

The Design and Implementation of HomeRF: A Radio Frequency Wireless Networking Standard for the Connected Home

JIM LANSFORD AND PARAMVIR BAHL, SENIOR MEMBER, IEEE

The HomeRF¹ Working Group (WG) is a consortium of more than 100 companies from the computer, telecommunications, and consumer electronics industries. This group has developed an open specification called the Shared Wireless Access Protocol-Cordless Access (SWAP-CA) that enables radio frequency (RF) wireless connectivity between a diverse set of devices and computing resources in and around a typical home. Built around an RF spectrum with worldwide availability, SWAP-CA includes operational support for both managed and ad hoc networks of devices. It combines and extends wireless networking and cordless telephony into a single unified protocol allowing mobile devices to communicate via both voice and data traffic simultaneously over the Internet and/or over the public switched telephone network (PSTN). For battery-operated devices, it includes a power management mechanism that ensures connection longevity. The technology has been specifically optimized for consumer applications and price points, and, consequently, the HomeRF WG has the broad backing of the major corporate stakeholders interested in enabling tetherless networking within the home.

Keywords—Home networks, power conserving protocols, quality of service, wireless LAN standards.

I. INTRODUCTION

Two major factors have presented a real opportunity for data networking within the home. The first is the explosive growth and usage of the Internet. The Internet has clearly revolutionized the delivery of information and entertainment to the home. The second is the emergence of sub-\$1000 powerful home personal computers (PCs). With these inexpensive devices, the barrier to getting on the Internet and rediscovering the utility of the PC is low enough to reach the vast majority of middle-income households.

Manuscript received December 15, 1999; revised August 8, 2000.

J. Lansford was with Intel Corporation, Santa Clara, CA 95052-8119 USA. He is now with Mobillian Corporation, Hillsboro, OR 97124 USA (e-mail: jim lansford@mobillian.com).

P. Bahl is with Microsoft Corporation, Redmond, WA 98052-6399 USA (e-mail: bahl@microsoft.com).

Publisher Item Identifier S 0018-9219(00)09782-6.

¹HomeRF is a registered trademark of the Home RF Working Group (HRFWG) [1]

However, consumers soon find that the PC/Internet combination, though very compelling, lacks some key attributes in terms of mobility and convenience of location compared with many of their traditional information and entertainment options, such as newspapers, magazines, television, videos, FM radio, DVD/CD/stereo, etc. The powerful home PCs (and the printers and peripherals attached to them) often end up turned off 20–22 hours per day while tucked into a bedroom or den corner where access is possible only within a 2–3 foot “bubble.” The major opportunity for networking in the home is, thus, to extend the reach of the PC and Internet throughout the home and yard and connect the resources of the PC and Internet with legacy home applications such as telephony, audio entertainment, and home control systems. Another opportunity is the sharing of resources (such as an Internet gateway or high-quality printer) amongst PCs in multi-PC homes.

With these issues in mind, several major stakeholders in the PC and wireless industry formed the Home RF WG in early 1997 [1]. The primary goal of this group is to enable interoperable wireless voice and data networking within the home at consumer price points. The group began by pooling market research data from the member companies to produce a *Market Requirements Document* on user expectations for a home networking wireless technology. This document was then used by the HomeRF technical team to specify design guidelines that best fulfilled the needs of the potential users (see Table 1). Technical proposals were solicited, submitted, and carefully evaluated and—with tremendous cooperation from the RF communications industry and the nascent wireless local area network (WLAN) community—the Shared Wireless Access Protocol-Cordless Access (or SWAP-CA) specification was created. The SWAP-CA specification describes in detail Layer 1 and Layer 2 of the International Standards Organization’s (ISO’s) Open Systems Interconnect (OSI) networking model. In designing these layers, the HomeRF technical team chose to reuse proven RF networking technology for data and voice communications and added simplifications where appropriate for home usage. With this approach, SWAP-CA inherited native support for Internet access via TCP/IP networking and for voice telephony via the public switched telephone

Table 1
Translation of Key Requirements Emerging from Market Survey to Design Decisions

Market (User) Expectations	Design Decisions (Phase 1)
Reasonable data rates	0.8 Mbps (standard), 1.6 Mbps (optional).
Support for multimedia traffic	Simultaneous support for 4 interactive voice sessions + multiple asynchronous data connections.
Cover a single family dwelling	Range: ~ 50 meters.
Usable battery life	Low power operation (nominal 100 mW) with explicit support for power management.
Global interoperability	2.4 GHz unlicensed spectrum with world-wide availability
Tolerance to interference	Frequency Hopping Spread Spectrum system
Protection against eavesdropping	Built-in Encryption support (optional)
Connection to multiple devices	Up to 128
Legacy software should "just" work	Plug-and-play with tight integration with TCP/IP and PSTN
New applications that improve lifestyle.	Family Communications – advanced messaging system etc. Entertainment/creativity – interactive games etc. Home Control – home security, baby monitoring, etc. Resource Sharing – Internet, printers, file sharing etc.
Affordable, sub-\$100 range	Reuse available technology, and reduce component cost wherever possible

network (PSTN) and voice-over-IP (VoIP). Additionally, because of this design approach the HomeRF WG made rapid progress in finalizing the specification and bringing it to market in a timely manner [2].

Today, the HomeRF organization consists of approximately 100 member companies representing the bulk of the PC, telecommunications, and consumer electronics industries. General information on the organization is available at [1]. The specification described in this paper started at a Rev 0.1 from a proposal made in late 1997 and was approved and published as Rev 1.0 in January 1999 [2]. As of this writing, the Rev 1.2 specification is available, which includes methods of bridging between a HomeRF network and wired networks such as HomePNA [3] and Ethernet [4].

In this paper, we present the vision, design, and implementation of the HomeRF networking standard. In particular, in Section II we present a sampling of applications and usage scenarios that motivated the development of this standard. In Section III, we present the architectural overview of the network and discuss some of its important features. In Section IV, we describe the medium access control (MAC) protocol. In Section V, we describe the physical layer (PHY) and radio design for SWAP-CA devices. In Section VI, we describe the software architecture, which allows SWAP-CA hardware to inherit legacy-networking applications and for creating new applications. In Section VII, we compare SWAP-CA with some of the other wireless networking technologies, and, finally, in Section VIII, we conclude with a discussion on the future of the HomeRF group and the evolution of the SWAP-CA standard.

