

## **HPNA 2.0 10Mbps Home Phoneline Networking**

Edward H. Frank and Jack Holloway  
Broadcom Corporation

### **Abstract**

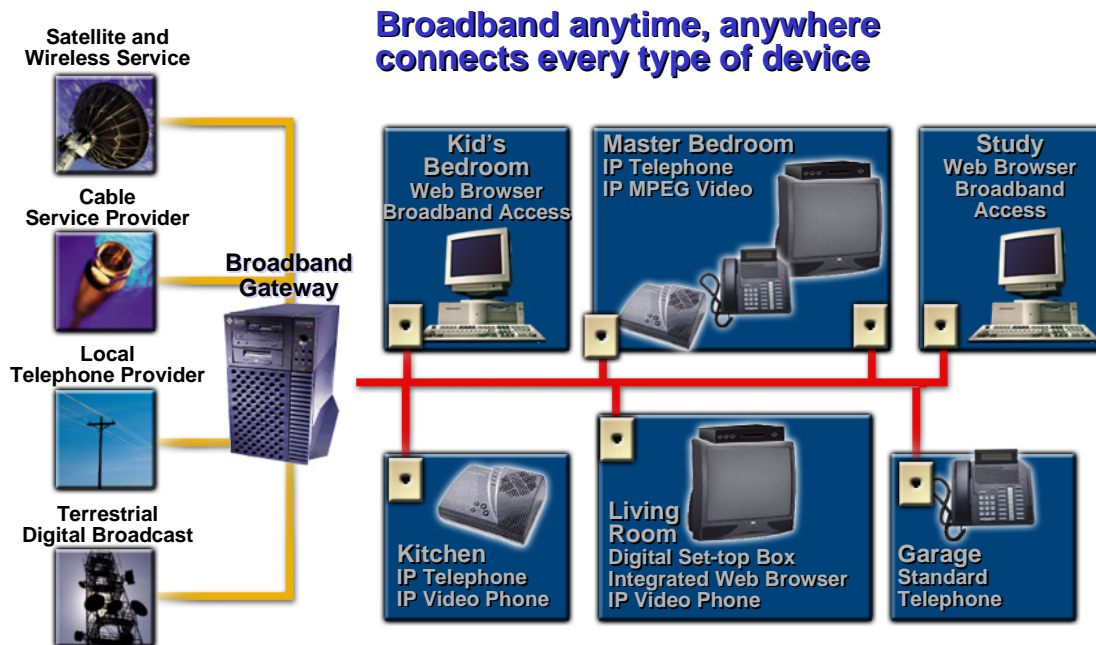
Computers are now in more than 50% of American homes, and about 20 million of these homes have at least two computers. The multiple PC home, together with the anticipated growth of Internet appliances, has created the need for a low-cost high-performance home networking technology. One approach uses the same pair of wires as the existing analog telephone service (POTS). Standardized under the auspices of the Home Phoneline Networking Association (HPNA), this technology is already in its second generation, operating at speeds up to 16 Mbps. Phoneline networking, unlike traditional Ethernet, must work robustly over a widely disparate range of transmission channels that have significant dynamic impairments.

### **1. Introduction**

We live in an age of ever-accelerating technological change. The signal event at the end of the Second Millenium was

almost certainly the explosion of the Internet. In 1995 there were 20 million users on the Internet, by 1998 there were 160 million. It is estimated that by 2003 there will be 500 million users worldwide, and over 14 countries will have more than 40% of their population on-line – countries that represent more than half of the world's GDP. Internet-based commerce has grown from essentially zero in 1995, to \$50B in 1998 and is projected to grow to \$1300B by 2003. There is an unprecedented level of investment in Internet-related business ventures – a direct consequence of the appreciation that the “new world order”, built on a wired information network, will profoundly effect the way we work and live.

What is less well appreciated is that the electronic dendrites of this network will extend beyond the PC to every electronic device within the home – connecting literally billions of devices (See Figure 1).



**Figure 1:** Connectivity in a networked home.

Traditional consumer electronics – television, stereo audio, telephones -- are already in the process of being redefined to use digital technology. In the new era, these devices will be designed with “the Network” built-in as a standard component, mirroring the absorption of the embedded microprocessor that occurred in the previous era. Network-connected devices will be smarter, easier to use, easier to maintain. The very nature of television, radio and the telephone will be transformed.

If every consumer electronic product will have an “Internet Inside” sticker, what connector will be used? How quickly will network-enabled products be adopted if consumers have to install network wiring and learn how to setup and administer a network? Will the home have to be “network-enabled” before these products can be used? The hard reality is that consumers don’t want to buy networks – but will be motivated to buy smart network-connected devices

that entertain, inform, educate, connect and increase convenience and choice. To initiate rapid market adoption, these devices will need to plug-in as simply as a telephone, with “no new wires”.

In the home, there are basically three existing wiring infrastructures that can be exploited: phone wiring, wireless and AC power wiring. It appears that all three will be used, with phoneline networks deployed first.

In 1998 the computer and semiconductor industries created an Alliance to select, promote and standardize technologies for Home Phoneline Networking (see HPNA 2.0 system [1] and <http://www.homepna.org>). This group has introduced a first generation 1 Mbps technology (based on a system developed by Tut Systems), and a second generation 10 Mbps technology based on a proposal from Epigram, Inc. (now part of Broadcom Corporation). Home Phoneline Networking is well

suited for the interconnection of broadband voice, video and data within the home. Industry reports estimate shipments of 1 million HPNA compatible interfaces by the end of 1999, and somewhere between 5-10 million interfaces by the end of 2000.

Networking over the existing home phoneline infrastructure suffers from many impairments (as do all no-new-wires physical media), namely high attenuation, reflections, impulse noise, crosstalk and RFI ingress/egress. These challenges must be overcome by a successful technology.

## **2. Requirements for Home Networking**

It is our belief that for a home networking technology to be successful, it must properly address the following major issues:

1. Leverage existing wiring infrastructure and be easy to install.
2. Leverage existing standards and interwork with common operating systems and software platforms.
3. Implement a quality of service (QOS) mechanism that provides low latency for

telephony and other voice applications; implement guaranteed bandwidth for streaming audio and video applications.

