



for any wire, anywhere in your home

PLC Neighboring Networks Interference

How G.hn solves PLC's Neighboring Networks interference problem

Author

John Egan

Marvell Semiconductor, Inc

With contributions from

Agustin Badenes Corella, Marvell Semiconductor

Chano Gomez Martinez, Marvell Semiconductor

David Thorne, BT

Erez Ben-Tovim, Sigma Designs

Marcos Martinez Vazquez, Marvell Semiconductor

Ravi Mantri, Metanoia Communications

Roy Silon, TangoTec

about HomeGrid Forum

HomeGrid Forum (HGF) merged with the HomePNA Alliance in May 2013, forming an industry alliance of over 70 members including some of the world's largest Service Providers, system manufacturers, and silicon companies. HGF promotes development and deployment of a single, unified, multi-sourced home networking technology, G.hn, over coax, phone wires, powerline, and plastic optic fiber while continuing to support the existing base of HomePNA deployments. HGF provides silicon and system certification through its compliance and interoperability testing programs to ensure that retail customers and service providers can have confidence in all G.hn and HomePNA products.

HGF members collectively provide an eco-system covering all aspects of the technology from Retailers to Service Providers, utilities to Smart Grid think tanks, system developers to test houses and silicon companies. Our goals include promoting the benefits of G.hn; enhancing G.hn technology to meet evolving industry requirements; ensuring interoperability, performance based on our certification program; and supporting the needs of Service Providers deploying G.hn and HomePNA technologies.

For more information on HomeGrid Forum, please visit our website at <http://www.homegridforum.org>



HomeGrid Forum Certified
G.hn Products
The mark of Certified
Compliance, Interoperability,
and Performance



HomeGrid Certified
HomePNA Products
The mark of Certified
Compliance, Interoperability,
and Performance

Note

THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE.

HomeGrid Forum disclaims all liability, including liability for infringement, of any proprietary or intellectual property rights, relating to use of information in this document. No license, express or implied, by estoppel or otherwise, to any proprietary or intellectual property rights is granted herein.

HomeGrid™, HomePNA®, HomeGrid ForumSM, HomeGrid Forum design, HomeGrid G.hn and G.hn-Lite Certified and Design, and HomePNA design are trademarks and/or service marks of HomeGrid Forum. All other names and logos are trademarks and/or service marks of their respective owners.

Copyright © 2013 by HomeGrid Forum. Specifications and content subject to change without notice.

	page
1 Overview	4
2 Neighboring Network Interference	5
2.1 PLC Signals Travel Between Units in MDU Buildings	5
2.2 Interference Patterns	8
2.3 Symptoms	9
2.4 Interference Mitigation Between Different PLC Technologies	10
2.5 Interference Between Networks of the Same PLC Technology	11
2.6 Deployment Considerations	14
3 Legacy Mitigation Approaches	14
3.1 Line Filtering	14
3.2 Power Reduction	14
3.3 Sharing the Wire—Coexistence for the same PLC Technology	15
3.4 Summary of Legacy Mitigation Technologies	17
4 G.hn's Neighboring Domain Interference Mitigation (NDIM)	18
4.1 G.hn NDIM Overview	18
4.2 G.hn NDIM Tools	20
4.3 Using the Tools	22
5 Comparing Same Technology Interference Mitigation	23
5.1 Legacy PLC Test Results	24
5.2 G.hn Test Results	26
5.3 Test Summary	26
6 PLC MIMO NN Mitigation in MDUS—a Unique G.hn Strength	27
7 Summary	28

1 Overview

The majority of the world's urban populations live in multi-dwelling unit (MDU) buildings. The need for high-speed networking within these often closely spaced units has resulted in a tremendous increase in the use of High Frequency¹ (HF) powerline communications (PLC). However, PLC networks within a given MDU building will be in close proximity to one another. Due to the nature of PLC signals going beyond the physical bounds of a living space, PLC signals from one network may be detectable on adjacent networks. These outside network signals may appear as valid PLC signals to the adjacent networks, which leads to substantial performance degradation. This cross-network interference is commonly known as Neighboring Networks (NN) interference. One PLC technology, G.hn's PLC mode, is unique in having advanced mitigation techniques that virtually eliminate the impact of NN interference.

NN interference is not only a problem for powerline systems. Any networking technology whose signals are not physically constrained² to a service area can also experience NN interference. Therefore, the same concerns apply to radio frequency (RF) based Wi-Fi networks, whose signals easily crossover between Wi-Fi networks. For PLC, signals can transfer to other PLC networks through inductive propagation or when the living units' power wires share common feeder lines.

PLC inter-network interference is a growing problem; for as the number of deployed PLC networks increases so does the probability of interference. Prior to G.hn, no comprehensive, robust solution was available to the NN problem. G.hn's Neighboring Domain Interference Mitigation (NDIM), defined in the global standard Recommendation ITU-T G.9961, specifically includes a set of techniques to address this. This paper describes the PLC NN problem, discusses previous largely unsuccessful attempts to resolve it, and introduces the NDIM enhancements to G.hn that overcome it.

Furthermore, there is a second set of concerns over closely deployed PLC networks related to their signals traveling outside the living space. If signals transmitted on one network can be received as valid PLC signals by users on other networks, there are issues of privacy and security. This is discussed in detail in a HomeGrid Forum White Paper, *Secure G.hn Networks*. G.hn security was defined to eliminate these problems so that even when G.hn PLC signals crossover to other networks, the user's data is secure from third parties decoding it, therefore content is protected and network access restricted.

¹ High Frequency, also denoted as Broadband or Wideband, PLC is used to denote PLC above 2 Megahertz. The term "PLC" is used throughout the document to mean HF PLC or when used as a generic term meaning all powerline communications according to the context of use.

² Not physically constrained means not restricted to a physical space or the network served area. As PLC and Wi-Fi signals can travel outside their service area, they are considered "unbounded" network technologies.

2 Neighboring Network Interference

A Service Provider will often deploy networks in many apartments in an MDU building. If these networks are PLC based and closely spaced, there is a high probability of substantial inter-network signal crossover. This crossover of signals is detected as interference in the local network when external (alien) PLC signals are on the wires. Such alien signals may just raise the noise floor on the line or appear as valid PLC signals. The problem is compounded when multiple Service Providers have PLC subscriber networks in the same building as well as by consumers installing their own PLC networks acquired through retail channels. Business rules and regional regulation generally preclude Service Providers from coordinating traffic in an attempt to mitigate the NN problem, therefore a means of automatically and equitably sharing the line needs to be used.

2.1 PLC signals travel between Units in MDU Buildings

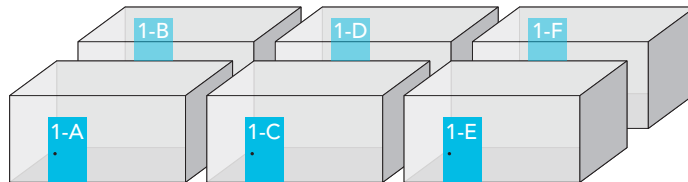
This section explains the underlying concept of neighboring networks interference.

PLC network nodes (transceivers) communicate using high frequency signals transmitted over a residence's power mains wiring. The signal power is generally sufficient to allow communication between nodes on any of the home's power sockets; meanwhile, PLC signals can also reach beyond the intended service area. Many PLC technologies transmit at the highest signal strength allowed, to overcome noise and ensure they can pass data at the maximum rate within their own network; however, this high power level means the signals are strong even when traveling outside their service area, which increases the NN problem for their neighbors' networks. The number of neighboring networks affected depends on the PLC signal strength, topology of the MDU wiring, and the attenuation between networks. It is quite common for PLC networks more than one floor away to detect signals from another PLC network.



» Figure 1: Three-story apartment building showing units

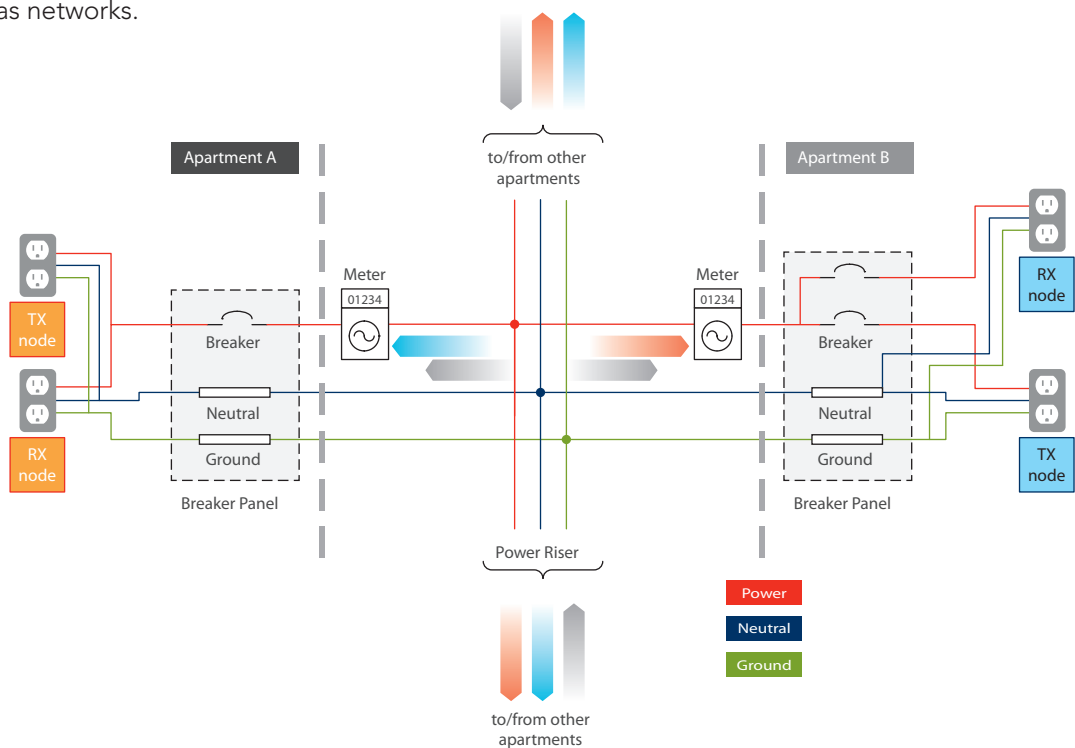
To assist the explanation, we have defined a typical MDU as a model. The simple figures (Figure 1 above and Figure 2 below) depict a three-floor apartment building with six living units on each floor. This example building is used in the following discussions to illustrate the concepts of neighboring networks' mutual interference, network performance impact, and mitigation techniques used by G.hn and legacy technologies.



