Home networking leaps to digital

Keeping up with the demand and proliferation of technologies and applications, home networking breaks the analog barriers.

By Guillaume Bichot, Louis Litwin and Kumar Ramaswamy

With the proliferation of digital technologies and digital content, today's home networking techniques are jumping on the digital bandwagon as well. The biggest driver of networking in the home is the Internet and broadband connectivity, and the desire to access the Internet from multiple points in the home.

Two options — wired and wireless

Present key technologies for home networking are:

• Wired Connectivity — Today's homes typically have power line, telephone wiring and cable wiring that can be used for home networking. However, access to these wires is restrictyed or inconvienient. Also, it is not desirable to expose to a neighbor (both for reasons of coexistent resource sharing and privacy) the data carried on a shared wire (such as the power line).

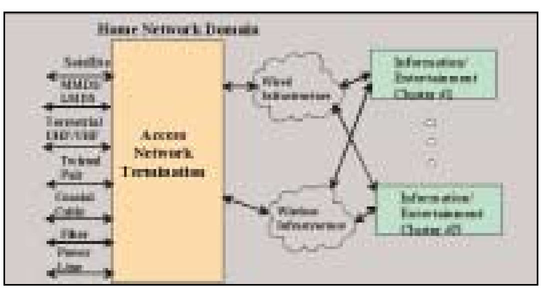

Figure 1. Logical Representation of the elements of a home network.

• Wireless connectivity — Wireless solutions are elegant in the fact that they do not require any modification to the existing infrastructure. Therefore, the installation and maintenance of this technology is easier.

However, wireless technologies tend to be more expensive than wired technologies for a given capability. With emerging high-speed silicon integration, this disadvantage starts to fade, making the case stronger for investment into wireless technologies for the home.

Key issues for digital home networking:

Digital home netoworking has many dependent and independent issues. Among them are:

• Network and access protocols — Network and access control protocols are an important element of the overall home networking system. These provide the rules, or the etiquette, for sharing the underlying resources among various devices in the network. There is a further problem of making sure that these rules are consistent among the various wired and wireless physical layers, which will be addressed later. Because these standards are developed in diverse industrial groups, ensuring network interoperability is often a challenge.

• Security, privacy and content protection — In addition to making sure that content is seamlessly available within the home it is also important to ensure that this is done securely and maintains privacy. It is also important, from the viewpoint of a content provider, that, whenever content is consumed, there must be a mechanism to monetize the transaction.

• Coexistence, interactivity and interoperability — Given the number of possible options on technologies and competing standards, it is important that these standards and technologies coexist. The average consumer would expect "plug and play" operation on most devices and applicances, hence it is important to guarantee this by design. In practice, however, due to the emergence of competing technolgies, this is often done as an afterthough.

Interoperability — a must have

Interoperability assumes direct interaction between devices to enable a certain functionality. It is easier to explain this with an example.

Consider a user who wishes to record a specific program on a digital video recorder from a television set that is connected to the recorder. The user starts with the interface of the television to pick the specific channel of interest. The next step would be to interact with the user interface of the recording device and request a recording. This is the interactivity paradigm, which should be replaced by interoperability. In the interoperability paradigm, the user simply requests a recording of a program (for example, using the user interface of the television set). The television set is responsible for coordinating the recording of the program in an available storage device anywhere in the network. This will clearly require an additional layer of sophisticated protocols among the devices involved.

Services for a home network

A typical list of services that may reside in a home network is shown in Table 1. It is obvious that several services with different constraints on delay and bit rates are possible. It must also be noted that the error rates at which these services need to be delivered are also different for each application.

