Quality of Service in the Home Networking Model

Abstract *– The market for home networking will soon see rapid growth. In addition to traditional data networking, this market will be driven by the desire of consumers to have access to multimedia audio, video, and gaming services. The Quality of Service (QoS) requirements these demands have put on home networking technologies has led to new standardization activities designed to deliver the QoS consumers will demand. In this paper we discuss the many ways in which QoS can be delivered, and then focus on the specific attributes of the HomeRF standard that enable it to deliver high QoS voice and multimedia services over a wireless home networking infrastructure.*

Introduction

The field of home networking is poised for rapid growth, and demand for home networking technologies will soon increase dramatically due to the impact of several factors.

Homes with computers are now becoming homes with multiple computers.

Although signs suggest "a plateau in PC penetration growth in the U.S.," this is not true for consumers who already have computers. "Repeat buyers accounted for 70% of personal computer sales in the second quarter of 1999, and – among those repeat buyers – almost 60% were purchasing an "additional" computer not to replace outdated units but to meet the increased demand of home users, according to statistics gathered by the Tech*Watch service of ACNielsen in New York City." [1]

The number of homes with a connection to the Internet continues to increase.

Data from many sources shows growth in both the penetration of computers and Internet usage. [2]

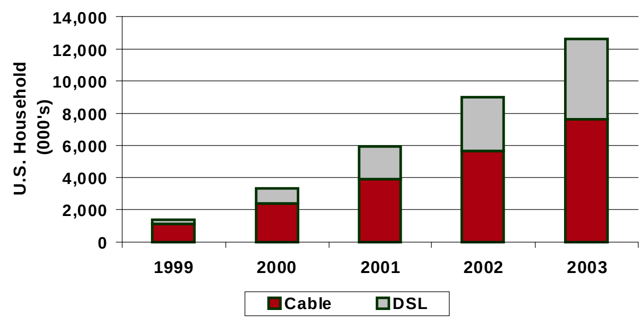

Figure 1: Broadband Households

It is the adoption of the broadband Internet lifestyle by multiple PC households that is driving the demand for home networking. One prediction for the growth of networked homes is shown below.[3]

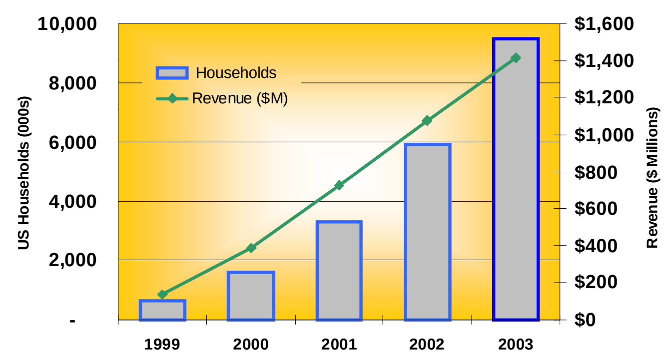

Figure 2: Networked Home Forecast 1999-2003

Home Networking Requirements

A wide range of requirements will be brought to bear by consumers as different technologies vie for success in the home networking market. These will include a no-new-wires installation and straightforward configuration and management. Of equal importance will be the requirement to support many types of services. Unlike

---

[1] http://www.mediainfo.com/ephome/news/newshtm/webnews/wt101399.htm

[2] Boyd Peterson and Karuna Uppal, the Yankee Group, "A Look Ahead: Home Networking in 2000." For related statistics see also "Falling Through the Net" at http://www.ntia.doc.gov/ntiahome/fttn99/, and DSL Internet Services Continue to Gain on Cable Modem, Business Wire, March 12, 2001.

[3] Ibid.

the enterprise environment, which is largely focused on data networking, the home environment will be characterized by the need to support services including:

- Toll quality voice communications;
- Internet surfing;
- Video and audio streaming;
- Traditional data networking;
- Internet gaming;
- Alarms, security, and other monitoring services;
- And many more.

The services that will drive the adoption of home networking, and some of the services that are most important to consumers, are shown in the following chart.[4] While Internet sharing is still the best-understood service, audio streaming and gaming applications are ranked very highly.

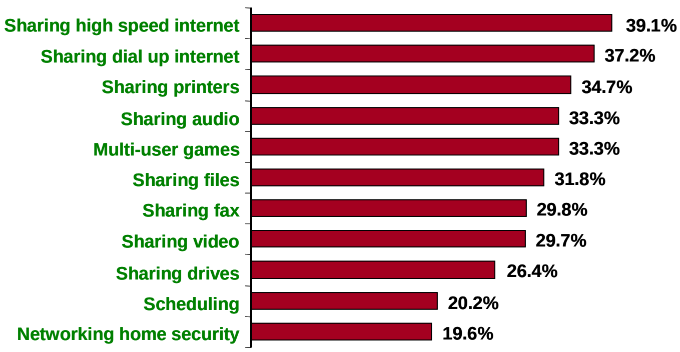

Figure 3: Applications Driving Consumer Adoption of Broadband

The Quality of Service (QoS) requirements for the streaming of audio and video over a home network are dramatically different from the requirements for the transmission of standard data. As with the Internet itself, many of the available home-networking standards were originally developed with “best-effort” traffic in mind. The demand for multimedia services over computer networks has resulted in the need to support a higher-level of QoS not originally included in these networking standards. In addition, wireless networks suffer from additional problems typically associated with open-air transmission such as interference. Hence there is an even greater need for enhanced QoS in wireless home networks. This paper discusses some basic networking protocols and the plans to include QoS provisions into existing and anticipated networking standards. Special attention will be paid to the QoS enhancements included in the latest release of the HomeRF specification.

[4] Ibid.

Quality of Service Basics

Before discussing QoS it is best to have some definition of the term. By one definition, QoS is: “A collective measure of the level of service delivered to the customer. QoS can be characterized by several basic performance criteria, including availability (low downtime), error performance, response time and throughput, lost calls or transmissions due to network congestion, connection set-up time, and speed of fault detection and correction.”[5]

A key element of this definition is that QoS can be determined by any one of a number of parameters, or any combination of those parameters. Equally important, QoS will not be defined in the same way for all services.

For the transmission of time sensitive information (voice, audio, video, etc.) over a network, Quality of Service is often defined in terms of the following parameters.[6]

- Bandwidth: The maximum data rate supported by a networking technology. Bandwidth indicates the theoretical maximum capacity of a connection, but as the theoretical bandwidth is approached, negative factors such as transmission delay can cause deterioration in quality.

- Latency: Delay in a transmission path or in a device within a transmission path. In a router, latency is the amount of time between when a data packet is received and when it is retransmitted.