» Figure 2: Layout of six units per floor

We are assuming that a Service Provider has installed a network using the same PLC technology in each tenant's unit. These networks will interfere with one another to an extent dependent on their relative physical location, potential for signal propagation between networks, and the electrical path the signals can take between networks.

The diagram below (Figure 3) shows how signals can crossover between networks via the in-building wiring. The Figure depicts meters in close proximity to apartments, which may or may not be the actual case, varying by building. Typically, there is 20 to 40 dB of attenuation between networks due to circuit breakers, meters, cable distances and topology ; however, this value will vary by individual nodes as well as networks.



» Figure 3: Typical PLC signals transiting between apartments over power lines

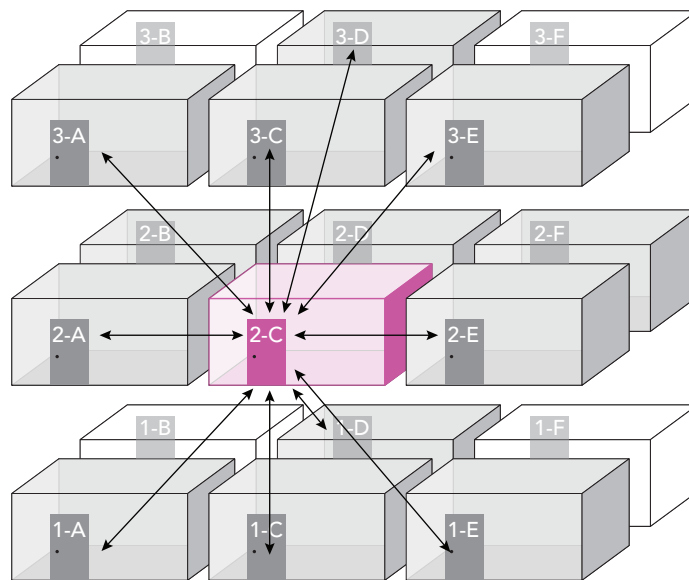
3 Per ETSI TR 102 269, the median attenuation @ 15 MHz between sockets in the same flat is 40 dB, while median attenuation @ 15 MHz between sockets in different flats is 60 dB.

For this example, we are assuming that power riser cables do not provide additional opportunities for signal propagation over and above those related to the network's physical location. From Figure 4 it can be seen that apartment C on floor 2 could be interfered by, and interfere with, the largest number of networks.

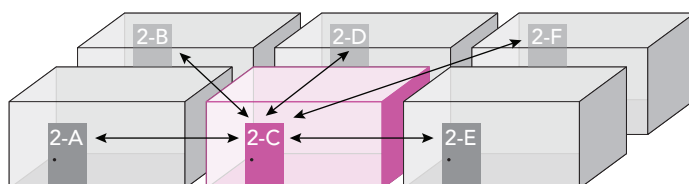
The interference from neighboring networks can be categorized as just raising the level of noise on the line (higher noise floor) or actually being detected as PLC signals in adjacent networks. When detected as PLC signals, they can be at either a high or low power level, depending on attenuation encountered. If a signal is detected in a local PLC network with enough signal strength to be detected as a PLC signal, yet with a low enough power level as to still allow communications in spite of it, this is considered "low level interference." If a signal from an outside network is of enough power to overcome local signals, this is known as "high level interference."

Low level interference may be 'ignored' to some extent by a local PLC network but throughput is impacted as it reduces the signal to interference and noise ratio (SNIR) of the effected nodes' signals. Also, local nodes can misinterpret these alien signals to be valid local signals, causing local signals to be lost. High power level NN interference cannot be ignored or effectively transmitted over, as it reduces the SNIR to unusable levels and blanks out local signals.

In the following figure (Figure 4) the unshaded living units denote units whose networks may or may not interfere in the real world with the network in apartment 2-C; but for simplicity, we are going to assume that their signals do not reach 2-C's network and vice versa.



» **Figure 4:** Interference between the network in 2-C and neighboring networks inter-floor



» **Figure 5:** Inter-network interference on the same floor between the network in 2-C and other units

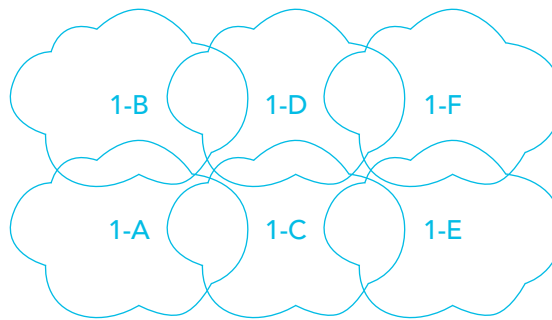
In Figure 4 and Figure 5 the arrows signify the probability for detectable inter-network interference occurring in relation to the network for apartment C on floor 2. In this case, there are 13 such networks (floor 1: units A, C, D, and E; floor 2: units A, B, D, E, and F; and floor 3 units A, C, D, and E).

2.2 Interference patterns

Each network in the MDU building experiences its own set of interference; therefore, each network has its own NN mitigation needs. The interference a network experiences is known as the network's interference pattern. Further, each node in each network has its own interference pattern, as each node has a distinct set of nodes from other networks that interfere with it and that it interferes with.

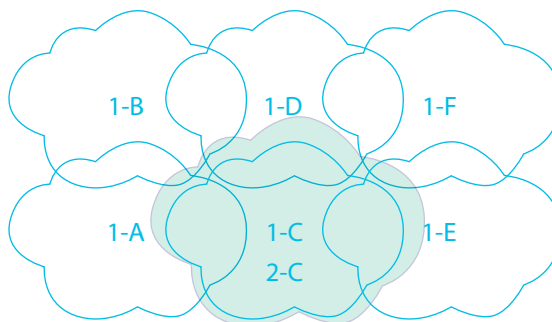
NN interference is time varying with respect to amplitude or even presence. A light switch turning on a circuit may power up a node in one network that then interferes with another network's nodes. On the other hand, a new network may power up for the first time, or an existing network power up after a power outage, dynamically changing (or creating) the other networks' interference patterns. The same applies to individual network nodes. These nodes not only experience NN interference when they are powered up, they also change the interference pattern for all networks and nodes that detect their signals. Further, it is possible that a node in a network may change its transmit power dynamically, depending on which node it is sending data to and line conditions.

The following diagram (Figure 6) shows the possible signal overlaps between networks on the building's first floor.



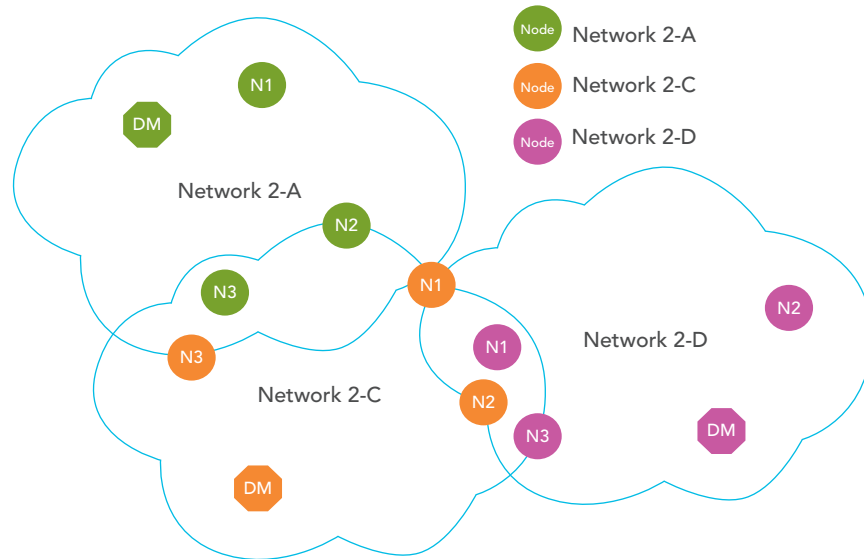
» Figure 6: Overlapping networks and their interference zones for the MDU's first floor

Adding network 2-C to this network view changes each network's interference pattern and complicates the process of mitigation (Figure 7).