II. VISION AND USAGE SCENARIOS

The HomeRF WG sees SWAP-CA as one of several connectivity options in the home of the future. The relationship of SWAP-CA with other connection options is shown in

Fig. 1. In this scenario, the main home PC is linked to an Internet gateway that might be a 56K, xDSL, or cable modem. This link may be a simple cable, a wired network connection, or a wireless SWAP-CA network connection. This main home PC would likely have a variety of built-in or peripheral resources such as a printer, a scanner, a CD drive, a DVD drive, etc. For most home PCs today and looking forward, it is likely that the universal serial bus (USB) [6] would be the bus of choice for many peripherals that do not need to be mobile or remote from the PC. For video applications, which require connection to devices such as camcorders and video players, the IEEE 1394 [7] is the expected choice. For sharing resources amongst multiple PCs, options such as conventional 10/100-BaseT Ethernet [4], home phone line networks [3], and AC power line networks [9] will exist. (This last option is particularly well suited for many home automation applications where very low data rates are acceptable.) As of this writing, there are no viable RF networking alternatives at consumer price points that can be used as a networking technology for home networks. The goal of the HomeRF WG is to fill this void.²

The networking vision of the HomeRF WG is shown in Fig. 1. The RF technology supports both a network of isochronous clients that are slaves to the main home PC and a network of asynchronous peer devices that is effectively a wireless Ethernet. In most cases, the system contains a *connection point* (CP), which is usually connected to the main home PC via USB. The isochronous clients, such as cordless telephones, wireless headsets, or remote input-output (I/O) devices to the home PC (a consumer personal information manager, or PIM), are always bound to the CP that assigns them guaranteed bandwidth for bounded latency communication. For asynchronous communications between devices,

²Very Fast Infrared [8] at 16 Mb/s is a reasonable "no-cable" choice for short-distance, line-of-sight communications.

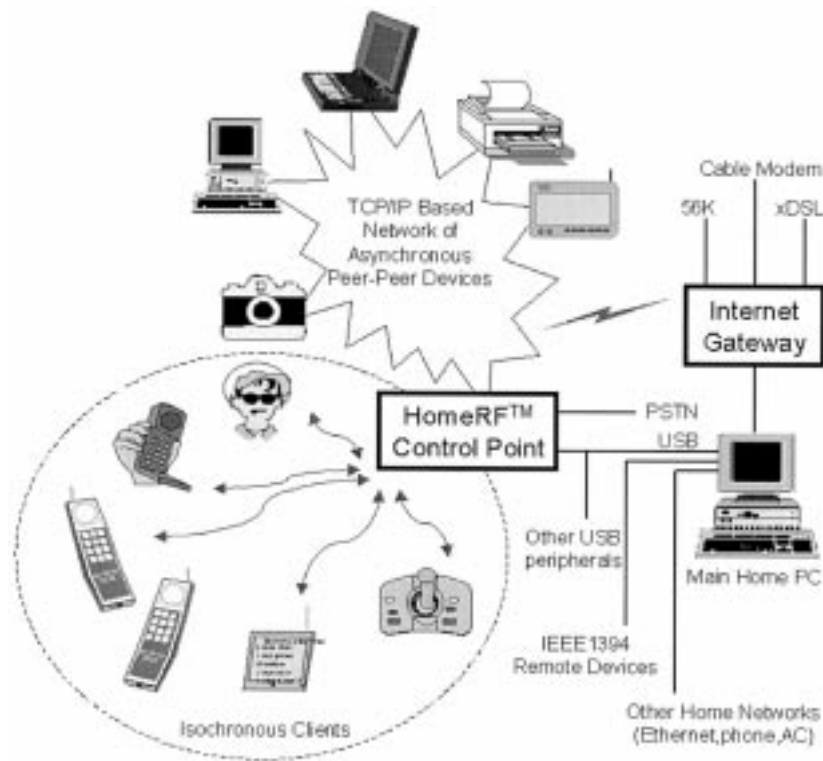


Fig. 1. HomeRF WG's vision for home networking.

the CP is not absolutely necessary. Asynchronous peer devices can communicate with each other directly, and the main home PC is just another peer device.

We now describe three distinct application areas that can directly benefit from the technology built by the HomeRF WG.

Our first example is a PC-enhanced cordless telephone. Today, there are no standards-based digital cordless telephones for consumer use in the United States, and interoperability of multiple vendors is not there. SWAP-CA defines a new standard for interoperable digital cordless telephones both in the U.S. and globally. Furthermore, the SWAP-CA specification includes a standard method for connecting the cordless telephone to the home PC software applications. Thus, many new enhanced features are possible. For example, for an inbound call, the caller ID information is sent to a PC application that looks up the caller's name and then routes the call to a specific handset instead of another number. For an outbound call, the PC interprets a spoken destination name (e.g., "Call Victor") through voice recognition, then depending upon the date/time determines the most likely number to reach the called person, and finally routes the call using the lowest cost option available at that time (e.g., using IP telephony). The handset could be used to pick up voice-mail selective to the user from the home PC call center. With voice synthesis, the handset could also be used to "listen" to e-mail. With more sophisticated application software, the handset could achieve PIM functionality by using voice or keypad I/O to store lists (i.e., "Add three quarts of milk to my shopping list") or control home automation features (i.e., "Turn the

temperature up three degrees"). All of these and undoubtedly much more creative features are possible because of the standard interoperable method of connecting to the home PC. The cordless handsets themselves are slightly different but not substantially more complex or expensive than the existing "dumb" cordless handsets sold in multimillion-unit volume today.

A second example is a mobile viewer appliance. This could take many forms but fundamentally consists of a color LCD display (like that of a notebook computer) with some limited input device (such as a pen) and a SWAP-CA radio network connection. Such a device could be either an extension of the home PC (like an X-terminal or terminal server client) or simply a web-browsing extension of an Internet gateway. In either case, the viewer communicates entirely through receiving and sending TCP/IP packets.

The third of many potential applications is resource sharing amongst multiple PCs in the same home. The resource to be shared could be a high-quality printer, a backup storage device, a file server, or an Internet connection. Clearly, these resource-sharing applications have received considerable attention from other home wiring-based networking alternatives. It is important to note that the market for HomeRF is not strictly multi-PC homes. Any home equipped with a modern home PC or an Internet gateway is a candidate for compelling mobile devices enabled by the SWAP-CA specification. Consider, for example, handheld devices that allow multiplayer gaming.

Thus, SWAP-CA is a hybrid in several ways; it is client-server between the CP and voice devices, but is peer-to-peer between data devices. The interactive voice


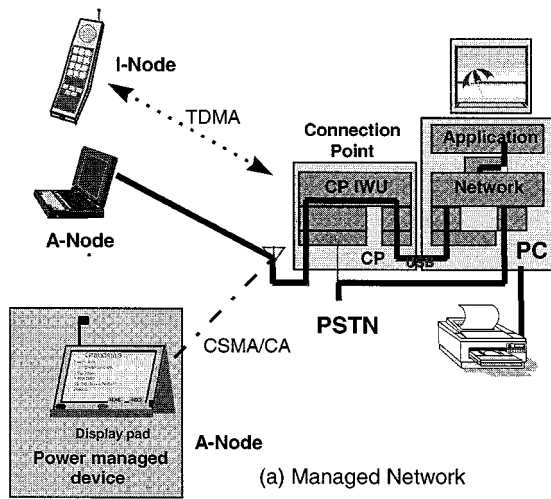


Fig. 2. SWAP-CA network topology flexibility.

transactions are circuit switched, but the asynchronous transactions are packet switched. It is precisely this richness that gives SWAP-CA the capability to be used broadly in the home; it is not designed to support hundreds of users doing similar things in an enterprise, but rather the variety of applications that occur in a residential setting.

In the next section, we describe the different device types, the network configurations in which these devices operate, and the management and operation of SWAP-CA networks.

