CONNECTING THE HOME WITH A PHONE LINE NETWORK CHIP SET

IN ADDITION TO SHARED INTERNET ACCESS FOR PCS, THE HOME NETWORK WILL CONNECT TO EVERY CONSUMER ELECTRONIC DEVICE. TO MAKE THIS POSSIBLE, WE MUST CONSIDER ROBUST SYSTEM REQUIREMENTS, HOME PHONE LINE STANDARDS, COSTS, AND IMPLEMENTATION OF A SUPPORTING ILINE10 CHIP SET.

Edward H. Frank
Jack Holloway

We live in an age of ever-accelerating technological change. The signal event at the end of the second millennium was almost certainly the explosion of the Internet. In 1995 there were 20 million Internet users; by 1998 there were 160 million. We could have 500 million users worldwide by 2003, with over 14 countries having more than 40% of their population online. Those 14 countries represent more than half of the world's GDP (gross domestic product).

Internet-based commerce has grown from essentially zero in 1995 to $50 billion in 1998, and industry observers expect it to reach $1,300 billion 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 affect the way we work and live.

More than half of US homes today have access to the Internet, primarily via dial-up voice modems. Multibillion-dollar investments are being made in cable modem and DSL (digital subscriber line) infrastructure to upgrade that access to broadband reception. By 2003 we expect one third of connected US homes to subscribe to a broadband always-on access service. Several new broadband digital wireless technologies—such as Fixed Wireless Local Loop (FWLL), terrestrial television broadcast, and satellite broadcast—are being developed and will increase the flood of digital information reaching the home. It is generally assumed that our 50-year-old voice telephony infrastructure will be replaced with an integrated broadband digital network carrying voice, audio, video, and data.

What is less well appreciated is that the electronic dendrites of this network will extend beyond the personal computer to every electronic device within the home, connecting literally billions of devices (see Figure 1, next page).

Connecting the home

Traditional consumer electronics (TV, 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 a communications 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, and easier to maintain, transforming today's TV, radio, and telephone.

Novel Internet appliances will be invented to access new services that are carried by the new-era network. As shown in Figure 2, IDC projects this product category to grow rapid-

0272-1732/00/$10.00 © 2000 IEEE

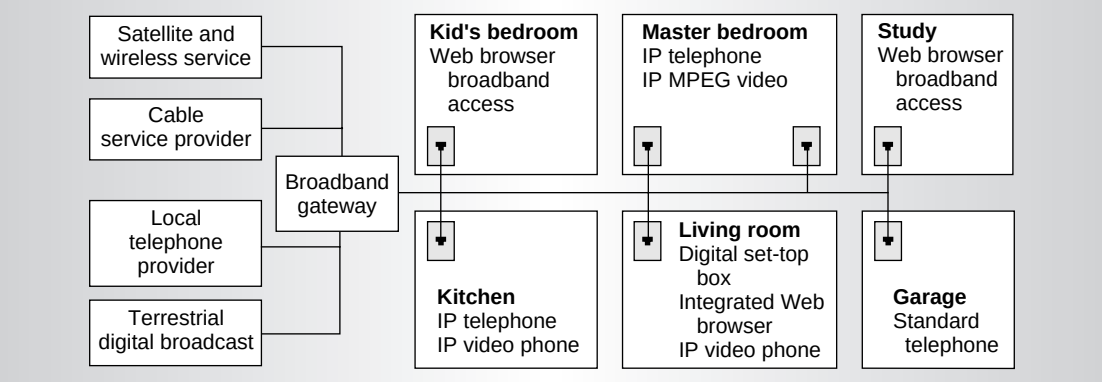

Figure 1. Connectivity in a networked home.

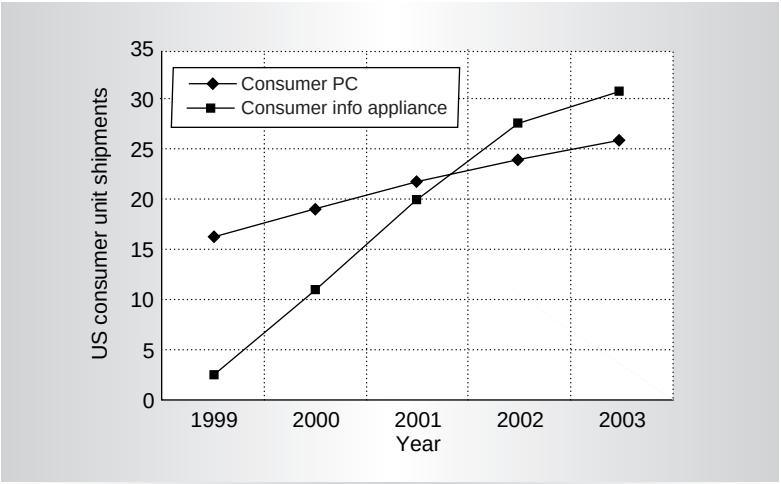

Figure 2. The growth in home PCs and Internet appliances. *Source: International Data Corporation (IDC), 1999*

ly, eventually far exceeding the number of PCs in the home.