» Figure 7: Network 2-C added with increased interference patterns

The following diagram (Figure 8) depicts how each node of a network may have its own unique interference patterns. Those nodes straddling another network's region may experience varying degrees of interference while a node such as N1 in network 2-D may experience higher interference from node N1 in network 2-C, and less interference from N3 in 2-C. Of course, the scenarios can become very complex with patterns overlapping patterns and nodes in and out of zones of interference. Further, interference might not be symmetrical, i.e. a node (e.g. N1 in 2-D) may cause interference to some nodes in 2-C, but experience interference from a greater number of 2-C nodes.



» Figure 8: Interference pattern of nodes

2.3 Symptoms

NN interference can be detected at network start up, or sometime later as it may develop over time as other networks' nodes are installed, or other networks' nodes' transmission power is increased to overcome line conditions.

Clearly, the first PLC network installed in an MDU building initially suffers no NN issues. However, as the number of PLC networks deployed in the building increases, neighboring network interference increases and service could deteriorate, with resultant service calls and consumer dissatisfaction.

The end user may experience NN interference as vastly reduced throughput, pixelated or frozen IPTV content, or even loss of service at specific nodes or even network wide.

Such NN interference can result in the dispatch of service personnel and, once the cause is determined, either the PLC network is removed or mitigation techniques are implemented.

Network interference mitigation techniques for legacy PLC technologies can be relatively expensive and only partially eliminate the problem, whereas G.hn PLC networks automatically handle interfering G.hn networks.

2.4 Interference Mitigation between Different PLC Technologies

When different PLC technologies are in close proximity they cannot inter-communicate under most circumstances. Technologies that do not intercommunicate and do not coexist tend to harm each other's signals to the point that, at best, only minimal network traffic gets through.

Under certain circumstances, specific PLC technology networks may be able to coexist by taking turns to use the wires for transmissions; however, this is a trade-off of 'time on wire' versus interference. The following table (Table 1) shows several PLC technologies where coexistence is possible, their effect on one another, and if they can or cannot coexist. Older PLC technologies do not coexist with other PLC technologies and therefore have not been included, for example, HomePlug (HP) AV. However HP AV's successor technology, based on IEEE 1901 FFT, does have coexistence specified as mandatory. The relatively new HomePlug AV2 specification includes the same inter-system protocol (ISP) as 1901 FFT, and is included in the Table, as it is anticipated that implementations of AV2 that are fully compliant with the HP specification will have the ISP coexistence properties.

Recommendation ITU-T G.9972 (also known as "G.cx") defines a means for G.hn to coexist with the two distinct IEEE 1901 PLC technologies using ISP.

» **Table 1:** Coexistence Attributes for Major PLC Brands

	1901 FFT/AV2 (Note 1)	HD-PLC 1901 Wavelet compliant (Note 1)	UPA	G.hn with G.cx (Notes 2, 4)
1901 FFT/AV2 (Note 1)	Yes (using ISP or legacy NN mitigation)	Coexists if 1901 FFT have ISP	No, will interfere with one another	Coexists with 1901 FFT if ISP present
HD-PLC 1901 Wavelet compliant (Note 1)	Coexists if both have ISP	Yes, if network controllers directly communicate	No, will interfere with one another	Coexists with 1901 Wavelet if ISP present
UPA	No, will interfere with one another	No, will interfere with one another	Yes	Note 3
G.hn with G.cx (Notes 2, 4)	Coexists with 1901 FFT if ISP present	Coexists with 1901 Wavelet if ISP present	Note 3	Yes

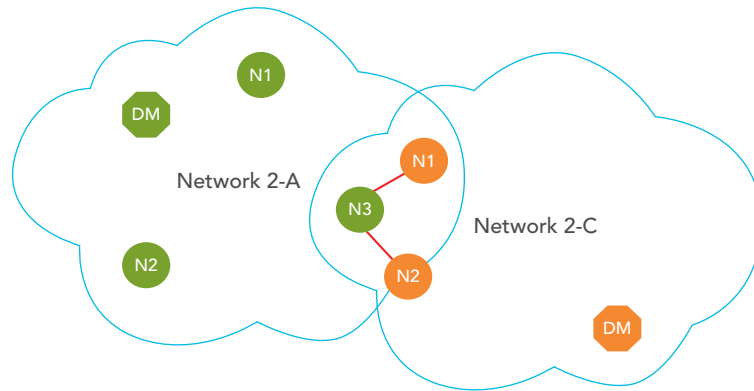
Notes:

1. The IEEE 1901 standard mandates that ISP be in all conforming 1901 products. To date, we are unaware of its being included in silicon that is claimed to be IEEE 1901 FFT compliant; therefore, coexistence problems between the two 1901 technologies and between IEEE 1901 FFT and G.hn with G.cx may exist.
2. G.hn devices require implementation of the optional G.cx technology to coexist with 1901 devices.
3. Marvell has announced a coexistence capability between their G.hn silicon and UPA-based devices.
4. Sigma Designs and Marvell have announced G.cx support in their respective G.hn silicon implementations with Sigma Designs also supporting coexistence with HP AV in their G.hn silicon

The coexistence functionality of IEEE 1901's ISP and G.cx works well in a relatively sparsely populated (i.e. small number of networks in close proximity) scenario. If the number of coexisting networks rises to more than two of the same type, there are potential coexistence issues due to limitations in ISP definition.

2.5 Interference between Networks of the same PLC Technology

Closely spaced PLC networks of the same type will also encounter interference.

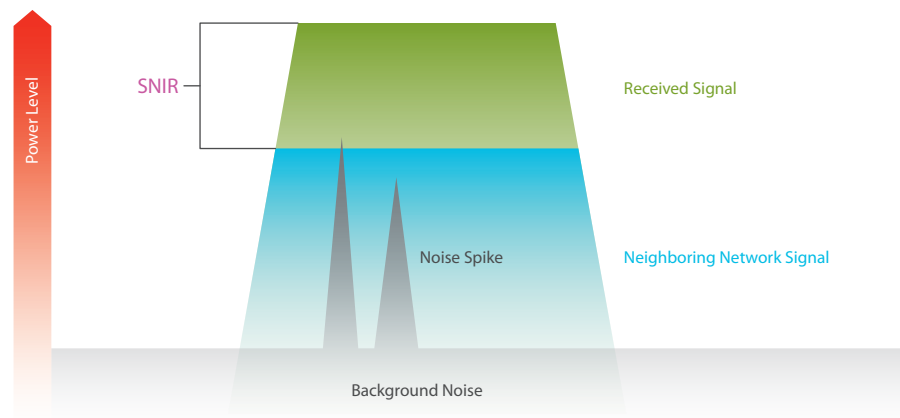


» Figure 9: Inter-network interference detection

In Figure 9 above, networks 2-C and 2-A are neighboring networks of the same technology.

Node N3 of network 2-A can detect signals coming from two nodes in network 2-C. Detection by N3 of the network 2-C nodes' signals may be as valid PLC signals or as raising the noise floor.

The following figure (Figure 10) depicts what could be the power level of received signals at N3 and how an even moderate level of neighboring network nodes' signals can affect the performance of N3. The SNIR is what determines the ability of a node to detect communications effectively from other nodes in its network. The lower the SNIR the less likely the node may receive error-free transmissions, or be able to receive transmissions at all.



» Figure 10: Moderate NN signal impact on SNR of node N3

2.5.1 PLC Interference Levels

When the NN PLC signals are attenuated to such a degree that they cannot be determined to be valid PLC signals, the local network's nodes only recognize such NN PLC signals as noise, raising the local nodes' noise floor and reducing the signal to noise ratio (SNR) of the local nodes, which may effect their throughput and ability to overcome other noise. With NN PLC interference at a power level such

that signals are only detected as noise, individual PLC technologies handle this in their designs with varying degrees of success.

The handling of PLC signals as noise is just as if the background noise was higher. It is only when the NN PLC signals are detected as valid PLC signals that true mitigation techniques be used. There are two classifications used for NN PLC signals detected as valid signals, “low level” and “high-level” interference. These terms correlate to relative signal power levels of the NN PLC signals.

2.5.1.1 Low Level Interference

When signals are detected as valid PLC signals, but at a substantially lower power level than local signals, this is considered “low level interference” (see Table 2). Low level interference can result in a very low SNIR for local signals, causing high error rates and reduced throughput. From a service perspective, this can result in pixelation, lost or frozen IPTV frames, and disconnected VoIP calls, with even total loss of connectivity experienced by nodes or even whole networks. The effect is worse during high network traffic periods, such as in the early evening when many are home watching the news or favorite shows.

Low interference is determined by several factors and is system design and implementation specific. For example, low interference may be considered generally when the difference in the power of the local signals versus neighboring signals is approximately 25 dB or more; however, this value is used here as a general example, only.

G.hn determines the mitigation steps to use on a per node, then per network basis; with low level interference mitigation automatically invoked using orthogonal preambles after detecting NN interference. See Section 4.2 for a discussion of orthogonal preambles and how this benefits network performance over a majority of NN interference cases.