III. NETWORK ARCHITECTURE AND OPERATION

The SWAP-CA network is designed to operate in the 2.4-GHz *industrial, scientific, and medical* (ISM) band. The 2.4-GHz band is an unlicensed frequency band that is available all over the world, and so SWAP-CA devices can operate globally. The standard is unique in the way it combines ETSI Digital European Cordless Telephony (DECT) standard [10]–[12], for carrying time-sensitive real-time traffic such as interactive voice, and IEEE 802.11 [14] and OpenAir [15], for carrying traditional data networking traffic such as file transfer.

We now describe the different types of devices that can exist in a SWAP-CA network along with their key operational features.

A. Device Types

There are four types of devices that can operate in a SWAP-CA network. These are:

- 1) *Connection point* (CP), which is the gateway between a SWAP-CA device, the PSTN, and a PC (possibly connected to the Internet).
- 2) *Isochronous node* (I-node), which is a voice-centric device such as a cordless phone and a walkie-talkie.
- 3) *Asynchronous node* (A-node), which is a data-centric device such as a handheld notepad and a personal digital assistant (PDA).
- 4) Combined asynchronous-isochronous node (AI-node).

The CP can either be a separate device connected to the main home PC typically via the USB connection, or it can be an integral part of the PC. It can also have a direct connection to the PSTN. The CP is capable of performing data transfers to and from SWAP-CA devices using both a contention-based and a contention-free protocol (Section IV). When configured to do so, it can provide power management services to both A-nodes and I-nodes (see Section IV-C).

Fig. 2(a) shows an example of a typical SWAP-CA network consisting of two A-nodes, one I-node, and a CP. One of the A-nodes is a power managed display pad whose communications traffic is managed by the PC so it can maximize battery life. Although not shown in this figure, the laptop A-node could also be power managed. As this figure shows, SWAP-CA has a unique ability among networking protocols to mix intense high-demand packet traffic with infrequent command and control traffic, and high-quality real-time voice traffic. The PC is an integral part of the SWAP-CA system, although peer-to-peer data networking is available even when the PC is inoperative. Every SWAP-CA device has a 48-bit IEEE MAC address, which is configured by the manufacturer prior to distribution.

B. Network Configuration

The SWAP-CA network is designed to operate in two basic modes: the network can be configured either as a *managed network*, as shown in Fig. 2(a), or as a *peer-to-peer ad-hoc network*, as shown in Fig. 2(b). In the managed network configuration, the network is under the explicit control of a CP, which functions as its gateway to other devices, the Internet, and the PSTN. Furthermore, in this mode the network can simultaneously support both real-time audio traffic, such as interactive voice, and nonreal-time data traffic, such as traditional TCP connections. In the *ad-hoc* network mode, the network provides traditional data networking support only and does not need a CP for proper operation.

It is possible for two or more SWAP-CA networks to co-exist at the same time, even when they are within range of one another. This is explained as follows.

C. Multiple Overlapping Virtual Networks

There are certain characteristics of RF networks that make them unique. For example, RF signals are not restricted to well-defined boundaries; consequently, a RF node can “hear” other RF nodes operating on the same frequency when they are within its range. This then opens up the possibility that two or more SWAP-CA networks can overlap when devices belonging to these networks come within range of one another. Scenarios like this are likely to occur in places such as an apartment complex where every apartment has its own SWAP-CA network, but since the apartments are close to one another devices in one apartment are within range of devices in a neighboring apartment.

Usually overlapping SWAP-CA networks do not interfere with each other because the physical layer of the SWAP-CA network uses frequency hopping spread spectrum (FHSS) modulation for transmission of packets. Since different networks are generally not synchronized and they use different frequency hopping patterns, the probability of interference between overlapping networks is low. Still, because there is a small finite probability that neighboring networks can sometime hop on to the same channel, SWAP-CA defines a 24-bit *network identifier* (NWID). The NWID is used to prevent “promiscuous” reception of packets from neighboring networks. Devices having the same NWID are part of a logical network and the NWID is present as part of every data packet in the network; devices only accept those packets that contain the NWID of their private network. Thus, the NWID is used to separate the different overlapping virtual networks located in the same area of coverage.

Depending on its configuration, a node may either use the NWID that is explicitly assigned to it by the user or alternatively derive a NWID from its MAC address. For an existing network, new nodes can learn the NWID as they join the network. For reasons just described, configuring the NWID explicitly is the recommended option to avoid the problem of nodes learning the NWID of neighboring networks and joining those.

In the following, we describe how nodes join an existing network and create a new one using a derived, learnt, and explicitly configured NWID.

D. Discovery and Creation of Networks

When a SWAP-CA node is turned on, it immediately enters into a *network discovery phase* in which it tries to determine if another node or CP is present within its range and if a network already exists that it can join. The node accomplishes its discovery phase by operating in a *passive scanning* mode. In the passive scan mode, the node listens on every channel,³ within its operational frequency band, for a specific amount of time, greater than a single *superframe* (see Section IV). “Listening” on a channel is accomplished by the physical layer, which forces the node to hop on a known scan pattern that is a good spread of the hopping patterns on all the

³The device has to conform to the rules of the geographic region in which it is operating. For example, in North America and Europe, the number of channels the MAC hops over is 75, in Japan it is 23, in France it is 35, and in Spain it is 27.

channels of the network. During a scan, the node receives all network packets regardless of the NWID or destination address. These packets are analyzed by the node’s MAC management module, and a decision is made on whether or not to join the network. The scan can be terminated as soon as the first network is found or when the management module so determines. When a network is found, the node synchronizes to the network by setting its hopping pattern to the discovered network’s hopping pattern and by recording the network’s NWID. The synchronization information is available on every data packet on the network and is used by the “joining” node to set its own parameters.

When the network identifier is known, a SWAP-CA node joins the known network by scanning all the channels for a specific NWID and then by locking on to the channel on which it finds the NWID. The scanning procedure is terminated as soon as the sought-after NWID is discovered. In case the node does not find a network with the particular NWID, it can either give up or start its own network. To start its own network, the node randomly selects an available hopping pattern, records the NWID as the network identifier, and starts transmitting synchronization signals. It is noted here that not all nodes are allowed to create a network. Specifically, a CP can create a managed network and an A-node can create a peer-to-peer *ad-hoc* network. I-nodes do not create their own network.

In a managed network, the CP is responsible for transmitting synchronization information, whereas in an *ad hoc* network, all A-nodes participate in synchronization. A-nodes share responsibility for beaconing by adding a random backoff to the scheduled *ad-hoc* beacon transmission time. The synchronization information⁴ contains the network’s hopping pattern and the dwell time⁵ and is transmitted as part of the beacon signal. When there are two or more devices in the network that are capable of CP functionality, the first one to join/create the network becomes the active CP while the remaining ones become passive CPs. In case a passive CP does not hear 50 consecutive beacons from the active CP, it assumes that either the active CP has gone off-line or has moved out of range and consequently creates its own network by transmitting synchronization information. If 100 or more simultaneous beacons are missed by the A-nodes in the network, they start operating the *ad hoc* network synchronization operation and create an *ad hoc* network. The entire process of scanning followed by either joining or forming a new network can on an average take about 1.5 s.

