ENVISIONING • EMPOWERING • EXCELLING

G.hn Home Networking

White Paper

Revision History

Document Number	Release Date	Change Description
002WPR00	06/26/17	Initial release.

Table of Contents

List of Tables

List of Figures

Introduction

This document gives an insight on the G.hn Home Networking technology, the market it is destined to, the performance reached, the standard compliance aspect, and the MaxLinear's G.hn offering.

G.hn Overview

G.hn is the globally recognized high speed technology based on *ITU-T standards* pioneering important changes for the broadband connectivity, industry deployments and users experience.

G.hn provides connectivity with excellent quality of service over any type of wire (powerline, coaxial, phone line, and Plastic Optical Fiber (POF)), creating a high performance backbone throughout the home that enables seamless multi-media distribution, without the cost and disruption of having to install any new wiring.

MaxLinear's G.hn chipsets implement this standard in an optimized way enabling high speed home connectivity whether embedded in existing devices such as broadband gateways and routers, HD Smart TVs and OTT Set Top Boxes, and any consumer electronic products or in stand-alone devices that require Plug & Play networking capabilities at Gigabit speeds.

G.hn offers also the flexibility of providing high bandwidth for Home Networking using both a wired medium and in combination with wireless, multiple wireless access points with G.hn as backbone.

The G.hn family of standards was defined and approved by the Telecommunication Standardization Sector of the International Telecommunication Union (ITU).

The family is basically comprised of the following set of recommendations:

- *G.9960. System Architecture and Physical Layer specification.* See page 6 and 9.
- *G.9961. Data Link Layer specification*. See page 13.
- *G.9963. Multiple Input – Multiple Output (MIMO) specification*. See page 34.
- *G.9964. Power Spectral Density (PSD) specification*. See page 36.
- *G.9962. Management specification.* See page 38.
- *G.9980. TR-069 Remote Management of CPE over Broadband Networks*. See page 38.
- *G.9972. Coexistence Mechanism for Wireline Home Networking Transceivers (G.hn powerline Coexistence with IEEE 1901 Systems).* See page 39.

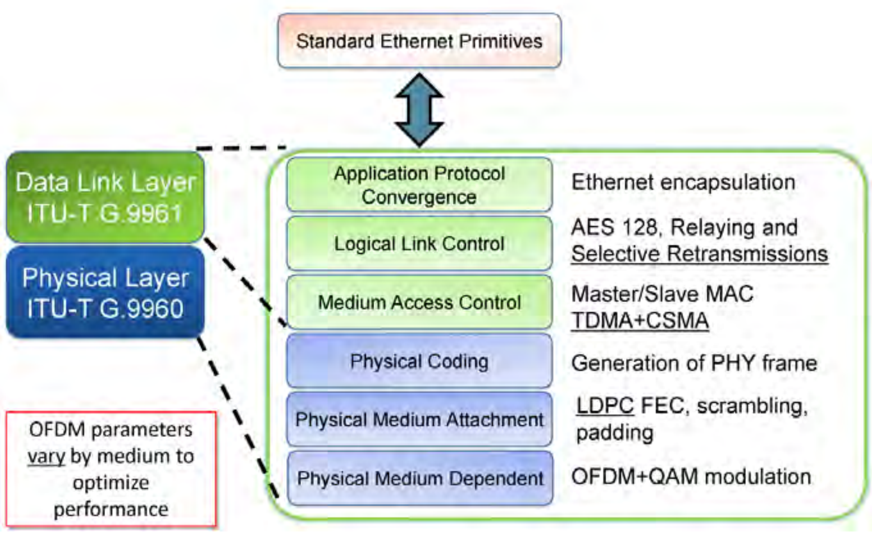

Figure 1: G.hn System Architecture Overview

Key Features of ITU-T G.hn Technology

- Single technology for all wires = powerline, coax, phone line, and Plastic Optical Fiber (POF)
- Multi-source silicon providers ecosystem
- Most optimum performance per medium:
  - Support for ITU-T G.hn baseband band plans 25 (PLC), 50MHz, 100MHz and 200MHz (coax and phone line).
  - Specific OFDM parameter set per medium to maximize throughput while minimizing latency
  - Up to 1500Mbps PHY rate over powerline and up to 2Gbps PHY rate over coax and phone.
  - MIMO techniques for powerline (based on G.9963) boosting PLC throughput.
  - Robust Communication Mode (RCM) supported for high noise environments.
- Support for multiple network access MAC schemes (TDMA, CSMA, token-passing) to accommodate application-specific requirements.
- Reliability and robustness
  - State-of-the-art LDPC Forward Error Correction
  - Enhanced ACK-based Selective Retransmission of PHY blocks to maximize impulsive noise mitigation.
  - Synchronization with the AC cycle for PLC medium.
  - Relaying function and auto-mesh networking for increased coverage and network resilience.
- Power management
  - Short term sleep = G.hn MAC enables power saving at MAC level thanks to schedule inactivity slots during a given MAC cycle (efficient-power mode L1).
  - Long term sleep = G.hn MAC enables long (up to 5 minutes) scheduled inactivity (low-power mode L2 and idle mode L3).
- Neighboring networks (Multi-network/neighboring domain interference mitigation) = G.hn standard mechanism achieves maximum aggregated performance in multi-dwelling units.
- Security
  - AES 128 bits encryption algorithm using CCMP protocol to ensure confidentiality and message integrity. Authentication and key exchange following *ITU-T Recommendation X.1035* enabling End-to-End encryption even when packets are relayed.
- QoS
  - Prioritized QoS with eight priority levels by selectable packet fields
  - “Per-stream” parameterized QoS
  - Admission control, congestion control and bandwidth limitation
- Power mask and notching allowing coexistence with radio systems

Applications

G.hn provides the flexibility and robustness to address many different scenarios and applications:

- **In-Home Networking** = Consumers are fueling the need for greater home network bandwidth requirements such as 4K HD TV.

  The consumer’s desire to seamlessly share applications across a growing number of digital devices makes In-Home Networking critical, overall in terms of supporting high bandwidth and reliable service.

  Wireline networking consists on reusing existing powerlines, coaxial cables, phone line, and POF to support this application. It is one of the easiest and most reliable ways to meet this market demand, offering a comprehensive solution that optimizes networking.

  Wireline can also help extend the Wireless networks by traversing the bandwidth requirements over existing consumer wiring without diminishing the integrity of the speeds needed to extend the Wireless foot print.

  The advantages of Wireline’s plug and play ease of use combined with Wireless’ ability to roam freely within a home make the next generation digital experience essential for having the ability to handle high bandwidth in an ubiquitous way.

- **Ultra Broadband** = MaxLinear’s G.hn extends broadband access from the fiber network termination in Multi Dwelling Units (MDU) or Distribution Points (DP) to users’ homes by taking advantage of the existing wiring. It offers significant end-to-end cost savings to operators around the world.

  Based on the ITU-T compliant MaxLinear’s G.hn chipsets, Time Division Duplexing (TDD) architecture and increased spectrum up to 200MHz OFDM, MaxLinear’s Ultra Broadband G.now platforms over phone line and Point to Multipoint (P2MP) solutions over coaxial cables are designed to support broadband access distribution for MDU topologies.

  MaxLinear G.now and P2MP solutions complement Fiber-To-The-Distribution-Point (FTTdp) deployments to transfer all the capacity of optical backbones directly to each home reusing cables.

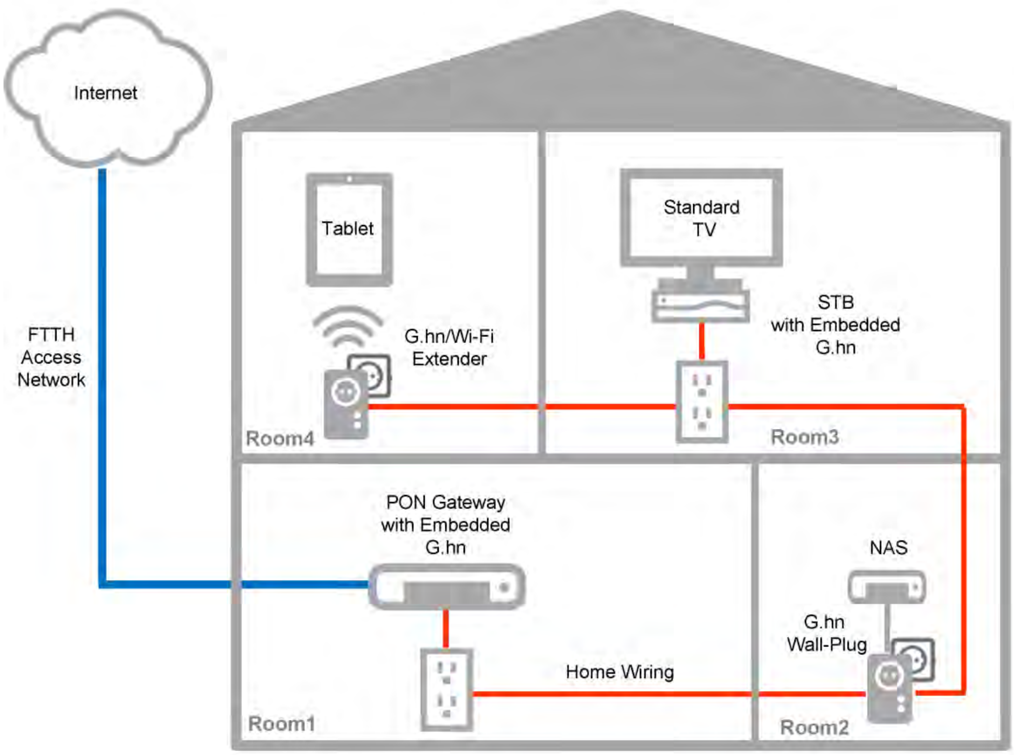

Figure 2: G.hn System Architecture Overview

G.hn System Architecture (ITU-G.9960)

A G.hn network is built upon one or more domains. A G.hn domain contains nodes connected to the same medium that as well belong to the same home network, and one of these nodes work as Domain Master (DM), and rest of nodes work as End Point (EP).

The DM is responsible for controlling operation of all nodes in the domain, including admission to the domain, bandwidth reservation, node resignation, and other domain management operations.

The DM avoids interference between nodes in a domain by coordinating their transmission time. This is simpler and more efficient than coordinating transmissions across several domains sharing the medium in the same home. However, it is still necessary when the medium is shared between neighboring networks, such as in many deployments over powerlines in high density population areas such as in Multi-Dwelling Unit (MDU) buildings. The user might have the ability to also establish multiple domains on the same medium, that is, by enabling different RF channels on coax.

In the DM failure, the DM function automatically passes to another node in the domain with DM capability.

Domains of the same network are interconnected by Inter-Domain Bridges (IDB) as shown in the following figure.

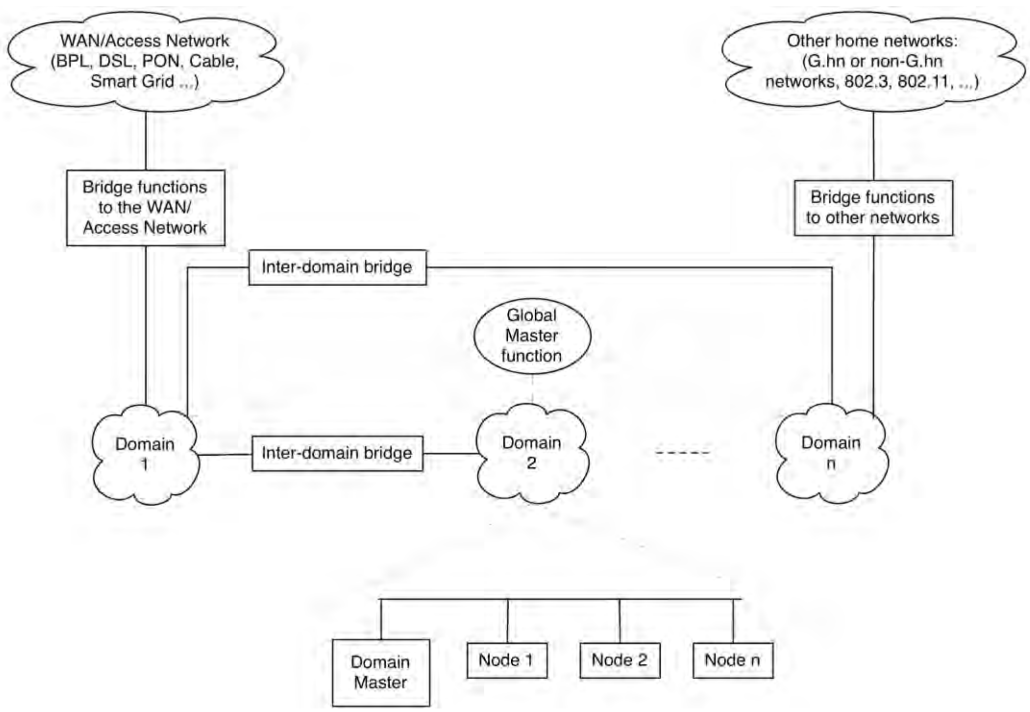

Figure 3: Home Network Architecture Reference Model

This allows nodes of any domain to communicate with any other node of any other domain in the network.

Any domain may be bridged to wireline or wireless alien networks (for example, DSL, PLC access, WLAN, other Home Networking technologies).

The network includes three domains: over coax, phone line, and powerline, each controlled by its DM. Alien networks are WLAN, USB2, Ethernet, and residential access network. A residential gateway bridges powerline and coax domains and bridges the G.hn network to alien networks.

Each G.hn node is configured to operate over the medium to which it is connected, and it can communicate directly with any other node of its own domain and, via IDB, with nodes of other domains.

Communications with nodes of alien networks, including the broadband access network, are through the residential gateway.

The following figure shows an example of residential Home Networking.