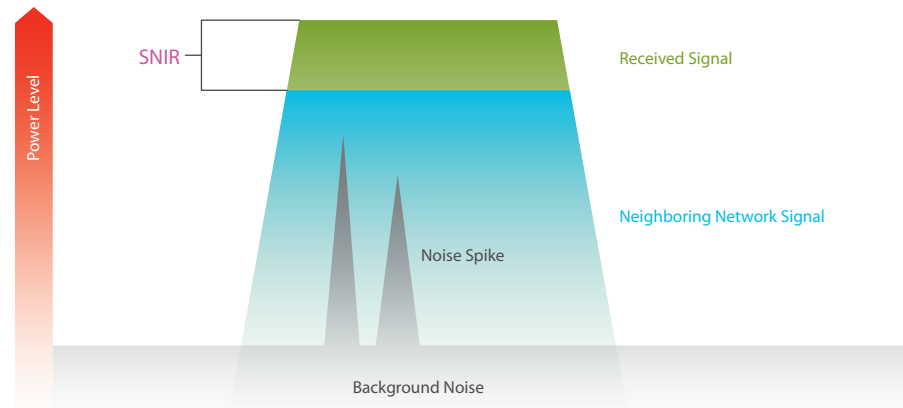
Depending on statistics, such as those in ETSI TR 102 269, between 85% and 95% of NN interference cases fall into the low level interference category. Resolving this case generally resolves the vast majority of instances.

2.5.1.2 High Level Interference

When the signals from neighboring domains are at a power level roughly the same as the local signals, this is known as “high interference.” High interference will result in loss of local PLC signals and disruption of service. High-level interference is defined to occur when the attenuation of alien PLC signals is roughly the same as local PLC signal attenuation, for example, the NN signal attenuation is 20 dB or less versus local signals. Local signals may attenuate 20 dB while transiting the premises’ cabling, thus a neighboring network with 20 dB attenuation of its signal would appear nearly the same as a local signal. Further, the 20 dB or lower SNIR deteriorates communications between nodes, reducing their throughput and QoS substantially.

The actual determination of low versus high interference is dependent on the system design and implementation, therefore, this 20 dB value will vary from design to design, but such design differences are not crucial as each node, based on its design, reports to its DM if it is experiencing high or low level interference based on the node’s capabilities and service requirements and then the DM determines what actions to take.

When interfering signals are at a high level, the interfered networks' nodes' ability to pass and receive data deteriorates and they must use some form of mutual mitigation, or they will operate with minimal throughput, poor QoS or even lose connectivity (see Figure 11). High level interference may be network-wide or only at specific nodes. When the interference is network-wide, it may cause PLC networks to crash and not connect. If at specific nodes, these may be unable to detect valid incoming signals and so may either drop out of the network, or generate a flood of messages causing a substantial loss of throughput, or even a network crash.



» **Figure 11:** High NN interference with reduced SNR

2.5.1.3 Summary of Inter-network Interference Levels

The following table summarizes the interference levels and what, if any, mitigation is possible.

» **Table 2:** PLC inter-network interference levels

Interfering PLC signals	Detected as valid PLC signal	Mitigation
Raise noise floor only	No	Handled as noise by PLC nodes according to their technology's capabilities to handle noise
Low signal power (low level interference)	Yes, but lower than local network signal power allowing enough SNR for intra-network communications	Depending on technology, legacy PLC has partial mitigation, using a time division scheme for time on wire
High signal power (high level interference)	Yes, with similar power to local network signal power	Most legacy networks cannot mitigate, or only use a time division scheme for time on wire

The level of interference amongst the nodes in any domain will vary according to their individual interference patterns. Some nodes may only detect a raised noise floor; others may detect low-level interference, while some may detect high-level interference. The ability to mitigate NN interference based on each node is crucial to maximize throughput and resolution of the problem while eliminating the negative effects of the interference.

2.6 Deployment Considerations

Legacy PLC technologies are unable to be upgraded to handle neighboring networks. As many installed networks are incapable of coexisting with other technologies or even like networks, there is an argument against backwards compatibility with them, as these legacy networks will always interfere with their neighbors.

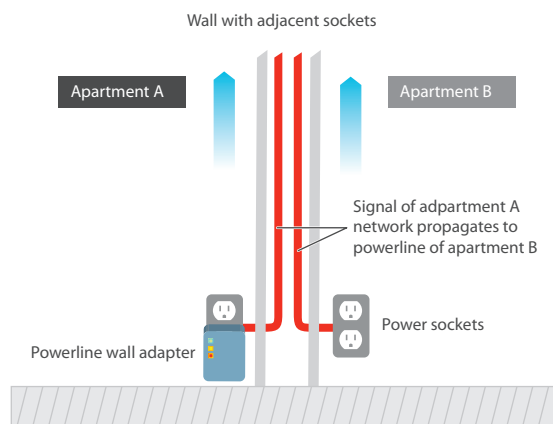
The best way for a Service Provider to avoid this is to decide on a single technology for PLC networks and one that can handle NN issues; thereby providing robust service to end users regardless of the number of other networks nearby. G.hn PLC is this technology.

3 Legacy Mitigation Approaches

PLC technologies other than G.hn have either no NN mitigation capabilities or very limited ones. This section discusses interference mitigation techniques used by non-G.hn PLC technologies.

3.1 Line Filtering

In some regions, powerline filters are manually installed to block the signals on the wire from transiting from one apartment to another. The filter placement is critical; it must be installed where the residence power lines enter/exit the unit. This is expensive in technician time and material's cost, with costs varying from region to region and depending on the type of mains filtering device installed. Moreover, it does not fully block the inter-network signals, as can be seen from Figure 12. PLC signals in closely located cables will still propagate through induction, e.g. from the node in apartment A to the wiring of apartment B. There is no way to prevent this coupling through filter use.



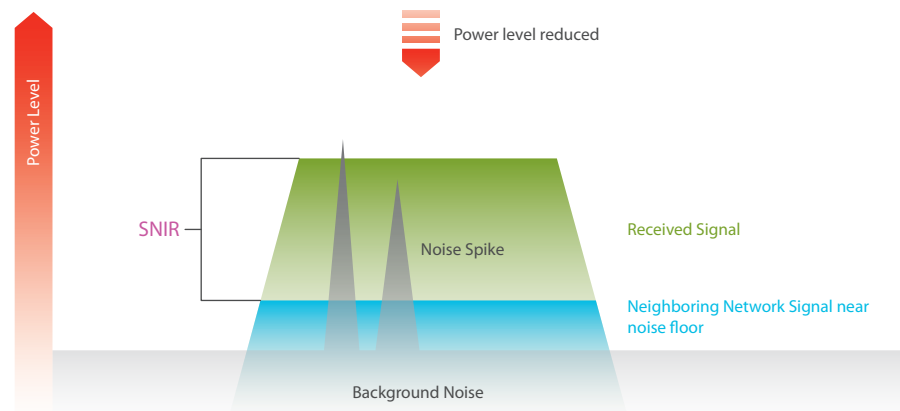
» **Figure 12:** Signal propagation from one apartment's network onto a different apartment's wires

3.2 Power Reduction

Another method is to reduce the power of each node's transmissions. While this somewhat reduces inter-network interference, it also reduces the SNR and SNIR for local PLC signals, and if not managed properly, there is a higher probability of noise interference, resulting in noise levels that can cause partial, intermittent or complete throughput failures.

The reason that power reduction works for very low interference scenarios and not for higher interference scenarios is that if a node causing interference reduces its transmit power too much, it cannot reach nodes in its own network. If it does not reduce its transmit power significantly then its transmissions may collide with neighboring network nodes and cause interference between networks. This is an acceptable approach if there is no significant NN interference and only very low noise levels are present. See Figure 13.

In practice, this technique is used for both high and low interference level scenarios, although it only has limited success in low interference scenarios and none at all in high interference scenarios.



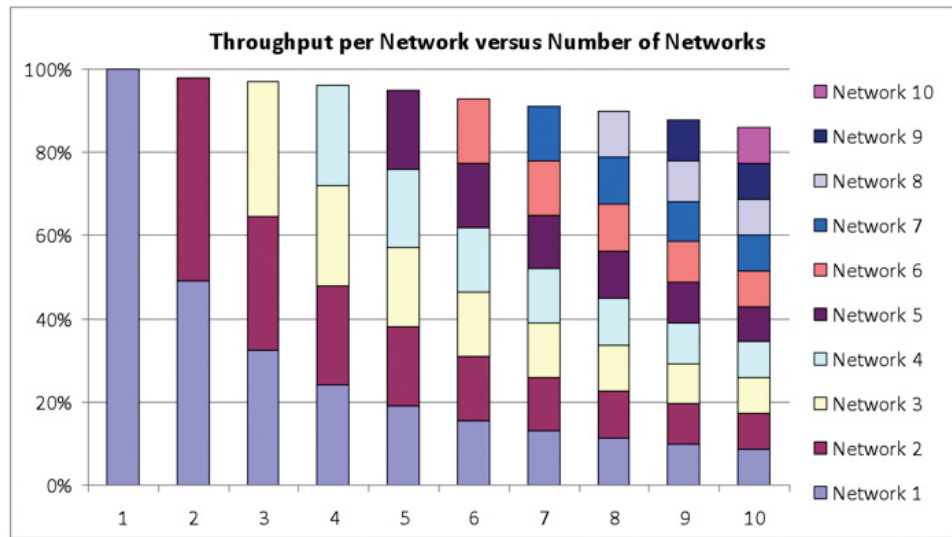
» Figure 13: Signal power reduction reduces interference, so reduces overall SNR and noise robustness

3.3 Sharing the Wire—Coexistence for the same PLC Technology

Some PLC technology vendors have adopted a set of “sharing the wire” approaches when inter-network interference is detected as valid PLC signals. To accomplish this, the networks establish a communication channel between their network master nodes. This channel is either direct between masters, or uses nodes as proxies, and requires that valid PLC signals be able to pass between networks. If signals are unable to be passed effectively, there is no coordination.