4. Be very robust and provide connectivity in essentially every home.
5. Support data rates in excess 10BASE-T Ethernet, and scale to 100 Mbps in a way that remains compatible with installed earlier generations.
6. Provide reasonable privacy at the physical layer. (Wireless and powerline require some level of encryption to achieve wired equivalent privacy.)
7. Be future safe, employing designs that are scalable and extensible so that users do not have to do “fork-lift” replacements when upgrading their networks in the future.
8. Be implementable with sufficiently low cost to allow inclusion as standard in a wide variety of products.

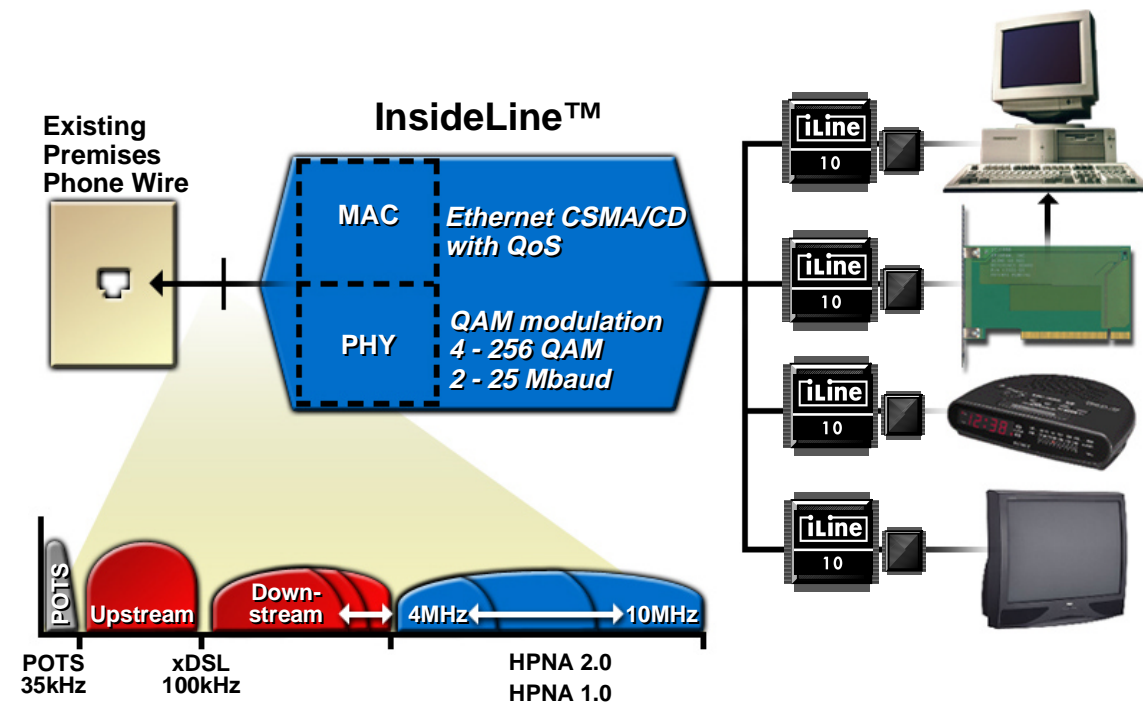
Table 1 summarizes how well the principal choices for home networking technology meet these criteria.

Parameter	HPNA 2.0	Wireless	Powerwire	Ethernet (Cat 5)
<b>1. Leverage existing infrastructure</b>	Good	Good	Good	Poor
<b>2. Leverage Standards</b>	Good (802.3 compatible)	Medium (too many standards <sup>1</sup> )	Poor (no standards)	Excellent
<b>3. QOS Support</b>	Good	Good to Poor (some standards have no QoS provision)	Unknown	Medium (Simple hubs don't support QoS. More expensive switches may.)
<b>4. Robustness</b>	Good	Medium	Unknown (Highly impaired channel)	Good
<b>5. Performance</b>	>10 Mbps, 100 Mbps next generation	1 to 11 Mbps Up to 50 Mbps at 5 GHz	Unknown (Highly variable channel capacity)	10, 100, 1000 Mbps
<b>6. Privacy of Physical Medium</b>	Good	Poor	Poor	Good
<b>7. Future-safe</b>	Good	Poor (too many standards, potential for interference)	Unknown	Good
<b>8. Cost</b>	Good	Medium (RF circuitry is harder to integrate)	Unknown (but should be comparable to HPNA)	Medium (low hardware cost, but higher cost if one considers installation of new wiring)

**Table 1:** Comparison of networking technologies.

---

<sup>1</sup> Several competing systems are under development and proposed for the unlicensed 2.4 GHz band (Bluetooth, HomeRF, 802.11b). The 2.4 GHz band has multiple sources of interference such as DECT phones and microwave ovens. The 5 GHz NII spectrum may also be used for home networking, using 802.11a or some other standard. Other standards and frequencies are proposed for systems to be used in Europe and Japan.



**Figure 2:** A view of the HPNA stack and spectrum

### 3. The HPNA 2.0 System

Figure 2 is an illustration of the HPNA 2.0 system from the point of view of network stack and frequency spectrum. The HPNA 2.0 system is a multi-point CSMA/CD packet network that supports unicast, multicast, and broadcast. As will be discussed in this section, and illustrated in Figure 7, it has the look and feel of Ethernet. However, it differs from 10BASE-2 and 10BASE-T in a number of respects. First and foremost, HPNA 2.0 places no restrictions on wiring type, wiring topology, or termination. Moreover, like 10BASE-2, but unlike 10BASE-T, HPNA 2.0 uses a shared physical medium with no need for a switch or hub. 10BASE-T on the other hand requires dedicated point-to-point CAT-3 or CAT-5 wires.

