize surround sound speakers).

MAC Data Plane

In the Data Plane, the MAC accepts MSDUs (e.g., Ethernet packets) arriving from the Convergence Layer and encapsulates them with a header, optional Arrival Time Stamp (ATS) and Check Sum to create a MAC Frame. The MAC Frames are then enqueued into the appropriate MAC Frame Stream. It is the MAC's responsibility to ensure that the MSDUs related to a given connection are delivered to the PHY in a timely fashion for transmission during the time allocated for the connection. For this purpose, it maintains individual queues for each connection's data, for each priority level of CSMA traffic and for each priority level of Control Messages.

Each MAC frame stream is divided into 512 octet segments each of which is encrypted and encapsulated into a serialized PHY Block (PB). As shown in Figure 4, the PBs are packed into an MPDU which is delivered to the PHY. The PHY transmitter applies forward error correction and places the resulting PPDU onto the powerline as described in the PHY section above.

As the receiver reconstructs the MSDUs, it selectively acknowledges the PBs; those that are not positively acknowledged are retransmitted during the next TXOP. The Selective Acknowledge (SACK) is an integral part of the TDMA allocation. When all the PBs composing an MSDU have been received correctly, the segments are decrypted and the resulting MSDU is passed to the Convergence Layer for delivery to the appropriate HLE.

Control messages are processed in an analogous fashion.

Since FEC and Selective Acknowledgment (SACK) are performed on relatively small blocks of data, the FEC is more robust and retransmissions are minimized. These two features contribute to HPAV's ability to operate at near channel capacity.

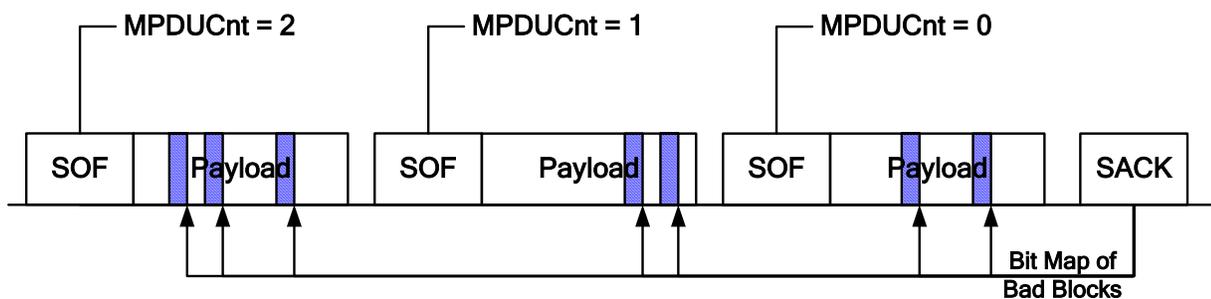


Figure 4 MAC Segmentation and MPDU Generation

Central Coordinator (CCo)

Each CCo controls an AV Logical Network (AVLN) which consists of several AV stations which all share a common Network Membership Key (NMK). The NMK and other details of security and encryption are discussed below. For now, it is sufficient to know that the NMK provides exclusive access to an AVLN so that the member stations can communicate in a private and secure environment.

As described above, the CCo provides bandwidth management services for the AVLN. These include admission control (determining whether to admit a new connection when it is requested). If the connection is admitted, the CCo schedules time allocations for the connection in the PCF Region. It manages this schedule via the beacon, which contains:

- the current schedule and the minimum number of Beacon Periods for which it will remain valid, and/or
- the new schedule and the number of Beacon Periods which will pass before it becomes valid.

When an AV Station is powered on, it listens to the medium. If it hears an existing AVLN, it will attempt to join it. If it does not hear an existing AVLN, it will form its own network by becoming a CCo and broadcasting a beacon. Eventually, another AV Station will be powered up and the two will associate and form an AVLN. (This is a highly simplified description; in reality, the HPAV specification provides for various cases of encountering more than a single AVLN, encountering HomePlug 1.0 devices, encountering other Powerline Networks and combinations thereof.)

The CCo attempts to learn the topology of its AVLN and of any neighboring AVLNs. To achieve this, each AV station broadcasts a Discover Beacon periodically (at a time allocated by the CCo in the non-persistent portion of the PCF Region). This Discover Beacon contains information about the station and the AVLN to which it belongs.