Note: G.hn overcomes this through the inter-network (inter-domain) messaging being more “robust” than normal traffic; this way the messages are more likely to be detected and useful to the receiving network. This means that, even in the case of low power PLC signals passing between networks, the G.hn networks will be able to co-ordinate.

This sharing by non-G.hn technologies is quite similar to the coexistence schemes between disparate technologies i.e., networks take turns on the wire. This has obvious drawbacks. For example if there are five networks vying for time on the wire, they average a maximum net throughput of 20% each, less any intercommunication overhead and throughput lost to CSMA collisions. This could be problematic where network performance is already low. For example, with only a maximum of 20% of network capacity available, a network may not have enough bandwidth to carry HD content effectively (see Figure 14).



» **Figure 14:** Average throughput per legacy network given number of networks interfering

Rather than simply having equal sharing of time on wire, an alternative approach is to do the sharing such that the high priority traffic gets through, e.g. by allocating more time on the wire to specific traffic with high QoS requirements. Therefore, one of the five networks could get more than its 20% fair share. This scheme presupposes that the networks are not set to claim more than their fair share. However, in such a scheme, a single network master acts as the central coordinator and could gain advantage over others. In this scenario a rogue network could take more than its fair share and this direct coordination between networks, with a network able to garner more time on wire for itself, may be contrary to business or regulatory rules.

Assuming the best case, optimal layer 2 throughput for a single legacy network is 250 Mbps (see Figure 18), and the service required for service to any node in the network is 35 Mbps layer 2 throughput, the maximum number of neighboring networks is 6, otherwise the metric is not met. Given that the 250 Mbps number is more theoretical (found in the lab) than real world, the likelihood is that a typical legacy PLC technology layer 2 throughput would be somewhere between 35 and 80 Mbps for one network. This means only one network alone or two networks interfering would meet the metric. Above two networks, the throughput is too low to meet service requirements.

Further, if the aggregate high priority (e.g. video) traffic across all networks is more than the available capacity, the networks default to a CSMA-based line access on a "first come, first served" basis. The devices needing to transmit have to listen to the wire, detect a quiet period of a specific duration, and start to transmit. This frequently results in data packets from different signals colliding on the wire, which in turn results in lost data, high latency, IPTV disruption, etc. The probability of collisions increases under more interfering networks and heavy traffic, resulting in more collisions, more retransmissions, and thus a deterioration of aggregate throughput, just when needed most. This method is not a good choice for video or TDMA traffic.

One legacy PLC technology has settled on an approach of centralized TDMA scheduling in an attempt to overcome this issue. Centralized TDMA reservation occurs when one network master instructs another network's master (or many network masters) to dedicate timeslots for the first network's TDMA traffic, which could lead to abuse, and reduces the time on wire for other networks. It does not work well in a multi-network interference scenario, or under heavy traffic.

Moreover, for non-priority traffic, the scheme is to allow all devices in all networks to vie for the wire using CSMA-CA line access during a specified period. This not only causes service problems, it also deteriorates the aggregate throughput as each time a device has a collision or must wait and listen for the wire to be quiet, there is loss of time for transmission, which equates to an overall throughput decrease.

Finally, there is the issue noted earlier regarding coordinating PLC networks from different Service Providers, and the associated business implications and regulatory issues. A rogue network could become the centralized controller and notify the other networks that all its traffic is top priority and thereby dominate the other networks, and capture substantial time on wire. In theory, the Service Provider with such rogue networks could get an advantage over the other PLC network providers. Of course, this could lead to all networks being rogues and misrepresenting their priority levels causing the situation to deteriorate into a pure CSMA scheme.

3.4 Summary of Legacy Mitigation Technologies

In summary, the inter-network interference issue between disparate legacy PLC technologies is likely to remain unresolved. Legacy PLC technologies do not support coexistence and only recently has some PLC silicon arriving on the market had any coexistence mechanism support. This means that Service Providers need to determine their best way to mitigate NN issues going forward.

The following tables (Table 3 and Table 4) summarize the issues raised when existing non-G.hn PLC networks based on like technology interfere in a NN scenario.

» **Table 3:** Low and High levels of interference versus Legacy Mitigation Techniques

	Low level interference	High level interference	Notes
Powerline filters	Works acceptably, but has relatively high cost in labor and materials	May work to some degree	Such filters cannot stop inductive propagation between wires from different units
Low transmit signal power	Works acceptably; however, disrupted by normal noise levels and when the number of interfering networks increases	Does not work very well even when every network has low levels. Also, levels may need to be so low as to make intra-network activity prone to errors or unsustainable	Limited value option when used by itself and reduced SNR/SNIR leads to data errors and service issues
Sharing the wire	May work acceptably for 2 or 3 networks, but problematic under heavy traffic and with more than 3 networks	May work acceptably for 2 or 3 networks, but problematic under heavy traffic and with more than 3 networks	CSMA rules time on wire under heavy traffic which is problematic for TDMA traffic. Priority traffic causes disruption and severe service impact, Rogue networks can gain too much time on wire
Combination of all three—filters, low signals, and sharing	Relatively high cost, low signal levels lead to disruptions and errors, while sharing reduces each network's time on wire	Relatively high cost, propagation still occurs, low signal levels lead to disruptions and errors, while sharing reduces each network's time on wire	The higher the number of networks, the worse the situation. After relatively few networks, does not work well

The following table summarizes the high-level pros and cons of the three legacy mitigation techniques.

» **Table 4:** Mitigation Techniques — Pros and Cons

	Pros	Cons
Powerline filters	Blocks PLC signals from leaving the premises via feeder wires	Does not stop inductive propagation, Installation is expensive and requires skilled personnel
Low transmit signal power	Reduces possible impact on other networks	Leaves own network subject to possible higher error rates and service issues
Sharing the wire	Allocates time on wire fairly, by priority, or using CSMA	Can be abused, unworkable under very high-density situations, defaults to CSMA where every packet battles for time on wire. Vastly reduced and uneven throughput per network causing services such as HD content to have problems

Note that all these legacy mitigation techniques are applied on a whole network basis, they cannot be confined to the small number of nodes which might be experiencing interference. The overall performance of the network may suffer as a result. Further, when one or only a few nodes in a network experience NN PLC interference, the network may not invoke mitigation, thus leaving the effected node (or nodes) to deal with the interference.

4 G.hn's Neighboring Domain Interference Mitigation (NDIM)

This section discusses the standard (mandatory) G.hn Neighboring Domain Interference Mitigation (NDIM) functions of all G.hn nodes. Please note that the term "domain" denotes a single G.hn network on a specific medium.

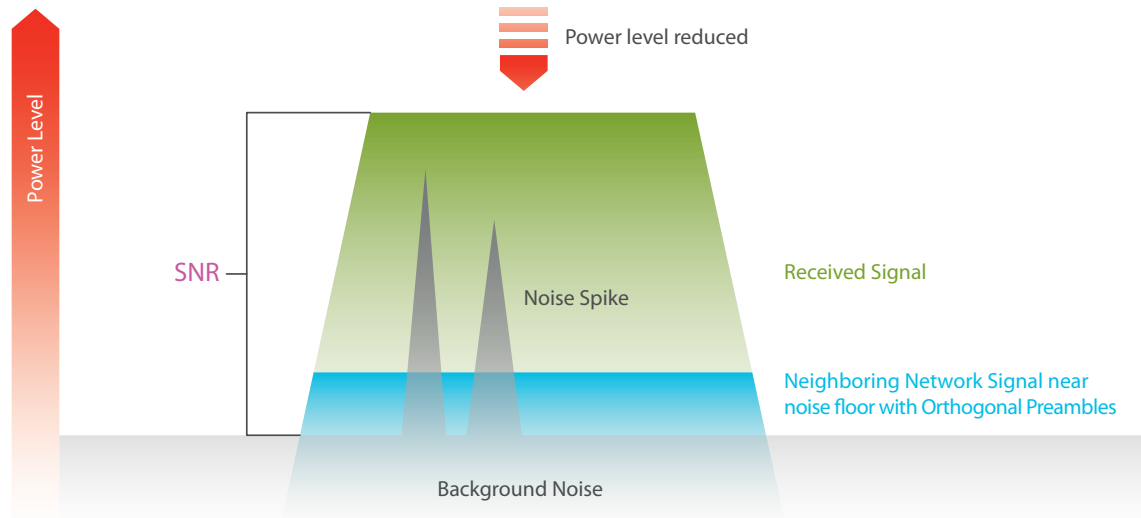
Among the ITU-T G.hn suite of standards, NN mitigation is in the core Recommendation ITU-T G.9961 standard.