- Jitter: The distortion of a signal as it is propagated through the network, where the signal varies from its original reference timing and packets do not arrive at its destination in consecutive order or on a timely basis, i.e. they vary in latency. In packet-switched networks, jitter is a distortion of the interpacket arrival times compared to the interpacket times of the original transmission. Also referred to as delay variance. This distortion is particularly damaging to multimedia traffic.

- Packet Error Rate: The rate at which the end user application receives a packet that differs from the packet as it was originally sent. This packet error rate may differ from the rate at which the medium causes packet errors, because mechanisms such

[5] A Quality of Service Glossary of Terms, http://www.stardust.com/qos/whitepapers/glossary.htm

[6] The definitions of these terms are also courtesy of A Quality of Service Glossary of Terms, http://www.stardust.com/qos/whitepapers/glossary.htm

as packet retry and error correction can be used to reduce this basic packet error rate.

Within the realm of video and audio streaming, bandwidth requirements for the transport of data streams can vary by as much as 3 orders of magnitude. A plot of some important services with their bandwidth requirements is shown below.

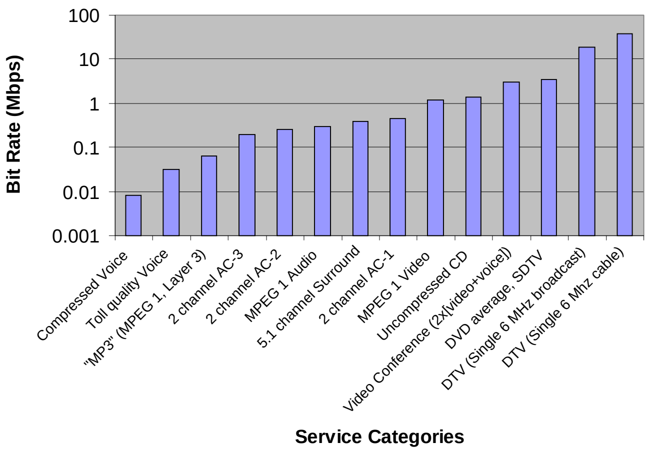

Figure 4: Services Ordered by Bandwidth Required

These services will have different requirements for latency, jitter, and error performance. For example, a two-way, real-time voice communication has a more strict latency requirement, but more tolerance for packet errors, than a high quality video stream. An illustrative table showing the different bandwidth, latency, jitter, and packet error rate requirements for various services is shown below.

Service	Payload Rate (Mbps)	Latency (ms)	Jitter (ms)	PER
High Quality Voice	0.064 × 2 streams	10	±5	10-3
Medium Quality Voice	0.008 × 2 streams	30	±20	10-3
Video Conference	1.5 × 2 streams	10	±5	10-5
HDTV	19.68	90	±10	10-5
SDTV	3	90	±10	10-5
CD Quality Audio	0.256	100	±10	10-5
High Speed Data	10	>100	>100	0
Medium Speed Data	2	>100	>100	0
Low Speed Data	0.5	>100	>100	0

Table 1: QoS Requirements by Service Type

Mechanisms to Guarantee QoS

Several different protocol variations can be used to guarantee QoS performance in networks. The general categories of resource reservation, priority mechanisms, and application control are discussed below.

Resource Reservation

When using resource reservation, a predetermined fraction of the network's resources (i.e. bandwidth) are reserved for the stream. The problem with resource reservation is that stream characteristics are difficult to know in advance. This may cause network under-utilization due to over-allocation of resources.

Priority Mechanisms

Priority mechanisms work by labeling data packets of each stream with different priorities. This allows network routers, or other devices, to treat packets differently based upon their priority.

The 802.1D standard is a layer 2 protocol that includes a prioritization scheme. There is no admission control procedure defined in 802.1D, but devices can use the priorities to determine which traffic get access to the medium first. 802.1D uses a 3-bit priority value, which allows for 8 different priorities. Table 2 lists the eight priorities and the associated traffic types.

Protocols such as Subnet Bandwidth Manager (SBM) allow the mapping of RSVP QoS services onto LANs.

Priority	Traffic Type
1	Background
2	Reserved
0	Best Effort
3	Excellent Effort
4	Controlled Load
5	Video
6	Voice
7	Network Control

Table 2: 802.1D priorities

Application Control

In application control the transmitting device adapts its sending rate to the current network conditions. If the network is congested the data rate will be slowed, which increases the chance of packets getting to the destination node.

Since this paper focuses on networking technologies, we will only address QoS mechanisms that operate in the MAC layer of the protocol stack, or below. Application control mechanisms, therefore, will not be part of the discussion.

QoS in Home Networking Technologies

Ethernet as a Basis

When discussing Quality of Service provisions among the various networking technologies, it is useful to begin with a baseline for comparison. A useful starting point is the Ethernet technology.

Xerox and DEC originally developed Ethernet as a way to interconnect machines without the use of a mainframe computer. This original protocol was adopted by the IEEE, and was adapted by that group in its creation of the 802.3, Carrier Sense Multiple Access with Collision Detection (CSMA/CD), protocol. Most people mean this protocol when they speak of Ethernet.

Ethernet uses a bus topology, in which only one computer at a time can send data. If more than one computer attempts to send data at the same time, the signals will “collide” on the wire, causing the data to become corrupt and creating data errors. The 802.3 standard defines a protocol for accessing the medium in such a way that these collisions are reduced.

The protocol defines a way for computers to listen to the network before they transmit any data. If there are no carrier signals present, the node will begin sending its data. This is the “carrier sense” mechanism. After the data is sent, the “collision detection” takes over. Collision detection is able to determine whether there was an error in sending the data. In the event that a collision does occur (two nodes can independently decide to send their data packets at the same time), the nodes wait for a number of time slots chosen at random and then attempt to resend their data. This “random backoff” mechanism helps to avoid a second collision when both nodes try to resend their data. A binary exponential backoff algorithm is used to calculate the backoff variable. The number that is generated from this algorithm is uniformly distributed in a range called the contention window. If after waiting the prescribed backoff time the transmission still does not succeed, the size of the contention window is doubled until some maximum size is reached.

The origin of Ethernet as a protocol used for sending data indicates its level of support for Quality of Service. In essence, it does not have any. All nodes access the medium using the CSMA/CD protocol and, therefore, all are subject to the time delays in packet delivery caused by these collisions. Ethernet, with its 10 Mbps and 100 Mbps physical layers is, therefore, an excellent technology for sending data packets (high data rate and loose requirements for latency and jitter), but it is less appropriate for the strict latency and jitter requirements of voice, audio, video, and gaming. And, of course, Ethernet is problematic in the home environment since it requires a separate wired infrastructure.

The Ethernet protocol, however, can be used to compare the other home networking technologies. Some technologies address QoS through the use of a reservation, guaranteed access, mechanism, while others opt for a priority mechanism. While it is not always possible to cleanly distinguish between the two, an attempt at such a grouping is shown below.