The hard reality is that consumers don't want to buy networks. However, they 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.

The home has three existing wiring infrastructures that can be exploited: phone line wiring, wireless, and AC power wiring. It appears that all three will be used, with phone line networks deployed first.

In 1998 the computer and semiconductor industries created the Home Phone Line Networking Alliance (HPNA) to select, promote, and standardize technologies for home phone line 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 joint proposal from Broadcom Corporation and Lucent Microelectronics). Home phone line networking is well suited for the interconnection of broadband voice, video, and data within the home since it offers data rates from 10 to 100 Mbps with good quality of service. Industry reports estimate shipments of 1 million HPNA-compatible interfaces by the end of 1999, and somewhere between 5 to 10 million interfaces by the end of 2000.

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

Home networking requirements

For a home networking technology to be successful, it must properly address certain issues. It must

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, and implement guaranteed bandwidth for

Table 1. Comparison of networking technologies.

Parameter	HPNA 2.0	Wireless	Power wire	Ethernet (Category 5)
Leverage existing infrastructure	Good	Good	Good	Poor
Leverage standards	Good (802.3 compatible)	Medium (too many standards)*	Poor (no standards)	Excellent
QoS support	Good	Good to poor (some standards have no QoS provision)	Unknown	Medium (simple hubs don't support QoS; more expensive switches may)
Robustness	Good	Medium	Unknown (highly impaired channel)	Good
Performance	>10 Mbps; 100 Mbps next generation	1 to 11 Mbps; up to 50 Mbps at 5 GHz	Unknown (highly variable channel capacity)	10, 100, 1,000 Mbps
Privacy of physical medium	Good	Poor	Poor	Good
Future safe	Good	Poor (too many standards, potential for interference)	Unknown	Good
Cost	Good (initial installation cost:<$100; new connected devices:<$50)	Medium (RF circuitry is harder to integrate)	Unknown (but should be comparable to HPNA)	Medium (low hardware cost, but higher cost when installing new wiring)

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

streaming audio and video applications;
4. be very robust and provide connectivity in essentially every home;
5. support data rates in excess of 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 power line technologies 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; and
8. be implementable with sufficiently low cost to allow inclusion as a standard in a wide variety of products.

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

Leveraging standards

The importance of leveraging standards cannot be overestimated. There are two issues to consider: 1) is there an accepted standard that guarantees interoperability between equipment from multiple manufacturers, and 2) does the system faultlessly support other networking standards and in particular the Internet protocol (IP) suite. Only protocols that have been extensively tested for many years in real use are likely to be ready for large-scale consumer deployment. Given the preponderance of IEEE-802.3 Layer 2 networking across the Internet infrastructure, HPNA has chosen a technology that uses 802.3 framing and Ethernet CSMA/CD (carrier sense multiple access with collision detection) MAC (media access control) behavior.

Quality of service

Anticipated home applications will drive the requirement for QoS support. The initial motivation for home networking is sharing resources

among multiple PCs such as Internet access, files, and printers. However, the ultimate applications that will dominate home networks are the transport of digital audio, digital video, and digital voice (IP telephony). Latency in voice connections must be controlled below 10 to 20 ms on the home network segment if voice quality is to be maintained. Streaming video and audio connections must receive an application-determined minimum bandwidth from the network.

Although the aggregate throughput rate of 10 Mbps for HPNA 2.0 is more than adequate for many application scenarios, burst loads presented by TCP transfers between PCs, without some QoS mechanism, would at times make the network unable to meet the latency and guaranteed bandwidth service requirements. Furthermore, bandwidth allocation within a given class of service should be fair. The traditional Ethernet MAC[2] layer exhibits a phenomenon known as packet capture,[3] which can result in long tail access latency distributions. The HPNA 2.0 MAC layer introduces eight priority levels and an improved collision resolution technique that eliminates packet capture.

Robustness

The primary difference between twisted-pair Ethernet and other technologies is the quality of the communications channel. Running over Category-5 cable, Ethernet encounters a channel that has a number of very nice properties. They include point-to-point communication, proper termination, a well-characterized channel response (both in terms of nulls and overall attenuation), and very low cross talk. In contrast, all of the no-new-wires media available for networking within homes have the problem that the communications channel can be severely impaired.

The ad hoc home wiring topology results in

- reflections and frequency-dependent channel transfer functions;
- uncharacterized and highly variable wire transmission parameters, especially at higher frequencies;
- telephone instruments on the same wiring that present a wide range of frequency-dependent impedances;
- POTS (plain old telephone system) signaling and ringing that can produce significant transients;
- impulse noise coupled from AC power wiring that is seen on many phone lines; and
- RF ingress and egress, particularly in the amateur radio frequency bands.

Little prior field survey has been done to statistically model these impairments.

Such less-than-ideal channels can have an impact on both robustness and performance. One way to overcome the challenge presented by what might be referred to as impaired media is to simply operate at very low data rates (when compared to the theoretical channel capacity) or over very limited distances. The X-10 system[4] used for control of lights and other devices on power lines adopts the low data rate approach.