4.1 G.hn NDIM Overview

NDIM provides a comprehensive toolbox with different tools for different interference scenarios, starting from low level to high level interference between neighboring networks. The specific tools are as follows:

- **Adaptive power control:** This technique works if the PLC interference power level is very low compared to local signals. This is similar to the technique supported in legacy PLC devices however with intelligent signal power management, where Power Spectral Density (PSD) levels are modified on a per node and per connection basis, as needed.
- **Orthogonal preambles:** This technique is appropriate for low level interference, enabling the local nodes to treat the NN interference as additional noise on the wire without significant impact on throughput (see Figure 20 and Figure 21 for actual throughput screen shots depicting this). The goal in this approach is to "ignore" neighboring transmissions by tuning the node's detectors for the preambles of the node's own domain-specific signals. Coupled with adaptive power control, this can effectively "hide" (separate) networks from one another (see Figure 15). Legacy PLC technologies do not have an ability to use a network-specific preamble to mask their signals. They cannot avoid

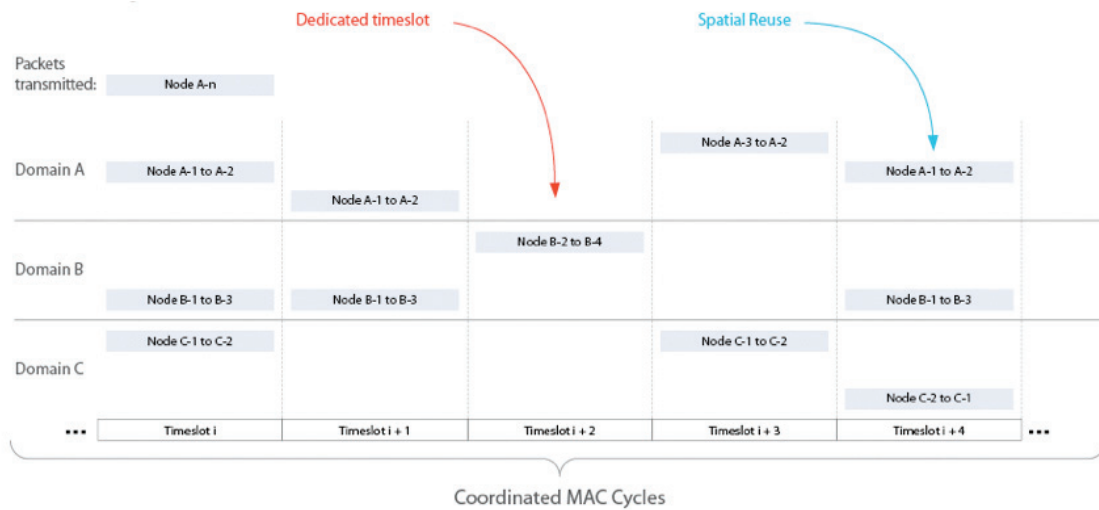
the interference and, by detecting and processing transmissions from neighboring domains, their nodes possibly miss the transmissions of their own network, increasing the negative impact of low level interference, therefore they resort to the methods discussed in Section 3 (above) and drastically reduce their throughput



» Figure 15: G.hn Orthogonal Preambles with adaptive power control masks NN

- **MAC cycle alignment and inter-domain coordination:** To counter high level NN interference, G.hn has tools defined to automatically coordinate network traffic between domains, without the legacy PLC approach of resorting to a centralized controller selected to make routing and time on wire decisions for multiple networks. Instead, G.hn allows for neighboring domains to:
 - Detect one another
 - Have a means of coordinating their timeslot activities in all networks so that their transmission allocation periods are aligned (these periods are known as MAC Cycles, which span 2 ac power cycles and consist of timeslots allocated for nodes to transmit in)
 - Provide a common signaling timeslot for domains to exchange interference and coordination information
 - Vector traffic in a fashion to avoid interference by each specific node

These techniques allow per node — per domain determination of time on wire (timeslots) so that spatial reuse is possible in the case of low or no interference for the transmitting and receiving nodes. Dedicated timeslots can be assigned when inter-domain interference from specific node(s) in one or more neighboring domains is high. This applies to both the transmitting and receiving nodes' interference patterns; so that both transmit and receive sides of the timeslot are taken into consideration when determination of spatial reuse is made. See Figure 16.



» Figure 16: Coordinated MAC Cycles with examples of dedicated and spatial reuse timeslots

4.2 G.hn NDIM Tools

The G.hn standards have several clauses that specify how to handle NN interference. These define a set of signaling steps, time on wire management algorithms, and coordination protocols that in aggregate provide tools for interference mitigation, from low-level interference to high-level interference. These tools address the scenario of high-density MDU environments where many networks cause mutual interference, where the number of networks is quite high, requiring a breadth of capabilities to enable mitigation of the aggregate interference.

The ability of G.hn nodes to dynamically adjust their signal levels is a standard feature of G.hn, it is applied in this case to reduce NN interference, yet was not added specifically for mitigating NN interference. Instead, it was defined to ensure as low an electro-magnetic 'footprint' of emissions from a G.hn domain was possible at all times. The fact that the PSD control is dynamic and can be lowered as part of the network coordination steps ensures a more complete mitigation of the NN interference from specific nodes.

The tools in G.hn to handle NN interference specifically are as follows (see Table 5 for a summary):

- 1 Domain specific seeds to generate orthogonal preambles. Orthogonal preambles differentiate domain signals and enable nodes from different domains to "ignore" signals that are not from their own domain. This is especially useful in low level interference scenarios where non-local PLC signals are detected and local nodes can set their filtering to ignore these alien signals. As stated earlier, this is used when the attenuation of neighboring signals at individual nodes is higher than the local signals. Domains do not have to coordinate the use of orthogonal preambles; this is an automated function of G.hn domains with nodes experiencing any interference. This step is used alone (low level interference detected even by a single node in a domain) or in conjunction with the coordinating steps following, as mitigation, while network-wide, is based on the interference experienced per node.
- 2 An inter-domain signaling window (IDSW) timeslot is defined to be in every MAC cycle. The IDSW contains a predefined inter-domain presence signal (IDPS), which allows fast detection of

neighboring domains. Any nodes in a G.hn domain can detect other neighboring G.hn networks' IDPS, which they report to their domain masters (DMs). DMs are the controlling entity for each domain, setting time on wire allocations and managing the network overall. DMs use IDPSs as a means to align their MAC cycles with neighboring domains. By having the signaling window at a known point in a MAC cycle, the IDPS can be detected readily by any node in any neighboring network.

- 3 Aligning MAC cycles among neighboring networks. By aligning MAC cycles, the neighboring networks are able to establish common times on the wire for coordinating transmissions from and for specific nodes and services. DMs automatically align MAC cycles among neighboring networks, even if the interference is low, in order to quickly deal with newly added nodes or domains that may cause high interference.
- 4 G.hn DMs exchange (or trade) timeslots as well as share them for either TDMA or CSMA traffic in a decentralized fashion. This is well defined and controlled and not under the control of a single entity. Centralized TDMA reservation control between networks is NOT supported in the standard, as this may enable rogue network controllers to gain more time on wire for their network and it may be contrary to business or regulatory rules. The G.hn approach prevents rogue networks claiming more than their fair share, ensuring equitable sharing of time on wire, without the need for exchanging customer data or allowing a single entity to gain control of all networks.
- 5 A common inter-domain communication channel (IDCC) established between domains for exchange of inter-domain management messages. IDCC is for management messages and interference measurements data exchange between networks. The exchange of information related to NDIM enables domain masters to manage traffic in order to maximize use of wire time and resources based on traffic scheduling using advanced algorithms. The interference pattern reported by each node enables each DM to determine the best allocation of time on wire for its own nodes. The scheduling algorithms are extensible, allowing for upgrade and modification to address NN situations in a more customized and specific fashion. This extensibility makes G.hn more future-proof than any other option.
- 6 DMs exchange the list of interfering and interfered with nodes, and a grid of allocations and forbidden zones for transmitters is created in a distributed fashion. This distributed means eliminates the "centralized" mode where one domain master might dictate the behavior of other domains. The key feature of this algorithm is that the grid of allocations is established for each node that suffers interference from nodes in neighboring domains or causes interference to nodes in neighboring domains. Spatial reuse is possible in most cases where specific nodes in different domains may be able to transmit or receive at the same time. Spatial reuse enables such concurrent non-interfering (or low interfering) traffic across all networks or a subset of networks, maximizing the aggregate throughput of all networks, while reducing the amount of time dedicated per network or node to a minimum. This is especially important in deployment scenarios that have overlapping and numerous NNs, where each network has its own set of high and low level interfering networks.

- 7 Based on (1 and 6), DMs assign their respective domains unique orthogonal preambles so no two interfering networks use the same preamble. This way, two or more networks can use the channel simultaneously whenever possible (e.g., when interference level is low) while other DMs may time-share the wire (e.g. when interference level is high). Furthermore, the grid of allocations is valid for a short period (about 4 seconds, per allocation period) during which new interference information is collected and the grid of allocations updated. This allows adapting quickly to interference from new nodes that join the G.hn networks, adapting to any changes in interference from existing nodes, or when new domains appear and change the interference patterns. The transition from one grid of allocations to another is seamless across domains using synchronizing counters.
- 8 Regular, non-DM nodes can act as detectors and proxies for their DM in finding and communicating with neighboring networks. The use of all nodes to detect neighboring domains ensures a comprehensive view of the interference pattern by the DM. Meanwhile, a DM may use proxy nodes to pass information to the interfering networks' DMs so that MAC cycles are aligned and time on wire is coordinated, if the two DMs cannot communicate directly.