E. Authentication and Privacy

As discussed in Section III-C, RF signals are not restricted to well-defined boundaries and consequently, unlike a wired network, an RF wireless network is difficult to secure. The transmission medium is open to anyone within range of the transmitter. Even when the physical layer of the network

⁴A beacon signal is a broadcast packet transmitted by a node (a CP in a managed network) periodically. It contains information that the network devices need for proper operation within the network.

⁵The dwell time is equal to the *superframe* period

is based on spread spectrum communication and different networks use different hopping patterns, the system is insecure as it is relatively easy for a malicious user to scan all channels and determine the hopping pattern and NWID of the target network. Thus, neighbors who receive RF signals from each other can conceivably “listen in” on each other’s conversations and intercept each other’s data packets. This is clearly undesirable. Consequently, an encryption mechanism that provides security and privacy is clearly needed and should be part of any wireless networking standard.

Data privacy and authentication in a SWAP-CA network is accomplished by using a well-established shared-key encryption algorithm. Notably, all I-nodes follow the per-session key security model defined in the ETSI DECT specification [12]. The authentication process in this model is split into a key generation process and the encryption process. The purpose for splitting this process is to support roaming of handsets between DECT base stations. SWAP-CA defines its authentication process in terms of the DECT security model.

A node indicates that it supports encryption to the destination node as part of the capability exchange process. Data packets are encrypted only if the destination node can decrypt the message. The encryption algorithm takes a 56-bit key and a 32-bit initialization vector and uses these to convert unencrypted data (called *plaintext*) to encrypted data (called *ciphertext*). The initialization vector is a combination of the packet sequence number and a hash of the 48-bit MAC address of the source node. The sequence number keeps track of the number of packets the node has encrypted and prevents against replay attacks. A magic number, a byte containing all zeros, is appended to every packet before encryption to allow the receiver to check, as it decrypts the packet, whether or not it has the right key. All A-nodes in the SWAP-CA network share a single common 56-bit key. The key may be entered through a *management information base* (MIB)⁶ or computed dynamically. A property of this algorithm is that the output ciphertext is of the same length as the input plaintext and the encryption algorithm is symmetric (i.e., decryption performs the same process as encryption). The core of the encryption algorithm is common to both asynchronous and isochronous data services.

All multicast and broadcast traffic is sent unencrypted since there is no guarantee that all A-nodes within range support the encryption option.

F. Compression

Since wireless networks generally have limited bandwidth, an often-stated design goal for these networks is to build a system that is bandwidth efficient. One way to improve bandwidth efficiency is to compress the data before transmission.

Data compression in a SWAP-CA network is optional and is left to the designer’s discretion as it provides a tradeoff between battery longevity and bandwidth. With compression, nodes can transmit more data in a given amount of time than

⁶An MIB contains information that is used to manage the operation of a SWAP-CA node. This information can be used by higher layer protocols to manage the node as well.

Table 2
MAC Behavior Depends on Device Type (M: Mandatory, O: Optional)

MAC Function	CP	A-node	I-node	AI-node
Synchronization within a managed network	M	M	M	M
Synchronization within a ad-hoc network	-	O	-	O
Power Saving	O	O	O	O
A-node power support	O	O	-	O
Power Management for multicast	M	-	-	-
Connection Management	M	-	M	-
TDMA access	O	-	M	M
CSMA/CA access	M	M	-	M
Service-slot access	-	-	M	O
CP assertion	M	-	-	-

nodes that do not compress their data; however, compression consumes power and, therefore, contributes to battery drain for portable devices. Another problem is that data compression involves the reduction of redundancy in the data. Consequently, any corruption of the data is likely to have severe effects and be difficult to correct. The point is that the network operator has to consider all these tradeoffs before deciding whether or not to enable compression.

The recommended compression algorithm for a SWAP-CA network is a lossless compression algorithm that uses a combination of the LZ77 algorithm [16] and Huffman coding [17]. The efficiency of this algorithm is comparable to the best currently available general-purpose compression methods. The data can be produced or consumed, for an arbitrarily long sequentially presented input data stream, while requiring very few resources in terms of processing power and memory [18].

As with encryption, compression is used only if the source node determines that the destination node is able to decompress the message. This is determined during connection-setup and capability exchange time. Compression is not used for multicast and broadcast traffic.

We now discuss in detail the medium access control mechanism that makes SWAP-CA unique, and a compelling technology for home networks.

IV. MEDIUM ACCESS CONTROL

The SWAP-CA medium access control protocol is derived from the ETSI DECT standard [10]–[12] and from the popular WLAN standards such as IEEE 802.11 [14]. It is optimized specifically for the home environment. The MAC is designed to work over, and to take advantage of, the frequency hopping radio subsystem. It includes a time division multiple access (TDMA) service for delivery of real-time isochronous data, and a carrier sense multiple access with collision avoidance (CSMA/CA) service for delivery of asynchronous data. The MAC’s behavior depends on the devices that are in use in the network. Table 2 contains the functionality supported by the MAC for different device types.

Some of the highlights of the MAC protocol are as follows.

- Simultaneous support for both voice and data traffic using a unique combination of TDMA and CSMA/CA access mechanisms.

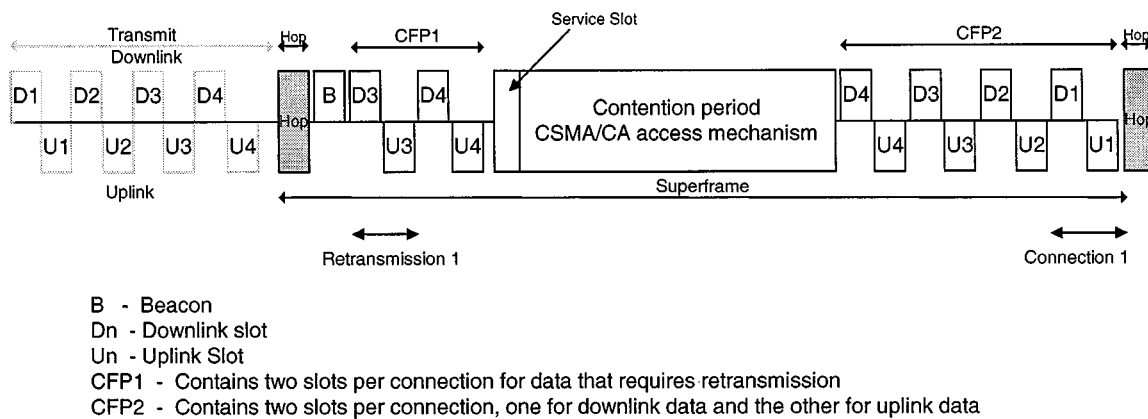


Fig. 3. SWAP frame description.

- Support for four high-quality, 32-Kb/s adaptive differential pulse code modulation (ADPCM) voice connections.
- Data throughput of about 1.6 Mb/s.
- Built-in power management capabilities for both isochronous and asynchronous nodes.
- Multiple levels of data security are provided—None/Basic (FHSS + NWID)/Robust (FHSS + NWID + 56-bit stream cipher)
- Support for multiple networks in the same physical area with a 24-bit network identifier.