Also, the data rates and constraints for Internet content are dictated by the type of service exercised. For example, if a video stream is being downloaded for immediate viewing, the delay constraints can be important, while they may be less important while viewing a Web page. Clearly, the design of the protocol and the digital communication technology itself is to be guided by the choice of intended appli-

Type of Service	Data Rate Range	Type of Communication	Delay Constraints?
Voice	8- 64 Kbps	Bi-directional	Yes
Hi-Fi Audio	64 Kbps- 4 Mbp/s	Unidirectional	None
Video Telephony	384 Kbps-1.5 Mbp/s	Bi-directional	Yes
Broadcast Video - Standard Definition	1.5 Mbp/s - 6 Mbp/s	Unidirectional	None
Broadcast Video - High Definition	12 Mbp/s - 18 Mbp/s	Unidirectional	None
Computer/Internet Data	Varies	Bi-directional	Depends on Application
Home Automation	Several Kbps	Bi-directional	Depends on Application

Table 1. Potential application services for distribution in a home network.

	Bluetooth	802.11b	802.11a
Max power level:	1 W	1 W	800 mW
Frequency band:	2.4 to 2.483 GHz	2.4 to 2.483 GHz	5.15 to 5.35 GHz, and 5.725, to 5.825 GHz
Data rate:	1 Mbp/s	1 to 11 Mbp/s	12 to 72 Mbp/s
Modulation:	FHSS	IR, FHSS, DSSS	OFDM
Range:	10 m, 100 m with external power amp	100 m	15 - 30 m

Table 2. Summary of several wireless digital home networking standards.

cations over different home networking technologies.

Elements of a home network

A logical representation of a home network domain is shown in Figure 1. Several possible broadcast and interactive mechanisms exist to access the outside world through the access network termination. It is also possible that some of these mechanisms function cooperatively. For example, a digital broadcast satellite and a digital subscriber loop (DSL) services.

As alluded to earlier, the existing wiring infrastructure in the home includes the power line, cable and twisted pair wiring. The power line medium is a harsh medium with large variations in dynamic characteristics. Furthermore, due to radiation regulations, the operational frequencies are restricted to below about 30 MHz, with limited power. The power line medium is likely to find a niche in data redistribution — especially from broadband modems — but is unlikely to be a solution for high-speed video streaming applications.

The twisted pair wiring in the home carries plain old telephone service (POTS) in the 0 to 4 KHz band. The frequency band up to 1.1 MHz is allocated for digital subscriber loop (DSL) applications. The Home Phone Network Allicance (HPNA) developed the technology to carry digital data in the home phone copper wires above the DSL band. But given the limited number of phone access points in an average home, this technology is not expected to play a major role in home networking.

Cable wiring for redistributing content also has a limited appeal due to the number of available cable TV outlets in the average home.

Furthermore, the channels that are allocated for television on the cable network are dependent on the local service provider. So, it is difficult to design a home network to share the cable resource with a service provider.

The physical layers

This section of the article will examine some of the physical layer aspects of wireless digital home networking, problems and issues faced when operating a wireless device in the home. It will describe some techniques for handling those problems. In particular, we will use examples from three popular wireless digital home networking standards: Bluetooth, IEEE 802.11b, and IEEE 802.11a.

Spectrum allocation

Most wireless digital home networks (WDHNs) operate in one of the frequency ranges that belong to the unlicensed industrial/scientific/medical (ISM) bands. The FCC has designated the ISM bands as license-free as long as devices stay below the specified transmit power level. In the United States, there are three sections of the ISM bands: 902 to 928 MHz, 2.4 to 2.483 GHz, and 5.725 to 5.850 GHz.

Most WDHNs operate in the 2.4, 5.7 and 5.15 to 5.35 GHz unlicensed unlicensed national information infrastructure (UNII) bands. Both the Bluetooth and IEEE 802.11b standards use the 2.4 GHz band and the IEEE 802.11a standard uses the 5.7 GHz band.

The wireless home environment

Many factors make the home a difficult environment in which to operate a wireless communications system. Interference from man-made sources is a major issue. Many devices in the home radiate energy (either intentionally or unintentionally) in the ISM bands.