Each station that hears the Discover Beacon adds the information it contains to a Discovered Station List (DSL). While building its DSL, if the station encounters a Discover Beacon from a station in a different AVLN, it adds the information about the other network to a Discovered Networks List (DNL). Periodically the CCo asks each station for its DSL and DNL and use the collected lists to compose a topology map.

The CCo uses the topology map it builds from the collected DSLs and DNLs to determine if there is another station in the AVLN that would make a better CCo than it. The criteria for making this decision, in order of priority, are:

1. User's Selection
2. CCo Capability
3. Number of discovered STAs in the Discovered Station List
4. Number of discovered AVLNs in the Discovered Network List

If the current CCo finds another station that would make a better CCo, it will negotiate a handover of CCo functions to the new CCo. Depending upon the capabilities of the old and new CCo's, the handover may or may not result in existing connections being torn down.

The CCo may also select another CCo-capable station to be its backup in case of failure. If the station accepts the backup role, it will monitor the AVLN and, if the CCo's beacon is not heard by any stations in the AVLN for a specified number of Beacon Periods, the backup CCo will assume the role of CCo and attempt to maintain the existing connections without disruption.

Since a station must be capable of communicating with the CCo in order to join an AVLN and establish connections, a proxy capability is provided to support stations that are hidden from (i.e., unable to communicate with) the CCo. This capability provides for the creation of a Proxy Coordinator (PCo) to repeat the beacon information in Proxy Beacons and to relay control messages between the hidden station and the CCo. Note that only control messages are relayed. The station must be able to communicate directly with any stations with which it wishes to establish a connection. The PCo also transmits a Proxy Beacon during each Beacon Period to convey scheduling and other information to the hidden station.

When all stations are idle, the CCo causes the AVLN to enter a power saving mode. In this mode, there is only a small CSMA Region (so stations can initiate communication) and a small PCF Region (just long enough for Discover and Proxy beacons). Stations must have their receivers on during these small regions to participate in the AVLN; they may turn their transmitters and receivers off for the remainder of the Beacon Period. This makes it easier for stations to qualify for Energy Star certification.

Convergence Layer

The Convergence Layer (CL) serves as the interface between the HLEs and the MAC in the Data Plane. It accepts data payloads through Service Access Points (SAPs) at the H1 Interface and processes them as needed prior to handing them off to the MAC through the M1 interface. The only Data SAP specified by AV is the Ethernet II-class stack. This stack supports packet formats as specified by IEEE 802.3 with or without IEEE 802.2 (LLC), IEEE 802.1H (SNAP) extensions, and/or VLAN tagging. Using the Ethernet format makes it easy for AVLNs to interface to other LANs.

Among the services the CL provides on the transmit side are classification and auto connection. If requested for a connection, the CL will also associate an Arrival Time Stamp (ATS) with the data payload. On the receive side, the CL provides (optional) smoothing and insures that the received MSDUs are delivered to the appropriate H1 Service Access Point (SAP). On both sides, it provides the Connection Manager sufficient information to monitor the level of QoS being provided by the connection.

When a connection is established, the CM provides the classifier with a set of rules that will enable the classifier to uniquely associate incoming packets with the connection. For example, a set of rules might specify the source and destination MAC addresses and the TCP source and destination ports of the connection.

The classifier examines each packet received at the H1 interface and attempts to match it with a connection using the classification rules that have been provided to it. If it finds a match it will label the packet with the Connection ID (CID) of the appropriate connection, otherwise the classifier will release the packet for transmission in the CSMA region at the appropriate priority level.

If the transmitting station supports the optional Auto Connect Service (ACS), all packets which are released by the classifier without being associated with a connection will be examined by the ACS which will evaluate the data flow(s) between a given source and destination and attempt to identify flows which are worthy of a connection. This evaluation and identification may be based on a mix of the following:

- Policies established by an HLE (or a manufacturer),
- Templates such as traffic associated with ports known to have a particular usage,
- Heuristics such as the volume and regularity of data which is being transmitted.

Until the ACS identifies a connection, it releases the packets for transmission in the CSMA period immediately upon completion of the packet's inspection.

If the ACS identifies a particular data flow as connection worthy, it behaves in a manner analogous to an HLE and asks the CM to set up a connection, providing classifier rules, etc. When the CM establishes the connection, the Classifier will start associating the packets with the newly established connection and the ACS will no longer see them. The ACS is, however, responsible for servicing the connection in the same way that an HLE would.