Performance

The HPNA 1.0 technology[5] uses a pulse position modulation (PPM) technique with a spectral efficiency of 0.16 bits/baud, resulting in a 1-Mbps data rate. The second-generation system targets an order-of-magnitude higher data rate.

History has taught us that higher network speeds are always better. In home networking several external influences persuade us that we require at least 10 Mbps. The common broadband access technologies such as ADSL (asymmetric digital subscriber line) and the DOCSIS (Data Over Cable Service Interface Standard) cable modem require home networks with data rates of 6 Mbps or more to share the access bandwidth. Moreover, applications such as multiple DVD streams or high-definition digital video make it easy to imagine that even 10 Mbps isn't enough. Therefore, the alliance designed the HPNA 2.0 system to achieve data rates up to 32 Mbps in approximately the same bandwidth as the HPNA 1.0 system and be forward compatible with future stations operating at speeds up to 100 Mbps.

To achieve high data rates over impaired channels, we need to use some form of complex modulation and a receiver that is adaptive to the frequency-dependent channel characteristics. Channel response is a function of the location on the network wiring of both the transmitter and receiver, so this adaptation is different for each unique transmitter-receiver pair. Channel response can also vary over a short time frame

when telephone instruments change impedance during switch-hook transitions and dialing. To work reliably over a wide range of possible home wiring configurations, we have chosen to design stations to be rate adaptive. Then when there isn't sufficient channel capacity on a given path for the full rate, connectivity can still be achieved at lower rates.

The fundamental design challenge for home phone line networking is how a receiver can determine—for each burst transmission seen on the channel—the equalization and demodulation that should be applied to recover the packet. And conversely, for each transmitter, what modulation rate is feasible when sending to a given destination station?

In HPNA 2.0, this is accomplished with a self-describing frame format with PHY-level signals that can directly control equalizer training and demodulation.

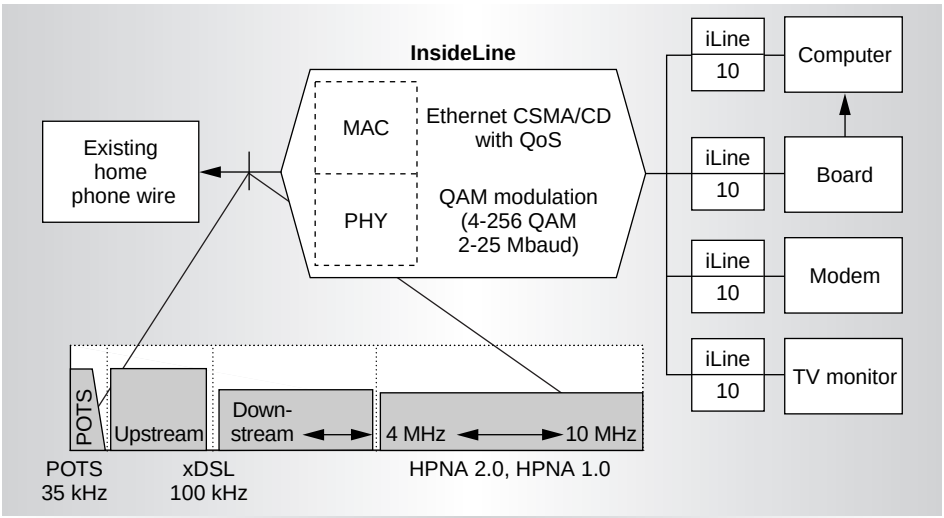

Figure 3. A view of the HPNA 2.0 stack and spectrum.

Physical layer privacy

Both power wire and wireless systems allow users to share the same physical medium. Therefore, they must have encryption at the link level to attain a degree of privacy equivalent to the phone line. This requires some user key configuration, which somewhat defeats the plug-and-play objective. On the other hand, the presumptive privacy of phone wires is not a substitute for true cryptographic security.

Future safeness

Once installed, home networks are likely to remain in place for many years. Worse yet, as home network interfaces become embedded in appliances, it may become almost impossible to replace them. Thus, a good home-networking technology ideally has built into the current generation a plan for interoperability with future generations.

Cost

Finally, there is the issue of implementation cost and complexity. As has become very well understood over the last 10 years by the computer and networking industries, volume is everything. With decreasing prices for computer equipment—especially for the home—a successful home-networking technology must be inexpensive, and it must ride Moore's law. Practically, the technology must be implementable in mainstream CMOS technology and require few additional external components.

HPNA 2.0 system

Figure 3 illustrates the HPNA 2.0 system from the point of view of the network stack and frequency spectrum; the system is a multipoint CSMA/CD packet network that supports unicast, multicast, and broadcast transmissions. 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. On the other hand, 10Base-T requires dedicated point-to-point Category-3 or Category-5 wires.

Physical layer

At the physical layer, the system is frequency division multiplexed on the same wire as used by the standard analog phone service and splitterless ADSL.[6] 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 makes 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.[7] Cross talk between phone lines 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.