Resource Reservation Technologies

IEEE 1394

The IEEE 1394a standard is an external bus standard that includes two types of channels, asynchronous and isochronous. The asynchronous channel is a best effort data channel that has no QoS mechanism associated with it. The Isochronous channel is a reservation based assured bandwidth channel.

The capacity of IEEE 1394a is typically 200-400 Mbps. Therefore, QoS guarantees on IEEE 1394 buses will likely be accomplished by reserving bandwidth on the isochronous channels. Because of the wide bandwidth of IEEE 1394a, the resulting “waste” of bandwidth is not considered a serious issue.

HiperLAN/2

HiperLAN/2 includes a centralized, controlled, (synchronous) Data Link Control (DLC) layer, which means that the access point (AP) controls how the resources are allocated in a MAC frame. Each mobile terminal (MT) requests capacity in future MAC frames when it has some data to send. The AP informs the MTs at which point in time in the MAC frame they are allowed to transmit their data. This time allocation dynamically adapts according to the request for resources from each of the MTs. The air interface is based on time-division duplex (TDD) and dynamic time-division multiple access (TDMA), i.e., the time-slotted structure of the medium allows for simultaneous communication in both downlink and uplink within the same time frame. Time slots for downlink and uplink communication are allocated dynamically depending on the need for transmission resources.

The connection-oriented nature of HiperLAN/2 makes it straightforward to implement support for QoS. Using connection control combined with error control, each connection can be assigned a specific channel, which satisfies the QoS requirements of bandwidth, delay, jitter, and bit error rate for the traffic flow.

Priority Mechanism Technologies

802.11

The 802.11 standard makes use of two different access mechanisms, the Distributed Coordination Function (DCF) and the Point Coordination Function (PCF).

The DCF uses CSMA during contention periods. The CSMA mechanism operates essentially as described above for Ethernet. The PCF mechanism operates by polling nodes to see if they have data. Nodes that have data to transmit can then proceed to transmit without contention. The PCF operates on top of the DCF. The point coordinator (PC), which is part of the access point, must gain access to the medium using the DCF access mechanism. Once it has gained access, a contention free period (CFP) begins immediately after the PC sends a beacon. During the CFP the polling operation begins, and the DCF mode is suspended.

The PCF already provides near-isochronous service, which makes it ideal for providing QoS. For this reason proposals for a QoS enhanced PCF have been presented as part of the 802.11e (QoS MAC) working group. On the other hand the PCF is difficult to implement, and is not found in any current products. For this reason there have also been many proposals for a QoS enhanced DCF. Since very little is solidified with the 802.11e standards group, we cannot predict what protocol the 802.11e group will eventually choose to provide QoS support.

However, it seems likely that the group will adopt some kind of prioritized DCF mechanism. The IEEE 802.11e subgroup adopted document IEEE 802.11-00/360r2 as its first draft in January 2001. The draft defines a set of MAC sublayer QoS parameters to be used to define QoS traffic. Some of the parameters have open issues and are subject to change in the future. These parameters, collectively called a traffic specification and applied to a traffic category (TC), are Traffic Type, Ack Policy, Delivery Priority, Retry Interval, Polling Interval, Transmit Interval, Nominal MAC Service Data Unit (MSDU) Size, Minimum Data Rate, Mean Data Rate, Maximum Burst Size, Delay Bound, and Jitter Bound.

A description of one of the possible priority mechanisms is sufficient to explain the concept. In one such proposal before the 802.11e group, multiple DCFs would operate at the same time. Each DCF would generate a backoff number out of a unique set of contention variables. Those variables are the contention window size and the contention offset. Higher priority packets will use smaller contention windows, and hence have a higher chance of getting lower backoff variables. The offset variable can be calculated such that there is no overlap in contention periods. This kind of prioritization mechanism is similar to the Priority Asynchronous Data Service defined in the HomeRF specification. This mode will be described in more detail below.

Phone line

ITU T Rec. G.989.1 specifies characteristics for Home Phoneline Networking transceivers. Specifically, the PHY and MAC layers are defined, although the PHY payload formatting is unspecified.

G.989.1 devices employ CSMA/CD contention to access the shared medium. The G.989.1 contention mechanism supports eight priority levels, implementing absolute priority among devices contending for access. These levels are labeled from zero to seven, with seven corresponding to the highest priority traffic. Devices are only permitted to initiate transmission of a frame during a Priority Slot with a number less than or equal to the frame’s priority level.

Generally, collisions occur between frames at the same priority level, but since devices may initiate transmission on a Priority Slot with a value less than the frame priority, it is possible for frames with different priorities to collide. Transmitting devices that detect a collision will cease transmission. Colliding devices resolve the contention via a random back-off mechanism. After a collision, three Signal Slots are defined prior to the Priority Slots. Contending devices choose a Signal Slot at random and transmit a Back-off Signal within the Signal Slot. Devices transmitting in the lowest number active Signal Slot then contend for access in the subsequent Priority Slot for its level; other devices increment a Back-off Counter. A successful transmission causes Back-off Counters to be decremented; those whose Counter values reach zero then contend for access.

Power line

The HomePlug Powerline Alliance is producing the “Draft HomePlug Medium Interface Specification.” This specification defines the functions, operations, and interface characteristics of the HomePlug 1.0 system for high speed networking using the medium of power line wiring (the HomePlug Power Line Networking System).

Medium sharing is accomplished by the CSMA/CA mechanism with priorities and a random back-off time following busy conditions on the channel. Prioritized access is achieved by a Priority Resolution Period in which stations signal the priority at which they intend to transmit, allowing only the highest priority available to continue in the contention process. A random back-off mechanism spreads the time over which stations attempt to transmit under busy conditions to reduce the probability of Collision, using a truncated binary exponential back-off mechanism similar to Ethernet.

Hybrid QoS Technologies

Bluetooth

Bluetooth defines two types of channels for transmission, synchronous and asynchronous. Synchronous channels (SCO, synchronous, connection-oriented link) are symmetric and provide a 64kb/s bi-directional connection between the Master and a specific Slave. Transmit and receive slots are sent periodically with a fixed interval between them. Most of the SCO slots are targeted to voice distribution, but one of them can carry both voice and low-speed data simultaneously.

Asynchronous packets (ACL, asynchronous, connectionless link) are sent on the slots left after SCO assignment. The slaves send information only after they receive information targeted to them from the Master. There are multiple types of ACL slots, which differ by their payload size and FEC protection. For the ACL channels, the BT specification defines the L2CAP layer that enables segmentation and re-assembly of packets as well as QoS services and connection establishment.

HomeRF

Unlike Ethernet, HomeRF uses two separate medium access protocols for its two primary services, voice and data. Voice packets are sent using a guaranteed access TDMA/TDD protocol. Data packets are sent using a wireless version of Ethernet; yet, even within this data protocol a powerful QoS mechanism has been implemented.