At the receiving station, the CL demuxes the received packets. If the packets are associated with a connection for which smoothing (a.k.a. de-jittering) has been requested, the CL will buffer the packets for the appropriate time so that all packets

are released to the HLE at a fixed interval after they arrived at the H1 interface at the transmitter, which the receiver knows from the ATS it received with the packet and the synchronized network clock.

On both ends of a connection, the CLs provide sufficient information to the CM that it can monitor the level of QoS being provided to a connection to insure that the guarantees are being met. The CM will take corrective measures specified by the CSPEC if there are any violations to the QoS guarantees.

HPAV Security

Admission control procedures ensure that only permitted devices are allowed into the AVLN. A station's ability to maintain multiple security keys allows it to participate in multiple AVLNs.

All data traffic and nearly all control traffic within the AVLN—the exception being a strictly limited set of control messages that simply cannot be encrypted—is secured by 128-bit AES encryption, providing a high level of security. This encryption uses the Network Encryption Key (NEK) and is performed on individual segments as the MPDUs are created. The NEK may be automatically and dynamically changed.

In order to join an AVLN, a station must obtain a Network Membership Key (NMK). If it already possesses an NMK it may join the network immediately; otherwise it must be provided with the NMK. This provisioning may occur in a variety of ways, including:

- Using the default NMK that is programmed into all AV stations. While this default NMK provides a seamless, plug and play experience for the user when the equipment is initially installed, it does not provide any privacy since it is known by every HPAV-certified station.
- The user can define and enter a Network Password (NPW) directly into a new station. This NPW is hashed to create the NMK, a 128 bit AES encryption key. The user must enter a NPW on at least one station to initially define the NMK for the AVLN.
- All AV stations are also programmed with a unique Device Access Key (DAK). The user may enter this key into any suitably programmed station already in the AVLN and that station will use the DAK to encrypt the NMK and broadcast it. Since only the new station has the DAK, it will be the only station that is capable of decrypting the broadcast message and so it and only it will receive the new NMK.
- Using asymmetric Public/Private Key encryption, the AV stations may provide the user the ability to join the new station to the AVLN without needing to remember or enter passwords. This may be as simple as having the user press a button or make a menu selection on the new station and on a station already in the AVLN.

When a station has the correct NMK and actually joins the AVLN, it will be given the current Network Encryption Key (NEK) which is used to encrypt data during segmentation in the MAC.

The design also permits encryption key management by higher layer Security and Authentication Standards such as 802.1x and EAP.

Multiple Networks

AV incorporates mechanisms to provide for the coordination of Neighboring Networks (NNs). Once detected, neighboring CCo's can cooperatively schedule transmissions in their own networks without causing interference in the other. In the case of multiple AVLNs, each CCo maintains an Interfering Network List (INL). The INL identifies all the nearby CCo's that the CCo can hear. Each CCo then communicates its INL to the other CCo's. HPAV requires that a CCo must recognize all of the INLs that it is aware of and not interfere with those networks.

When the CCo in an AVLN discovers another CCo, it attempts to coordinate with the discovered CCo. It negotiates for a beacon slot within the beacon region—which may need to be expanded to provide room for another beacon slot—and for any time it requires for its PCF Region. When the negotiations are completed, both CCo's synchronize a schedule change which results in each AVLN having a space for its own PCF Region, which is identified by its neighbors as a Stayout Region in their beacons. All of the NNs share a common CSMA region. An example is shown in Figure 5.

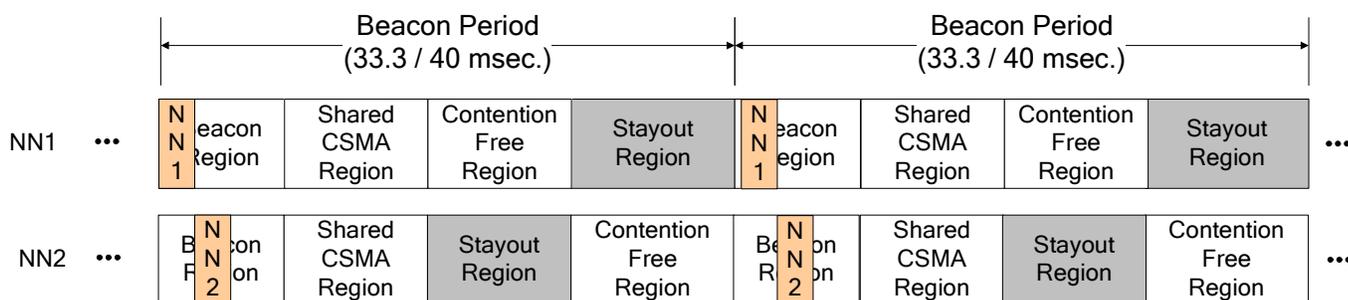


Figure 5 Neighbor Network Coordination

When multiple AVLNs are coexisting, each has a bandwidth quota, i.e., a portion of the Beacon Period that it uses for its PCF Region. The quota is defined by policies which are suitable for the regulatory region in which the AVLN is operating. The default quota is equal shares.