#### Physical Layer

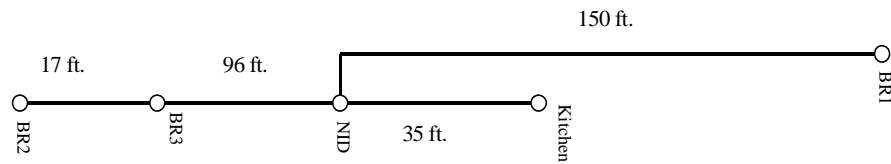
At the Physical layer, the system is frequency division multiplexed on the same wire as standard analog phone service (POTS), as well as other splitterless ADSL[2]. Analog telephony uses the low part of the spectrum below 35 kHz. ADSL (both G.Lite and G.Heavy) use spectrum up to 1.1 MHz.

HPNA selected the 4 to 10 MHz band for several reasons. The lower limit of 4 MHz was chosen to make it feasible to implement the filters needed to reduce out-of-band interference between HPNA and splitterless ADSL. After modeling several thousand representative networks with capacitive telephones and common wire lengths, it was determined that the spectrum above 10 MHz was much more likely to have wider and deeper nulls caused by reflections [3]. Crosstalk between phonelines increases

with frequency, and the analog front end is harder to implement at higher frequencies. The particular choice of 4 to 10 MHz only overlaps a single Amateur Radio band (40 meters), which simplifies ingress and egress filtering.

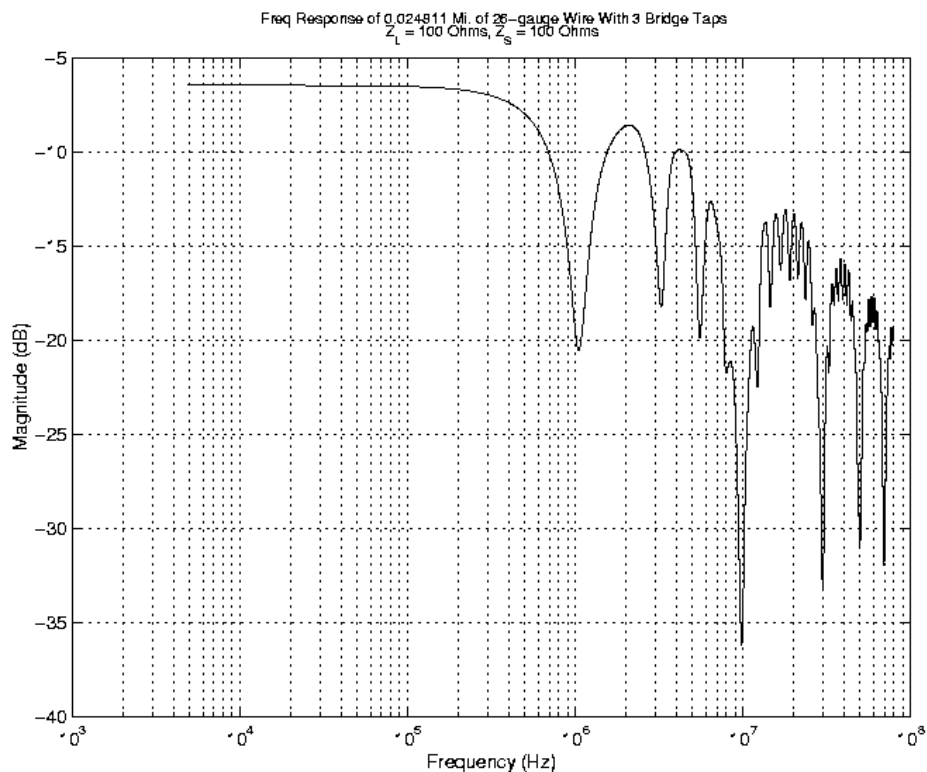
The no-new-wires media that are available for networking within homes

have the problem that the communications channel can be severely impaired. The nature of the impairments is illustrated in Figure 3 and Figure 4, which shows the kind of channel response one might find on a typical phone wire loop inside of a house.



4

**Figure 3:** Topology of a simple home network



**Figure 4:** The channel response of a simple home phonewire network.

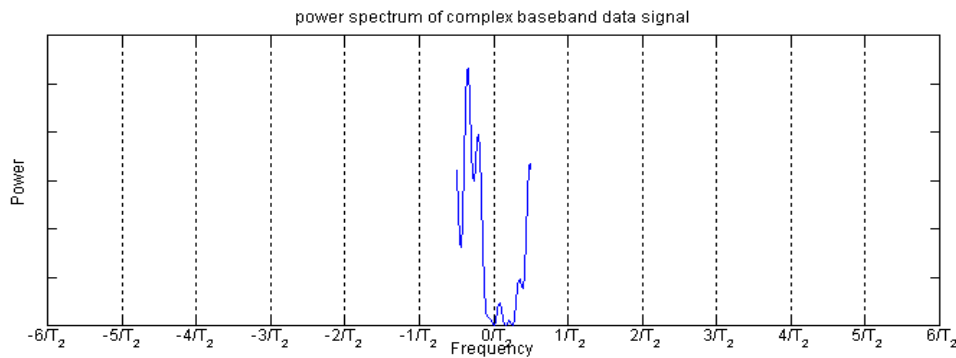
The wiring topology found in homes is ad-hoc resulting in reflections and frequency dependent channel transfer functions; the transmission parameters of the wire used is uncharacterized and highly variable, especially at higher frequencies; telephone instruments on the same wiring present a wide range of frequency-dependent impedances; While HPNA 1.0 uses Pulse Position Modulation as mentioned above, HPNA 2.0 uses quadrature amplitude modulation, both to get more throughput in the same bandwidth, as well as to achieve greater robustness. However, due to fact that the channels may have very deep nulls, and multiple nulls in band, two techniques are used to improve robustness. The first technique is to be rate adaptive. Instead of having a fixed number of bits per symbol, a transmitter may, on a packet by packet basis, vary the packet encoding from 2 to 8 bits per symbol. The second technique is to use spectral diversity as discussed below