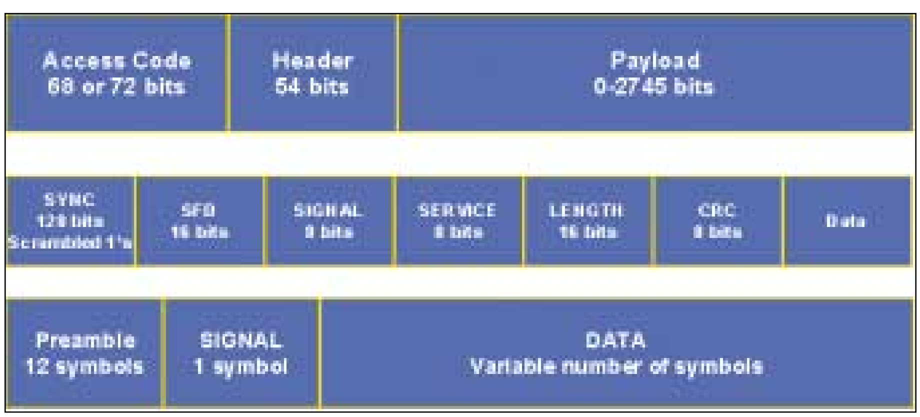

Figure 2. Burst formats for Bluetooth (top), IEEE 802.11b, and 802.11a.

Examples include interference from microwave ovens, cordless phones, amateur radio transmissions and radar. Because several WDHN devices operate in the ISM bands, there will also be interference in areas where multiple networks are being operated with devices from different standards.

Because the home is essentially an enclosed box, WDHN devices must cope with frequency-selective channels caused by a multipath environment. The effect of a frequency-selective channel is intersymbol interference (ISI). ISI occurs when the in-home reflections cause several delayed versions of a signal to arrive at the receiver. Typical delay spreads of an indoor channel are between 100 to 200 ns within a room and greater than 300 ns in hallways. These delay spreads can be even greater when using WDHNs in office or factory environments. The dynamic environment of the home, such as moving people and animals, as well as moving objects such as ceiling fans, means that the in-home channel is rapidly time-varying as well.

Another problem faced when operating in the home is that of path loss. Severe signal attenuation occurs when trying to pass the signal through the walls of a home at high frequencies. This attenuation can range from about 1 dB for plywood to more than 20 dB for metal walls. Thus, the material that a house is made out of can greatly affect the useful range of a WDHN device.

ISM band power limits

The FCC limits the transmitted power sent to the antenna in the ISM bands to 1 W. Spectral density limits also restrict the power density in any 3 kHz bandwidth region to 6.3 mW. By placing these requirements on the ISM band usage, the FCC prevents a user from severly degrading the performance of neaby wireless devices.

To stay below these limits, many WDHN radios, such as Bluetooth and IEEE 802.11b, use spread-spectrum (SS) modulation where the transmitted signal spectrum is significantly wider than the original data bandwidth. Instead of SS modulation, IEEE 802.11a uses orthogonal frequency division multiplex (OFDM) modulation (OFDM is described in detail later).

Basics of spread-spectrum

Two basic types of spread-spectrum modulation exist. Frequency-hopping spread spectrum (FHSS) is used by Bluetooth. The Bluetooth spectrum is divided up into 79 channels of 1 MHz each. After transmitting a burst on one channel, the transmitter hops to another channel and transmits the next burst at that frequency. The Bluetooth radio hops over these channels at a rate of 1.6 khops/s. By transmitting at relatively high power for only a short period on any given frequency, the signal's average power is kept below the FCC's limit. In addition, the radio is able to hop around interfering signals. Thus, a narrowband interferer will only be a problem during the short time that the radio is using that 1 MHz channel.

IEEE 802.11b radios support both FHSS and direct sequence spread spectrum (DSSS). In contrast to FHSS, a DSSS radio sends a signal that is spread over a wide bandwidth, but the spectral density at any given frequency is very low. The signal in DSSS is spread by multiplying the data signal by a wide bandwidth spreading sequence. Each bit in the spreading sequence is known as a chip. IEEE 802.11b uses an 11-chip Barker code as the spreading sequence for its 1 and 2 Mb/s modes.