Fig. 3 illustrates the protocol framing when the network is configured as a managed network.⁷ SWAP-CA defines a *superframe*, which is a periodic division of time. A superframe contains two contention-free periods (CFP1 and CFP2) and a contention period. The channel access mechanism used during the contention-free periods is TDMA and during the contention period is CSMA/CA. The duration of the superframe is fixed, and a synchronization mechanism allows the nodes of the network to agree on the superframe timing.

SWAP-CA devices employ a frequency hopping spread spectrum modulation for transmission of packets (details in Section V). All nodes in the network hop together using the same hopping pattern, and they hop to a different frequency at the start of every superframe. By hopping to a different frequency channel every superframe period, the dwell-time of the devices on a portion of the spectrum in which there may be high interference is reduced. This, in turn, reduces packet errors and packet loss, thus, improving the overall throughput and reliability of the system.

The start of a superframe is marked by a *connection point beacon* (CPB) signal. The CPB is used for multiple purposes, including: 1) maintaining network synchronization; 2) controlling the format of the superframe; 3) managing the network during the contention-free period by indicating when each node should transmit and receive data. The CPB can include a list of active voice connections (and, therefore, slot assignments), retransmission slot assignments for the current superframe, connection status information, and paging information; and 4) providing power management services for isochronous and asynchronous nodes to maximize the battery

⁷In *ad hoc* networks, the SWAP-CA MAC reduces to a CSMA/CA MAC.

life of portable devices (see Section IV-C). Slot assignment and synchronization information does not change on a per frame basis, so if a node misses a beacon it uses the information contained in the most recent valid beacon. All connection and paging status requests and information are repeated until the receiver acknowledges them.

Real-time voice traffic is tightly integrated into the MAC. This is discussed as follows.

A. Voice Encoding and Transmission

The SWAP-CA standard specifies a 32-kb/s ADPCM codec, defined in CCITT Recommendation G.726 [13], as a baseline voice encoding mechanism. I-nodes process 20 ms segments of 14-bit linear PCM audio samples, sampled at 8 KHz. These samples are companded and encoded to a sequence of 4-bit ADPCM codewords before being packetized and queued for transmission in a chronological order. Encoded voice packets are transmitted during the contention-free period.

The contention-free periods are divided into a number of pairs of fixed length slots, two per voice connection. The first slot in each pair is used to transmit voice data from the CP to a node (downlink), and the second is used to transmit voice data from a node to the CP (uplink). CFP2 at the end of the superframe is used for the initial transmission of the voice data, while CFP1 at the start of the superframe is used for the optional retransmission of any voice packet that was not received or was incorrectly received in the previous dwell. Each voice packet transmitted by an I-node includes in the packet header a piggyback acknowledgment (ACK) of the last voice data message received by the node, i.e., in the uplink packet the voice node ACKs the downlink packet sent by the CP. This allows the CP to determine prior to a hop which voice data transmissions were lost, and determine whether or not retransmissions are required. Retransmissions are advertised in the beacon at the start of the next superframe; each voice data packet is retransmitted only once.

The time between the first voice packet transmission during CFP2 and its retransmission in CFP1 is a function of the voice codec and is fixed at 20 ms. This provides an acceptable performance with respect to latency. The length of the dwell period is equal to a single voice data message



- Sender selects a slot (backoff counter) and then decrements the counter whilst the medium is clear
- Medium must be free for a DIFS period before the backoff counter is decremented.
- Example shows transmission of a packet in slot 5.

Fig. 4. CSMA medium access procedure.

containing 20-ms segments of ADPCM data (640 bits), which is equivalent to an extended DECT B-field plus 56 bits of control data, and equivalent to the DECT A-field plus some additional addressing information [11]. With a 20-ms superframe, SWAP-CA can support up to four voice connections simultaneously with full retransmission possibility.

Robustness against interference is achieved by choosing CFP2 for the initial packet transmission and the following CFP1 for any retransmission. Thus, the system provides both frequency and time diversity. This is particularly important given the potentially noisy environment in which the protocol operates.

At the end of CFP1, there is space reserved for a *service slot*. I-nodes use the service slot to request a connection from the CP. Since there is only one service slot, it is possible for two or more nodes to transmit at the same time and for their transmission to collide. Each management message is explicitly acknowledged by the CP in the CPB, and if there is no ACK, a node performs a random backoff across a number of dwell periods before resending the message. The average time to connect can vary since CFP1 varies, depending on number of retry slots. CFP2 is more constant, but always depends on the number of active handsets. Normally, the connection setup will succeed at the first possibility. Subsequent attempts to use the service slot will take place on average 160 ms apart. CFP1 varies, depending on number of retry slots. CFP2 is more constant, but always depends on the number of active handsets. Normally, the connection setup will succeed at the first possibility. Subsequent attempts to use the service slot will take place on average 160 ms apart.

B. Data Transmission

For data traffic, a CSMA/CA access mechanism is used during the contention period of the superframe. With this scheme, the protocol provides efficient data bandwidth even with concurrent active voice calls. Peak effective user throughput of up to 1 Mb/s is possible under lightly loaded conditions in the 1.6-Mb/s mode. Furthermore, data transfer between nodes can occur even when four voice calls are active simultaneously.

The CSMA/CA protocol attempts to avoid collision by adopting listen-before-transmit (LBT) etiquette—sensing transmissions on the radio medium and starting transmission when there is no such activity. The mechanism is similar to Ethernet (IEEE 802.3 [4]), enabling easy integration with an existing TCP/IP protocol [19] stack within a host platform; the main difference with Ethernet is the slotted contention mechanism and the addition of MAC level ACK of unicast packets. Fig. 4 illustrates how the medium is accessed during the contention period.

Access to the wireless medium is controlled through the use of interframe spacing (IFS) time intervals. The IFS intervals are mandatory periods of idle time on the transmission medium between frames. Two IFSs are specified in the standard: Short IFS (SIFS) and DIFS (the name is taken from the IEEE 802.11 standard [14]). SIFS is the time it takes for the PHY layer to get ready for a transmission or reception. It is constrained by the physical layer's (PHY) performance and dictates the minimum time separating two independent MAC transmissions on the channel. *Slot-time* (or *SlotDuration*) is also constrained by the PHY layer. It is the time taken by the PHY layer to estimate the state of the wireless medium as idle or busy with certain accuracy. DIFS is one *SlotDuration* longer than a SIFS period. When the channel is free for a DIFS time period, it indicates to the node that the previous data exchange is complete and that it may contend for the channel.