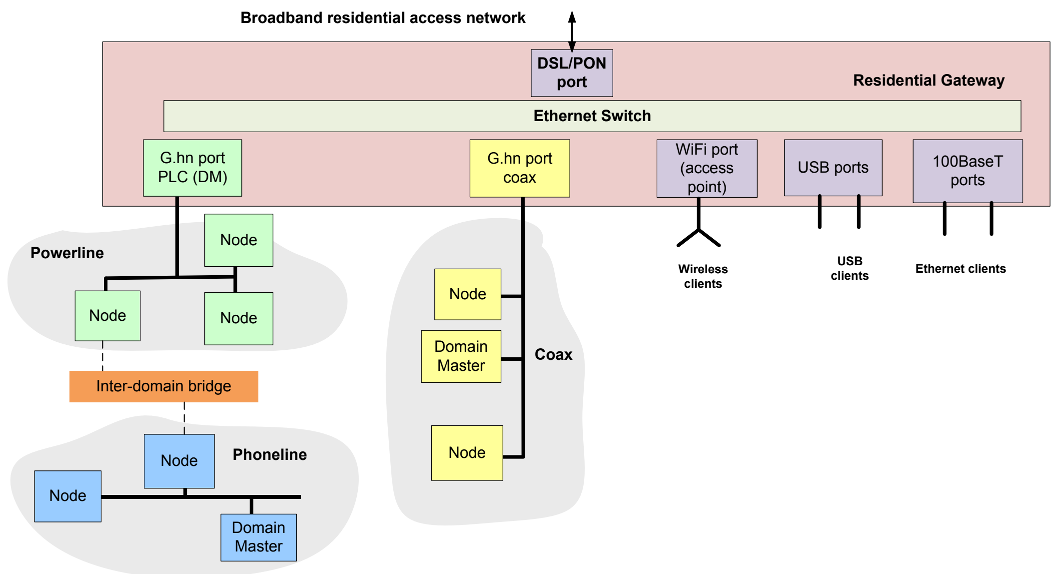

Figure 4: HN Topology Associated with Residential Access

Nodes in a domain can communicate with other nodes directly or through one dedicated relay node called the domain access point (centralized mode), or both directly and through one or more relays. Nodes that are hidden from the DM are coordinated via a DM-proxy node assigned by the DM. Mutually hidden nodes can communicate via relays.

Protocol Reference Model

The protocol reference model of a home network transceiver is shown in Figure 5.

- The Physical Layer is specified in *G.9960 recommendation*
- The Data Link Layer (DLL) in specified *G.9961 recommendation*.It includes three main reference points:
  - Application interface (A-interface) = it is user application protocol specific (for example, Ethernet). All intermediate reference points are independent of the type of medium and are defined as functional (logical) interfaces in the terms of functional flows and logical signals.
  - Physical Medium-Independent Interface (PMI) = it is both independent of medium and application. It is defined as a functional interface in terms of functional flows and logical signals.
  - Medium-Dependent Interface (MDI) = it is a physical interface defined in the terms of physical signals transmitted over a specific medium (powerline, coax, phone line) and mechanical connection to the medium.

The DLL is composed of the following functional sublayers:

- The Application Protocol Convergence sub-layer (APC) provides an interface with the Application Entity (AE), which operates with an application-specific protocol (for example, Ethernet). The APC also provides the data rate adaptation between the application entity and the G.hn transceiver.
- The Logical Link Control (LLC) sub-layer coordinates transmission of nodes in accordance with requests from the DM. It is responsible for establishing, managing, resetting and terminating all connections of the node within the domain. The LLC also facilitates QoS constraints of the flow, defined for its various connections.
- The MAC sub-layer manages access of the node to the medium using various medium access protocols.

The PHY sublayers are:

- The Physical Coding (PCS) provides data rate adaptation (data flow control) between the MAC and PHY and encapsulates transmit MPDUs into the PHY frame while it also adds PHY-related control and management overhead.
- The Physical Medium Attachment (PMA) provides encoding of PHY frame content for transmission over the medium.
- The Physical Medium Dependent (PMD) modulates and demodulates PHY frames for transmission over the medium using Orthogonal Frequency Division Modulation (OFDM). By implementation, the PMD may include medium-dependent adaptors for different media (coaxial, powerline, and phone line).

The layers above the data link layer (above the A-interface) are beyond the scope of G.9960.

	Layers and Interfaces
A	A Interface
APC	Application Protocol Convergence
LLC	Logical Link Control
MAC	Media Access Control
PMI	PHY MAC Interface
PCS	Physical Coding
PMA	Physical Medium Attachment
PMD	Physical Medium Dependent
MDI	Medium Dependent Interface

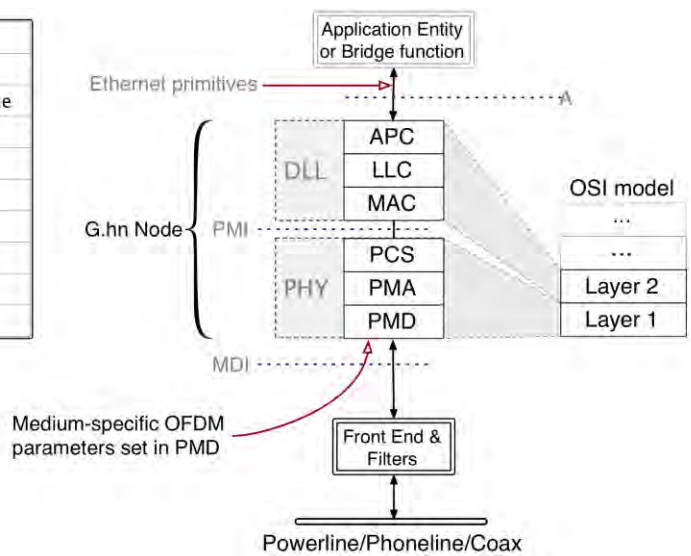

Figure 5: G.hn Functional Model

G.hn Physical Layer (ITU-G.9960)

Modulation

G.hn adopted windowed OFDM with the following programmable set of parameters to address different types of wiring:

- Number of subcarriers, N = 2n, n = 8–12
- Subcarrier spacing, Fsc = 2k × 24.4140625 KHz, k = 0 (powerline), 1 (phone line) or 3 (coax)
- Central frequency Fc
- Window size

The values of media-dependent parameters were selected taking into account the following considerations:

- Subcarrier spacing is selected based on channel characteristics
- Modulator design is significantly simpler if:
  - All values for the subcarrier spacing (Fsc) are a power-of-two multiples of a basic spacing
  - All values for the number of subcarriers (N) are a power-of-two
  - All values of sampling frequency are dividers of a common reference frequency
- Same values of subcarrier spacing and sampling frequency as used by legacy technologies simplify implementation of dual mode devices.

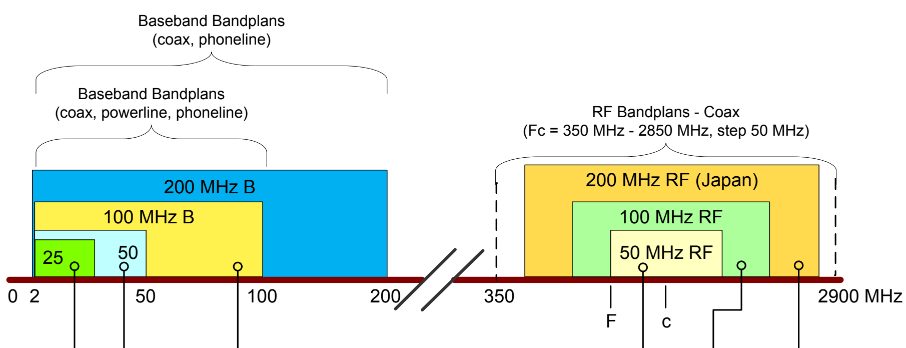

Medium	25 MHz LC	50 MHz B	100 MHz B	200 MHz B		50 MHz RF	100 MHz RF	200 MHz RF
Powerline	1024	2048	4096					
Phoneline		1024	2048	4096				
Coax		256	512	1024		256	512	1024

Figure 6: G.hn Profiles

G.hn defines baseband bandplans and RF bandplans as shown in above figure. Note that coax RF 200MHz profile is only included in the Annex for Japan.

For each particular medium and bandplan, G.hn defines only a single set of OFDM parameters so that overlapping bandplans use the same sub-carrier spacing. This rule, plus a unified per medium default preamble structure and PHY frame header, facilitates interoperability. The number of subcarriers used in each bandplan depends on the media and varies from 256 to 4096. There are also eight selectable values for the payload Cyclic Prefix (CP) length: k × N/32, k = 1, 2 …8.

G.hn defines flexible per subcarrier bit loading in the range between 1 and 12 bits. Gray-mapping is used for all constellation points of even-bit loadings and for almost all constellation points of odd-bit loadings.

The bit loading for each connection can be negotiated between the transmitter and receiver, providing sufficient flexibility to adopt channels with wide ranges of frequency responses and noise PSDs.

Functional Model of the PHY Layer

The functional model of the PHY is presented in the following figure.

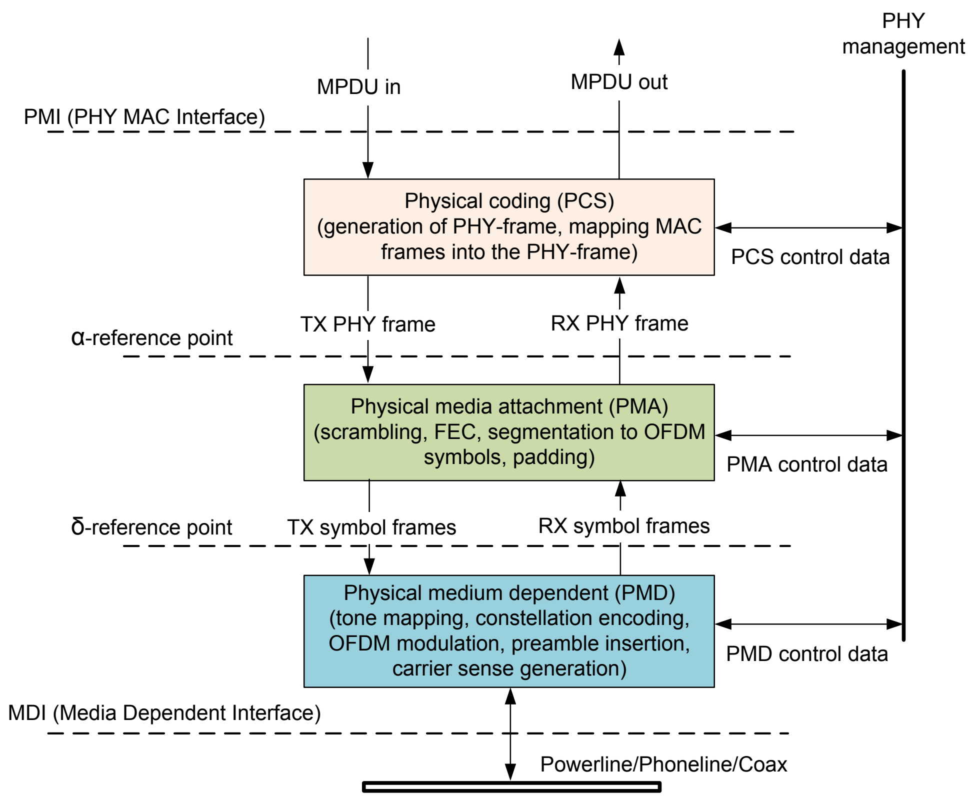

Figure 7: Functional Model of the PHY

The PMI and MDI are, respectively, two demarcation reference points between the PHY and MAC layers and between the PHY and the transmission medium.

Internal reference points δ and α mark separation between the PMD and PMA sublayers, and between the PCS and PMA, respectively.

In the transmit direction, data enters the PHY layer from the MAC via the PMI in blocks of bytes called MAC Protocol Data Units (MPDUs). An incoming MPDU is mapped into a PHY frame in the PCS, scrambled and encoded in the PMA, modulated in the PMD, and transmitted over the medium using OFDM modulation with relevant parameters for each medium. In the PMD, a preamble is added to facilitate synchronization and channel estimation in the receiver.

In the receive direction, frames entering from the medium via the MDI are demodulated and decoded in the PHY layer. The recovered MPDUs are forwarded to the MAC layer over the PMI. The recovered PHY-frame headers are processed in the PHY to extract the relevant frame parameters.

PHY Frame

A transmit frame (PHY frame) consists of a:

- Preamble
- Header
- Additional Channel Estimation (ACE) symbols (optional field, its presence is frame type dependent)
- Payload

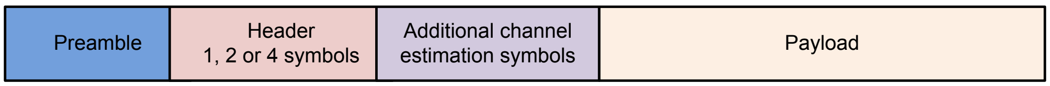

Figure 8: Format of the PHY Transmission Frame

- The Preamble and ACE symbols are added to the PHY frame in the PMD sublayer.

  The preamble does not transport any user or management data and is intended for synchronization and initial channel estimation.

  Preamble can be modulated using default seed or a domain specific orthogonal seed based on Domain ID (DOD). This mechanism is used to reduce level of interference in case of neighboring networks scenarios. For further details see Neighbor Domain Interference Mitigation (NDIM).
- The Header and Payload each contain an integer number of OFDM symbols.

  The PHY frame header carries settings of all programmable parameters related to the payload, such as guard interval, bit loading, and FEC parameters. The parameters of the PHY frame header are unified and specific per medium to ensure interoperability and selected to allow reliable detection of the header over noisy channels even without preliminary channel knowledge.

  The length of the Payload may vary from frame to frame; it may be of zero length.

  The Payload includes one MPDU composed of one or more LLC Protocol Data Unit (LPDUs). Each LPDU carries a segment of the transmitted data, a header identifying the carried segment, and the CRC to detect errored LPDU for selective retransmission. For the Payload, different coding parameters and bit loading can be used in different frames, depending on the channel/noise characteristics and QoS requirements.

Scrambling and FEC Scheme

All data starting from the first bit of the Header up to the last bit of the Payload shall be scrambled with a pseudo-random sequence generated by the linear feedback shift register (LFSR) with the polynomial $p(x)=x^{23}+x^{18}+1$.