To send information, the transmitter takes either one or two bits and uses those values to differentially modulate the spreading sequence using either binary phase-shift keying (BPSK) or quadrature phase-shift keying (QPSK). Thus, the information is conveyed in the phase of the spreading sequence.

The higher date rate modes of 5.5 and 11 Mb/s use a different type of modulation known as complementary code keying (CCK). In addition to using two bits of the symbol to perform differential QPSK modulation, the remaining bits of the symbol are used to choose from a set of either four or 64 complex orthogonal spreading sequences. This selection is performed on a symbol-by-symbol basis. Thus, in CCK modulation, information is contained in both the phase of the spreading sequence and in the choice of the spreading sequence used. By using a wide bandwidth signal, a DSSS signal will interfere with all frequencies in the channel. However, the low spectral density means that the interference will be small.

The DSSS radio has the drawback that it must always transmit on the entire channel as opposed to having the flexibility to hop around other interfering signals.

OFDM and ISI

WDHN radios use different techniques to cope with the severe frequency-selective channels that occur in the gigahertz bands within the home. FHSS systems such as Bluetooth are able to hop around channel nulls so that such nulls will only distort the received signal for a brief amount of time. IEEE 802.11a radios take a different approach by using OFDM. The 20 MHz spectrum is broken down into 64 sub-bands of 312.5 kHz each. The user data are modulated onto 48 subcarriers and four subcarriers are used to carry known pilot sequences that aid in signal recovery. The remaining band-edge and DC subcarriers are zeroed out to reduce the effects of co-channel interference.

An advantage of using a multicarrier modulation such as OFDM is that it takes 64 symbols that have a short symbol duration and it combines them all into a single OFDM symbol. This OFDM symbol has a symbol duration 64 times as long as the original symbol duration. In single-carrier systems, the time span of the channel is typically much longer than the symbol period. Hence, ISI occurs due to interference from a large number of adjacent symbols. However, the long symbol period in OFDM means that the symbol duration is longer than the timespan of the channel. The ISI in an OFDM system is due to the previous symbol, and this ISI will only affect the first few samples of a given OFDM symbol.

To combat the effect of this ISI a cyclic prefix is used. The cyclic prefix is formed by taking 16 samples from the end of the OFDM symbol and prepending a copy of those samples in front of the symbol. Because ISI is only due to the previously transmitted symbol, only the cyclic prefix will be corrupted. Because the cyclic prefix contains redundant information, it can be discarded at the receiver and thus the effects of ISI are essentially thrown out.

Another type of interference known as intrasymbol interference occurs when delayed versions of a given transmitted symbol cause the symbol to interfere with itself at the receiver. Intrasymbol interference is corrected by using a frequency-domain equalizer that has a single complex equalizer tap for each subcarrier.

Bursty communications

Unlike a broadcast system such as digital television, WDHNs are bursty communications systems. Because the receiver only has a short burst of data to work with, it must quickly perform functions such as carrier and timing recovery and channel equalization to be able to successfully recover the entire burst of data.

To aid the receiver in this process, preambles are prepended to the beginning of a burst. The preambles contain known sequences that the receiver can use to synchronize itself prior to the reception of the user data part of the burst. The preambles will also contain information about the length of the user data payload as well as the data rate of the payload. To ensure good reception of the preamble, it is typically transmitted using the most robust format available in a given system. Many systems send the preamble using BPSK modulation and rate 1/2 convolutional coding.

Applications

The three standards used as examples are suited for different applications. Bluetooth is typically used as a cable replacement solution to wirelessly connect devices that were previously connected using a cable.

An example is connecting a laptop to a cell phone. Bluetooth's low cost and low power consumption allow it to be integrated in a variety of products that need a low data rate (1 Mb/s) and short range (10 m) wireless link. IEEE 802.11b's higher data rate (11 Mb/s) and longer range (100 m) make it better suited for use as a wireless data network.

A typical application uses the 802.11b link to give devices such as laptops and PDAs mobility while still retaining the ability to connect to a wired data network or a phone line. IEEE 802.11a radios provide a high data rate (72 Mb/s) over a range of between 15 to 30 m. The high throughput of this link allows audio and video to be wirelessly sent throughout the home or office.