While HPNA 1.0 uses PPM, HPNA 2.0 uses quadrature amplitude modulation (QAM), both to get more throughput in the same bandwidth and to achieve greater robustness. However, because the channels may have very deep nulls, and multiple nulls in band, two techniques are used. The first technique, as mentioned earlier, is adapting the modulation rate. 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. A packet header is always encoded at 2 bits per symbol, so that every receiver can demodulate at least the packet's header.

The system uses a fixed 7-MHz carrier frequency and can operate at either 2 Mbaud or 4 Mbaud with modulation encodings of 2 to 8 bits per symbol. The base symbol rate is 2 Mbaud. At this rate, the system has a peak data rate ranging from 4 to 16 Mbps, though overhead reduces the actual throughput the system can achieve. In practice, to achieve performance equivalent to 10Base-T Ethernet, a packet must be sent at 6 bits per symbol.

Frequency-diverse QAM

Unfortunately, the nature of channel nulls can be such that even rate adaptation 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 (that is, one with which the equalizer can't cope) in the band, the system will fail to operate. At its 2-Mbaud rate, HPNA 2.0 implements a modified version of QAM invented by Eric Ojard called frequency-diverse QAM (FDQAM).

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 filter's width, the output signal has two redundant copies of the baseband signal. Thus, the signal is frequency diverse, motivating the name FDQAM.

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 FDQAM performance versus that of QAM on arbitrary channels is more complicated, and this analysis is not included here. However, on channels with a low signal-to-noise ratio, or SNR, where a large part of the spectrum is severely attenuated, FDQAM works robustly in many cases where uncoded QAM would fail. Such channels are common on home phone lines. Unlike most other methods of handling severe channels, FDQAM does not require the transmitter to have knowledge of the channel characteristics. This simplifies the protocol and enables 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 20 Mbps.

Frame format

Figure 4 shows the frame format on the wire. The frame begins with a known 64-symbol preamble. The preamble supports robust carrier sensing and collision detection, equalizer training, timing recovery, and gain adjustment.

Following the preamble is a frame control field, the first part of which is an 8-bit frame type. Frame-type = 0 is shown in the figure, where other codes can be assigned for future system frame formats. Following the frame type is an 8-bit field that specifies the modulation format (bits per symbol, for example). There are other miscellaneous control fields in frame control including an 8-bit CRC header. 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-

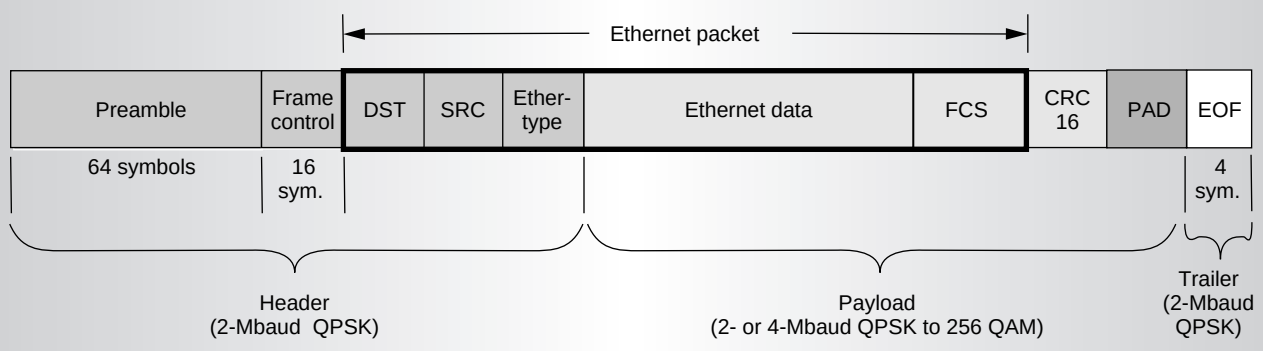

Figure 4. HPNA 2.0 frame format.

Mbaud, 2-bits/symbol rate so that any station able to demodulate a packet can do it at this encoding. Thus, even if the payload is encoded at a rate or bits/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 (RRCF) to the sender, asking it to reduce the number of bits/symbol or the symbol rate.

In practice, the system starts out sending at 2 bits/symbol unless the receiver sends an RRCF, asking for future packets to be sent at higher data rates. Several algorithms can be used to determine when to send RRCFs and to estimate the channel capacity using an approximate SNR and bit error statistics. The rate adaptation algorithm can optimize the rate used when sending to multicast and broadcast groups.

Media access control

As mentioned earlier, HPNA 2.0 is a CSMA/CD system, just like the 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.

For example, three nodes (N1, N2, and N3) could contend for access to the network. Node N2 is transmitting a voice-over-IP (VoIP) packet. Initially, N0 accesses the wire and transmits a frame (TX). 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 back-off 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 back off. In this manner, the queuing discipline can become very unfair for N2. If N0 and N1 are PCs 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 is to introduce different access priorities in which the VoIP station uses a higher priority than best-effort file transfer traffic. HPNA 2.0 accomplishes this by organizing the time following the interframe gap into an ordered series of priority slots. See Figure 5.