The FEC scheme consists of a systematic QC-LDPC-BC encoder and a puncturing mechanism with five code rates (1/2, 2/3, 5/6, 16/18, and 20/21) and three block sizes of 21 (Header), 120 bytes and 540 bytes (Payload). Three parity check matrices are used for code rates 1/2, 2/3, and 5/6, whereas the other two high code rates are obtained by puncturing the rate 5/6 code. The range of FEC parameters together with bit loading capabilities are designed to fit the retransmission scheme.

Physical Medium Dependent (PMD) Sublayer

In the transmission direction, the Tone mapper splits the incoming symbol frames of the Header and Payload into groups of bits and associates each group of bits with a specific sub-carrier onto which this group shall be loaded. The constellation encoder converts each group of incoming bits into a complex number representing the constellation point for this subcarrier. The unused sub-carriers are modulated by a pseudo-random bit sequence.

The OFDM modulator converts the stream of the N complex numbers at its input into the stream of N complex valued time-domain samples. After adding the Preamble, the transmit signal is up-shifted by Fus (up-shift Frequency). The real part of the resultant signal is transmitted onto the medium (powerline, coax, or phone line). Parameters of the Preamble are determined by the PHY management and depend on the type of the transmitted PHY frame.

Frames are sent out on the medium with inter-frame gaps.

In the reception direction, the incoming frames from the medium are demodulated and decoded. The recovered symbol frames are transferred to the PMA over δ-reference point. The Preamble is processed and preamble data is transferred to the PHY management entity.

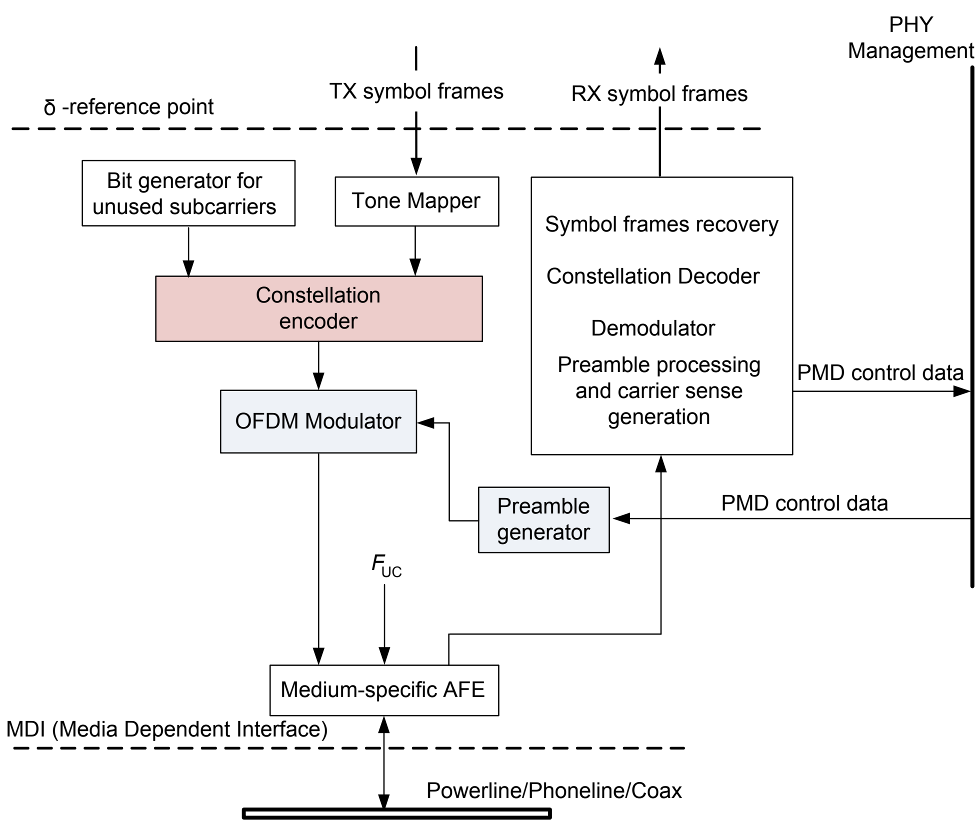

Figure 9: Functional Model of the PMD

G.hn Data Link Layer (ITU-G.9961)

G.9961 specifies reference models and functionality for all components of the Data Link Layer (DLL) of G.hn home network transceivers designed for the transmission of data over phone line, coaxial, and powerlines:

- Contention free TDMA, contention based CSMA, and STXOP “token passing” medium access control
- Security and confidentiality, including authentication, end to end encryption and key management procedures.
- Hidden nodes with data relaying
- Parameter based and priority based QoS
- Internal and external management communications
- Unicast and multicast retransmission based on selective acknowledgement and frame based acknowledgement protocols
- Bidirectional transmission
- Power saving mechanisms
- Network management procedures:
  - Network initialization procedures
  - Admission control
  - Node authentication and encryption key assignment
  - Connection management
  - Channel estimation
  - Bandwidth reservation and flow control
  - Topology maintenance and routing mechanisms
  - Automatic Mesh Network set up and re-establishment
  - Recovery procedures after Domain Master (DM) failure

Functional Model of DLL

The functional model of Data Link Layer is shown in the following figure.

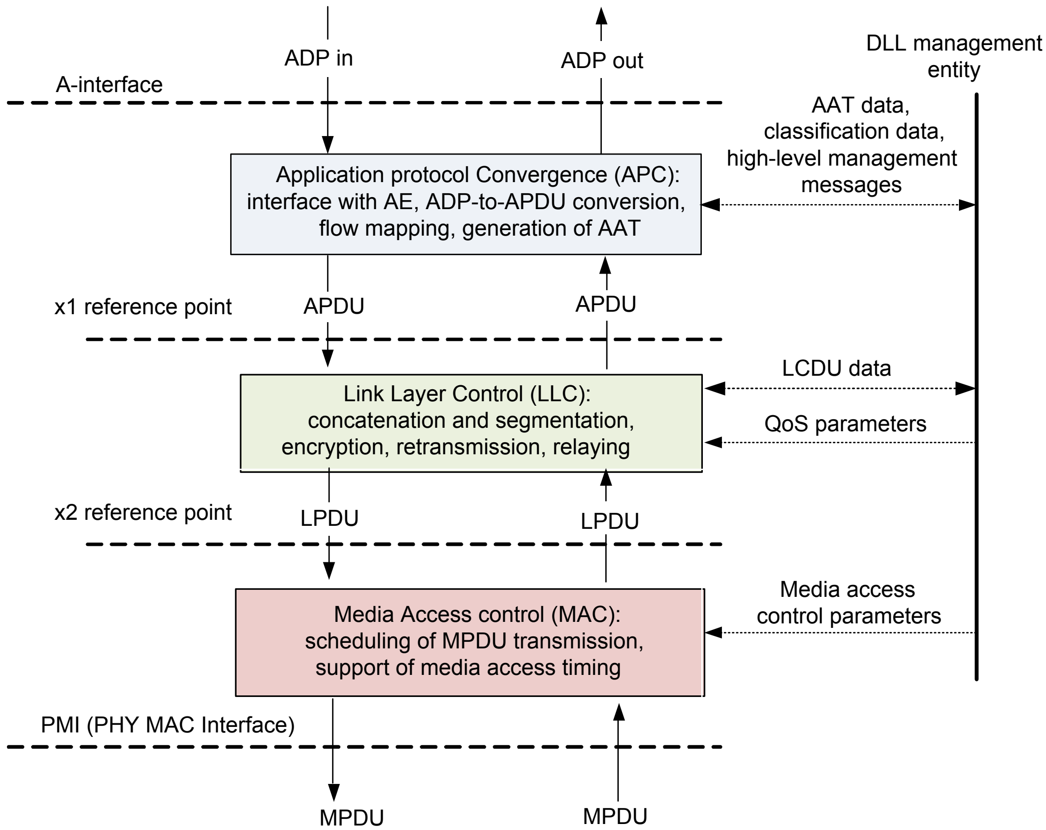

Figure 10: Functional Model of the DLL

The Application interface (A-interface) is the demarcation point between the Application Entity (AE) and the Data Link Layer (DLL).

The Physical Medium Independent (PMI) interface is the demarcation point between the DLL and the Physical Layer (PHY).

Reference points such as x1 and x2 show logical separation between sublayers: APC and LLC, and LLC and MAC, respectively.

In the transmission direction, Application Data Primitive (ADP) sets enter the DLL from the AE through the A-interface. Every incoming ADP set meets the format defined by the particular application protocol; for an Ethernet type AE, the ADP set has one of the standard Ethernet formats. Each incoming ADP set is converted by the APC into APC Protocol Data Units (APDUs), which include all parts of the ADP set intended for communication to the destination node or nodes. The APC also identifies ADP classification primitives (for example, priority tags), which can be used by the LLC to support QoS requirements assigned to the service delivered by the ADP. Furthermore, the APC is responsible for establishing flows of APDUs between peer APCs and assigning one or more queues for these flows according to the classification information associated with each APDU.

The APDUs are transferred to the LLC across the x1 internal reference point, which is independent of both application and medium. Additionally, LLC receives management data primitives from the DLL management entity intended for LLC control frames, which are mapped into Link Control Data Units (LCDU). The LLC is responsible for establishing flows of LCDU (control frames) between peer LLCs.

In the LLC sublayer, the incoming APDU and LCDU are converted into LLC frames and might be encrypted using assigned encryption keys. LLC frames are subject to concatenation and segmentation. Segments are transformed into LLC Protocol Data Units (LPDUs) by adding an LPDU Header (LPH) and CRC. LPDUs are then passed to the MAC across the x2 reference point. The LLC is responsible for retransmission and relay operations too.

The MAC sublayer is responsible for concatenating LPDUs into MAC Protocol Data Units (MPDUs) and then conveying these MPDUs to the PHY in the order determined by the LLC (scheduling, using number of transmission queues) and applying medium access rules established in the domain by the DM.

In the reception direction, MPDUs from the PHY enter the MAC across the PMI together with associated PHY frame error information. The MAC disassembles the received MPDU into LPDUs, which are passed over the x2 reference point to the LLC. The LLC recovers original APDUs and LCDUs from the LPDUs, performs decryption (if required), and conveys them to the APC and LLC management entity, respectively. In the APC, ADPs are generated from the received APDUs and conveyed to the AE.

The LLC is responsible for the decision regarding LPDUs received with errors (PHY frame errors). It decides whether to request retransmission of errored LPDUs (and generates the ACK response to assist retransmission), or to discard these errored LPDUs.

The functionality of the APC, LLC, and MAC is the same for all types of medium. However, some of their functions and control parameters may be adjusted for efficient operation of the G.hn transceiver over each particular medium.

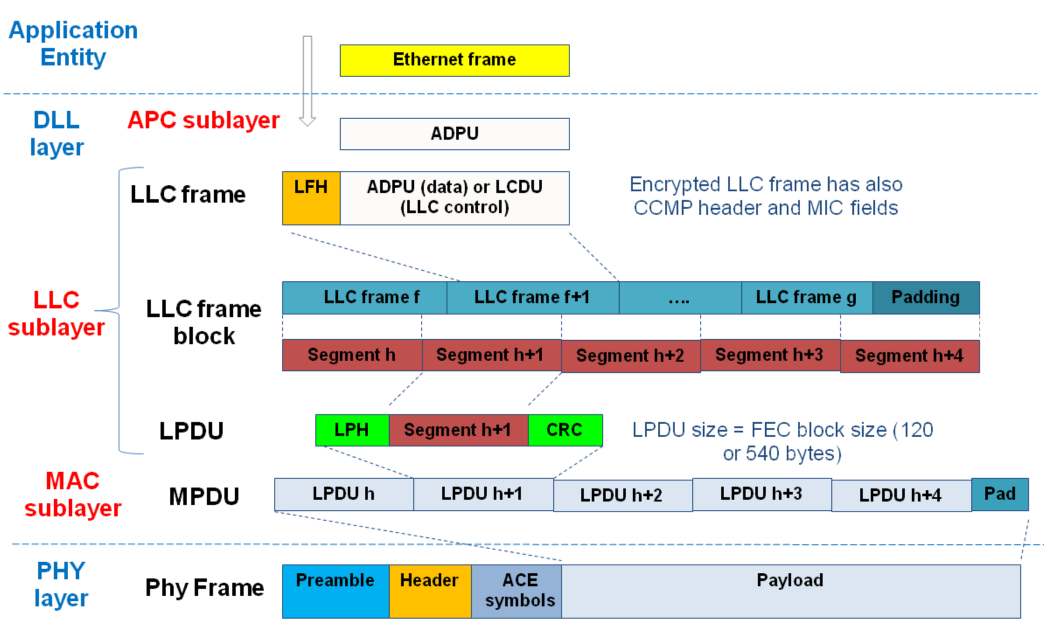

Figure 11: Data Plane in G.hn

MAP and TXOP

G.hn defines a synchronized media access scheme (for example, transmissions in the domain are coordinated by the DM and synchronized with the MAC cycle).

The time period of synchronized access is called a MAC cycle.

In G.hn PLC mode, the MAC cycle can be synchronized with the mains — to cope with periodically time varying behavior of channel response and noise caused by electrical devices and appliances plugged into the powerline.

Each MAC cycle is divided into time intervals associated with Transmission Opportunities (TXOPs) assigned by the DM for nodes in the domain. The DM assigns at least one TXOP for it to use to transmit the Media Access Plan (MAP) frame, which describes the boundaries and parameters of the TXOPs assigned for one or several following MAC cycles.

The assignment of TXOPs over multiple MAC cycles protects against MAP erasures by impulse noise. Other TXOPs are assigned by the DM to nodes requesting to transmit application data (for example, video services, data services, VoIP).

All nodes in the domain synchronize with the MAC cycle, read and interpret the MAP, and transmit only during the TXOPs assigned to them by the DM. The DM sets the order, type, and duration of TXOP based on requests from nodes and available bandwidth. The assignment of TXOPs is dynamic, and can change from cycle to cycle to better adapt to medium and network changes and to node requirements. The DM may also assign TXOPs over multiple MAC cycles.