## Frequency Diverse QAM

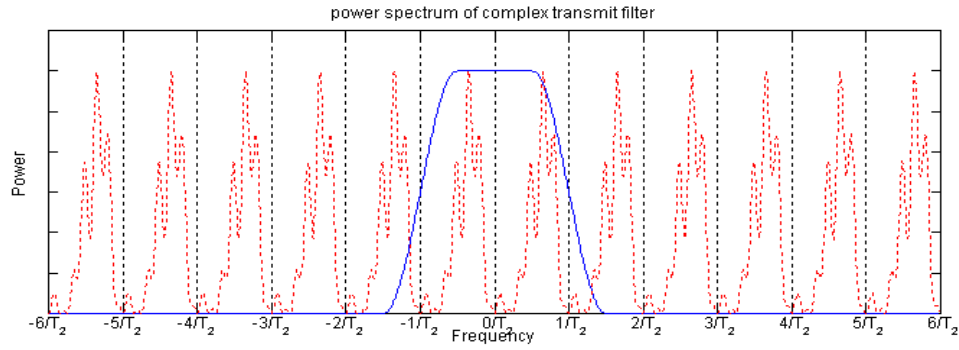
Unfortunately, the nature of channel nulls can be such that even rate adapting down to 2 bits per symbol is not sufficient to guarantee that the packet can be received.

In a traditional QAM system, if there is an extreme null (i.e. one with which the equalizer can't cope) in the band then the system will fail to operate. At its 2 Mbaud rate, HPNA 2.0 implements a modified version of QAM invented by one of our colleagues Eric Ojard, called Frequency Diverse QAM (FDQAM) [5]. While a full discussion of FDQAM is beyond the scope of this paper, Figure 5 through Figure 8 illustrate the basic concept.

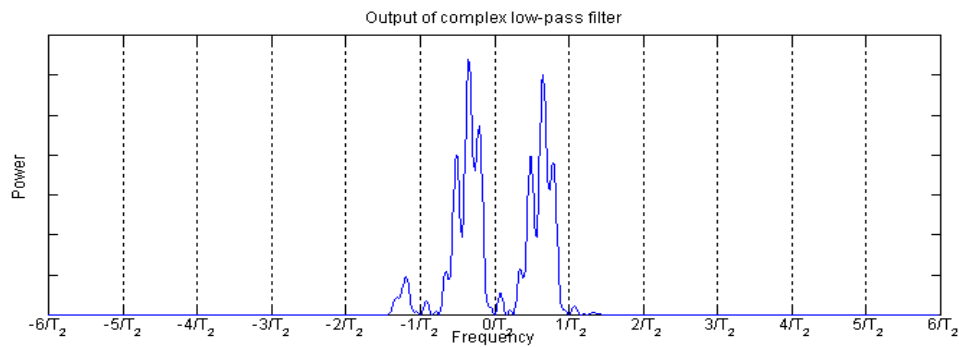
In a traditional QAM system, a single copy of the baseband signal is sent and received. Because in FDQAM the baud rate is less than half the width of the filter, the output signal has two redundant copies of the baseband signal, as shown in Figure 5. Thus, the signal is frequency-diverse, motivating the name FDQAM.



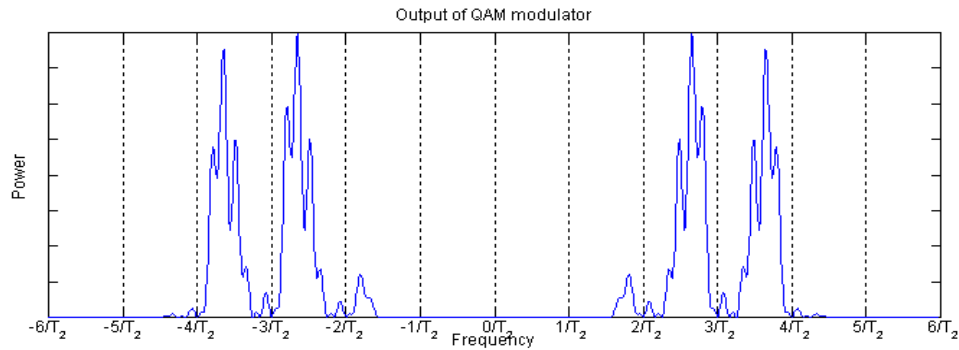
**Figure 5:** Spectrum of complex base-band signal.



**Figure 6:** Spectrum of upsampled complex baseband signal, overlayed with low-pass filter.



**Figure 7:** Output of frequency-diverse low-pass filter.



**Figure 8:** Output of FDQAM modulator.

Intuitively, it's easy to see that on channels where half of the spectrum is nulled out, one copy of the signal will still make it through. Quantifying the performance of FDQAM versus QAM on arbitrary channels is more complicated, and this analysis is not included here. It can be shown, however, that on channels with low SNR where a large part of the spectrum is severely

attenuated, FDQAM works robustly in many cases where uncoded QAM modulation would fail. Such channels are common on home phonelines. Unlike most other methods of handling severe channels, FDQAM requires no knowledge of the channel characteristics by the transmitter, simplifying the protocol and enabling robust



performance over time-varying channels.

In cases where the channel nulls are not particularly deep, HPNA 2.0 allows for a higher performance 4 Mbaud mode, which achieves peak data rates up to 32 Mbps, and throughput above 20Mbps.