When N0 finishes transmitting, all stations on the network with a 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, and penalty. Each uses a different priority level.

Within a given priority level, HPNA 2.0 uses a new algorithm for collision resolution. Each station keeps track of a back-off level and after a collision, randomly chooses to increment the back-off level by 0, 1, or 2. During a collision resolution cycle, stations

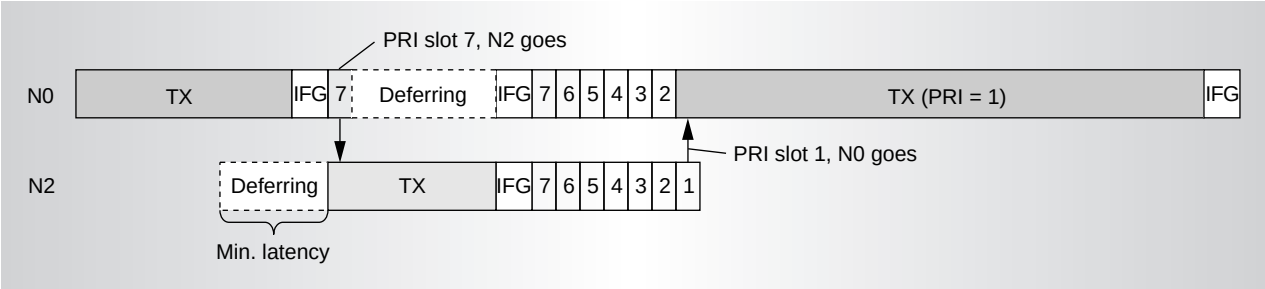

Figure 5. Priority example.

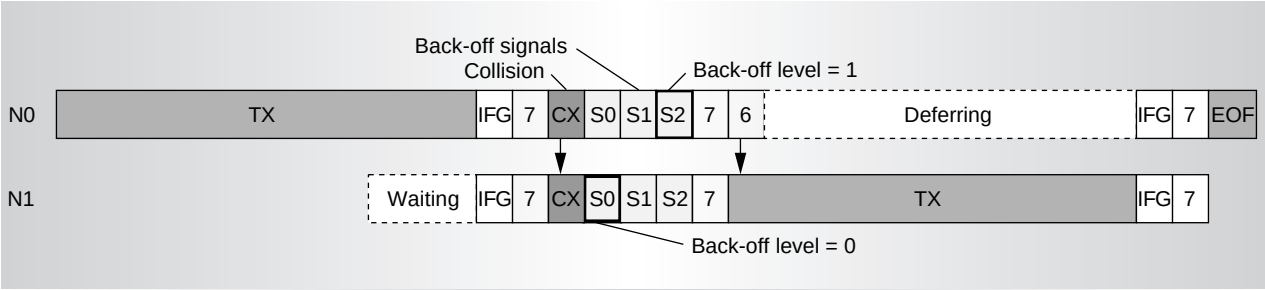

Figure 6. HPNA 2.0 collision resolution algorithm.

incrementally establish a partial ordering. Eventually, only one station remains at the lowest back-off level and gains access to the channel.

In the example shown in Figure 6, N0 and N1 enter into a collision resolution cycle. N0 randomly chooses to increment its back-off level by 2, and N1 by 0. To optimize the partial ordering, eliminating null levels, stations send a special signal immediately following a collision. The signal reflects the back-off 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 back-off level by 1 because it saw the back-off signal from N1 in S0, but no station indicating in S1.

In practice, even on saturated networks, HPNA 2.0 behaves very well, and unlike traditional Ethernet does not exhibit the capture effect. As the offered load increases, Ethernet experiences delays of hundreds of frame times for several percent of transmission attempts. When compared with the Ethernet delay, the HPNA 2.0 access latency distribution reveals a negligible tail beyond twice the minimum time to service each active station, even at high offered loads.

Link layer protocols

Another impairment is a problem for all home networks using no-new-wires technology: impulse noise. On phone wires, 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 operates only on a single segment of the network, it is very effective in hiding packet erasure from TCP/IP, as demonstrated in Figure 7.

Finally, note that 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 is traffic on the wire, in which case a reduced number of frames may be sent.

iLine10 chip set

The iLine10 two chip set implements a combined MAC/PHY approach for both HPNA 1.0 and HPNA 2.0. The major components of a complete PCI NIC board are

- the BCM4210 MAC/PHY chip,
- the BCM4100 analog front-end chip,
- a magnetics module for phone line isolation and protection,
- a serial PROM with MAC address and other configuration, and
- a crystal.

BCM4210 MAC/PHY chip

We implemented the primary component of the system—the BCM4210 MAC/PHY chip—using standard cells together with compiled blocks for registers and memory. Approximately half the chip is dedicated to the PHY signal-processing functions. The frame processor implements MAC layer functions such as CRC processing, address filtering, and wake-on-LAN. Figure 8 shows the major blocks of the BCM4210; Table 2 describes its parameters.