» **Table 5:** G.hn NDIM versus Low and High Levels of Interference

	Low level interference	High level interference
1 Domain specific preamble seeds	Used to ignore signals from other G.hn networks and treat them as noise	Used whenever possible to ignore signals from other G.hn networks and treat them as noise
2 Common Signaling Window (IDSW) and Common Channel	Used to detect signals and receive messages from other G.hn networks and attempt to align the MAC cycle with them	Used to detect signals and receive messages from other G.hn networks and attempt to align the MAC cycle with them
3 Align MAC cycles	Enables rapid adjustment to changing interference patterns, especially when new nodes appear with high interference patterns	Used to enable spatial reuse and manage time on wire by exchange of coordination messages
4 Decentralized TDMA reservation	Not needed	G.hn domain masters, in a distributed fashion, exchange information and allocate timeslots between their respective domains for TDMA or CSMA use without the need for a centralized controller function, as used by legacy technologies
5 Coordination Protocol	Not used	Used for coordinating time on wire and passing lists of nodes with interference pattern data
6 Lists of nodes interfering and interfered with	Not used	Enables spatial reuse and establishing time on wire for dedicated use
7 Allocation periods	Not used	Helps DMs maintain an accurate record of interference patterns per node for timeslot assignments and inter-domain coordination
8 Nodes as detectors and proxies	Used to continuously detect and characterize interference, and attempt to contact neighboring G.hn networks	Used to continuously detect and monitor interference pattern for any changes and act as proxy between DMs

4.3 Using the Tools

Each domain master determines the interference level on a per node basis in their domain, and automatically invokes the appropriate set of mitigation techniques to be used.

Upon detection of any interference, the domain master will set that domain's preamble to a unique setting to eliminate any low level interference for its nodes and any neighboring domain nodes (orthogonal preambles).

If a neighbor network's inter-domain presence signal is detected, the domain aligns its MAC cycle to the neighbors' cycles, thus enabling any inter-domain coordination.

If even one node in a domain is experiencing a high level of interference, then the domain takes the mitigation steps for high-level interference i.e. coordinate the time on wire for the nodes experiencing high interference, and for those that are generating high interference into other domains.

Frequently, domain masters cannot communicate with one another directly, nor detect all of the surrounding domains, thus all nodes in domains are proxies for NN detection and those with the best links to other domains, as inter-domain communication nodes. This use of proxies ensures that a domain, as discussed in Section 4.2 above, with even a single node experiencing interference is able to detect the interference, report this to its DM, which in turn sets the preambles, communicates with the other domain(s) and coordinates traffic.

5 Comparing Same Technology Interference Mitigation

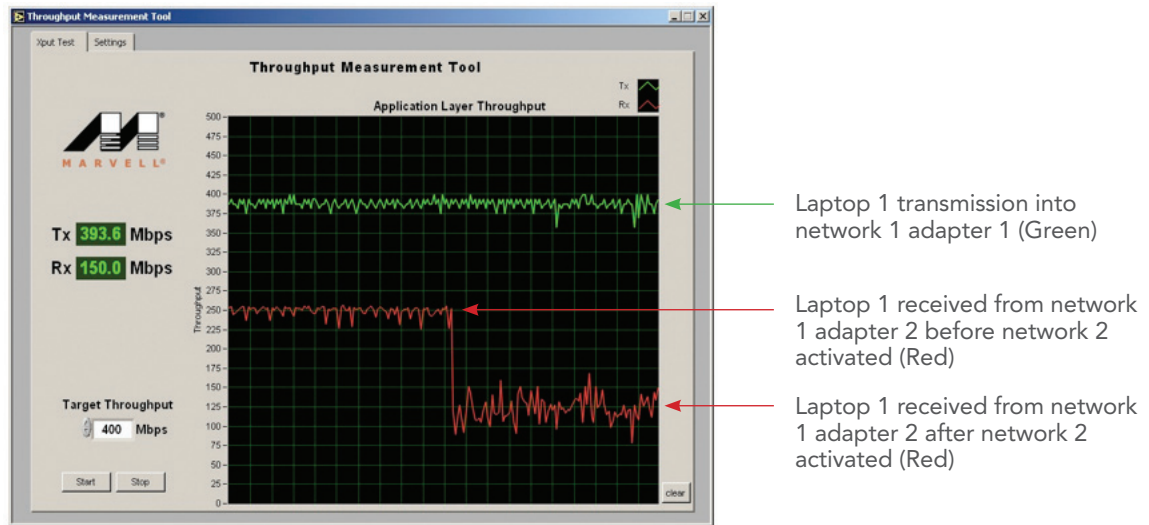
This section provides a comparison between a legacy PLC technology and G.hn in regards to mitigating low level NN interference between PLC networks of the same technology. As stated earlier, low-level interference is the majority of NN interference encountered; therefore, this test was made to compare the low-level interference approaches and results of G.hn PLC mode and a legacy PLC technology currently available in retail channels.

The test was done in a fashion that can be repeated by any interested parties. The test set up is according to a standard methodology followed by PLC test labs. The tests were run several times with adapters moved to different sockets and pairings to ensure that no variables existed to effect results.

The table below describes the test set up shown in Figure 17.

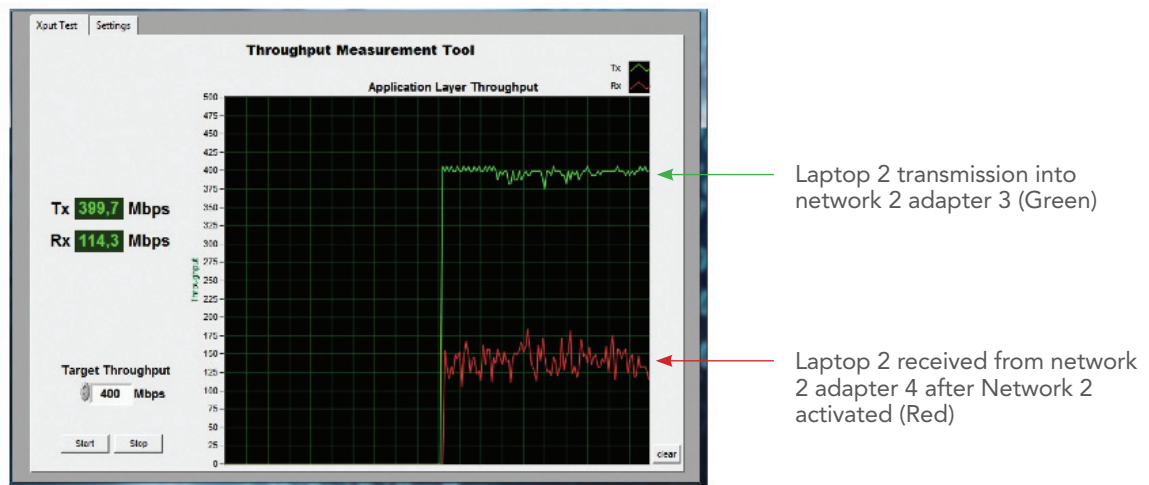
» **Table 6:** PLC NN Interference Test Set Up

Item	Description	Notes
Laptop 1 & 2	Traffic Generator and Recorder per network	Generating UDP Ethernet traffic at approximately 435 Mbps over each network, from one modem to the other in the respective networks. Each laptop has two Ethernet interfaces at 1 Gbps per interface
PLC Adapters	G.hn modems from Marvell or legacy PLC devices	Either all PLC modems were G.hn or legacy for both networks. G.hn modems are off the shelf evaluation kits from Marvell for their G.hn product. Legacy PLC modems acquired through the retail channel, mid-2013
Adapter Differences	G.hn adapters running at 80 MHz SISO mode, while the legacy PLC devices were running at 68 MHz SISO mode	While frequency ranges and throughput differ between the adapter types, this is not relevant to the test results which were only for verifying operation under NN interference, not to showcase the better throughput of G.hn
Attenuation intra-network	20 dB between adapters, which is typical for lab tests	Each network set at 20 dB between its adapters by use of two 10 dB attenuation units
Attenuation inter-network	Using a variable device, attenuate signals by 40 dB, nets a 60 dB drop between neighboring nodes which results in a 40 dB differential between local and alien node signals	To provide a typical signal drop between networks, creating a low interference level. Note that inter-network signals arrive at the other network's nodes at equal strength
Filtered power	AC power filtered coming into the test set up to ensure no outside interference causes variation in results.	
PLC to Coax splitter	Used to route the PLC signals from the AC wiring over coax to the attenuation sources and the other adapter in the domain as well as the far network's adapters	This is standard test equipment for evaluating PLC performance.



» Figure 18: Legacy PLC Network 1 Input and Resulting Throughput under Test

The following is the screen shot of the input to (green trace) throughput (red trace) for network 2 of the legacy adapters from before (no activity) until after network 2 started up.