To address different applications, two main types of TXOP are defined:

- Contention-Free TXOP (CFTXOP) defines medium access based on time-division medium access (TDMA): only one node can transmit during this TXOP. It targets services with fixed bandwidth and strict QoS (that is, video) needs.
- Shared TXOP (STXOP): A TXOP for which access is defined amongst a group of nodes. An STXOP is divided into a number of Short Time Slots (TSs), representing an opportunity for one or more nodes to start transmitting. Each TS inside an STXOP is identified by its ordinal position within the STXOP. An STXOP may contain two types of TS:
  - Contention-Free Time Slot (CFTS): Each CFTS is associated with a single node and a single flow/priority, identified by the node's DEVICE_ID and FLOW_ID/priority, which is exclusively allowed to transmit within that TS. If a node has nothing to transmit it passes the opportunity to next node.
  - Contention-Based Time Slot (CBTS): A CBTS is not associated with a specific node and its associated priority. This TS can be used by all the nodes or by a group of nodes indicated by the DM node in the MAP. CSMA/CA is used during this TS for medium access with a priority resolution scheme for ensuring highest priority traffic contends for the medium before lower priority traffic.
- An STXOP can be composed of only CFTSs, only CBTSs, or both CFTSs and CBTSs. An STXOP that is composed of CBTSs only is denoted as CBTXOP.

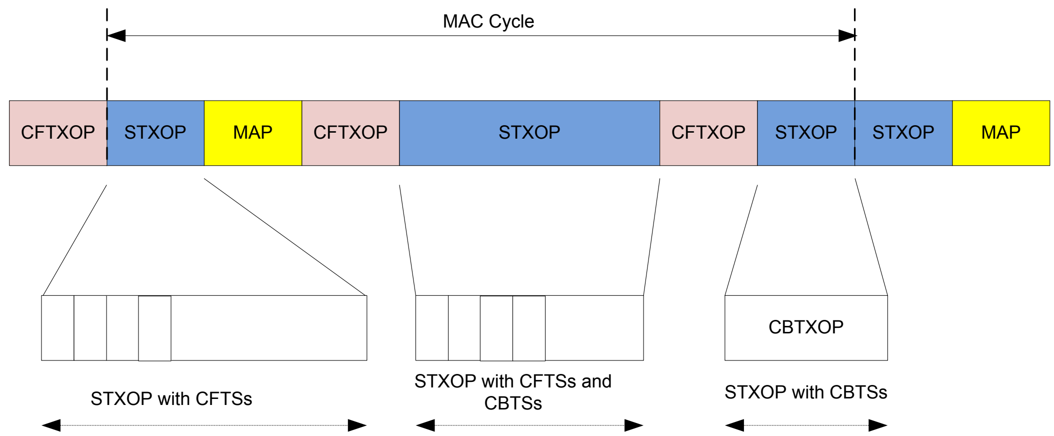

Figure 12: Example of a MAC Cycle Structure

The DM partitions the MAC cycle in CFTXOP and STXOP. The partition is done in accordance with service requirements of network nodes and domain scheduling decisions. At least one MAP is sent each MAP cycle in a dedicated CFTXOP assigned by the DM.

Medium access within STXOPs shall be performed using CFTSs and CBTSs. Each STXOP may contain zero, one or more CFTSs, each assigned to a given node. Similarly, each STXOP may contain zero, one, or more CBTSs, and each CBTS is assigned to several nodes potentially contending for this CBTS.

The type, placement and duration of TXOPs within a MAC cycle and the order of the TSs inside a STXOP are assigned by the DM according to internal scheduling decisions that are vendor specific.

An example of STXOP with CFTS and sequential TS rule assignment ("token passing") is shown in the following figure.

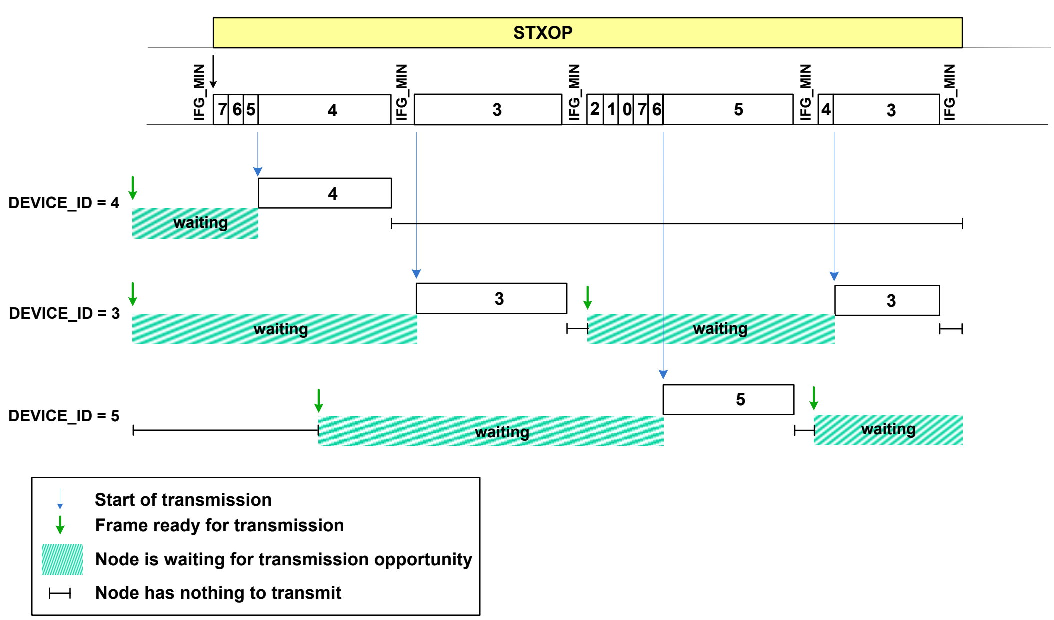

Figure 13: STXOP with CFTS and Sequential TS Assignment Rule

Nodes with data to be sent are using their assigned TS to do it, ending their transmission with an Inter-Frame Gap (IFG_MIN). And the nodes that do not have data to transmit are just letting pass the TS.

CBTXOP Operation

On the other hand, transmission during CBTXOP is arranged by contention periods. At the beginning of a contention period, each contending node indicates the priority of the frame it intends to send using Priority Resolution Slots (PRS) signaling. PRS signals determine nodes with frames of highest priority: only these nodes are allowed to contend, while all others back off to the next contention period. The probability of collision between the selected nodes is reduced by a random pick of back-off time before a node's transmission slot inside the contention window. From the beginning of the window, all selected nodes monitor the medium (by carrier sensing). If the medium is inactive at the slot picked by the node, the node transmits its frame; otherwise, it backs off to the next contention period. Next contention period is issued in the following CBTS.

To facilitate virtual carrier sensing, every frame indicates its duration in the frame header. Also, Request-To-Send (RTS) and Clear-To-Send (CTS) messages, similar to *IEEE 802.11*, are defined to reduce time loss in case of collision and improve operation in the presence of hidden nodes.

Use Case Example

An example of use case of several kinds of traffic within a four G.hn Home Networking nodes is shown in the following figure.

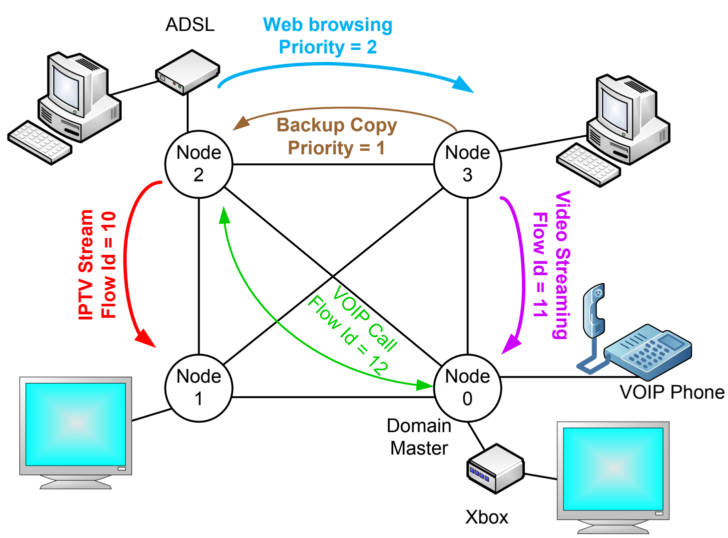

Figure 14: G.hn Network Use Case Example

This figure shows how G.hn nodes can manage these traffic types in terms of TXOP assignation depending on their bandwidth, latency and jitter requirements. Node #0 is the DM. Attached to it there is an Xbox and a VOIP phone.

The Xbox provides video services to a TV set connected to it. Another TV set is connected to node #1. An ADSL gateway/router, connected to node#2, bridges the network to an external WAN.

Two computers are also connected to the network, one through the ADSL gateway/router and another one connected via node #3.

Two video streaming, two data, and one VoIP traffic flows are shown in the above figure.

The following figure shows two consecutive MAC cycles on the network where the two streaming videos (flow id 10 and 11) and one VoIP call (flow id 12) are assigned to three different CFTXOP during the two MAC cycles, which duration depends on their service requirements (videos need more bandwidth than VoIP call).

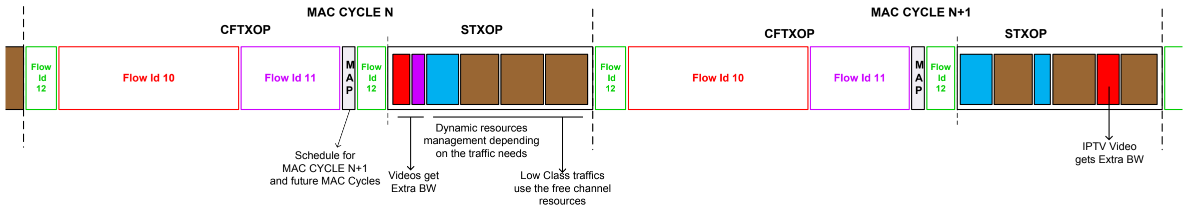

Figure 15: MAC Cycle Distribution of G.hn Network Use Case Example

One STXOP per MAC cycle is assigned for the rest of traffics, in this case for Web browsing and backup copy transmissions. Although, some TS of this STXOP are also used by video streaming traffics because they need extra bandwidth due to some temporary data rate peaks in the video flows (Variable Bit Rate videos) not covered by assigned CFTXOPs.

Security

Since G.hn is intended to operate over shared media, such as powerline and coax, its threat model includes two kinds of threats:

- **External threat** implies an attacker capable of eavesdropping on transmissions and sending frames within the network, but with out-of-network access credentials.

  G.hn defines an authentication procedure based on the Diffie-Hellman algorithm and the Counter with Cipher Block Chaining-Message Authentication Code algorithm (CCM), which uses AES-128.

- **Internal threat** is from a legitimate user of the network who has an illegitimate interest in the communications of another user or access to a specific network client. In case of hidden nodes, communications between two particular nodes may pass through a relay node, raising the potential for a man-in-the-middle threat.

  Internal threats typical for public installations, G.hn defines pair-wise security: a unique encryption key is assigned to each pair of communicating nodes and is unknown to all other nodes. Pair wise security maintains confidentiality between users within the network and builds another layer of protection against an intruder that has broken through the network admission control. Besides, end to end encryption avoids man-in-the-middle in the relays, as G.hn messages remain encrypted when relayed.

  The expected grade of security in G.hn is the same as or stronger than that defined in the most recent specification for WLAN IEEE 802.11n.

In both cases the goal is to protect against attackers with reasonably powerful computing resources but no access inside operating nodes.

Security Steps

G.hn security functions are based on controlled entry of nodes into the domain, Security Controller (SC) authentication of the nodes before communications is enabled, encrypted communications between nodes using Peer-to-Peer keys, and periodic re-establishment of authentication.

When a new node wants to join a domain, it must first register with the domain's DM. The DM verifies the node has a valid identity and is allowed to join the domain. Once this verification is done, the DM registers the node. This registration means that the node is a member of the domain; however, it is not allowed yet to communicate with any other node.

A newly registered node must then authenticate with the SC. The SC is a function that may be a separate function in the domain, co-resident with the DM function, or it may be distant from the domain. Once a node requests authentication from the SC, the SC validates the node's identity and, if the node is allowed into the domain according to the logic and records the SC may access, then it performs a series of steps with the node to establish a secure link with the node. This secure link enables the SC to pass encryption keys to the node for use when it communicates with other nodes. At this point, the node is allowed to communicate with other nodes.

When two nodes want to establish a communications flow between them, they query the SC for a set of keys that is unique to their specific flow. No other nodes have access to these keys, thus the flow between the two nodes is secure both from other nodes in the domain as well as external threats.

Periodically, the SC re-authenticates the nodes of the domain, authenticating the nodes and creating new keys for communicating between it and each node. Any node failing re-authentication is set as deregistered from the domain and unable to communicate with other nodes.

Strengths

G.hn's security architecture is composed of six major elements:

- Authentication = Procedure to ensure that only authorized nodes/users can exchange data.
- Key distribution = Procedure to distribute secret keys to authenticated devices in the network.
- Data encryption = Procedure to encrypt data with a secret key only able to be determined by the receiving node.
- Data integrity = Procedure to ensure that messages are not altered after they have been transmitted.
- Message uniqueness = Messages cannot be retransmitted by a third node in an effort to produce a "replay attack".
- Man-in-the-Middle (MITM) Protection = A relaying node within a domain cannot decrypt messages.