On transmission, the system transfers packets via DMA over the host interface into an on-chip frame buffer. Each packet is augmented with a preamble, frame control, FCS, CRC16, and padding in the block labeled framing. All but the initial part of the frame is scrambled and encoded into symbols, where the number of bits/symbol depends on frame position and rate adaptation. The symbols are upsampled and modulated into a QAM-like signal and then filtered. The receiver reverses the functions performed by the transmitter.

The carrier sense and collision detection blocks are critical DSP functions. These operations must be performed on signals with

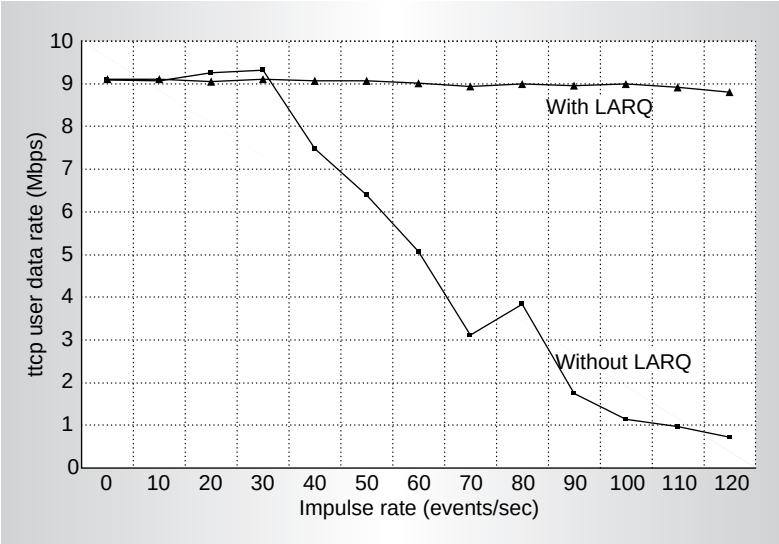

Figure 7. User-level throughput versus impulse noise events/sec.

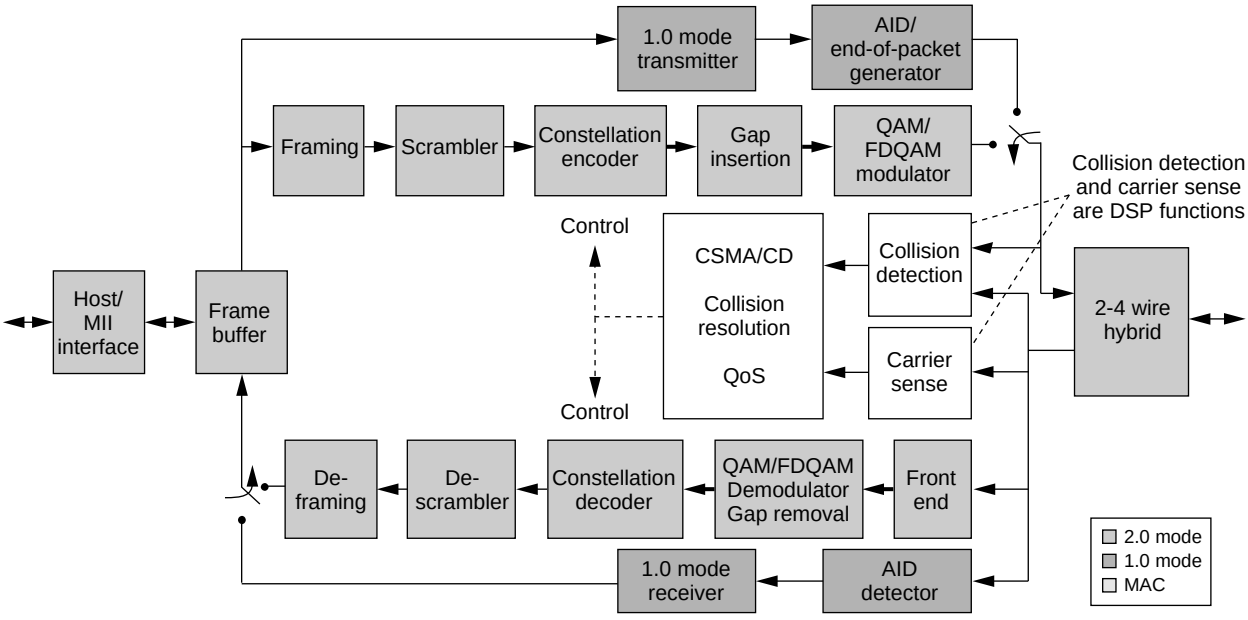

Figure 8. Block diagram of the BCM 4210.

widely varying frequency-dependent attenuation, in real time.

We chose a high sample rate of 32 Msamples/sec to eliminate the need for precise analog domain filtering. Approximately 500 Moperations/sec are required for receive filtering, carrier sensing, and collision detection, and this led us to prefer a hardwired data path. Each of the chip's subfunctions has its own multipliers and adders. To minimize area, we kept the precision of each multiplier-accumulator to the minimum required for that step of the algorithm. The transceiver operates in two modes, native 2.0 and backward-compatible 1.0, so that both HPNA 1.0 and HPNA 2.0 can interoperate on the same phone wire.