The access procedure is designed to provide fair access to the wireless channel for all nodes by using a contention window and backoff counter as shown in Fig. 4. Before any node transmits a packet, it selects a backoff counter (a number of contention slots) and then starts "listening." When the medium has been clear for a DIFS period, it decrements its backoff counter for each free contention slot. When the backoff counter expires, the node transmits the message. Whenever the medium is busy, the countdown is suspended and only resumes when the medium has been free for a DIFS. This backoff mechanism reduces the probability of collision, and performing a backoff before transmission ensures that responses from multiple nodes responding to a broadcast message on an otherwise idle network do not all collide. If a retransmission is required because of a collision

Table 3
Typical Values of Some Attributes Used by the MAC Layer

Attribute	Typical Value
Superframe duration	20 msec
SIFS	142 usec
DIFS	309 usec
Maximum beacon period duration	1278 usec
TDMA slot-pair duration (2*TDMA slot duration +SIFS)	2364 usec
Maximum unicast MSDU length	1500 octets
Maximum time to spend transmitting/receiving a data packet	100 msec.
Maximum contention window size	64 slots
Minimum contention window size	8 slots
Persistence of events in the CPB	3
Time after which node can go into power saving mode	1 sec
Maximum time between checking CP beacons by a PS-node	4 sec
Maximum time between power management requests	60 sec
Maximum number of dwells I-nodes backoff for service slot	16

DA	SA	Ethertype	SC	MSDU	Status

DA - IEEE 48 bit MAC address of Destination
 SA - IEEE 48-bit MAC address of Source
 Ethertype -16-bit etheretype value
 SC - Service Control (Encryption, compression, data rate)
 MSDU - MAC Service Data Unit
 Status - success or expired

Fig. 5. Basic MAC frame for the asynchronous data service.

or transmission failure the size of the collision window is increased from an initial value of eight exponentially up to a maximum 64 to avoid congestion.

Typical values for parameters used by the MAC are shown in Table 3.

The basic structure for the MAC data frame exchanged between peers is shown in Fig. 5. The *Ethertype* used in SWAP-CA is compatible with the Ethernet Etheretype, and its value is managed by the IEEE. An example of a popular Etheretype, for Internet protocol, is 0x0800. The asynchronous data service supports a *MAC service data unit* (MSDU) with sizes of up to 1500 bytes. All MAC frames are protected by a 32-bit CRC.

Reliability of unicast data packets is guaranteed through an ACK/retry mechanism; however, it is not guaranteed for multicast data packets. Also, the order of unicast data packets sent between two SWAP-CA nodes is generally preserved. The MAC is not allowed to gratuitously reorder data packets; however, reordering is possible when the device is in power-saving mode.

When there is no CP present, data nodes can create an *ad-hoc* network in which control of the network is distributed between all the nodes. When there are no active voice connections, the CSMA/CA period expands to occupy the entire superframe, with the exception of the hop and beacon, thus, maximizing the network data throughput.

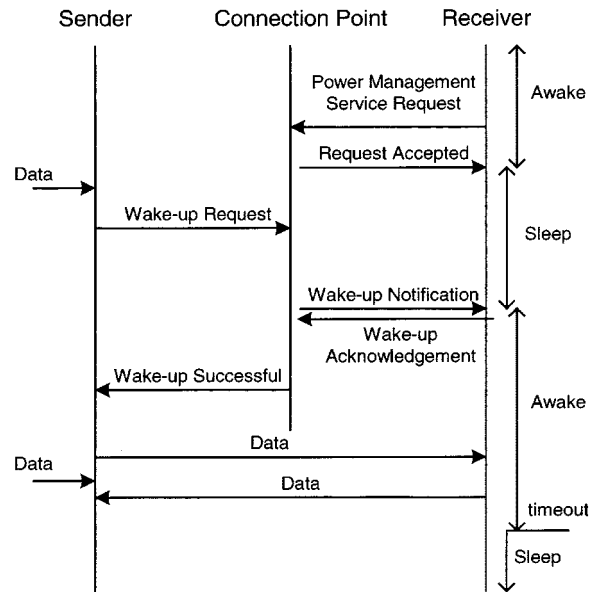


Fig. 6. Power management of TDMA nodes in a SWAP-CA network.

The MAC contains special support for saving power. This is explained in detail in the following section.

C. Power Management Services

It is a well-understood fact that battery energy is a limited resource that is not expected to increase in potential more than 30% in the near future [20]. Consequently, for portable battery operated devices, this resource needs to be used carefully and efficiently. A wireless network transceiver can typically use anywhere from 15%–30% of the power from a typical portable computer, the display being the only part that consumes more power [21]. Most wireless network interface cards that are in the market today consume close to 12x the power of a standard 10 Mb/s Ethernet card, and the battery longevity for wirelessly connected portable devices is reduced by as much as 60%, i.e., from a three-hour operation down to 72 min. It is therefore mandatory for any wireless standard to support power management functions and be power-aware. So, one of the primary design goals that influenced the design of SWAP-CA MAC was to provide power management services. We discuss how this service is incorporated in the standard.

Power management can be enabled only when the network is configured as a managed network and the CP is a key component of the power management services in a SWAP-CA network. Both A-node and I-nodes can save power by enabling the power management services within the CP.

For I-nodes, the procedure for power management is straightforward. During an ongoing connection (e.g., an active voice call) the I-node that has registered with the CP as being a power saving node powers on, initially only for the duration of the CPB, to receive slot assignment information. It then powers down until its assigned slots are due. When not in an active connection state, I-nodes power up every N dwells, where the system designer chooses N as a compromise between power saving and speed of

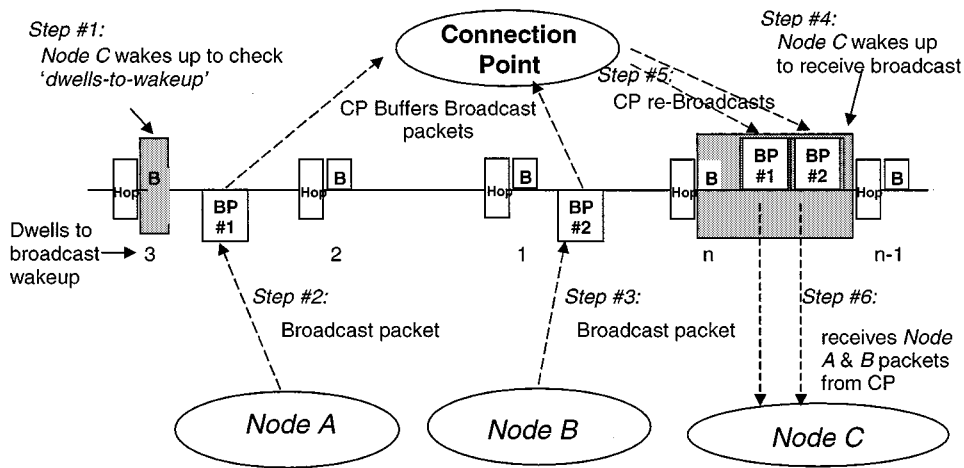


Fig. 7. Power management of CSMA nodes for broadcast messages.

response to a new connection. In every beacon the CP broadcasts a list, or directory, of the nodes, which have traffic pending for them. Power management for I-nodes is illustrated in Fig. 6.

For A-nodes, the behavior of the power management service depends on the type of packet being transmitted. Fig. 7 illustrates the process of sending broadcast messages to power-saving asynchronous nodes. The process is described with the help of an example in which two asynchronous nodes (*Node A* and *Node B*) send broadcast packets, and *Node C*, a power saving node, is a receiver of these packets:

- Node C* powers up and receives a “dwells-to-wake-up” counter of three, which is broadcast in the CPB. The maximum value of the “dwells-to-broadcast” counter is a system design parameter that allows the designer to trade off latency, CP buffer size, and broadcast reliability against battery life. The CP maintains a count-down to the next dwell when power saving (PS) nodes would wake up.
- Node A* transmits a broadcast message, which is received and stored by the CP.
- Node B* transmits a broadcast message, which is received and stored by the CP.
- Node C* wakes up when its dwells-to-wake-up counter reaches zero. It checks the CPB and sees an indication from the CP that packets are available for it and consequently stays powered up.
- The CP transmits the buffered broadcast messages during the dwell.
- Node C* receives the broadcast messages and then goes back to sleep mode.