G.9961 security services have the following characteristics:

- Encryption based on AES-128 and CCM/CCMP
- Advanced Authentication and secure admission of nodes into a domain, based on *ITU-T Recommendation X.1035*.
- Key Management, including generation, secure communication, update, and termination of encryption keys.
- High confidentiality and integrity of all transactions, due to point-to-point authentication and unique encryption keys for unicast and multicast communications.
- Support of secure (encrypted) operation over nodes that relay messages to other nodes: relay nodes do not possess encryption keys of the relayed frames and do not pass relayed messages above the LLC sub-layer.
- Simultaneous operation of distinct, separately secured domains on the same medium.
- Procedures for setting up a secure network that can be self-contained in its processes or open to secure communications with a Security Controlling entity that is remote from the domain.
- Periodic and automatic re-registration, re-authentication, and encryption key updates.
- Replay protection

Scheduled Inactivity (Power Saving)

Nodes inform to the DM about their inactivity schedule, and DM does not allocate resources (time on wire to transmit or receive) to a node during its inactivity schedule. Inactivity periods can be within a single MAC Cycle (short inactivity scheduling) or a number of MAC Cycles (long inactivity scheduling). The rest of nodes are notified by the DM when another node is in inactive state. Message traffic to the inactive node can be buffered until the node is active once again. Therefore, power consumption is reduced thanks to the inactive periods. L1 (effective-power mode) takes advantage of short inactivity scheduling, and L2 (low-power mode) and L3 (idle mode) from long inactivity scheduling

The following figure shows short inactivity scheduling, where D and E nodes are only active to receive MAP and to listen to non-forbidden TXOP, the rest of MAC cycle both nodes are inactive.

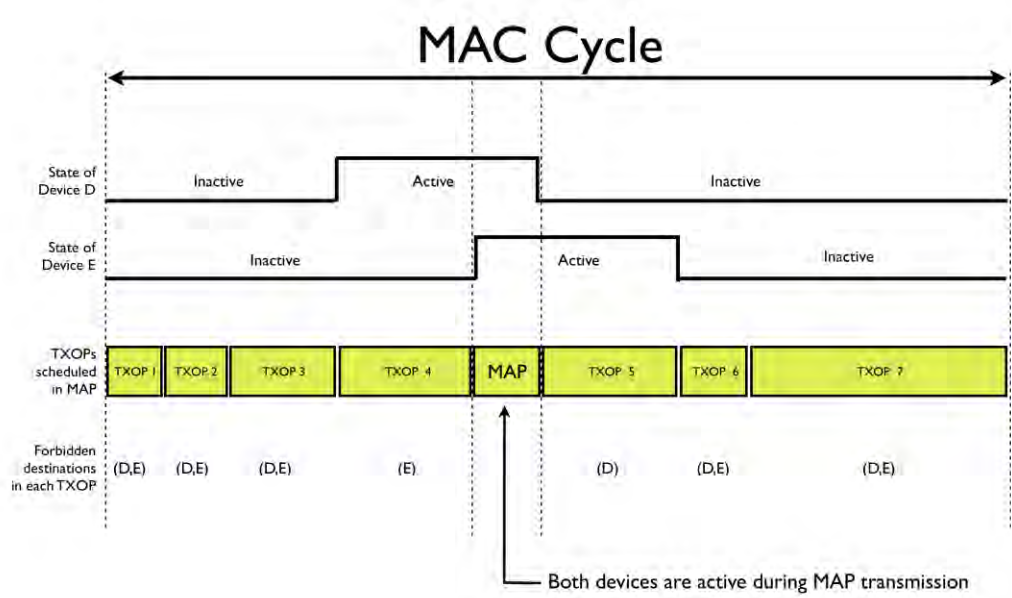

Figure 16: Scheduled Inactivity during a MAC Cycle (Short)

Bidirectional Traffic

Bidirectional transmissions between two nodes can be used to improve throughput and minimize latency of bidirectional traffic, such as TCP traffic with acknowledgements.

The defined bidirectional mechanism is only applicable to nodes communicating directly (for example, not via a relay node).

In the case of bidirectional transmission, the node originating the bidirectional traffic and the destination node exchange special frames: a Bidirectional Message (BMSG) frame and a Bidirectional Acknowledgement (BACK) frame. Both BMSG and BACK transport data, and in the case of acknowledged transmissions, also an acknowledgement on the recently received frame.

An exchange of BMSG and BACK frames forms a bidirectional frame sequence that shall last strictly inside the boundaries of the particular TXOP or TS assigned in the MAP for the node sourcing the bidirectional transmission. For an acknowledged bidirectional transmission only immediate acknowledgement is allowed.

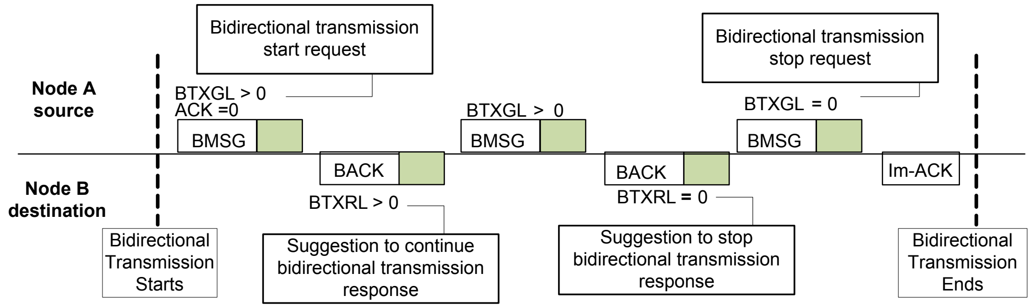

Figure 17: Example of Bidirectional Transmission

Functions of a G.hn End Point Node

End Point (EP) node is that node which currently is not a DM. The following points describe the functions of an EP node.

MAC Cycle Synchronization and Synchronized Transmissions

An EP node is not allowed to transmit until it detects the DM of the domain and it synchronizes with the MAC cycle indicated by the DM in the MAP or RMAP.

After synchronization with the MAC cycle, the EP node may register with the DM using the admission procedure.

After registration and authentication, the EP node must operate according to the medium access rules, and must strictly follow the TXOP and TS assignments in each MAC cycle advertised by the DM via the MAP message.

MAP and Messages Relay

- MAP relay = In some mediums there is a chance that a node may not receive direct transmissions from the DM (it is hidden from the DM). In order for such hidden node to be able to synchronize with the MAC cycle, another EP node shall relay the MAP upon the DM request, for instance, generate and send a Relayed MAP (RMAP) frame. There is MAP relaying for two purposes:
  - Registration of Hidden Nodes = A node that intends to join the domain may not detect the Default MAP (MAP-D) frames transmitted by the DM because the node is hidden from the DM.

  In order for such a node to register with the DM, another EP node (that is not hidden from the registering node) shall transmit MAP-D frames at the DM's request.

  The DM shall specify in the transmitted MAP message a TXOP descriptor to schedule the transmission of the RMAP-D and to specify the relay nodes to allow registration of new incoming nodes into the domain.
  - Operation of Registered Hidden Nodes = The DM ensures that every node already admitted to the domain can receive either the MAP or an RMAP in every MAC cycle.

  When the DM learns that at least one of the nodes in the domain is hidden from the DM it must designate one or more nodes to relay the MAP in every MAC cycle.

  A node is designated if the DM allocates a TXOP or TS to it to send an RMAP.
- Messages relaying = In some cases it might occur that some of the nodes are hidden from others, and may be or may not be hidden from the DM.

  To allow communication between hidden nodes, other nodes act as data relays.

Bandwidth Reservation

To support bandwidth reservation for flows and to manage flows that require QoS, EP nodes shall support a Flow Signaling Protocol that is used to establish flows with particular QoS parameters, modify them, or terminate them.

Routing of Application Data Primitives (ADPs)

Each node shall inform the domain about the nodes of the common domain that it has detected.

Each node can have one or more applications associated with its AE (above it's A-interface).

Each application is identified by a unique 6-octet MAC address.

Each node shall maintain the full list of addresses associated with applications above its A-interface. This list is called local Address Association Table (AAT). Each node shall also maintain the list of addresses associated with the AEs of other nodes in the domain. This list is referred to as a remote AAT. Each node provides its local AAT to the DM and to other nodes of the domain using topology management messages.

When a node receives an ADP from the A-interface, it checks its AAT to determine if the ADP is intended for the node itself (local in-band management message) or for an AE associated with another node.

Broadcast of LLC Frames

To facilitate broadcast of an LLC frame, every node shall obtain the Broadcast Routing Table (BRT). The BRT of a particular node contains a list of destination nodes (list of DEVICE_IDs), to which this particular node must relay the broadcasted APDU or LCDU. This list depends on the source from which the broadcasted APDU or LCDU was received. It is intended for the node to create multicast groups or use PHY unicast transmissions or PHY broadcast transmissions to reach the destination nodes indicated in the BRT (the DID, Destination ID, of the PHY frame could be a DEVICE_ID, or a MULTICAST_ID, or a BROADCAST_ID).

Reporting of Detected Neighboring Domains

All nodes must send information to the DM and to other nodes of the domain about any and all detected neighboring domains (those neighboring domains whose signals are detected as valid G.hn signals and not noise).

Functions of a G.hn Domain Master

A DM capable node is a node that, as well as supporting all of the required capabilities of an EP node, it is also able to assume the role of a DM.

A DM capable node must support all of the functions specified in the following points.

Only one node is allowed to act as a DM for a domain at any given time. The rest of the nodes within the domain are managed by this DM. If a DM fails, another node of the same domain, capable of operating as a DM, takes over the function of the Domain Master.

The DM's node shall perform medium access control using the same medium access rules as for EP

Domain Master Selection Protocol

The selection of DM is done automatically among all nodes which have the capability of being DM. DM selection protocol provides the mechanism to select one single DM in presence of several nodes that have the capability to become DM.

During the initialization, G.hn nodes shall search for a MAPs or RMAPs with the same Domain Name as the configured in their “Target Domain Name” parameter during t0 (typically 10s). If during t0 a MAP or RMAP is received, the node starts the procedure to join to the domain.

If after t0 the node has not received any MAP or RMAP, the node starts after t1 (random time between 0 and 1 second) a new domain becoming the DM of the created domain and start transmitting MAP frames. If during t1 time a MAP or RMAP is received, the node aborts the process to create a new domain and starts the procedure to join that domain as an EP node.

In case a failure of the DM is detected and no backup DM exists, a node capable of becoming DM shall take the role of DM after waiting for t1 without receiving MAP frames from other node.

Standard also defines DM ranking capabilities and a handover procedure of the DM role to a more capable node depending on that ranking.

Network Admission

To join a G.hn network, all nodes shall first be registered with the DM using the Network Admission protocol. Normally non registered nodes are able to receive successfully the MAP frames only if the MAP is transmitted in the Default MAP (MAP-D) format. Therefore the DM must send periodically MAP-D messages in addition to Active MAP (MAP-A) transmissions to enable registration. MAP-A is a MAP modulated with specific settings of the domain that are announced in the MAP-D. If a node does not have direct communication with the DM, because it is hidden from the DM, this node can still register and become part of the network using Relayed admission.

For registration, a unique Registration Identifier (REGID) is assigned to every node before its installation. REGID is only intended for registration and may be communicated unencrypted. The value of the REGID must be equal to the MAC address of the node.

The registering node shall identify the domain it wishes to join by comparing the domain name information in the received MAP-D frames, with the target domain name provided to the node by the user (to distinguish his network from neighboring networks), through TR-069 ACS message, preset by the manufacturer or service provider or obtained during the first registration.

- If the domain operates in non-secure mode, a node that successfully registers with the DM can communicate with other nodes in the domain.
- If the domain operates in a secure mode, a registered node shall also follow authentication procedures. After authentication, the node becomes a member of the secure network and is in a position to communicate with any other node in the domain.

If no MAP frame meeting the target Domain Name Identifier (DNI) is found, the node that is not capable to act as a DM may continue searching for the target DNI. A node that is capable to act as a DM creates a new domain.

The DEVICE_ID of the registering node shall be set to zero. After registration is complete, the DEVICE_ID shall be set to the value assigned by the DM. A node cannot establish connections until it has been assigned a DEVICE_ID other than zero.

Bandwidth Management

The DM must be capable of allocating TXOPs and TSs to different nodes, user priorities and service flows. These allocations should be such that nodes transmitting within the assigned TXOPs and TSs should meet priority constraints for priority traffic and QoS bandwidth, latency and jitter constraints specified in the Traffic specification (TSpec) for the established service flows. This should work even in the presence of neighboring domains operating in the same medium.

The output of the scheduling process is the MAP. The MAP is sent every MAC cycle and defines the TXOPs and TSs allocated to node(s), user priorities and service flows in the next MAC cycle(s). The DM should maintain state information concerning the scheduling of medium resources in the domain and shall control the admission of new service flows and the allocation of medium resources.

Admission control of new service flows should guarantee that the minimum QoS requirements for existing services are not violated. The DM shall service requests to add and remove service flows and requests to change service flow characteristics. If a request is made to add a new service flow and the QoS requirements specified in the TSpec cannot be accomplished, the DM denies admission of the new service flow and a denial of service status is returned to the requestor. Be aware that denial of service flow establishment means that no QoS guarantees can be given to a particular service flow. In this case, medium access may still be performed on a priority-basis within STXOPs (CBTSs).

The TSpec describes the set of parameters, characteristics, and expected QoS related to a particular flow.

The TSpec may be provided to the node by its associated client before the flow is established.

The TSpec may include any of the following QoS attributes:

- Traffic Priority
- Maximum information Rate (MIR)
- Maximum Traffic Burst
- Committed information Rate (CIR)
- Tolerated Jitter and Maximum Latency
- Grant Interval
- Polling Interval
- APDU Size

Changes in line conditions should be detected by channel estimation. If the line conditions change and the transmitting node are forced to use a lower bit loading for an admitted service flow, it shall notify the change to the DM. The DM should then reallocate medium bandwidth reservations to account for the change (flexible TS sizing based on channel status). If there are not enough bandwidth resources to support the low bit loading, the DM may decide to reduce the allocations of one of the current active flows by updating its TSpec attributes or the DM may finish one flow or more in order to release the needed bandwidth resources.

If the service flow data rate is changed at the A-interface, the originating node shall initiate a new procedure with the DM to update the TSpec attributes of the service flow according the new service flow parameters.

Synchronization to an External Source

Synchronization of the MAC cycle to an external source by the DM is optional when operating over phone line or coax media. Over powerline medium connected to AC mains supply, the AC cycle shall be considered as the external source, with a nominal cycle frequency of 50 or 60 Hertz, the DM shall synchronize the MAC cycle to the powerline cycle every 2 Hertz.