Link layer protocols

The link layer is the layer between the physical layer and the network layer according to the open software interconnection (OSI) layering model specified by the International Standard Organization (ISO). The physical layer allows the bit coding and transport of the information over the wireless medium. The link layer provides protocols to enable communications between two devices interconnected through the medium.

When the medium is shared among several interconnected devices, the layer is split into two sub-layers. The lowest sub-layer (directly over the physical layer) is called medium access control (MAC) and it controls and arbitrates the access to the shared medium. The highest sub-layer is called logical link control (LLC) and provides the protocols to ensure the communication between two devices directly connected to the medium. Typically, the service offered by the LLC sub-layer, and hence by the link layer, is either a reliable (or acknowledge) service where the receiver acknowledges each data frame or packet, or a non-reliable service where a data frame or packet is sent without guarantee that it has been received by the destination device.

Wireless link

From the link layer point of view, a wireless link is no different than a wired except for the issue of reliability. The bit error rate may be over 10^{-3} which can increase packet loss and implies dedicated strategies such as forward error correction (FEC) and/or retransmission.

Another issue is the way a device has to be connected to the medium. In a wired medium, the physical connection consists of plugging the device into the medium. In a wireless medium, the device has to automatically perform certain operations called scanning and joining to recognize (and be associated with) the network. The wireless nature of a link implies some inherent problems such as the hidden node problem.

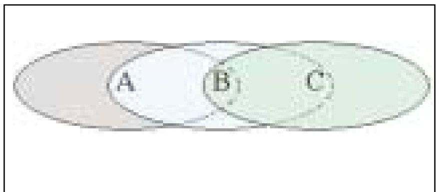

Figure 3. Illustration of a hidden node.

A wireless link is characterized by the maximum range within two devices can communicate. Consider the case in which three devices (A, B and C) are located in such a way that A and B are in range together, and B and C are in range together, but A and C are not in

	IEEE 802.11	Hiperlan2	Bluetooth
Network mode	Ad-hoc Infrastructure	Ad-hoc Infrastructure	Ad-hoc
MAC	Distributed: CSMA/CA	Centralized: scheduler	Centralized: polling
QOS	No	Yes: unicast + multicast	Yes
Multimedia	Near isochronous service	Fixed slot allocation	Circuit switched service Near isochronous service
Handover	Inter-base stations protocol not specified	Inter-base stations protocol not specified	
Security	Basic	Ok	Ok
Error correction	Retransmission	Retransmission FEC (home extension)	Retransmission FEC
Misc	Interferences within the 2.4 *GHz* band Mounted attacks against the security barrier	Difficult to implement the Scheduler	Limited set of active device within the scatter-net Interference within the 2.4 *GHz* band Limited interactions between slaves in a scatter-net

Table 3. Wireless LAN comparison.

range (see Figure 3).

A starts to transmit information to B. C senses the medium and assumes that the medium is free (A being out of range). It consequently starts to transmit to B as well. The consequence is interference at B. The MAC protocol needs to provide some means to solve this problem.

Another important issue is power management. There are a few fundamental issues to be resolved in relation to power management. The first one is related to the power of the transmitter that has to be adjusted according to the distance between the transmitter and the receiver. The physical layer usually performs this function. The second function relates to the power consumption of the devices. The mobile device is battery operated. Therefore, it must be power efficient. The MAC sub-layer may provide some mechanisms to control and limit the power consumption by allowing the device to be in "sleep" mode and to be awakened when some message is intended specifically for it.

The final important issue is that of security. Wireless links are prone to eavesdropping. Also, unauthorized users could access the network unless there are explicit mechanisms that prevent them from doing so. Consequently, the link protocol usually has some security mechanisms such as authentication and encryption.

Wireless architecture

Two types of architectures are commonly found in wireless local area networks (LANs). The first, called the infrastructure mode, implies that some device which forms the network infrastructure is required. Such a device is called a base station (or access point). This device provides some additional intelligence and memory resources to extend the wireless interconnections and to help coordinate bandwidth management.