The host interface includes a complete 33-MHz, 32-bit PCI master and slave, as well as a variety of other slave-only interfaces that allow the BCM4210 to be used both in standard PCs and a variety of embedded applications. The DMA engine supports a descriptor-based architecture similar to what is found in other Ethernet MAC chips. The frame processor implements perfect filtering for up to 64 MAC addresses and an imperfect filter for 2,048 addresses. The chip set is PC99 compatible, implementing four 128-byte wakeup patterns and dissipating less than 375 mA when in D3 sleep state.

Table 2. Implementation parameters for the BCM4210 MAC/PHY chip.

Parameter	Description
Process	0.25-micron CMOS
Package	144-pin TQFP
Gates	250,000
Memory size	80 Kbytes
Clock speed	64 MHz
Voltage	3.3 V (core and pads)
Power	500 mW
Die size	4.5 mm × 4.5 mm
Equivalent transmit (operations/sec)	0.7 billion
Equivalent receive (operations/sec)	2.2 billion

BCM4100 analog front end

Figure 9 is a block diagram of the BCM4100 analog interface to the phone wire. The receive path dominates the active chip area. We carefully designed the layout and wire placement to minimize cross talk between digital and analog blocks. Table 3 summarizes the key implementation parameters.

One of the challenges of phone line networking is that due to both FCC part 68 requirements for maximum signal levels on phone lines and avoiding audible noise in phones caused by accidental envelope detection of the HPNA signal, the transmit power is severely limited. The HPNA 2.0 system has a peak signal level of approximately 750 mV peak to peak, and a PSD (power spectral density) of –74dBm/Hz.

Because the system is intended to operate with almost 40 dB of insertion loss, the minimum received signal can be less than 10 mV. This amount of insertion loss, combined with the relatively modest transmit voltage, means that significant care was necessary not to introduce additional noise in the analog front end.

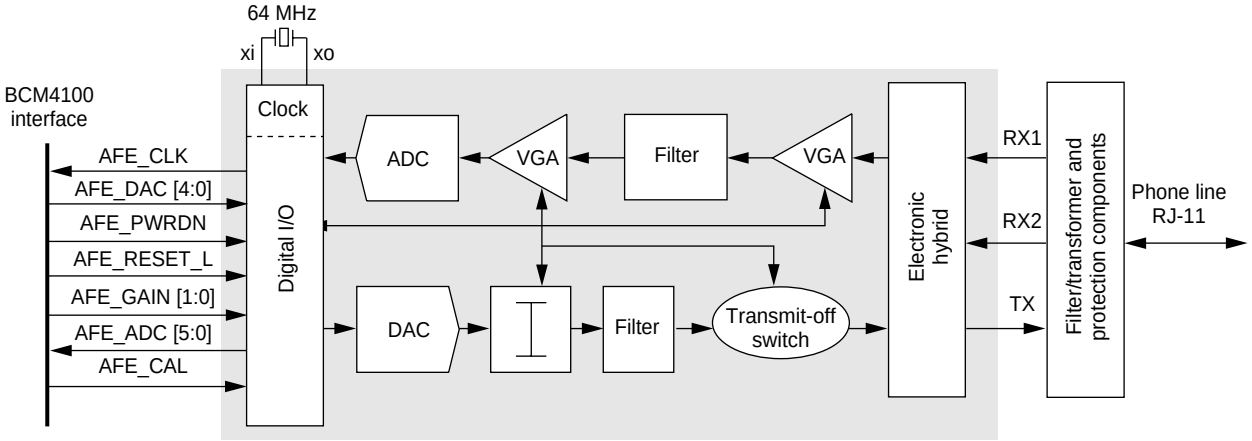

Figure 9. Block diagram of the BCM4100.

For example, we determined that the combination of the DAC and the active transmit filter generated several dBs of noise, whether or not the transmitter was active. As a result we added a low-noise output switch on the transmit path that decouples the on-chip circuitry when no packet is being transmitted.

The ADC uses a pipelined architecture and achieves an SNDR of 57 dB, or an equivalent number of bits (ENOB) of almost 9.2 bits. The dynamic range required led us to a design that uses automatic gain control. The analog front end, or AFE, contains a variable gain amplifier that is controlled digitally by the BCM4210 as a function of an estimation generated during the preamble processing.

Table 3. Implementation parameters for the BCM4100 analog front-end chip.

Parameter	Description
Technology	0.35-micron, double-poly CMOS
Package	48-pin TQFP
Voltage	3.3 V
Power	500 mW
DAC resolution	10 bits
ADC resolution	10 bits
SNDR (ENOB)	57 dB (9.2 bits)
Sample rate	32 Msamples/sec

Magnetic and serial PROM

The magnetics module contains a transformer that provides electrical isolation and a hybrid function. Although several different implementations are possible, for performance reasons we use a split winding design. The external transformer together with a small number of passive components performs some low-pass filtering to remove AM radio, POTS, and ADSL signals that may be on the wire.