Routing and Topology Management

In G.hn there are two modes for routing and topology management:

- Centralized Routing and Topology Management (CRTM) Mode. Every node in the domain transmits their topology information to the DM. The DM calculates the routing table for all nodes of the domain and sends the calculated routing table and topology information to all nodes in the domain. After that, the nodes in the domain use the routing table received from the DM.

- Distributed Routing and Topology Management (DRTM) Mode. Each node calculates its own routing table using the standard algorithm indicated in the MAP by the DM and distributes this topology information to all nodes in the domain.

A domain may operate in CRTM or in the DRTM mode. The DM shall advertise the operation mode in the MAP. All nodes admitted to the domain shall use same routing and topology management mode which was indicated by the DM.

A node shall indicate its capability to support the DRTM mode during the registration process and in the topology update message.

The DM shall use CRTM mode if not all nodes in the domain are capable of calculating routing tables.

The DM and the backup DM shall be capable of gathering and keeping information describing the network topology.

The DM shall broadcast updates of its topology information by transmitting a message after it receives any update. In case the update of the topology received by the DM requires update of the routing table, the DM shall calculate a new routing table and a new Broadcast Routing Table (BRT), and transmit those to all nodes in the domain. Thus, the DM maintains a routing table that contains routing information from each node to all other destination nodes in the domain. This routing table contains N × N elements, where N is the number of nodes in the domain.

Backup Domain Master

Each DM-capable node is able to act as a backup DM.

The role of a backup DM must be assigned to a DM-capable node by the acting DM.

Only a single node acting as a backup DM or none may be assigned at any given time.

The backup DM shall take the role of the DM only in case a failure of a DM is detected.

Selection of the DNI and the DOD

When a node starts its role as a DM, either at initialization or as a result of a handover, it shall set the values of Domain ID (DOD) and Domain Name Identifier (DNI) and use them beginning from the first transmitted MAP.

The DM, prior to sending its first MAP frame, shall monitor the DODs of all visible neighboring domains and pick the DOD randomly between 1 and 15, excluding values already used by neighboring domains. The value DOD=0 is reserved by ITU-T. Furthermore, DM shall compute the DNI using the default value of the hash key and using the value from the 32-byte Domain Name. Furthermore, the DM shall monitor the DNI of all visible neighboring domains and re-compute the DNI using one of the alternative hash keys if the same value of DNI is discovered.

Each node that joins the domain shall verify the DNI using the hash key indicated in the DNI_KeyID field of the MAP Header.

Addressing Scheme

Node Identifier

The following three node identification parameters are used:

- DEVICE_ID = after a DM has successfully registered the node to the domain, it must assign a unique DEVICE_ID for that registered node and communicate the assigned DEVICE_ID to the node. The assigned DEVICE_ID (an 8-bit unsigned integer with valid values in the range from 0 to 250) shall be used until the node is resigned from the domain.

- MULTICAST_ID = it must be generated autonomously by the nodes creating PHY multicast groups for the multicast transmission. MULTICAST_IDs only apply to multicast transmissions among nodes communicating directly (for example, not via a relay node). A node shall generate a unique MULTICAST_ID for each PHY multicast group that it creates (8-bit unsigned integer with valid values in the range from 1 to 254).
- BROADCAST_ID = it is a MULTICAST_ID with a fixed value of 255 and must be used for broadcast transmission only to all nodes.

The same node can be identified by its unique DEVICE_ID, by several MULTICAST_IDs and by the BROADCAST_ID, which is the same for all nodes. The node identifier shall be used to identify the assignment of TXOPs and TSs within STXOPs to the nodes in the MAP and to identify the source and destination identity of a PHY frame (SID and DID).

***Table 1:* Definition and Valid Values of Node Identification Parameters**

Parameter	Valid Values	Description
DEVICE_ID	0	The ID used by a new node joining the network before it is assigned a unique DEVICE_ID by the DM. The DM shall not assign the DEVICE_ID = 0 to any node admitted to the network.
	1 to 250	IDs reserved for assignment by the DM to nodes admitted to the network.
	251 to 255	Reserved by ITU-T
MULTICAST_ID	0	Reserved by ITU-T
	1 to 254	IDs reserved for assignment for PHY multicast traffic
BROADCAST_ID	255	A special value of MULTICAST_ID reserved for broadcast traffic

Flow Identifier

A node may source multiple flows where each flow is identified by a FLOW_ID and considering that a flow is a unidirectional stream of data between two nodes related to a specific application (or characterized by specific QoS requirements).

The FLOW_ID is uniquely assigned by the node originating the flow when the flow and its associated data connection are established, and is released by the same node when the flow is terminated. A FLOW_ID must be represented by an 8-bit unsigned integer. Valid values of FLOW_ID are in the range of 8-254. A flow is uniquely identified in the domain by the tuple SID and FLOW_ID.

Retransmission and Acknowledgement Protocol

The retransmission and acknowledgment protocol specifies acknowledgment (either immediate or delayed) by the receiver. The transmitter must indicate in the header of the transmitted frame, using the RPRQ field, whether Imm-ACK, delayed-ACK, or no acknowledgement is required.

Retransmissions are made at PHY block level (LPDU frames) in a selective manner, where Maximum ACK Window Sizes are:

- 512 for data connections with LPDU size of 540 bytes
- 564 for data connections with LPDU size of 120 bytes
- 16 for management connections with LPDU of size 540 bytes
- 32 for management connections with LPDU of size 120 bytes

There are different procedures specified to acknowledge unicast PHY frames and multicast PHY frames. It is also possible to request for ACK retransmission.

Unicast Acknowledgement

When Imm-Ack is required for a unicast frame, the acknowledging node shall follow the reception of a frame with an Imm-ACK frame after the end of the frame that has requested the Imm-ACK.

In case of delayed-ACK, the transmission of the acknowledgement shall be deferred to a TXOP or TS assigned to the receiver. If an Imm-ACK request is received prior to transmission of the delayed-ACK, the deferred acknowledgement shall be sent in the requested Imm-ACK.

Multicast Acknowledgement

With multicast acknowledgement a frame addressed to a group of nodes is acknowledged by one or more nodes of the group using acknowledgment frames that are transmitted in predefined time slots that immediately follow the frame requesting acknowledgement response. Each Mc-ACK frame slot is uniquely assigned to a single destination node from the multicast group that acknowledges the multicast frame. In addition, a NACK signaling time slot may follow Mc-ACK slots, if requested by sender, and in this case all destination nodes of the multicast group that are not assigned a Mc-ACK slot in which to respond, shall indicate reception failure by transmitting a NACK signal in the NACK signaling slot.

The reception failure shall be declared if one or more errors were detected in those LPDUs of the received frame that were not received correctly in the previous transmissions; in case errors are detected only in the LPDUs that were already received correctly in previous transmissions, the frame should be considered as received rightly.

All destination nodes of the multicast group that received at least one Mc-ACK frame but did not receive the original multicast frame corresponding to this Mc-ACK frame must transmit NACK signal in the NACK signaling slot.

Channel Estimation

The channel estimation protocol describes the procedure of measuring the characteristics of the channel between the transmitter and the receiver nodes. The procedure involves initiation of channel estimation, transmissions of PROBE frames, and selection of parameters.

Channel estimation can be done in two phases:

- Channel discovery – initial channel estimation
- Channel adaptation – subsequent channel estimation to adapt to a changing channel

The protocols used for channel discovery and channel adaptation can be started either by the transmitter or by the receiver. The main part of the channel estimation protocol is always initiated by the receiver (receiver-initiated channel estimation). The transmitter can request the receiver to initiate channel estimation (transmitter-requested channel estimation).

During the initiation process, the transmitter and receiver jointly determine input parameters for channel estimation such as channel estimation window (a fraction of a MAC cycle over which channel estimation should be executed), and parameters for the PROBE frame.

The receiver selects the BAT_ID associated with the Bit Allocation Table (BAT) to be updated. This BAT_ID is used for an identifier for a particular channel estimation process throughout the rest of the process.

Once the channel estimation process is started, the receiver may request the transmitter to send one or more PROBE frames. The receiver can change parameters of a PROBE frame at each request. In case the receiver requests a PROBE frame without specifying its parameters, the transmitter sends the PROBE frame using parameters previously selected by the receiver. The receiver is not required to request PROBE frames if it chooses other means such as MSG frames or PROBE frames transmitted to other nodes to estimate the channel. The protocol provides numerous options to speed up the channel estimation process for faster channel adaptation.

The receiver ends the channel estimation process by sending the outcome of channel estimation to the transmitter. This includes at least the following parameters:

- BAT
- FEC coding rate and block size
- Guard interval for payload
- PSD ceiling

The receiver may cancel the channel estimation process without generating new channel estimation parameters.

Connection Management Protocol

Connection management protocol is in charge of negotiating the establishment, release and reset of unicast and broadcast connections that could be intended for data and management. Connections should be established by the transmitter and released by the transmitter or receiver, and it shall be established before exchanging any segment associated with that connection.

At a given time there must be multiple open connections between a transmitter and a receiver in a network:

- None or one management connection (for LCDUs exchange)
- Between zero and eight prioritized data connections (for APDUs not mapped to flows and LCDUs mixed with APDUs)
- Between zero and 247 data connections associated to service flows

Connections are unidirectional and it is supported to establish connections with or without level 2 ACKs.

The FEC block size chosen by transmitter for a connection shall be indicated in the PHY frame header when connection is established.

Multicast Binding Protocol

The multicast binding protocol must be used to create PHY multicast groups and manage multicast transmissions identified by (SID, MULTICAST_ID) tuples among nodes communicating directly (For example, not via a relay node). PHY multicast groups share same BPC table. DLL multicast specifies routing of multicast LLC frames in the network.

A transmitting node of a multicast stream may initiate the multicast binding protocol upon detecting the presence of a multicast source (that is, when IGMP Query or multicast traffic is transmitted by the multicast source) when there are nodes that requested to receive the multicast stream (that is, via IGMP join message). In case the transmitter initiates the multicast binding protocol, it shall compute the BATs to be used for the multicast stream based on the BATs reported by the receiver nodes that requested to receive this multicast stream. The transmitter shall then determine the number of multicast groups and the assignment of receivers to each multicast group.

There are procedures defined to initialize a multicast group for a new multicast stream (see figure below where multicast binding protocol messages are marked in blue arrows), maintenance of multicast binding information, and termination of a multicast group.

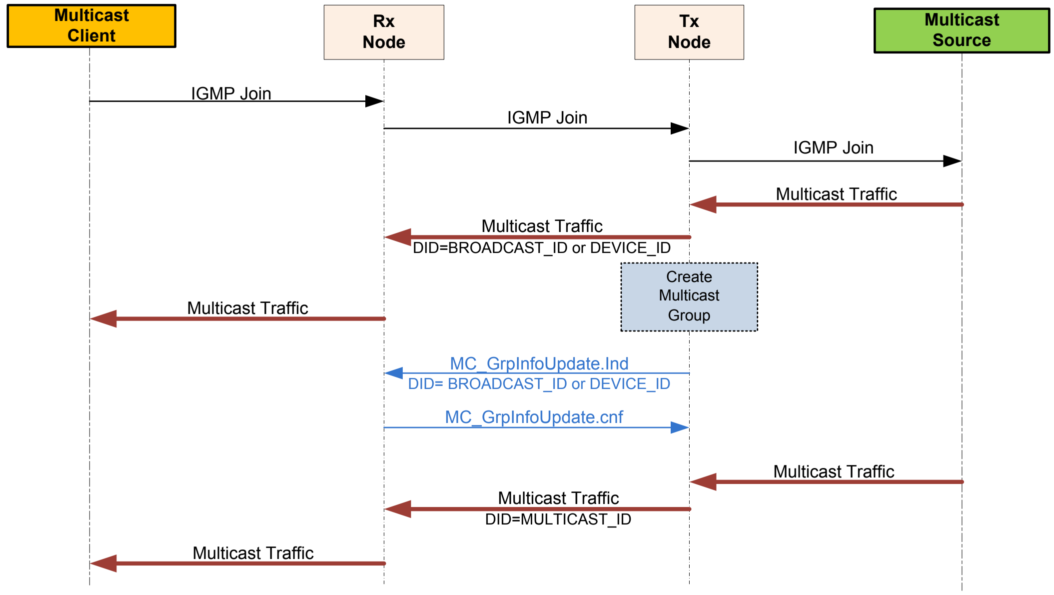

Figure 18: Example of Initializing a Multicast Group

Neighbor Domain Interference Mitigation (NDIM)

This functionality is intended for powerline and maximizes overall performance of several neighboring G.hn PLC home networks in multi-dwelling units through coordination or isolation mechanisms.

A G.hn node should be capable of detecting the presence of other neighboring domains operating over the same medium (either directly or via information sent by other nodes in its own domain). Once a neighboring domain is detected, a DM should take all needed actions to mitigate interference with the neighboring domains' nodes including coordinating with them while guaranteeing that QoS requirements for existing service flows are maintained.

The NDIM procedure includes the following mechanisms:

- A signaling mechanism for detection of neighboring domains
- A mechanism to align the MAC cycles of all the domains that are present on the same medium and can detect each other, independent of the level of interference among their nodes
- A routine mechanism allowing nodes to update the information of the neighboring domains that they can detect
- A mechanism specifying the usage of a near-orthogonal signal for generating and detecting the preamble, PR, INUSE and NACK signals in specific TXOPs. This mechanism may be applied in cases where the interference from neighboring domains is low. Based on this mechanism the transmissions from other neighboring domains are treated as noise
- A mechanism allowing coordination between domains of transmissions from and to nodes that can interfere with each other

An NDIM cluster is defined as a group of domains that have their MAC cycles aligned. A cluster must be uniquely identified by its 48-bit cluster identification (ClusterID).