A network may contain several base stations. The base stations may be interconnected through either a wired or wireless link. In this mode, a wireless device is said to be attached to a base station, meaning all transmissions coming from or to be sent to this device go through the base station. The base station acts as a relay and can buffer some data when the destination device is temporatily unavailable. If the wireless device is mobile, one base station may perform a handoff to another base station seamlessly. The process or protocol that provides this function is usually implemented within the wireless link layer and is called handover.

The second network architecture is called ad-hoc mode. In that mode, no infrastructure device exists such as a base station. Two devices may establish a communication relationship autonomously when they are in range of each other. It is also possible for a device be identified as a controller after an automatic initialization to control the medium access for instance. However, it is still in an ad-hoc mode because any device can potentially play that role.

MAC protocol

Two types of MAC protocols exist. The first type, called a distributed MAC, includes the protocols where the control of the access to the medium is distributed. This implies that all devices participate in that control and that there is not a central controller. Thus, all contender devices have the same chance to gain access to the medium. The second class of MAC protocols are called centralized MAC protocols in which one device controls and manages the access to the medium.

Wireless LANs

IEEE 802.11 specifies one MAC layer and several different physical layers. The LLC sub-layer is already specified by the IEEE 802.2 group and used over Ethernet (802.3) for instance. The obvious consequence is that the interconnection between an IEEE 802.11 sub-network and an Ethernet LAN, for instance, is natural. IEEE 802.11 defines two network modes, an ad-hoc mode and an infrastructure mode as defined above. IEEE 802.11 defines the service offered by the distribution system that consists of the set of interconnected access points (through either wired or wireless links), but doesn't specify the protocols to ensure handoff between these base stations.

The MAC protocol is carrier sense multiple access/collision avoidance (CSMA/CA) and it is a distributed protocol. All stations, regardless of architecture (ad-hoc or infrastructure), have the same status and the same priority regarding the medium access. The method in itself is based on a random process where a station that wants to acquire the medium first senses the medium. If some activity is detected (the medium is busy), the station has to wait a certain time corresponding to a certain number of time slots randomly chosen. After this waiting time, the station attempts to get access again. If it is still busy, the station has to double the time it waited before this attempt and so on.

This process is known as the binary exponential back-off mechanism. If no activity is detected during a certain gap then the medium is seen as free and the station may start to transmit. The basic transmission is called a transaction and it consists of sending a data packet to the receiver and waiting for the acknowledgement from the same receiver. To avoid the hidden node problem, the transaction may contain a short handshake before sending the data packet. This handshake is called request-to-send/clear-to-send (RTS/CTS) and it allows all other stations, in the range of the receiver (which will send back CTS) to detect the upcoming transaction. Embedded within each packet of the transaction (RTS, CTS, DATA, ACK) is a duration parameter, which informs all other stations waiting to acquire the medium about the duration of the transaction. This is the collision avoidance part of the MAC protocol.

Hiperlan2 also specifies two network modes: ad-hoc and infrastructure. The stack includes (from bottom to top) the physical layer, a data link control (DLC) layer comprising the MAC sub-layer, the radio link control (RLC) sub-layer and some convergence layers.

A convergence layer gathers specification to adapt other link protocols such as ATM, IEEE 802.3 (Ethernet), and IEEE 1394 over the Hiperlan2 protocol stack. The MAC protocol is a centralized one. A

Continued on page 58

Continued from page 38

station has the responsibility to manage/schedule the access to the medium. This station is the base station for the infrastructure mode and the central controller (decided on automatically by auto-configuration) for the ad-hoc mode. The time is divided into frames. Each frame gathers a set of time slots that correspond to transport channels. A transport channel contains either signaling information or data currently transmitting between the controller and a station. A station that wants to use the medium sends a signaling message (resource request) to the controller. The controller processes the request, allocates some radio resources (a set of time slots) to the requester and sends back the signaling response message. The station can then use these dedicated time slots to send data.