### Frame Format

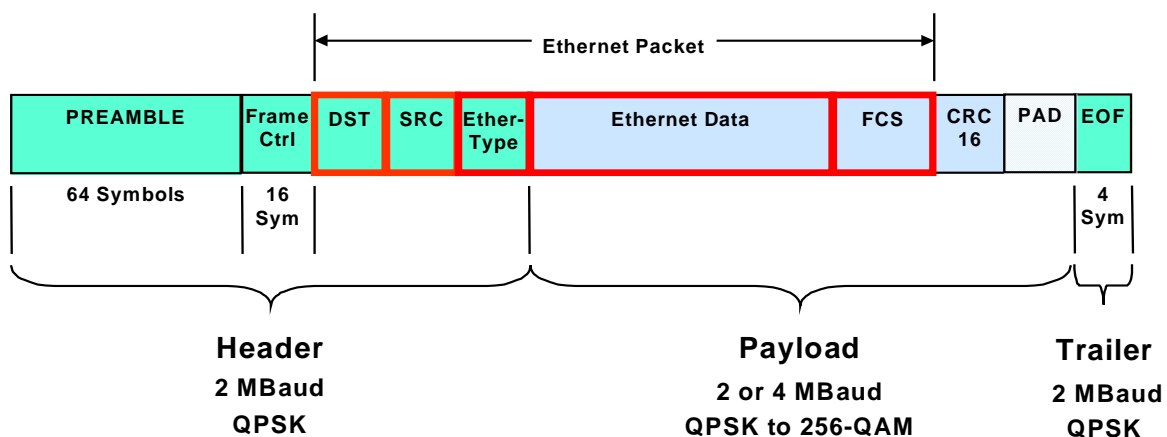
Figure 9 shows the frame format on the wire. The frame begins with a known 64 symbol preamble. The preamble is used for several purposes:

- Robust Carrier Sensing and Collision Detection
- Equalizer Training
- Timing Recovery
- Gain Adjustment

Following the preamble is a Frame Control field. The first part of the FC is an 8-bit frame-type. Frame-type=0 is shown, where other codes can be assigned for frame formats used by future systems. Following the frame-type

is an 8-bit field that specifies the modulation format (bits per symbol for example). There are miscellaneous other control fields in Frame Control including an 8 bit header CRC. The remainder of the packet is exactly an 802.3 Ethernet frame followed by CRC16, padding and EOF sequence. The CRC16 covers the header and payload, and reduces the undetected error rate for severely impaired networks.

Key to operation is that the first 120 bits of the frame are sent at the most robust 2 Mbaud, 2 bits per symbol rate. The reason for this is that if it is possible for any station to be able to demodulate a packet, it will be at this encoding. Thus, even if the payload is encoded at a rate or bits per symbol that the receiver can't demodulate, it will be possible to demodulate the header. In this situation, the receiver sends a Rate Request Control Frame to the sender asking it to reduce the number of bits per symbol or the symbol rate.



**Figure 9:** HPNA 2.0 Frame Format.

## Media Access Control

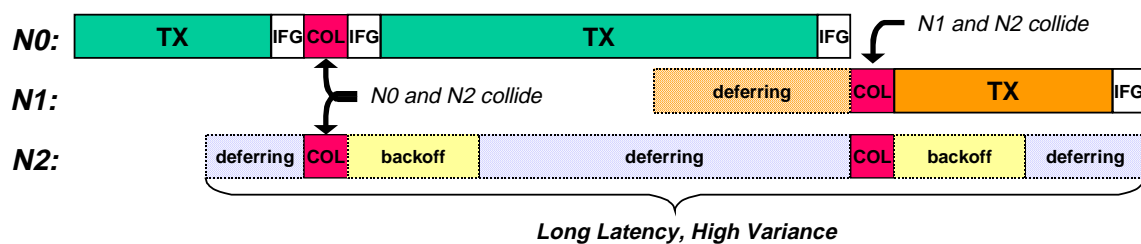
As mentioned above, HPNA 2.0 is a Carrier Sense Multiple Access (CSMA) with Collision Detection (CD) system, just like standard IEEE 802.3 Ethernet. HPNA 2.0 introduces eight levels of priority and uses a new collision resolution algorithm called Distributed Fair Priority Queuing (DFPQ).

Voice telephony requires a low-latency network service, and streaming audio or video applications require a guaranteed bandwidth service. With the MAC in Ethernet, there are no real service guarantees, as shown in Figure 10.

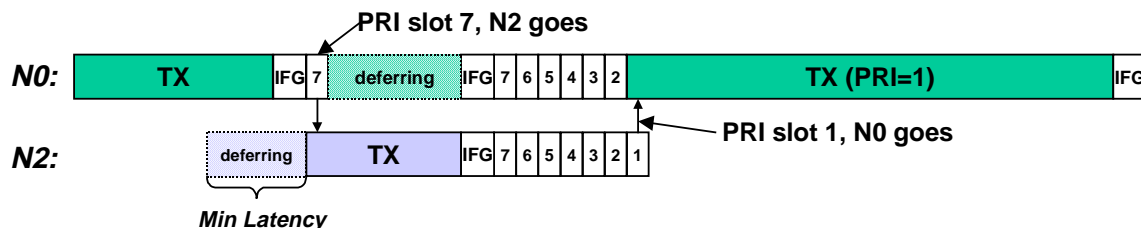
In this example, three nodes (N1, N2, and N3) are contending for access to the network. Node N2 is transmitting a voice-over-IP (VoIP) packet. Initially, N0 accesses the wire transmits a frame (TX), and during this transmission N2 has packetized a voice sample and is ready to transmit, but must defer to N0. At the end of the first transmission, N0 has a second packet ready to send, and

when N2 and N0 contend for access (resulting in a collision) N2 by chance chooses a longer backoff interval than N0. N0 gains access again and transmits. During this time, another station N1 becomes active, and starts deferring waiting for N0 to finish. Now, when N2 attempts to transmit it collides with N1. A possible outcome is that N1 succeeds in the collision resolution and N2 further increases its backoff. In this manner, the queuing discipline can become very unfair for N2. If N0 and N1 are PC's engaged in file transfers, they can generate enough traffic loading on the network to cause errors in the VoIP service operating on N2.