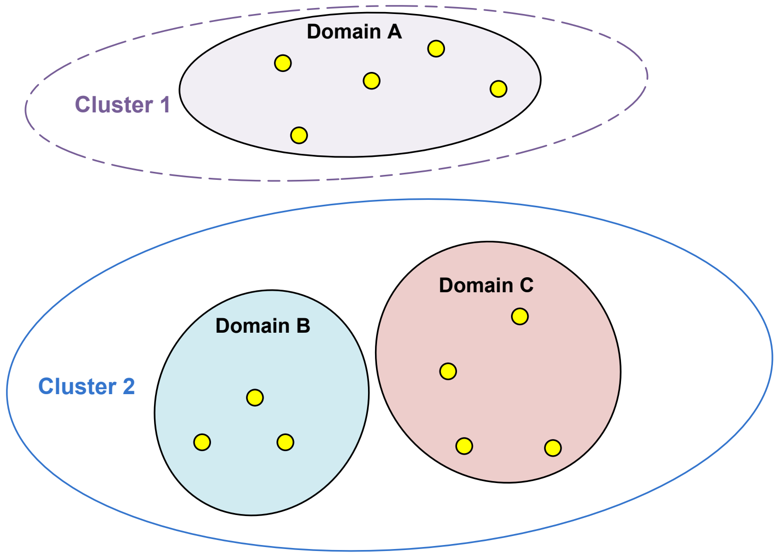

Figure 19: NDIM Clusters

To accelerate the processes of detection of neighboring domains and MAC cycle alignment, the DM shall reserve some time of the MAC cycle for Inter-Domain Communication Channel (IDCC) as depicted in the following figure.

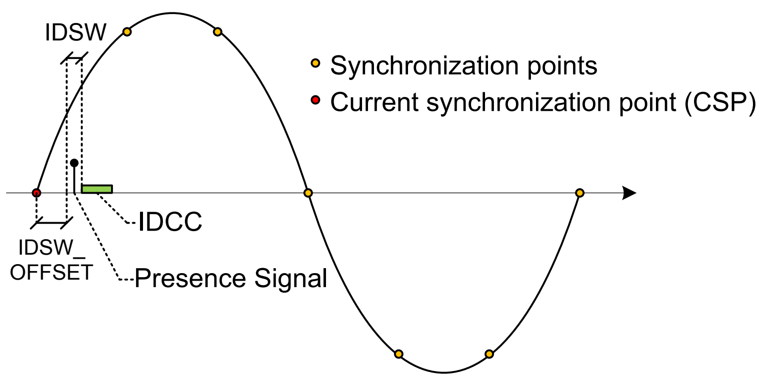

Figure 20: Position of Signaling Information in the MAC Cycle

Each cluster shall use one of these synchronization points as a reference for the position of the Inter-Domain Signaling Window (IDSW) and the Inter-Domain Communication Channel (IDCC). This selected synchronization point is called the Cluster Synchronization Point (CSP).

The IDSW defines a region of the MAC cycle that shall be used by all the nodes of a cluster for transmitting their Inter-Domain Presence Signal (IDPS). IDSW shall be present once every MAC cycle. The IDPS shall be generated within the IDSW by all nodes of the cluster domains to announce their presence to potential Neighboring Domains. The IDCC is used by every DM of the cluster to exchange inter-domain messages with the others domains' DM.

MAC Cycle Alignment Procedure

It consists of synchronizing the MAC cycles of two different clusters or a cluster and a domain by:

- Creating a single cluster including all the nodes and domains that belong to the original clusters.
- Choosing a unique CSP with associated IDSW and IDCC, to be used by all the nodes belonging to the newly created cluster for IDCC.
- Assigning a unique ClusterID for identifying the resulting cluster.
- Setting AC mains cycle synchronization to the same for all domains in the cluster.

Routine Maintenance Procedure for a Node to Follow After MAC Cycle Alignment

Once MAC cycle alignment has finished and only one single cluster has been created, there is a routine maintenance procedure for a node to keep track of the status of the neighboring nodes that interfere with it. For this, the DM shall publish in the MAP the DEVICE_ID of the node that shall contend to signal its presence in the IDCC in the following MAC cycle. This procedure allows finding any new nodes in the cluster that may interfere with the monitoring node or to keep track of the neighboring nodes that disappeared from the medium.

Communication between Neighboring Domains

An IDM message can be transmitted directly between nodes from different domains through the IDCC or between Domain Masters of different domains.

- Node to node communication through IDCC: it can be of two types, either unicast node to node communication or broadcast node to multi-node communication.
- Transmitting IDM messages from DM to Neighboring Domain Master. A DM selects a proxy to communicate with a given neighboring domain DM. It embeds the LCDU message and sends it to the chosen proxy in the neighbor domain via its own proxy using the unicast node to node communication.

Neighbor Domain Coordination Mechanism

- Input to the coordination algorithm consists of information about:
  - Nodes causing interference to nodes in neighboring domains
  - Nodes suffering from interference coming from nodes in neighboring domains
- For coordination, only nodes causing and suffering high interference are considered. Nodes with low interference can simply use different seeds for signals like preamble, INUSE and NACK, for interference mitigation.
- All nodes in the domain which are not interfering are allowed to transmit anytime during the MAC cycle.
- All nodes in the domain are allowed to transmit anytime during the domain’s assigned 1/16th of the MAC cycle. This assigned time position in the MAC cycle depends on its DOD.
- Permission for transmission for a node during the rest of the MAC cycle, for instance 14/16th of the MAC cycle, since 1/16th is allocated to reserved DOD 0, depends on the number of domains that this node needs to coordinate with and is based on a pre-defined rule that all domains follow.
- The default permissions for transmission for a node that are based on the interference information and pre-defined rule, can be changed by exchange of time resources between coordinating domains.
- Additional constraints could be applicable in some situations to deal with the time allocation corresponding to G.9972 in case of presence of one or more IEEE 1901 neighboring networks.

Intra-domain receptions for nodes that suffer interference from neighboring domains can be protected if the transmitting nodes of the domain are aware of the constraints applicable to the receiving nodes of the domain and transmit only during the time duration in which the receiving nodes are protected. This may be enforced by the DM by using appropriate scheduling policy or by communicating each node's time allocation information to all other nodes in the domain.

MIMO (ITU-G.9963)

MIMO refers to a technology that has the ability to use more than one transmit path and more than one receiver path, in other words, multiple inputs to the path (or channel) referencing the transceiver's transmitters and multiple outputs from the path (or channel) referencing the transceiver's receivers.

Examples of commonly used MIMO technologies are typically found in the wireless space, with the I*EEE 802.11n standard* defining MIMO options for WiFi transceivers.

ITU-T G.9963 (G.hn-MIMO) defines an enhancement to the operation of G.hn in PLC mode to add MIMO functionality.

In Wireline, the use of multiple transmitters and/or receivers depends on the number of electrical paths between the two communicating devices. Each path requires two wires to function.

In the Home, there is the use of 3-wire electrical sockets in many regions. These enable G.hn-MIMO enhanced devices to send and receive signals over all three wires as two logical circuits.

In Wireless, MIMO technologies enable better transmission/reception in crowded environments, while moving, and under increased noise conditions. Further, when conditions are favorable, MIMO transceivers are able to establish links at extended distances versus Single Input, Single Output (SISO) devices.

In powerline communications, the use of MIMO brings several improvements over standard SISO powerline transceivers.

- First, there is the ability to increase coverage in the building before the need for relays arises.
- Second, the ability of MIMO signals to cross over to other phases in the electrical wiring increases coverage and performance for many areas of the building.
- Third, there is the improved throughput as MIMO is based on a highly optimized communication scheme that sends spatially multiplexed signals over each port, with embedded self-noise cancellation techniques.

Introduction to G.9963

G.9963 specifies the basic characteristics of MIMO enhancement to G.hn Networking transceivers capable of operating over premises powerline wiring. It provides a high level introduction to the required technology and additions and modifications to recommendations G.9960 and G.9961 needed to define a MIMO enhanced G.hn networking transceiver.

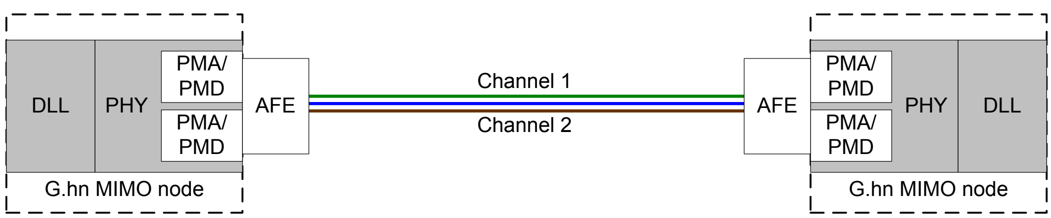

Figure 21: G.hn PLC Nodes with MIMO Linked by Multiple Channels or Paths

G.hn-MIMO transceivers are able to transmit and/or receive over three powerline conductors (phase, neutral, and ground) using more than one Tx and/or Rx port. It provides a substantially increased data rate, greater noise immunity, and enhancing the connectivity (for example, service coverage) of the Homenetwork.

G.9963 includes:

- The PHY functional models of MIMO transceivers;
- Descriptions of the modifications (changes and additions) needed in the PHY and DLL sections relative to G.9960 and G.9961 Recommendations (for example, transmission and reception are mapped to multiple TX and RX ports in the PMD sublayer).
- The means by which both G.9960/G.9961 and G.9963 transceivers interoperate when communicating on the same wires (heterogeneous G.9960/1 and G.hn with G.9963 domain).
- The means by which transmissions from G.9963 transceivers do not degrade performance of G.9960/G.9961 transceivers when operating on the same wires.

A G.9963 transceiver is defined to be fully compliant with G.9960 and G.9961 recommendations, which ensures interoperability.

Power Spectral Density (ITU-G.9964)

G.9964 specifies Power Spectral Density (PSD) mask requirements for transmission over phone, power, and coaxial lines, as well as the allowable total transmit power into a specified termination impedance. It also provides a set of tools to support reduction of the transmit PSD, as well as a list of radio frequency bands to be notched out in powerline and phone line profiles.

Some examples of PSD limits and transmitted power levels for main powerline, phone line, and baseband coax profiles are shown in figures below.

For more details about rest of profiles refer to the G.9964 specification.

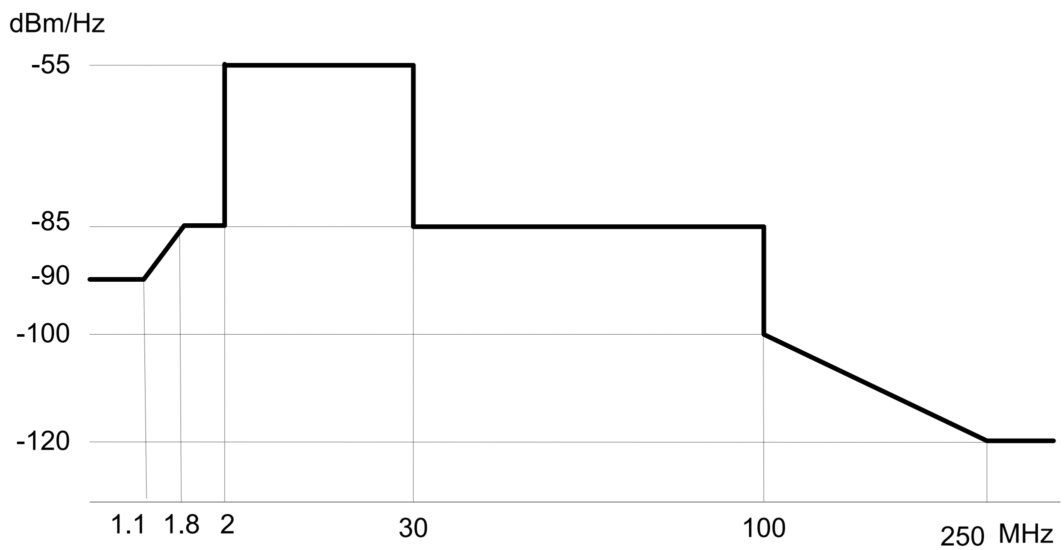

Figure 22: Transmitted PSD Limits for 100MHz Powerline Profile

Note: Notches are not shown.

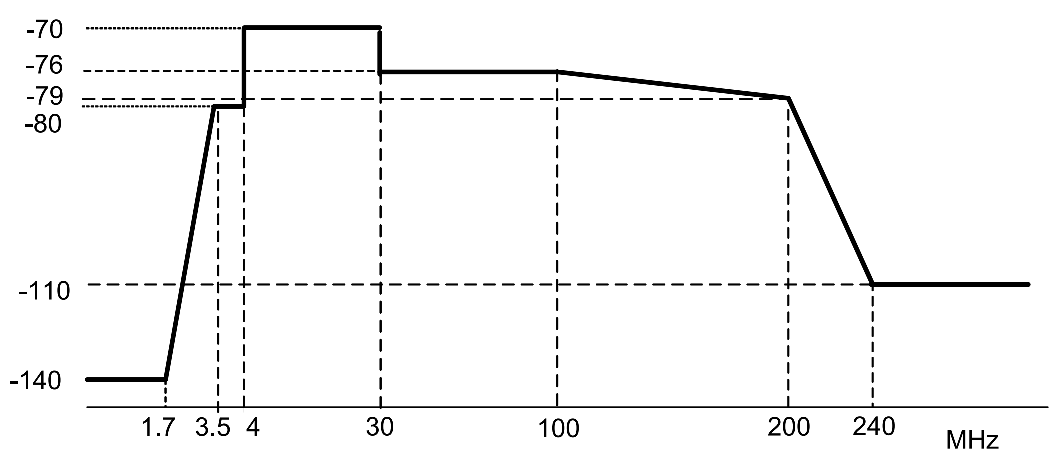

Figure 23: Transmitted PSD Limits for 200MHz Phone Line Profile

Note: Notches are not shown.