Fig. 8 illustrates the process for sending unicast messages to power-saving asynchronous nodes. Once again, we explain the process with the help of an example: *Node A* sends a unicast message to *Node B*, which is a power-saving node. The wake-up pattern of power-saving nodes (*Node B*) is controlled by a “wake-up” flag transmitted in the CPB upon request from the sending node. The example below describes how this feature is used.

- Node B* powers up and checks the CPB. It does not see a wake-up signal from the CP, concludes that no packets are waiting for it, and powers down.
- Node A* sends a request to the CP to “wake up” *Node B*.
- The CP asserts the wake-up flag for *Node B* in the CPB.
- Some time later (dependant on the designer’s tradeoff of power-saving versus latency) *Node B* wakes up, checks the CPB, and receives the wake-up flag. Consequently, it stays powered on for the dwell.
- Node B* and *Node A* transfer data using the normal CSMA/CA access method.
- Node B* powers down after the exchange of final message is complete.

Some general comments follow: The unicast and multicast state machines are independent; consequently, a power-saving A-node can operate unicast, multicast, or both types of power-saving procedures. The multicast power management service is completely transparent to the sender; the CP automatically resends all multicast packets regardless of their source. A SWAP-CA network may contain both PS and non-PS nodes simultaneously. The CP stores the address of the PS nodes as a list of nodes to which it provides power management service. To guarantee proper network operation, PS nodes periodically re-request power management services from the CP to guard against CP disappearance from the network (when it is stopped, or when a higher functionality CP takes over) and against CP’s holding on to resources assigned for a PS-node.

V. PHYSICAL LAYER AND COMPONENTS

The physical layer specification for SWAP-CA was largely adapted from the IEEE 802.11 FH [14] and OpenAir [15] standards with significant modifications to reduce cost while maintaining more than adequate performance for home usage scenarios. The SWAP-CA PHY layer provides the transmission and reception of data packets in the 2.4-GHz ISM band, using a 2-level frequency shift key (2-FSK) [22] modulation scheme, for 0.8-Mb/s raw data rate performance, and an optional 4-FSK, for 1.6-Mb/s raw data rate performance.

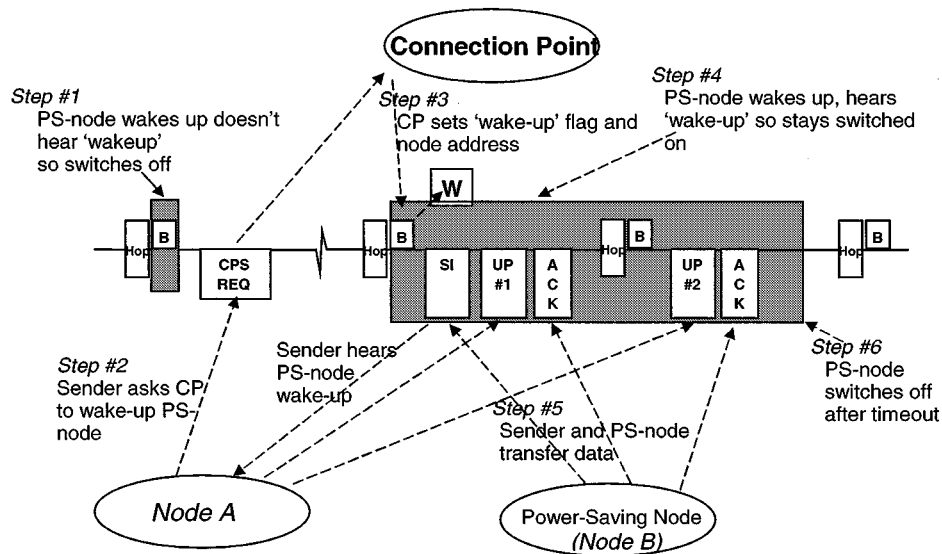


Fig. 8. Power management of CSMA nodes for unicast messages.

Some key features of the SWAP-CA physical layer specifications are as follows.

- Transmit power—Up to +24 dBm (or nominally 100–250 mW⁸)
- Low power transmit mode (optional)—between 0 and +4 dBm (for portable devices with limited peak current capability).
- Receiver sensitivity in 2-FSK (or 0.8 Mb/s mode)—less than –80 dBm; in 4-FSK (or 1.6 Mb/s) less than –70 dBm.
- Hopping time—300 μ s (to allow conventional synthesizers to be used).
- Transceiver turnaround time—134 μ s (very easy to achieve with existing synthesizers).
- Adjacent and alternate channel filtering—no requirement.

The combination of the transmit power and the receiver sensitivity represent a typical range that should easily exceed 50 m in most home environments. In the optional low-power mode, reliable indoor range is expected to be 10–20 m (which covers the bulk of the interior of most homes). Although the PHY layer design is quite similar to the IEEE 802.11 FH and OpenAir design, the SWAP-CA requirements impose substantially lower cost constraints for three reasons. First, the required sensitivity limit is relaxed by about 10 dB. Second, the greatly relaxed channel filtering specification causes dramatically less intersymbol interference due to filter group delay variations in the passband. And third, the SWAP-CA packet headers for 4-FSK add a special training sequence to allow optimum slicing threshold values to be determined for the changing propagation environment. Thus, for usage within most homes, the 1.6-Mb/s data rate is really available with SWAP-CA and adds virtually no cost to the 0.8-Mb/s solution.

⁸The nominal transmit power level is defined as the power delivered to the antenna averaged between the start of the first symbol and the end of the last symbol in the physical layer packet header.

Although the hopping time is easy to meet, the transceiver turnaround time creates challenges for many conventional RF transceiver architectures and components. This low transceiver turnaround limit is essential for SWAP-CA to provide low latency performance in a mixed voice and data network in the presence of microwave ovens and other interference sources. Fortunately, increasing levels of integration and speed in complementary metal oxide semiconductor (CMOS) circuits now make it possible to build very fast switching channel synthesizers capable of this requirement by adapting technology previously used in precision instrumentation.

In fact, the entire SWAP-CA PHY-layer specification has been written specifically to accommodate very low-cost, single-chip implementation in CMOS technology. A typical system partitioning is shown in Fig. 9. For many of the digital devices envisioned by the HomeRF WG, the digital MAC baseband portion of the component solution can be integrated into a large application specific integrated circuit (ASIC) already in the device. At $\sim 30K$ gates for the SWAP-CA data core, this is extremely low cost in the sub- 0.25μ CMOS era. The modem functionality can interface to the digital baseband via a very simple serial interface (with no analog quantities). The modem and RF functionality can all be integrated into a single mixed-signal CMOS integrated circuit (IC) as shown because of the specific technical requirements on filtering and modulation chosen by the HomeRF technical team. Note that it probably does not make sense to integrate the RF front-end functionality such as the low-noise amplifier, the power amplifier (if present), the antenna switches, and the band-select filter onto the CMOS IC even though technically feasible. This is because the semiconductor die area for the front-end functions is typically much less than 5% of the rest of the modem (hence, low-cost already) and the overall power consumption performance is driven largely by optimizing these functions in detail.