Like Ethernet, the data service is built on the CSMA protocol. In a wireless system, however, it is very difficult for a node to both send packets and receive them at the same time, so collision detection is not used. Rather, “collision avoidance” is used. Collision avoidance is very similar to collision detection in that, when a node fails to receive a positive acknowledgement that its packet was received, it will wait a random number of time slots before attempting to resend the packet in order to avoid collisions with other nodes. This is known within HomeRF as the Asynchronous Data Service.

For data with more strict latency requirements, like lower quality voice, or streaming audio or video sessions, HomeRF has also implemented a Priority Asynchronous Data Service, in which selected data packets gain priority access to the channel while using the CSMA protocol.

Finally, HomeRF uses a TDMA/TDD protocol for the transmission of its high quality voice service. TDD (time division duplex) allows for a full duplex communication, and TDMA (time division multiple access) provides guaranteed access to the medium once per frame for each voice connection.

The HomeRF protocols and QoS provisions are discussed in greater detail below.

HomeRF Quality of Service Support in Detail

HomeRF provides QoS for various services using a variety of different protocols. Each of these protocols is used by one or more of the technologies described above. HomeRF’s use of isochronous channels for the delivery of voice is similar to the Bluetooth synchronous service and the QoS mechanism described for IEEE 1394 and HiperLAN. The use of a priority asynchronous service is similar to what is being proposed for use in 802.11, phone line, and power line networking. HomeRF’s isochronous and priority asynchronous modes are described in detail below. In addition, an enhanced physical layer technique known as hopset adaptation is also described. Hopset adaptation enhances the quality of time-sensitive applications in the presence of a static, persistent source of interference.

Effects of Wireless Transmission on QoS

Wireless networks incur certain problems that make the handling of QoS more difficult than for wired solutions. The major factor is packet loss. The perceptual quality of an audio stream drops significantly as packet loss reaches 20% for non-adaptive application, even if retransmission techniques are used.[7]

---

[7] Chen, Tsuwei, Geria, Mario, Kazantzidis, Manthos, Romanenko, Yuri, Slain, Ilya, “Experiments on QoS Adaptation for Improving End User Speech Perception Over Multi-hop Wireless Networks”, International Conference on Communications, June 1999, quoted in “Quality of Service (QoS) for Streaming Audio Over Wireless LANs”, Jason S. Flaks, 2001.

HomeRF operates in the 2.4 GHz ISM band, under the FCC's Part 15 rules for unlicensed devices. This is a busy section of the radio spectrum. Bluetooth and 802.11b operate there as well, as do microwave ovens, amateur radios, military systems, medical systems, microwave lighting, etc. Unlike the Internet, where packet loss is primarily due to network congestion, wireless networks often suffer from packet loss due to interference. The most common culprit in the 2.4 GHz band is the microwave oven, although any of the other users, including other wireless networks or 2.4 GHz cordless phones, can be a source of disruption.

One of the most common ways of handling packet loss is through input rate regulation. Since interference in a wireless network sometimes causes bandwidth reduction, and other times can be completely disruptive, traditional methods of rate regulation will be ineffective in many instances.

Interference can also indirectly affect other QoS variables such as delay and jitter. For this reason QoS enhancements to wireless LAN technologies must consider the adverse effects of interference.

HomeRF Frame Structure

In order to understand the QoS support in HomeRF, it is important to understand the basics of the HomeRF frame structure. As mentioned above, HomeRF combines an asynchronous, CSMA protocol for best effort data with an isochronous, TDMA protocol used for the delivery of high quality voice.

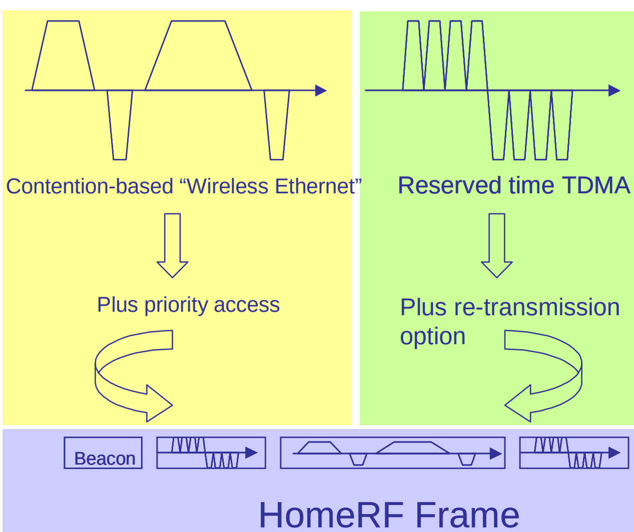

Figure 5: Basics of the HomeRF MAC Frame Structure

The contention protocol acts essentially as described above for Ethernet. However, during certain periods (the "Contention Free Periods") time is reserved for active voice calls. That time is broken into specific time periods, used for the uplink and downlink of the voice connections. When there are no active voice connections on the HomeRF network, that time is recaptured for use by the contention-based data protocol. This is detailed below.

While HomeRF is based on a 20 ms frame structure (with one hop every frame), HomeRF moves to a 10 ms subframe structure (with one hop every subframe) whenever there are active voice connections. The shorter frames provide decreased latency and increased interference immunity, hence enhanced QoS, to the voice connections. The longer frames have less overhead from the hopping procedure and, therefore, higher throughput for the data connections. The basic structure is illustrated in Figure 6.

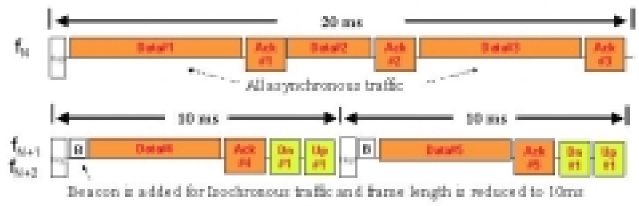

Figure 6: HomeRF frame structure both with, and without, active voice connections (Not to scale)

HomeRF Highest Level of QoS: The Guaranteed Priority Voice Service

The upper part of Figure 6 shows what the HomeRF frame looks like when only data is present on the network. The frame is 20 ms long, and all of the data accesses the medium using a contention-based access protocol. The lower part of the frame shows the frame structure when a voice communication is present. In this case, the part of the frame available for contention-based access is reduced, and part of the frame is set aside for TDMA access for the voice. The voice connection is full duplex using time division duplex (TDD), which explains the Dn#1 and Up#1 areas in the lower part of Figure 6. These are the downlink and uplink parts of the voice communication. Use of TDMA for the voice access methodology explains how HomeRF provides for low latency voice communications on the network. Low bit error rates are provided by adding a retry mechanism for the voice frames to the robust, frequency hopping, physical layer.