The DLC sub-layer includes some signaling protocols for power control, scanning and joining, and to set up some logical connections where a certain quality of service may be negotiated between the station and the controller. This QOS implies the reservation of time slots for several frames in order to provide an isochronous service.

The Bluetooth protocol stack is composed of the physical layer, the baseband sub-layer (MAC) a logical link control (LLC) sub-layer and a set of application-oriented protocols. Bluetooth proposes only the ad-hoc network architecture mode. However, this may be divided into two different sub-modes. The first defines a peer-to-peer relationship between two stations, whereas the second defines some one-to-multiple stations relationship to provide multicast service. The latter mechanism forms a so-called “scatter-net”.

A controller named the master has the responsibility to control the medium access. Any Bluetooth device has the ability to be a master. In fact, the master is the device that initiates the connection during the scanning/joining phase. In case of a scatter-net, the master may control as many as eight “slave” stations simultaneously.

The MAC protocol is a centralized one in which the master controls the medium access by a polling method. The time is divided in time slots. The even slots are dedicated to the master and the odd slots to the slaves. The master polls any slave station sending data and/or signaling information; the slave sends back its response and/or other signaling/data information through the following time slots. The master then polls the following station and so on.

Bluetooth specifies two services: an isochronous service (based on the reservation of dedicated slots) where up to three isochronous connections may be established between a master and one or several slaves, and an asynchronous service allowing only one connection between the master and its slaves. The master poll interval may be negotiated providing a certain level of quality of service (QoS).

Comparison

Table 3 points out the differences, drawbacks and advantages of the three protocols focusing on audio/video services within home networks. The security drawback for IEEE 802.11 is related to the size of the encryption key which is too small and the authentication mechanism which is not mutual. Hiperlan2 and Bluetooth both specify an FEC mechanism that might be useful for video transmission.

With its lower bit rates and lower power, Bluetooth is intended to be primarily a solution for a lower speed data connectivity between devices in a room. It is not expected to provide complete home coverage which may limit its use, especially for video based applications. The master/slave model has an impact on how the application is designed that implies some limitations regarding interactions between stations.

Conclusions

Wireless Home Networks are driven by applications. The primary driver for the first wave of applications is the sharing of broadband access to the home. There are several solutions from the corporate LAN space that can fit in home applications bearing in mind that home applications are cost sensitive. Given the relative complexity of the protocols and the diverse products in the home space, it is also a challenge to guarantee interoperability between devices and services. It is expected that as digital video content becomes more prevalent and available with acceptable business models for distribution, home networks will be needed to make the content available throughout the home in a cost-effective manner. Wireless technologies are expected to coexist with existing home wiring based technologies. As they become more cost-effective, it is expected that home electronics will be sold network-enabled out of the box.

RF

About the author

Guillaume Bichot is a Senior Member, Technical Staff, Corporate Research with Thomson Multimedia. He holds a degree of Network & Computer Science from the University of Pierre et Marie Curie (Paris VI). In 1996 he joined Thomson Multimedia to work on software development around the High Performance Serial Bus IEEE1394. Contact him at: *bichotg@tce.co*

Louis Litwin is a Member of the Technical Staff with Thomson Multimedia Corporate Research where he is currently working on 3G CDMA technology for mobile applications. Mr. Litwin received his M.S. degree in Electrical Engineering from Purdue University and his B.S. degree in Electrical Engineering (summa cum laude) from Drexel University. Contact him at: *LitwinL@tce.co*]

Kumar Ramaswamy is General Manager, Corporate Research, with Thomson Multimedia. He has additional responsibility for spearheading Thomson's worldwide efforts on Digital Home Networking. Dr. Ramaswamy holds a BE (Honours) from REC, Trichy, India, an MSEE from Vanderbilt University, MS (Mathematics) and PhD degrees from Rensselaer Polytechnic Institute. Dr. Ramaswamy has won the 2001 HKN Award for Outstanding Young Electrical Engineer. *swamyk@tce.co*