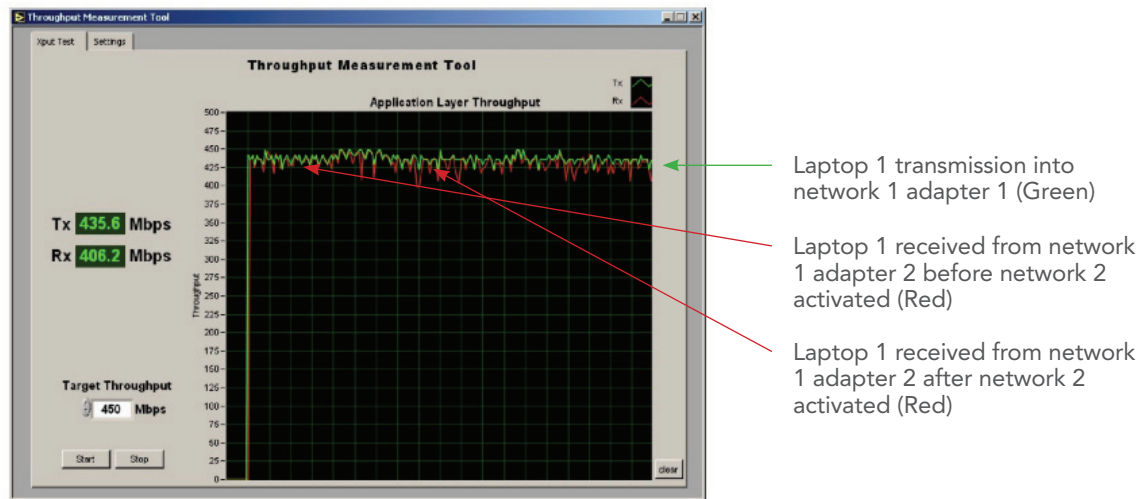
» Figure 19: Legacy Network 2 Input and Resulting Throughput under Test

Note: as explained in Section 3.3 and depicted in Figure 14, the output of legacy PLC networks decreases in proportion to the number of interfering networks, in this case the two networks drop to just under 50% each, netting less than a single network's throughput overall.

5.2 G.hn Test Results

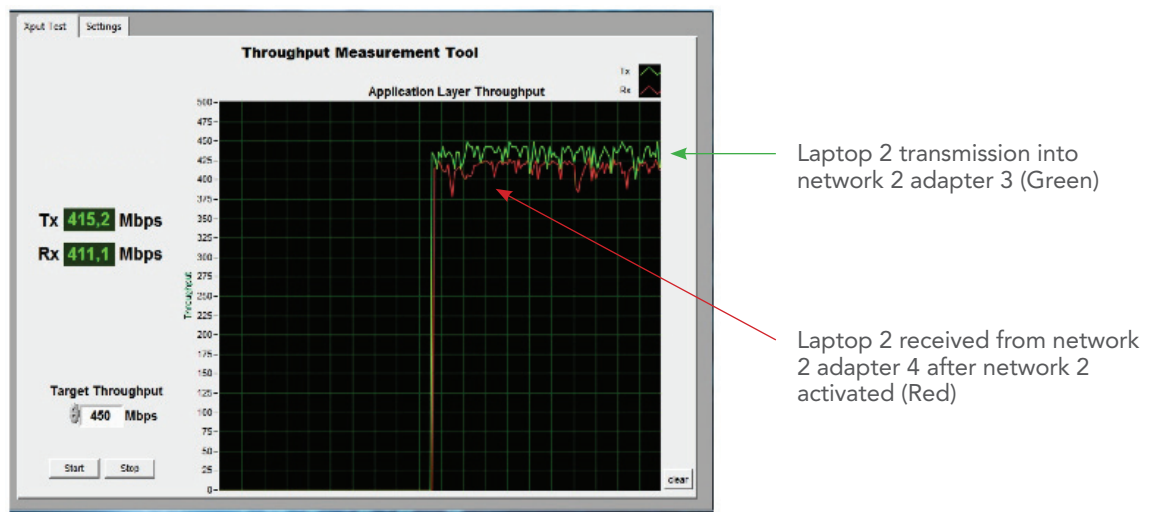
Using the same test equipment, set up and procedures, the Marvell G.hn units were tested.

Figure 20 depicts the input to network one (green trace) and the resulting output (red trace) from adapter 2 from before and after network 2 was activated.



» Figure 20: G.hn Network 1 Input and Resulting Throughput under Test

Figure 21 depicts the input (green trace) and output (red trace) of network 2 from before until after it was activated.



» Figure 21: G.hn Network 2 Input and Resulting Throughput under Test

5.3 Test Summary

The results clearly show the “sharing the wire” nature of the legacy PLC technology versus the “hidden” networks approach of G.hn NDIM with orthogonal preambles. With the legacy PLC technology, the throughput of each network drops to approximately the throughput of the

individual network throughput divided by the number of interfering networks. In this case, with two interfering legacy networks, the throughput drops to under half for each. While with G.hn, the effect is undetected.

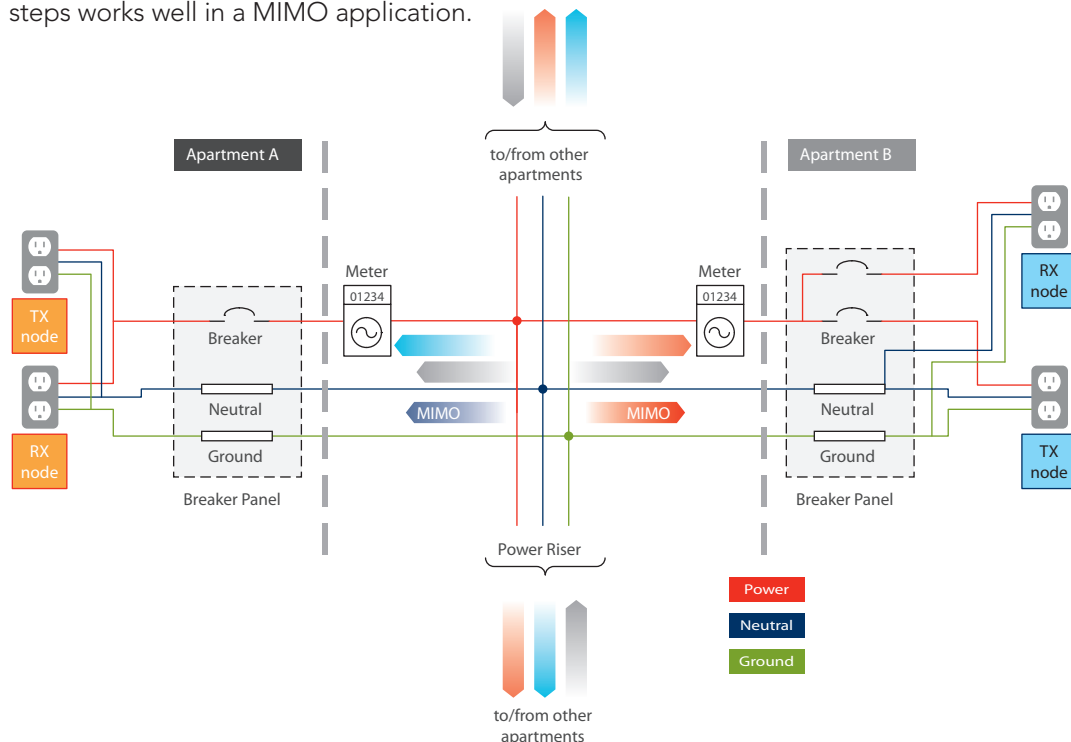
6 PLC MIMO NN Mitigation in MDUs — A Unique G.hn Strength

Multiple Input, Multiple Output (MIMO) is a means of using multiple transmit paths and receive paths when 3-wire cabling and sockets are used in residences. See the HomeGrid Forum White Paper, *MIMO for PLC*, for further information on MIMO. MIMO may actually exacerbate NN interference in that it may enable PLC signals from one unit to reach another unit with a stronger overall signal. The following Figure (Figure 22) depicts the paths each type of electrical service wire may take between two apartments. The power leads have higher attenuation than other lead types due to transiting breakers and meters. Further, the neutral wire, used as a return path for single input, single output (SISO) PLC has substantial noise. Meanwhile, earth ground, which is used as the MIMO second channel's return path, has less noise and therefore allows for a higher PLC signal level inter-unit.

While this figure is similar to Figure 3, it depicts the MIMO signal paths between living units.

With legacy PLC signal transmission over the power lead, there is signal attenuation due to the breakers and meters. This attenuation, in most cases, is not enough to eliminate inter-unit network interference, however.

As discussed earlier, pre-G.hn PLC technologies may require the user or service provider to install blocking filters as well as a time division methodology to handle neighboring network interference. The filters are relatively expensive and may not resolve the matter fully; moreover, neither of these steps works well in a MIMO application.



» Figure 22: PLC signals crossing between MDU apartments, showing MIMO signal paths

In a MIMO scenario, which increases the overall throughput of its network, PLC signals use the Ground wires, encountering less noise, therefore arriving at other living units with relatively higher signal strength. MIMO therefore increases the likelihood of substantial interference between PLC networks.

G.hn's NDIM functions enable G.hn PLC networks to operate in MIMO mode and still overcome NN interference, including coping with the increased interference due to MIMO in MDU buildings.

7 Summary

Neighboring PLC networks can interfere with one another to a point at which the service deteriorates to a level where the user experience becomes unacceptable. The ITU-T G.hn suite of standards specify a unique set of tools to allow neighboring G.hn domains (networks) to mitigate interference automatically and effectively, thus protecting the end users and providing a high quality of service and user experience.

The NDIM approach is very future-proof as it establishes a means for algorithm upgrades and other software-based modifications to meet specific challenges and evolving market needs.

While other PLC technologies have tried to retrofit mitigation techniques, even using external filters, these have been largely unsuccessful. G.hn has mitigation in its core standards and successfully mitigates NN interference today, and has the ability to extend this functionality to meet future or regional requirements.