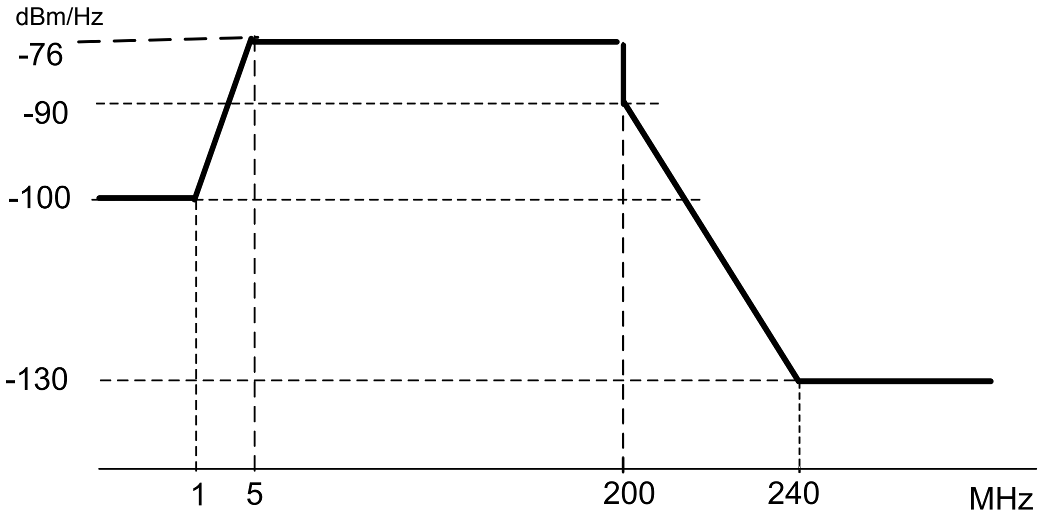

Figure 24: Transmitted PSD Limits for 200MHz Coaxial Profile

In G.9964 annexes there are some lists of international radio amateur, international broadcast, aeronautical and radio astronomy frequencies which might be notched out. For some frequencies dynamic notching is allowed and in other cases they must be permanently notched. Normally notches to be applied depend on regulation of each region.

Remote Management TR-069 and Management Plane (ITU-G9980/62)

G.9962 specifies the Management Plane for Physical and Data Link Layer of G.hn Home Networking nodes. This includes the management, architecture, protocols and common management parameters. It defines:

- Reference model (see Figure 25) and architecture for management layer.
- Protocols defined in G.9980, which are focused on TR069 remote management and needed for device configuration, checking status, diagnostics and security.
- G.hn management objects and parameters are included in *Broadband Forum's TR181 data model* (See https://www.broadband-forum.org/cwmp/tr-181-2-11-0.html#D.Device:2.Device.Ghn).
- Support for LCMP (Layer 2 Configuration and Management) protocol for configuring G.hn nodes.

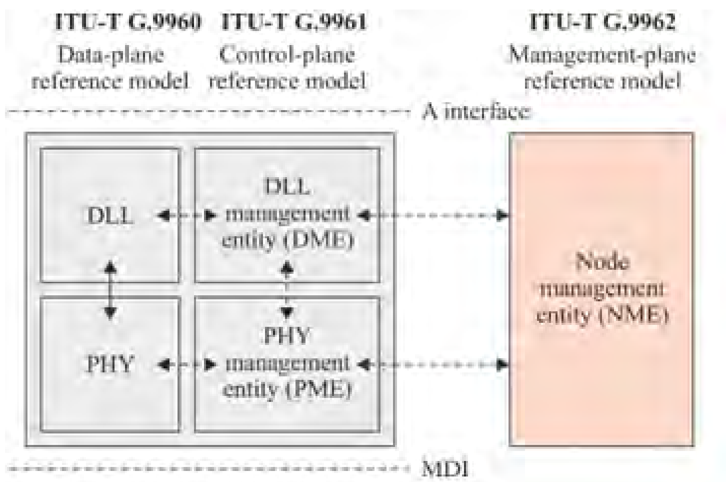

Figure 25: Management Reference Model

TR181 Issue 2 Amendment 11 specifies "Device.Ghn" object. This is an interface table for G.hn supported CPE, and it includes information related to G.hn interface itself in "Device.Ghn.Interface" object and G.hn diagnostics configuration and results in "Device.Ghn.Diagnostics".

Usually a G.hn node also includes other interfaces and assocciated objects such Ethernet, IP, DHCPv4 and v6, DNS, GatewayInfo, Time, etc, depening on application. Besides, each vendor can add their own vendor specific parameters to handle their particular configuration options and status information.

Coexistence Protocol between G.hn and IEEE 1901 over Powerline (ITU-G.9972)

The G.9972 Coexistence Protocol mitigates interference between G.9960/1 nodes and *IEEE 1901* networks. G.9972 provides Management Message Parameter List (MMPLs) for the G.9960/1 coexistence related management messages.

When mitigation using G.9972 is unnecessary, G.9972 provides a management message that is communicated between G.9960/1 nodes of the domain to cease transmission of G.9972 signals.

Powerline Communication devices may suffer interference from and create interference to neighboring powerline networks when operating over the same frequency range.

Therefore, when there is a chance that multiple non-interoperable powerline technologies are simultaneously using the same powerline cables in the same frequency range, it is strongly recommended that both G.9960/1 and neighboring devices use G.9972 to avoid performance degradation.

G.hn transceivers implementing the G.cx protocol coexists with 1901 devices that have the ISP functionality activated.

G.9972 uses the simple and flexible Inter-System Protocol (MMPLs), which allows other technologies to coexist with G.hn devices.

The ISP protocol adds a TDMA structure to the medium, efficiently allocating communication time to G.hn and non-G.hn devices.

Furthermore, the protocol empowers nodes with the capability of detecting when it is possible to transmit simultaneously to other nodes in their network (spatial reuse) without causing harmful interference and thus increasing overall network throughput.

MaxLinear's Offering

MaxLinear's G.hn development package is composed of the following elements:

- Silicon Products
- Spirit Firmware
- Reference Designs
- Software Tools

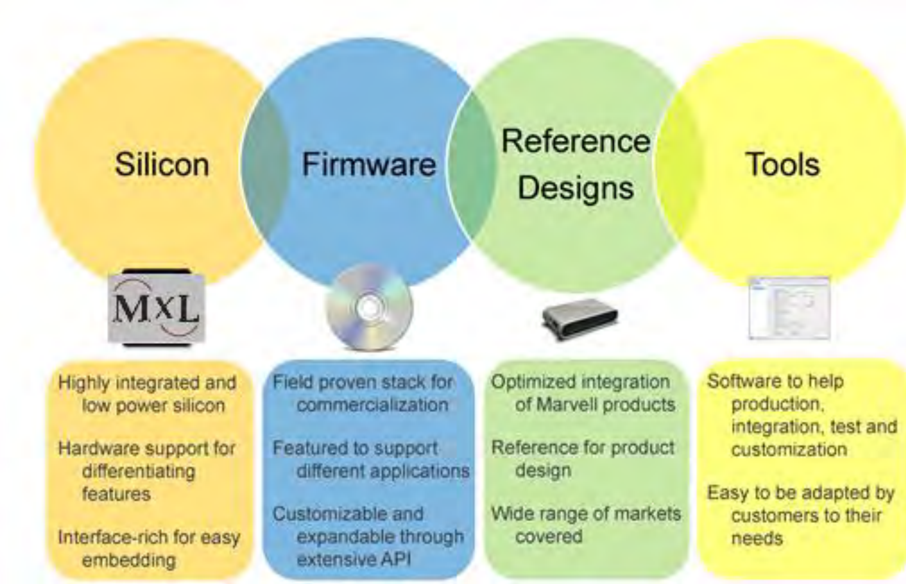

Figure 26: MaxLinear Turn Key Offering

Silicon Products

MaxLinear's silicon offering provides highly integrated and low power devices implementing G.hn specification and differentiating features with easy management and customization possibilities. Silicon is structured in a chipset comprised by a Digital Baseband (DBB) processor and an Analog Front End (AFE).

Digital Baseband Processor

MaxLinear's DBB processors are intended to support wireline communications using existing cables (powerline, coaxial, and phone cables).

They implement the ITU-T G.hn Standard providing Gigabit level and beyond throughputs. These devices contain industry-leading features and best in-class security, and directly interface to companion AFE devices, highly-integrated chips perfectly adapted to powerline, coaxial, and phone line media.

MaxLinear's DBB processors are low-cost integrated circuits optimized to enable high-speed video and data communications over coax, phone, or powerlines. Their high performance and low power consumption make them ideally suited for standalone Gigabit Ethernet adapters and embedded communication solutions in devices such as gateways and set top boxes as well as in combination with existing WiFi access points.

The devices include a G.hn data path, which implements a G.hn PHY as specified in ITU-T G.9960 and ITU-T G.9963, as well as hardware functions of the G.hn DLL layers specified in ITU-T G.9961, including the AES-128 encryption engine.

DBB processors also contain an embedded microcontroller subsystem to implement the software functions of the MAC and DLL layers, as well as management functions as specified by ITU-T G.9962. This subsystem includes peripherals such as timers, interrupt controller and serial UART and SPI controllers (in addition to the CPU and processor memory).

Regarding data connectivity, they incorporate rich set of interfaces such as high speed SGMII, 2500Base-X or RGMII.

For further information refer to the latest Digital Baseband Processor Data Sheet available in MaxLinear Watchdox.

Analog Front End

MaxLinear's AFEs are programmable, high-performance silicon devices implementing transmission and reception channels in charge of amplification, filtering, and driving the corresponding line signal perfectly adapted to powerline, coaxial, and phone line media.

In addition to the transmission and reception paths, AFEs contains a biasing circuit and a register block controlled from the digital interface.

Their operating mode is highly configurable to balance performance and power consumption based on the requirements of bandwidth, power density, and physical medium.

For further information refer to the latest *Analog Front End Data Sheet* available in MaxLinear Watchdox .

Firmware

The embedded microprocessor in the Digital Baseband Processor runs MaxLinear's time-proven Spirit firmware that provides a rich set of communication and management applications as well as a flexible API that developers can use to build their own code. It also.offers a high degree of customization of existing features. External processors are only required for very specific applications. The interconnection to external application processors can be achieved using SPI, SDIO, MII, and UART interfaces.

Spirit firmware dictates how G.hn nodes intercommunicate. Depending on the application, requirements are different. In order to match different scenarios, MaxLinear provides 3 different versions of firmware:

- Spirit Home Networking (HN) Firmware = Targeted to support Home Networking applications where peer to peer connectivity is required. Additional features specially tailored to these scenarios are also included.
- Spirit G.now Firmware = Targeted to support phone line in-building scenarios where point to point links share the same bundle presenting crosstalk. Specific features are included to handlethis crosstalk to maximize performance.
- Spirit Point to Multipoint (P2MP) Firmware = Targeted to support coaxial in-building scenarios where multiple end points share the same cable. Special MAC schemes are included to maximize the efficiency of the network.

Every Spirit Firmware release includes the following deliverables:

- Documentation = Data sheets, Test reports describing how the release has been validated, Validation Reports where the results of validation are described and release notes with a description of new features included.
- Spirit binary files = Default flash images for all supported products in a firmware release.
- Product Configuration Kit (PCK) = This software package is composed of a GUI and some auxilary software tools to encode configuration parameters and build flash images. PCK facilitates to ODMs and OEMs an easy way to customize production firmware images.
- Software Development Kit (SDK) = Set of compilers, libraries and tools running in Linux. Used for advanced firmware customizations such integrating a TR069 client, modifying drivers, changing LED behavior… Programming is in Posix C.

For further information refer to the corresponding Spirit Firmware package in MaxLinear Watchdox.

Reference Designs

MaxLinear's G.hn design package includes reference designs for developing different products based on G.hn silicon, mainly Ethernet to G.hn adapters, in all the variants of media covered (powerline, coaxial and phoneline),

MaxLinear also provides application designs and application notes for developing other G.hn solutions to complement Reference Designs such G.hn WiFi Extender, Embedded Modules, Multi-Medium implementations, etc.

The Reference Design has been thoroughly validated for EMC and thermal performance, and can provide an advanced manufacturing design foundation for manufacturers seeking to minimize their time to market.

Reference designs include the following documentation:

- Data sheets
- EVK user guide explaining how to start initial evaluation of equipment based on the reference design
- Bill of Materials (BOM).
- Schematics, layout and pick and place files for all circuit boards included in the reference design.
- A set of design guidelines and criteria for selecting alternatives to critical components and how to select region or country specific plug options.
- EMC report for the reference design explaining the EMC testing that has been undertaken on the design and the results achieved

Reference designs also include a number of special features to simplify product manufacturing and thus reduce production costs whilst ensuring performance parameters.

- All designs are RoHS Compliant and Halogen Free
- Components are chosen to ensure good global availability
- Built-in Self-Test (BIST) features are included

Complementary application design and application notes are fully functional but not thoroughly validated. It Includes schematics, BOM and layout. A data sheet or some basic document may be included. RoHS compliance has been verified

For further information refer to the latest *Reference Designs User Guides* available in MaxLinear's Wachdox.

Software Tools

Software tools are provided within the G.hn design package to assist engineers and evaluators during manufacturing, testing, and configuration of MaxLinear G.hn devices.

These tools are:

- Production Tool Kit (PTK) = It contains all the applications required for preparation of the manufacturing line, execution of production test and configuration of MaxLinear's G.hn chipset based final products.
- Spirit Configuration Tool (SCT) = It is a Java application that manages the configuration of the G.hn nodes of a network from a PC connected to one of the nodes. For doing so, it uses the Low Complexity Management Protocol (LCMP), an ITU standard layer 2, and included in Spirit firmware. It includes also a SNR viewer option that depicts the Signal to Noise (SNR) ratio and Channel Frequency Response (CFR) between pairs of G.hn nodes within same domain. This is very useful to test the impact of attenuation and noise on communication.
    - Additionally, Spirit firmware provides a Web interface that allows user access to the same configuration parameters available with the SCT. It is accessible from any Web browser. The only requirement is that the PC must be in the same IP subnet as the G.hn nodes.

For further information refer to the corresponding Software Tools folder in MaxLinear's Wachdox.

MaxLinear, Inc.

5966 La Place Court, Suite 100
Carlsbad, CA 92008
760.692.0711 p.
760.444.8598 f.

www.maxlinear.com