Voice packets that are blocked due to interference in one subframe are resent in the next subframe, less than 10 ms later. This is illustrated in Figure 7.

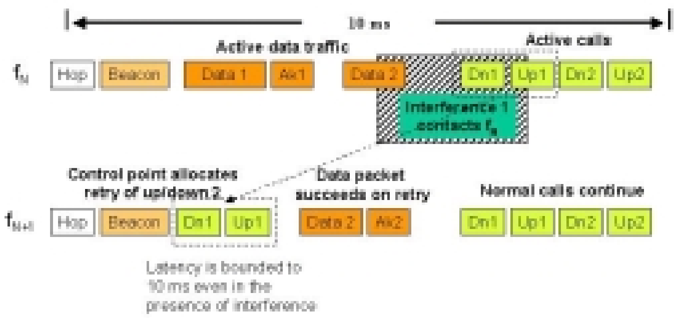

Figure 7: HomeRF provides a unique mechanism to retry voice packets

As shown in Figure 7, if interference appears in the frequency range and time such that HomeRF voice packets experience interference, those packets can be sent again no more than 10 ms later. Since the packets are retried after a frequency hop, the frequency environment will be entirely different, thus increasing the probability of a successful retransmission. This probability is further increased by the use of hopset adaptation, discussed below.

The reason for using a guaranteed access mechanism for voice connections rather than a contention based access mechanism is best illustrated by a comparison of the latency issues. Voice communications should have less than 125 milliseconds of end-to-end delay in order to be considered toll quality. This delay will be made up of many sources, both fixed and variable, including coding delay, buffer delay, and delays incurred by transporting the voice over a wide area network such as the PSTN, satellite links, the Internet, or some other packet network.[8] Because of these many components of delay, the home network, acting as the access point (equivalent to a wire in a simple POTS network,) should contribute only a very small delay (on the order of 30 milliseconds) to this total.

In a contention-based access mechanism, any packet needing to be transported must first contend with other packets waiting to be sent. Contention results in packet collisions (lost packets) that need to be retried, and variable delays in predicting when any given packet will be able to access the network. It is to reduce the size of these delays that HomeRF uses dedicated time periods (contention-free periods) to transport voice communications.

There are two main contributors to packet delay in a network. First, simple queuing delay results as more and more packets are presented to the medium. If the rate at which data is being presented to the medium exceeds the medium's ability to transport that data then some of the packets will be forced to "wait in line", thus causing a delay between their creation and their delivery. Second, contention on the network will result in collisions and the requirement to send packets again. Even in a two-node system, with one node sending packets and the other node receiving those packets, there will still be collisions on the medium that will contribute to packet delay. Those collisions are the consequence of the acknowledgement packets (ACKs) that many protocols, like TCP/IP, send from the recipient of the data packets to the sender of the data packets in order to ensure reliable transmission. For a real-time service like two-way voice communication, delays resulting from these sources can quickly become intolerable.

The effect of interference on these delays can be significant. In order to estimate it, we make use of a MAC layer simulation. This simulation includes packet creation, queuing, collision, an exponential backoff protocol, the creation of MAC and TCP acknowledgements, and the ability to simulate packet errors.

A typical delay profile resulting from the simulation is shown in Figure 8. In this case the traffic offered to the network was about 50% of the maximum channel rate, and no packet errors were included. The 99th percentile delay is shown. Whenever the delay is referred to in the following analysis, the 99th percentile delay on this delay distribution is what is meant.[9]

---

[8] An excellent discussion of the delays incurred in the transport of packetized voice is presented by Eric Larson and Steve Nikola of Motorola, "Voice Technologies for IP and Frame Relay Networks", available at http://www.mot.com/MIMS/ISG/mnd/papers/voice_technologies_for_ip_and_frame_relay_networks.html.

[9] For 99% of the packets, the delay is less than this value. This point on the delay profile was used under the assumption that packets experiencing too much delay in a real-time application would be dropped, and so would appear as packet errors. Therefore, the 99% delay point is roughly equivalent to the addition of 1% packet error to the end user. This is a reasonable target packet error rate for a voice system.

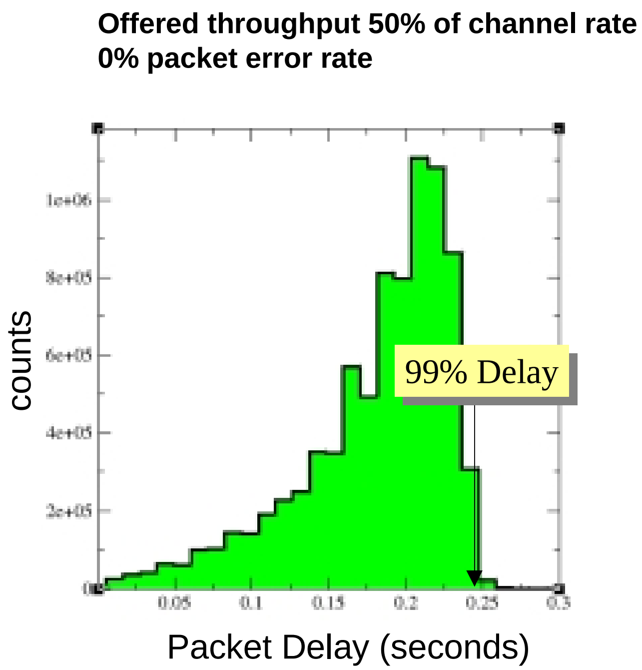

Figure 8: Delay profile for packets in a contention-based network

As can be seen in Figure 8, the simulation shows that a typical contention based network will contribute significant packet delays as the traffic offered to the network increases. Note that even with no packets lost due to interference, increased network traffic (a large file being downloaded, a video stream being established, a file being sent from one device to another or to a printer, etc.) can drive the delays attributable to a home network using a contention-based access mechanism to beyond 100 milliseconds as the offered throughput approaches the achievable data throughput (which is less than the physical layer rate), saturating the system. HomeRF, on the other hand, completely removes the voice communication from the contention mechanism. This keeps the delay for voice packets bounded to approximately the frame size of 10 milliseconds.

As also shown in Figure 9 the effect becomes increasingly significant as interference affects the network, causing more packet retries and collisions in the contention process. In the case of a 20% packet error rate delays quickly approach 300 milliseconds for high system loading. Neither of these effects (increasing traffic or increasing packet errors) will have an impact on the HomeRF voice service, however. HomeRF's TDMA access protects the voice traffic immune from network loading, and the hopset adaptation with subframe retries (see below) protects the voice traffic from network interference.