The primary purpose of the serial PROM is to store the 48-bit Ethernet MAC address. In addition, it contains information such as whether hardware link integrity is enabled by default, the Cardbus Information Structure, and PCI subsystem vendor IDs. Software can read and write the serial PROM to ease manufacturing. A locking mechanism prevents certain key data from being overwritten in the field.

Verification environment

Because of the complexity of these chips, particularly the BCM4210, it was necessary to implement a very robust verification environment. There are two key elements of this environment. The first is a channel modeling system that was developed based on measuring actual homes, combined with theoretical models for parameters such as wire gauge and wire balance. This channel model allows us to generate "cooked" packets. That is to say, it allowed us to create a set of digital samples corresponding to a packet that traversed a particular channel from one node to another.

The other key element is a C language implementation of both the transmitter and receiver. In its simplest mode, the verification system could be used to transmit a packet, which is then cooked using the channel model and demodulated using the receiver. In practice, the transmitter implementation is much simpler than the receiver. Using the channel modeler, we could generate millions of packets (thousands of packets cooked by thousands of network topologies) to verify that the receiver was operating correctly. In addition we hand-generated several worst-case packets that had, for example, worst-case peak-to-average signals, timing offset, and so on.

The next important use of the C implementation of the system is that it is possible to use it as a "soft-PHY." In this mode, together with a hardware analog front-end board that plugged into a PC, we could actually send and receive packets at a reasonable rate, prior to silicon availability. By using Pentium MMX instructions to augment certain inner loops, we could demodulate at rates greater than 4 Mbps on a single-processor 400-MHz Pentium. This soft implementation allows us to verify that the system operated correctly, debug the upper layers of the drivers and the link layer protocols such as LARQ, and perhaps most importantly give realistic and impressive demonstrations of the system.

We also used the C implementation in debugging the VHDL software used to implement the BCM4210. With great care and effort, we matched the C and VHDL implementations on a cycle-by-cycle basis in both time and arithmetic calculations on major block boundaries. We viewed any discrepancy between the two implementations as a bug. Though the general transmitter and receiver algorithms used in both the C and VHDL

implementations were the same, their details were very different. The C implementation processes a block of samples at once, while the VHDL implementation is a sample at a time in the hardware pipeline. By matching these two implementations at block boundaries, we convinced ourselves that the hardware correctly implemented the DSP algorithms.

Finally, the verification environment provided the means to connect multiple instances of either hardware or software simulations. This let us verify that the MAC was functioning correctly.

Home networking is now a reality. Using advanced signal-processing techniques and high-density CMOS makes it possible to transmit data over existing media, such as in-place telephone 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 become 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. MICRO

Acknowledgments

The work described here is the result of the efforts of many people. In particular we recognize Ruben Alva, Neal Castagnoli, Alan Corry, Mike Dove, Greg Efland, Nino Ferrario, Hariprasad Garlapati, Ali Hariri, 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, Jeff Vadasz, Chris Warth, Larry Yamano, and Chris Young.

References

1. *HPNA 2.0 Specification,* Oct. 1999; available to members of HPNA, see http://www.homepna.org.
2. R. Metcalfe and D. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Commun. ACM,* New York, N.Y., Vol. 19, No. 5, July 1976, pp. 395-404.
3. M.L. Molle, *A New Binary Logarithmic Arbitration Method for Ethernet*, Tech. Report CSRI-298, Computer Systems Research Institute, Univ. of Toronto, Apr. 1994.
4. X-10 FAQ, ftp://ftp.scruz.net/users/cichlid/public/x10faq.
5. *HPNA 1.0 Specification,* Nov. 1999; available to members of HPNA, see http://www.homepna.org.
6. W. Chen, *DSL: Simulation Techniques and Standards Development for Digital Subscriber Line Systems,* Macmillan Technical Publishing, Mar. 1998, ISBN 1578700175.
7. *Spectrum Selection for Home Phone Line Networking,* ITU-T Study Group 15, NG-101, Nuremberg, Germany, 2 Aug. 1999.

Edward H. Frank cofounded and is executive vice president of Epigram, now the Home Networking Division of Broadcom Corporation. Previously, he served as vice president of engineering for NeTpower, Inc., and was a Distinguished Engineer at Sun Microsystems. At Sun, he worked on the Green Project, which created Oak, the precursor to Java, and was coarchitect of several generations of SPARCstation workstations. Frank received his PhD from Carnegie Mellon University, and his BSEE and MSEE degrees from Stanford University. He holds approximately 15 issued patents.

Jack Holloway cofounded and is the chief technology officer of Epigram, now the Home Networking Division of Broadcom Corporation. Earlier, he served as vice president of broadband at MicroUnity, a closely held mediaprocessor developer, and as vice president, broadband technology for Bolt Beranek and Newman. At BBN, he started an ATM product division called Lightstream, which was eventually acquired by Cisco Systems. Holloway also cofounded Symbolics, Inc., and for several years was a principal research scientist at MIT and codirector of the MIT VLSI Laboratory.

Direct questions to Jack Holloway, Broadcom Corporation, 870 West Maude Ave., Sunnyvale, CA 94086; h@epigram.com.