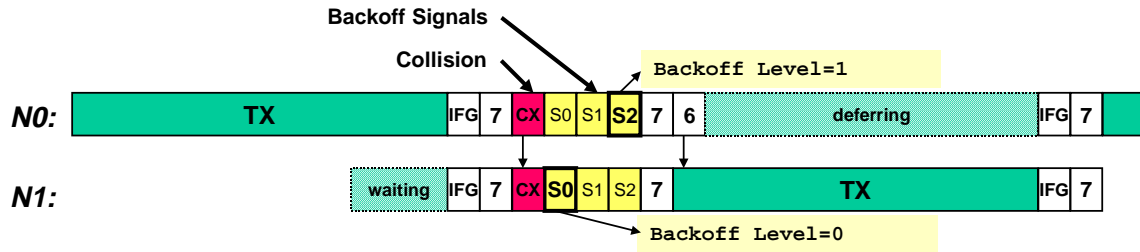
One solution to this problem is to introduce different access priorities, where the VoIP station uses a higher priority than best-effort file transfer traffic. HPNA 2.0 accomplishes this by organizing the time following the inter-frame gap into an ordered series of priority slots.



**Figure 10:** An example of access latency unfairness



**Figure 11:** Priority



**Figure 12:** HPNA 2.0 Collision resolution algorithm

Now in the example shown in Figure 11, when N0 finishes transmitting, all stations on the network that have lower priority than 7 wait, while N2 begins to transmit (without collision). After N2's transmission, no stations have traffic with priority higher than 1, so N0 again gains access to the channel with its next transmission.

Access priority lets software define different service classes, such as low-latency, controlled-bandwidth, guaranteed-bandwidth, best-effort, penalty, etc., each using a different priority level.

Within a given priority level, HPNA 2.0 uses an algorithm for collision resolution where each station keeps track of a *Backoff Level* and after a collision randomly chooses to increment the backoff level by 0, 1 or 2. During a collision resolution cycle, stations incrementally establish a partial ordering—eventually only one station remains at the lowest backoff level and gains access to the channel.

In the example shown in Figure 12, N0 and N1 enter into a collision resolution cycle. N0 randomly chooses to increment its backoff level by 2, N1 by 0. To optimize the partial ordering, eliminating null levels, stations send a

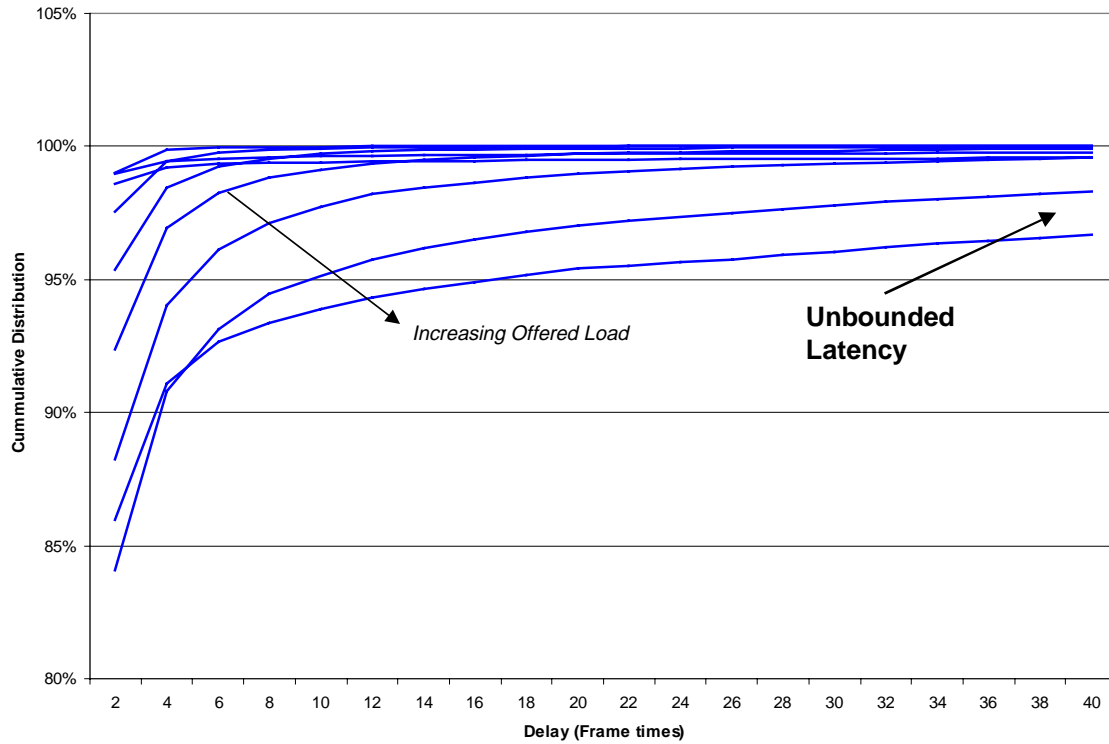
special signal immediately following a collision which reflects the backoff increment chosen (0 and 2 in the example shown). All stations observe these signals and perform a distributed computation to calculate the new (partial) ordering. In this case, N0 increments its backoff level by 1 since it saw the Backoff Signal from N1 in S0 but no station indicating in S1.

In practice, even on saturated networks, HPNA 2.0 is very well behaved, and unlike traditional Ethernet does not exhibit the capture effect. The relative performance of DFPQ and Ethernet are shown below.