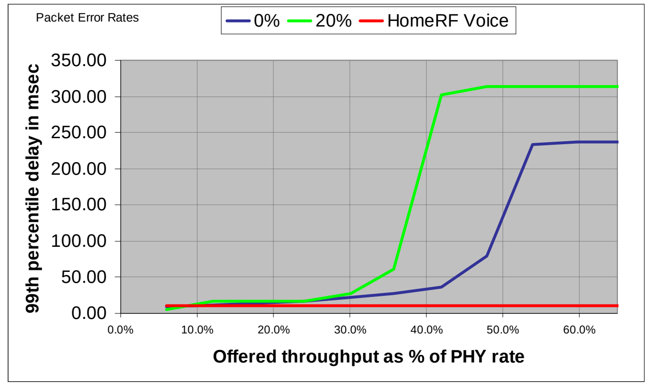

Figure 9: Simulated 99% Delay for Packets in Contention-Based Access Network Compared to HomeRF Voice Packet Delay

HomeRF Intermediate Level of QoS: The Priority Asynchronous Data Service

Between the strict latency requirements of voice communications and delay tolerant data communications lies the realm of streaming multimedia. In these applications a continuous video and/or audio stream is sent from one point to another via a network. For example the stream may come from the Internet to a home computer, stereo, or monitor, or it may come from a home storage device (like a computer) to a set of speakers or a headset also within the home.

Streams are usually not two-way communications, so latency is not their primary problem. The bigger problem for streams is jitter. In a two-way communication like voice, latency translates into the time between one person saying something and the other person hearing it. Extensive latency leads to gaps in the conversation that quickly become unacceptable. (See the discussion of voice, above.) In a one-way stream, however, latency appears only as an overall delay in the time that the stream begins. Jitter, on the other hand, refers to the time difference in the transmit time of packets in the stream. If some packets are delayed, this translates into "stutter" in the stream, so that the audio or video appears to pause while waiting for additional packets to arrive.

Jitter can be avoided by the use of "jitter buffers", which store a segment of the stream before displaying it so that if any packets are delayed the buffer will still have data that it can display. This is a method of trading latency (filling the jitter buffer adds latency) for reduced jitter. The other tradeoff of this method, however, is the need for enough memory in the display device to buffer the stream. This adds cost and size to the product.

Therefore, in order to support streaming connections without requiring extensive buffering at the client devices, a network should be able to provide some sort of quality of service (QoS) to streaming sessions in order to keep the latency below about 100 milliseconds and the jitter below about ±10 milliseconds. HomeRF does this by way of the Priority Asynchronous Data Service.

The basic Priority Asynchronous Data mechanism is shown in Figure 10.

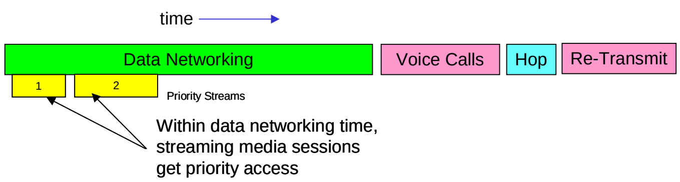

Figure 10: Principles of the HomeRF Priority Asynchronous Data Mechanism

In a standard contention-based data networking protocol like Ethernet, described above, packets contend for access to the channel ("the airwaves") by selecting a random number. The unit with the lowest number gets to transmit first. If two units select the same number and transmit at the same time, their transmissions will "collide", and both will be forced to retransmit. Both the random access method and the occurrence of collisions and retransmissions lead to latency and jitter. The existence of interference on the network only compounds these problems.

HomeRF's Priority Asynchronous Data service solves these problems in the following way. A streaming session is assigned its access number when the session is established. Those numbers are assigned based on priority. For example, in Figure 10 the stream with the higher QoS requirements would be assigned to position 1, and the stream with lower requirements would be assigned to position 2. HomeRF supports the assignment of up to eight simultaneous streams. The contention-based data networking protocol is then adjusted so that the random numbers selected by the asynchronous devices do not include the numbers already assigned to the streams. By removing the random access and the possibility of collision, the latency and jitter performance of the streaming sessions can be very strictly controlled.

To illustrate the difference between the packet delays that can be experienced in a contention-based data network versus those from the Priority Asynchronous Data service, the MAC simulation discussed earlier has been used in the following way. A scenario involving six users has been simulated. Four simulated users generate traffic at a constant rate of 1 Mbps, and two other simulated users generate traffic at a rate of 64 kbps. The 64 kbps data flows can be thought of as MP3 audio streams, for example. First the simulation was run using only the asynchronous data service for all users. A packet error rate of 10% was also included. The results are shown in Figure 11.

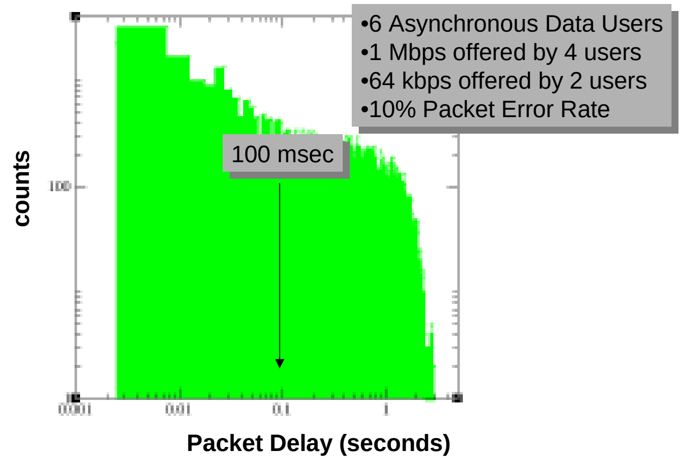

Figure 11: Delay Profile for 4 Asynchronous Data Users

This figure shows clearly that many of the packets (more than 70% in this example) have delays in excess of 100 milliseconds. Some are much longer than that. These delays, as mentioned before, are due to random access times, collisions, and retries. These delays will be present in any of the data connections, because each user's data will contend for access the channel in the same way.

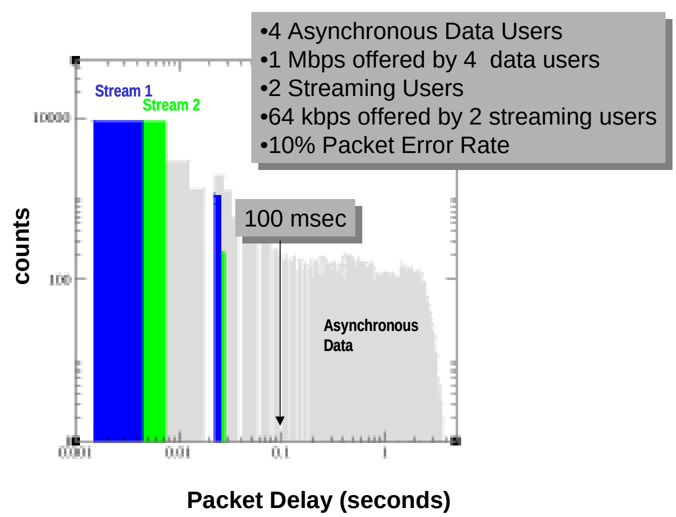

Figure 12: Delay Profile for Asynchronous Data Mixed with Priority Asynchronous Data

The Priority Asynchronous Data service leads to a very different delay profile. The simulation was next run using the Priority Asynchronous Data service for the two simulated MP3 connections. The effect is shown in

Figure 12. The priority streams show a very short delay profile, less than 30 milliseconds at the maximum, while the contention-based asynchronous data shows packets with significant delays.