After allowances are made for the Beacon Region and for a minimal Shared CSMA region, a CCo may allocate as much of the remaining bandwidth as is available to service its connections. It may even exceed its quota if the bandwidth is available. If an AVLN is using more than its quota—which depends upon the number of NNs—it must relinquish bandwidth down to its quota if requested by its neighbors, even if it has to reconfigure (squeeze) or release connections in order to fit within its quota. It is not required to drop below its quota or to release bandwidth to another NN if the release would cause that NN to exceed its quota.

Coexistence

HomePlug 1.0 Coexistence

The AV PHY enables coexistence and interoperability with HomePlug 1.0 devices. The specification requires coexistence but interoperability is optional. Coexistence means that AV devices are capable of the low-level communications with 1.0 devices needed to share the medium, but not necessarily the ability to communicate payload data. An optionally interoperable device has the ability to communicate payload data with 1.0 devices.

Coexistence is achieved with the use of preambles that all devices (AV and 1.0) can use for synchronization, as well as the addition of 1.0 Frame Controls. The additional Frame Controls and 1.0 coexistence mechanisms are only activated when one or more 1.0 devices are detected.

BPL Coexistence

HPLV employs BPL Coexistence through one of two methods: Coexistence of Services, and Coexistence of Technologies. Coexistence of Services provides an efficient, integrated extension of services, while Coexistence of Technologies allows simultaneous use of the powerline by differing technologies.

The Coexistence of Services method uses TDM with beacon signaling and messaging to coordinate the in-home and BPL networks. Network coordination allows flexible time allocation and re-use. Flexible allocation provides increased efficiency and throughput for both networks by allowing either network to utilize unused time in the other network. In addition, both networks can communicate allowing service-level integration from providers to in-home devices

The Coexistence of Technologies method uses FDM to allow differing technologies to coexist. While allowing unique technologies to share the powerline, it lacks the ability to share unused bandwidth with the other network(s).

Conclusion

In this paper, an overview of HomePlug AV has been presented. An overview of the architecture and some details of each of the functional blocks have been presented. In order to get complete details and access the specification, any company may join the HomePlug Alliance. Instructions on how to do this can be found at <http://www.homeplug.org/en/join/index.asp>.

Glossary

Acronym	Meaning
ACS	Auto Connection Service
AES	Advanced Encryption Standard
AFE	Analog Front End
AGC	Automatic Gain Controller
ATS	Arrival Time Stamp
AVLN	HomePlug AV Logical Network
BPL	Broadband over Powerline
BPSK	Binary Phase Shift Keying
CCo	Central Coordinator
CF	Contention Free
CID	Connection ID
CL	Convergence Layer
CM	Connection Manager
CSMA/CA	Collision Sense Multiple Access/Collision Avoidance
CSPEC	Connection Specification
DAK	Device Access Key
DNL	Discovered Networks List
DSL	Discovered Station List
EAP	Extensible Authentication Protocol
FDM	Frequency Division Multiplexing
FEC	Forward Error Control
FFT	Fast Fourier Transform
HDTV	High Definition Television
HLE	Higher Layer Entity
HPAV	HomePlug AV
IFFT	Inverse Fast Fourier Transform
INL	Interfering Network List

Acronym	Meaning
MAC	Medium Access Control
MPDU	MAC Protocol Data Unit
MSDU	MAC Service Data Unit
NEK	Network Encryption Key
NMK	Network Membership Key
NN	Neighboring Network
NPW	Network Password
OFDM	Orthogonal Frequency Division Multiplexing
PB	PHY Block
PCF	Persistent Contention Free
PCo	Proxy Coordinator
PHY	Physical Layer
PPDU	PHY Protocol Data Unit
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
SACK	Selective Acknowledge
SAP	Service Access Point
SDTV	Standard Definition Television
SNR	Signal-to-Noise Ratio
SOF	Start of Frame
STA	Station
TCC	Turbo Convolutional Code
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TXOP	Transmit Opportunity
VLAN	Virtual LAN