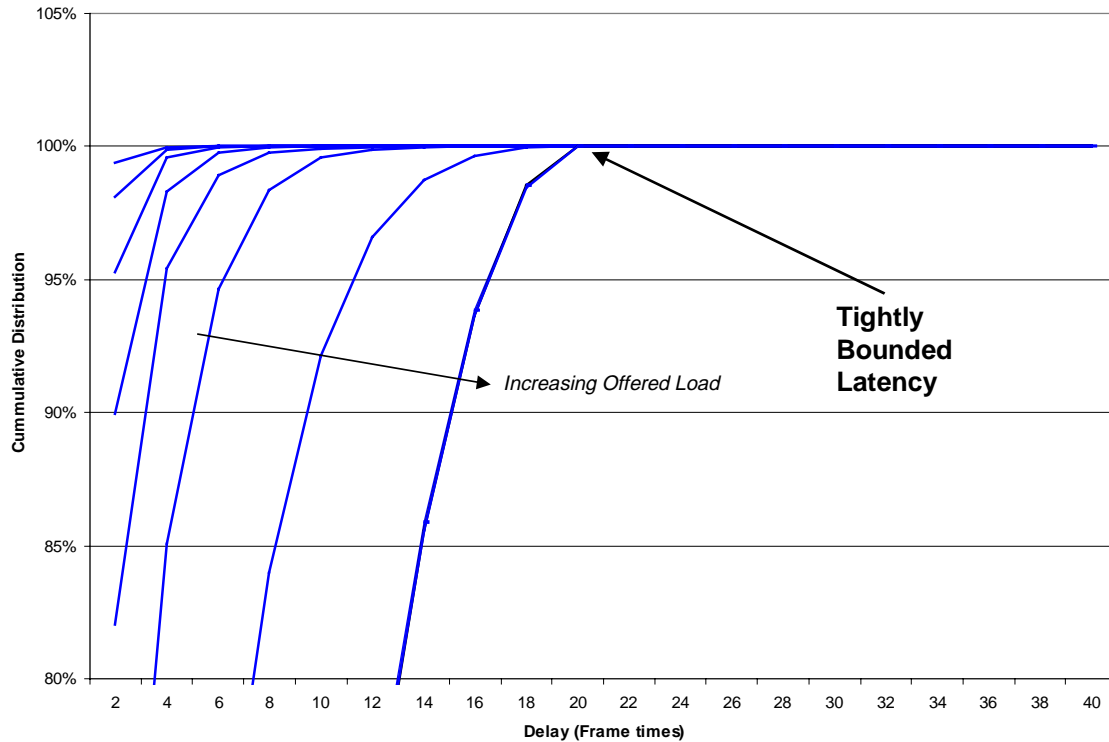
Figure 13 shows the distribution of access delay for Ethernet for offered loads ranging from 30% to 120% of network capacity. The simple simulation shown assumes a Poisson traffic model, 10 stations and a uniform frame size of 1500 bytes. Delay is indicated in units of frame transmission time, where a perfect queuing discipline would have a maximum delay equal to 10 frame times (the number of stations). As the offered load increases, Ethernet experiences delays of 100's of frame times for several percent of transmission attempts.

In Figure 14 the HPNA 2.0 access latency is shown. By comparison with Ethernet there is a negligible distribution tail beyond 20 frame times, even at high offered loads. It should be noted that the

results shown are for contention within a single priority level. High priority traffic will see on the average less than one frame time of delay when contending with lower priority traffic.



**Figure 13:** Ethernet Access Delay Distribution



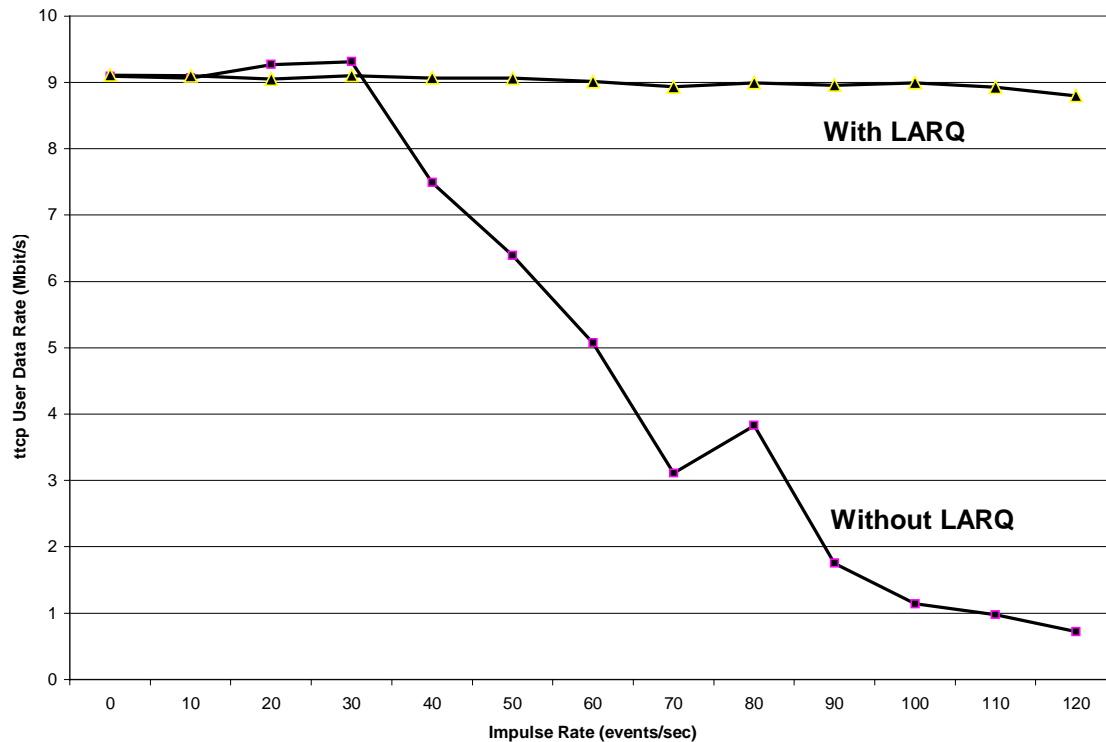
**Figure 14:** HPNA 2.0 Access Delay Distribution

### Link Layer Protocols

One impairment not mentioned previously that is a problem for all home networks using “No new wires” is impulse noise. On phonewire impulse noise exists due to phone ringing, switch hook transitions, and noise coupled from the AC power wiring. Fortunately, the impulses tend to be short, and destroy only a single packet. While there are coding techniques that might reduce the number of packets destroyed by impulses, we have chosen to use a fast retransmission mechanism we call Limited Automatic repeat ReQuest (LARQ). Because LARQ is implemented (in software) at layer 2,

and because it is only on a single segment of the network, it is very effective in hiding packet erasure from TCP/IP, as shown in Figure 15.