Examining the effect of added interference on these results further illustrates the power of HomeRF's Priority Asynchronous Data service. In this case we simulate the same six users as above, four sending 1 Mbps of data using the asynchronous data service, while the other two send 64 kbps using the Priority Asynchronous service. In Figure 13 we show the 99% delay in the delay profile of the packets as a function of channel interference, or packet error rate.

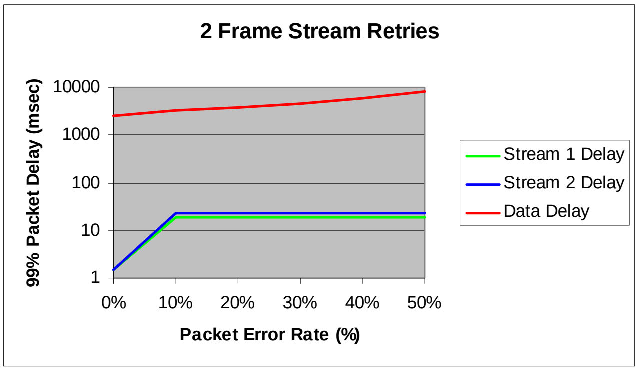

Figure 13: 99% Packet Delay versus Packet Error Rate. Streams have a 2 Frame Retry Limit

As is shown in this figure, increased interference causes the contention-based data packets to be retransmitted, resulting in further delay, further collisions, more delay, and so on. While such large delays may be acceptable for data transport (where getting all of the packets delivered is the primary goal), they are clearly too large for any streaming multimedia service. The results using the HomeRF Priority Asynchronous Data service, however, are much different. With guaranteed access to the channel and a two frame retry limit, the stream packets have a delay profile that remains bounded at about 30 milliseconds.

For streaming connections, a packet that is delayed too long by interference becomes outdated and should be discarded. Since, in this simulated implementation, the stream packets will not be retried more than one time, another good measure of the stream QoS is the percentage of the stream's packets that are delivered before being discarded. For the two streams simulated here, the related measure of "undelivered packets" is shown in Figure 14.

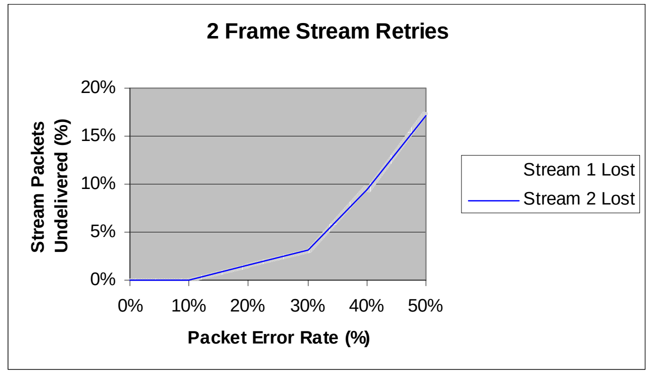

Figure 14: Undelivered Packets for Streams with 2 Frame Retry Limit versus Packet Error Rate

Together, Figure 13 and Figure 14 show that the Priority Asynchronous Data service is able to guarantee excellent QoS, even in the face of severe interference. Figure 13 demonstrates, for example, that with 30% packet error rate a contention-based data stream will show packets with several seconds of delay. With that same packet error rate, however, Figure 13 and Figure 14 show that two streams can deliver their 0.64 kbps streams with less than 30 milliseconds of delay, and with fewer than 5% of the packets undelivered.

Even this performance can be improved upon for applications that can tolerate more delay. This is because the number of frames over which stream packets can be retried can be increased in order to reduce the undelivered bit percentage. In the current example, increasing the retry limit to 4 frames completely eliminates the undelivered stream packets (100% of stream packets are delivered at all packet error rates up to 50% PER) while only increasing the stream packet delay to a maximum of 60 milliseconds. The resulting 99% delays are shown in Figure 15.

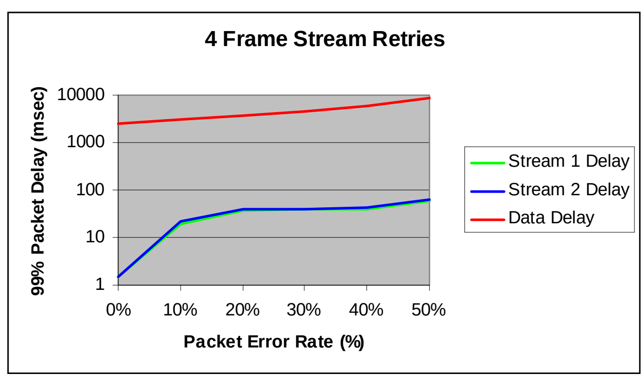

Figure 15: 99% Packet Delay versus Packet Error Rate for Streams with 4 Frame Retry Limit

Figure 16 demonstrates that the QoS for the streaming sessions can be improved by permitting more delay in exchange for higher QoS.

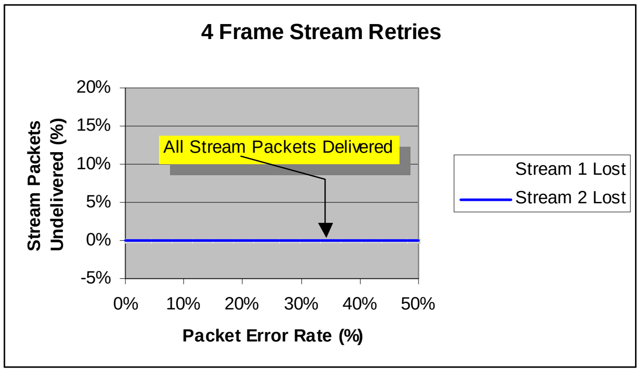

Figure 16: QoS for Streams versus Packet Error Rate for Streams with 4 Frame Retry Limit

In the examples above, the Priority Asynchronous Data service appears very similar to the isochronous voice service in that the streams have the same QoS performance (latency and undelivered packets.) This is due to the choice of stream bandwidth used in these simulations. Using other choices of stream bandwidth it is possible to see the different QoS performances of the different stream positions, as in the discussion of Figure 10.