Finally, it is worth mentioning, the HPNA 2.0 implements a link integrity mechanism, which can be implemented either in hardware or at low levels of a software driver. The virtue of link integrity is that it provides a quick and easy way for the end user to determine if the network has basic connectivity. Link integrity frames are sent once per second, unless there traffic on the wire, in which case the number of frames sent may be reduced.



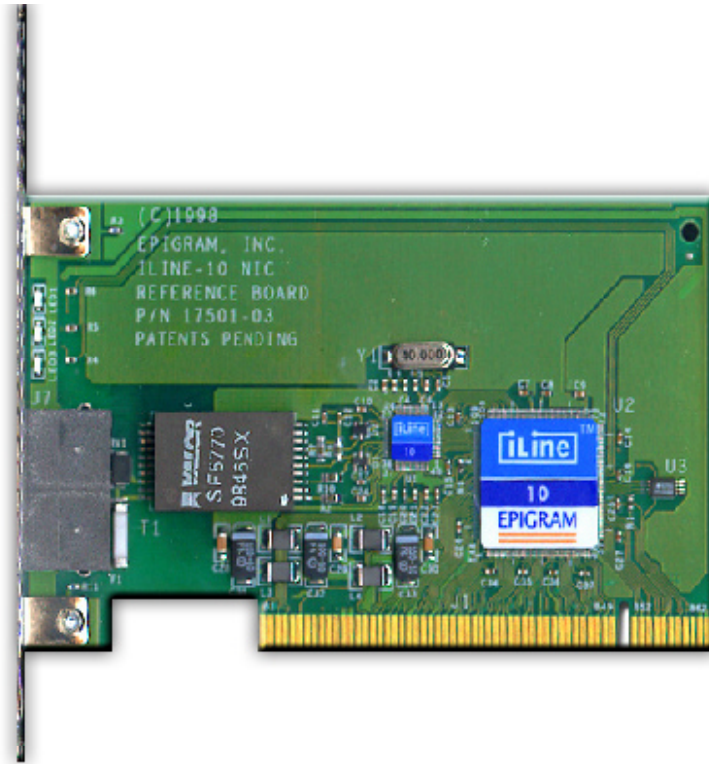
**Figure 15:** User level throughput vs. Impulse noise events/sec

#### 4. Example Implementation

Figure 16 shows the high integration at the chip-level and the low complexity at the system-level that can be achieved with HPNA 2.0, in this case using Broadcom's iLine-10 chipset, which combines the MAC and PHY for both HPNA 1.0 and HPNA 2.0. Chipsets from other vendors should have similar characteristics. The major components are:

A. The BCM4210 MAC/PHY chip  
(the larger iLine chip)

- B. The BCM4100 Analog Front End chip (the smaller iLine chip)
- C. Magnetics module for phoneline isolation/ protection (large black module)
- D. Serial Prom with MAC address and other configuration
- E. Crystal



**Figure16:** Broadcom's iLine10 HPNA2.0 NIC reference board.

## Conclusion

Home networking is now a reality. In 1999 estimates are that over 1 million home networking nodes were shipped. Estimates range as high as 10 million home networking nodes will ship in 2000. Using advanced signal processing techniques and high density CMOS it is possible to transmit data over existing media, such as in-place phone wire, at rates once considered impossible. Equally important, the cost of these solutions is such that the chips can be built into a wide variety of computers and internet appliances. Just as the Microprocessor has become an essential component of every digital device, we anticipate that a communications element, which we call the Internet-Chip or I-Chip, for short, will be come an equally essential element in every digital

system over the next several years. Consumers will come to expect that the devices they buy have an I-chip in them. By the year 2005, if not sooner, consumer devices that can't communicate in the Home LAN will be obsolete.

## Acknowledgements

The work described in this paper is the result of the efforts of many people. In particular we would like to recognize Ruben Alva, Neal Castagnoli, Alan Corry, Mike Dove, Greg Efland, Nino Ferrario, Hariprasad Garlapati, Ray Hayes, Mark Kobayashi, Harrison Kuo, Jim Laudon, Joe Lauer, Guillermo Loyola, Tracy Mallory, Rich McCauley, Subrat Mohapatra, Tushar Moorti, Walter Morton, Neal Nuckolls, Eric Ojard, Jay Pattin, Kevin Peterson, Henry

Ptasinski, Tim Robinson, Dane Snow, Bill Stafford, Jason Trachewsky, Chris Warth, Larry Yamano, and Chris Young.

### **Biographies**

Edward H. Frank (ehf@broadcom.com) was co-founder and Executive Vice President of Epigram, now the Home Networking Division of Broadcom Corporation. Prior to Epigram, Dr. Frank was Vice President of Engineering for NeTpower, Inc., and was a Distinguished Engineer at Sun Microsystems, where he worked on the Green Project, which created Oak, the precursor to Java, and where he was co-architect of several generations of SPARCstation workstations. Dr. Frank has a PhD from Carnegie Mellon University, and BSEE and MSEE degrees from Stanford University. He holds over 15 issued patents.

Jack Holloway (h@broadcom.com) was co-founder and Chief Technology Officer of Epigram, now the Home Networking Division of Broadcom Corporation. Prior to Epigram, Mr. Holloway was Vice President of Broadband at MicroUnity, a closely held mediaprocessor developer. Mr. Holloway was a Vice President, Broadband Technology for Bolt Beranek and Newman, where he started an ATM product division called Lightstream, which was eventually acquired by Cisco Systems. Mr. Holloway was a cofounder of Symbolics, Inc, and for several years was a Principal Research Scientist at M.I.T. and co-Director of the M.I.T VLSI Laboratory.

---

[1] *HPNA 2.0 Specification*, October, 1999, available to members of Home PNA, see <http://www.homepna.org>

[2] Chen, W., *DSL: Simulation Techniques and Standards Development for Digital Subscriber Line Systems*, Macmillan, 1998.

[3] "Spectrum Selection for Home Phoneline Networking", ITU-T Study Group 15, NG-101, Nuremberg Germany, 2 Aug 1999.

[5] "Frequency Diverse Quadrature Amplitude Modulation", by Eric Ojard, in preparation.