For example, if four streams of 1 Mbps each are simulated and each stream is limited to a two-frame retry limit, the undelivered packet distribution is shown in Figure 17.

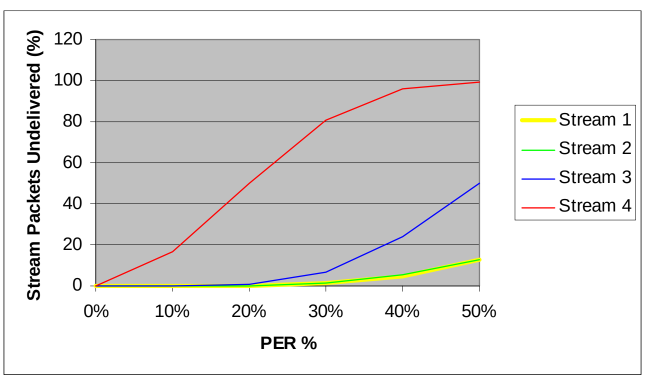

Figure 17: Undelivered Packet Performance for 4 Streams of 1 Mbps each with a 2 Frame Retry Limit

As this figure clearly illustrates, the first two streams deliver more of their packets than do the last two streams. This is because, with these required bandwidths, it is possible for the first two streams of use up most of the available network resources leaving the later streams with no opportunity to deliver their packets.

Adaptive Hopping Mechanism

Hopset adaptation is used to minimize the impact of long-term, static interferers. An example of such an interferer is a microwave oven or a frequency static wireless LAN. Under the FCC's frequency hopping rules for the 2.4 GHz band hopping systems using hopping channels of at least 1 MHz must occupy at least 75 MHz of the 83.5 MHz available in the band. Therefore, there is no opportunity to avoid entirely the frequency range occupied by a wideband static interferer. HomeRF takes advantage of the fact that communications that are blocked by interference can be re-transmitted on the next hop. HomeRF's hopset adaptation mechanism ensures that two adjacent hops will not both be within a frequency range where interference has been identified, up to a very wide interference bandwidth.

This mechanism works as follows. When an "interference range" is identified, the hopset is examined to find out if two consecutive hops are both within the range. If such a pair of hops is located, an attempt is made to switch one of those hops with a hop that is outside the interference range. Therefore, though the full set of 75 hops is still used, the hopping sequence now virtually guarantees that an "interfered with" hop will be followed by a hop without interference. This technique is very powerful, and leads to no consecutive "bad" hops in the presence of an interference of up to 31 MHz. This is shown in Figure 18.

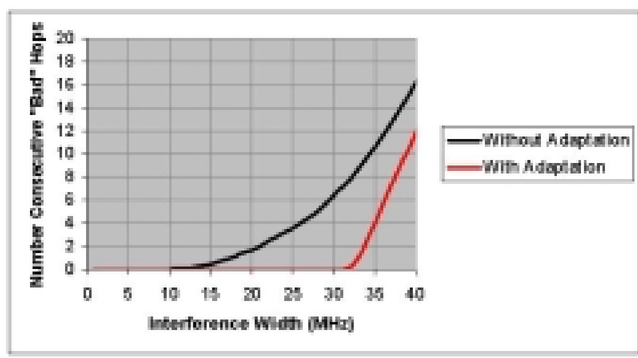

Figure 18: Performance of the HomeRF Hopset Adaptation Algorithm

The power of this algorithm is that in the presence of a wideband, static interferer (like a microwave oven or DSSS wireless LAN, for example) HomeRF can virtually guarantee that a hop that is interfered with will be followed by a hop that is free from interference. Because of HomeRF's retry mechanism for voice and stream packets (see the discussions above), this leads to very robust HomeRF voice and streaming media communications in the presence of this type of interference.

An estimate of the interference immunity of HomeRF in the presence of a microwave oven is shown in Table 3. The result, as shown, is that the expected bit error rate is significantly less than 1%.

Probability of a hit by microwave oven interference (Assumed characteristics: interference covers 20% of the band for 50% of the time.)	Raw $P_{hit} = 10\%$ (This is the extreme failure rate that would be experienced by some WLAN systems.)
HomeRF time/frequency diversity reduces this probability (Results in acceptable voice and stream quality)	Independent trials result in $P_{2\ hits} = P_{hit} \times P_{hit} = 1\%$
HomeRF also uses hopset adaptation to further reduce this probability in the presence of persistent interferers (Results in excellent voice and stream quality)	The probability of a second hop into the interference zone is reduced further. N_i = channels with interference N_T = total channels $P_{2\ hits} \sim P_{hit}^2 \times (N_i / N_T) << 1\%$

Table 3: Voice Packet Failures in the Presence of a Microwave Oven

The beneficial effects of adaptive hopping apply to both the HomeRF voice service and to the Priority Asynchronous Service. In both of these services it is likely that a packet will be retried only a limited number of times before being discarded permanently. This packet loss will lead to reduced QoS. The adaptive hopping mechanism increases the probability of a successful transmission on the hop following a hop in which a packet transmission was unsuccessful.

HomeRF Future Plans

As the future of home networking unfolds, HomeRF will evolve to meet the changing needs. Anticipated changes involve increased bandwidth to support more simultaneous services and services with higher bandwidth requirements, and enhanced services such as roaming.

As is shown in Figure 4, there are many services on the horizon (video conferencing, standard and high definition television broadcast) that the current generation of wireless home networking technologies including HomeRF are ill equipped to support.

In its next generation equipment HomeRF will be able to support 20 Mbps, or greater, at the physical layer. Unlike other attempts to achieve these data rates in the 2.4 GHz band the HomeRF migration path to higher data rates requires no changes to the current FCC rules. There are many proposed methods to achieve these rates, including the use of FSK modulation with a higher symbol rate, or the use of linear modulation techniques in order to achieve higher bits per symbol transmission. The technology choice will be finalized in 2001 with product expected in 2002.

The higher bandwidth availability will be used to support more QoS connections (more voice calls, more streaming calls) as well as to provide higher data rate connections both for the Asynchronous and Priority Asynchronous Data services.

In addition to increased bandwidth, HomeRF will enhance other components of its offering. For example, the latest version of the HomeRF specification describes a roaming protocol for use with the Asynchronous Data service. The next generation of HomeRF will extend this roaming protocol so that both the isochronous and priority asynchronous services will be able to use it without a noticeable degradation in QoS as the connections are transferred from one Connection Point to another.

Conclusion

The demand for home networking will be driven by more than the desire to transfer data more quickly. Multimedia applications will be an important component of the service offerings, and these demand QoS guarantees beyond the capabilities of most data-centric networking technologies. HomeRF – with its mix of isochronous, priority asynchronous, and asynchronous services – is well suited to handle these services, both now and into the future.