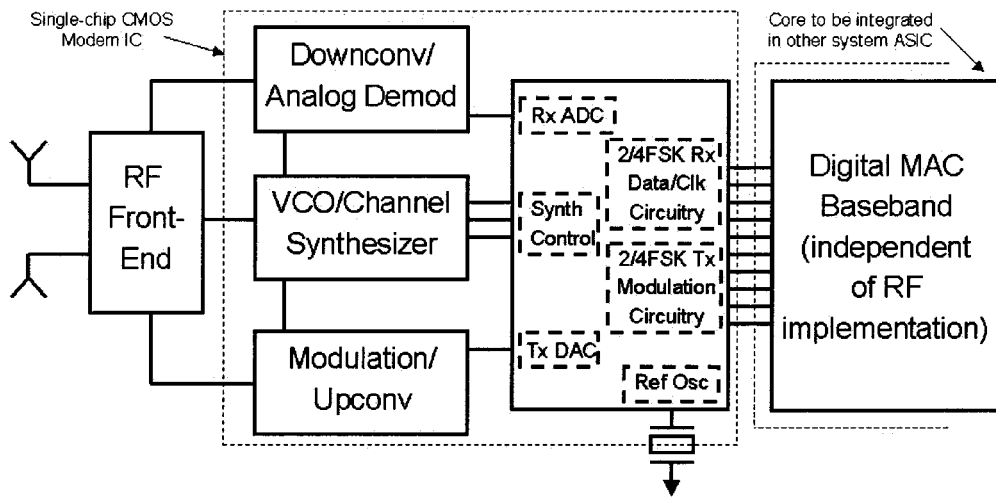


Fig. 9. Partitioning of RF modem and digital MAC component sections.

Table 4
Typical Values of Some Attributes Used by the PHY Layer

Attribute	Typical Value
Time to sense the channel	27 usec
Time to stabilize on new frequency	300 usec
Rx / Tx turnaround time	134 usec
Time to transmit single PHY symbol	1.25 usec

With such high levels of integration and optimized front-end, the RF modem section overall cost in multimillion-unit volume should be well below \$10 (similar to the situation today with DECT equipment) while the digital MAC section approaches “zero” other than IP royalties. Note that while this seems extraordinarily low compared to today’s common perception of RF data being hundreds per node, it is still very expensive compared to the very low cost of IrDA transceivers or USB controllers. Thus, cost remains a significant issue to making HomeRF a “throwaway” item in every electronic device. But consumers have consistently shown with voice that they will pay extra for personal mobility. Even today, cordless phones are significantly more expensive than corded phones, yet much more popular. If consumers begin to value mobility within the home for Internet-based content the way they do today for PSTN-based content, then the present cost projections for enabling SWAP-CA in devices should not be a serious barrier. Table 4 contains some typical values used by the SWAP-CA PHY layer.

VI. SOFTWARE ARCHITECTURE

Due to the widescale availability of Microsoft’s Windows operating system and its use as a *de-facto* standard operating system for home PCs, the HomeRF WG sought to streamline operation of SWAP-CA devices on these systems. We now briefly describe how SWAP-CA devices operate in the context of a main-line operating system.

A SWAP-CA enabled node connected to a PC exposes three distinct types of interfaces: asynchronous data, isochronous data, and management. Fig. 10 illustrates these types of interfaces. The interface exposed to legacy data-

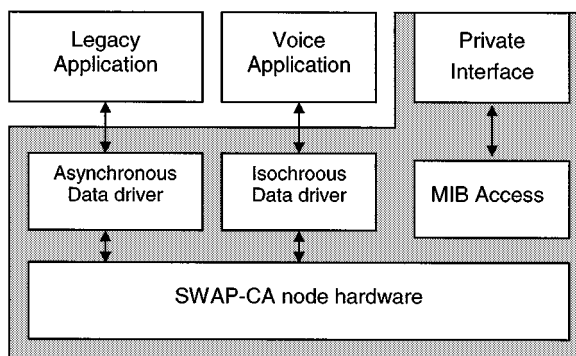


Fig. 10. Interfaces exposed by a SWAP-CA device.

centric applications and voice application is a standard part of the Windows operating system. Information about the device not covered in the standard programming interface is exposed by a private interface that may be provided by the hardware vendors to application writers as needed.

Windows defines a *network device interface specification* (NDIS) [24] for programming networking hardware. The NDIS software library, which is the implementation of this specification, performs many of the functions that are common to device drivers of networking hardware, including synchronization, device mapping, interrupt handling, event logging, etc. Furthermore, it provides a standard interface for higher level protocols and applications, which can be agnostic of the underlying hardware. Manufacturers of network adapters write an *NDIS miniport* driver that provides functionality specific to their hardware. Miniports of a given media type can be used by higher level protocols knowledgeable about that media type with no further modifications. Fig. 11 illustrates how the various software modules fit together. The shaded blocks are provided as a standard part of the operating system and the hardware vendor writes the miniport device drivers. Miniport drivers that are written to conform to the NDIS specification are guaranteed to work with the Windows operating system.

NDIS exports two distinct interfaces—a *connectionless interface* (used by broadcast media such as Ethernet) and

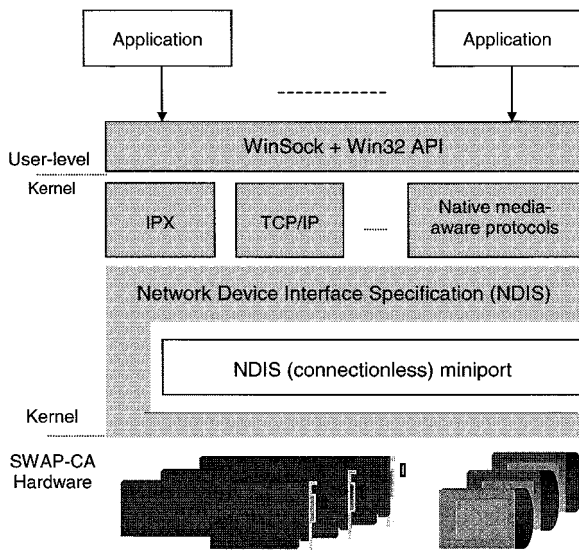


Fig. 11. SWAP-CA A-node driver architecture.

a *connection-oriented interface* (used by media that have explicit connections between endpoints, such as in ATM [25] network connections). Hardware manufacturers who produce A-node SWAP-CA devices are required to write a connectionless miniport device NDIS driver only that declares itself as a member of the Ethernet media type. To higher-level protocols and applications, SWAP-CA A-nodes are then indistinguishable from normal Ethernet adapters, allowing Ethernet-knowledgeable applications to immediately function with SWAP-CA devices. This inheritance of applications was an important design goal.

Hardware manufacturers who produce I-node devices create a miniport driver that declares itself as a member of the Ethernet media type and exposes a connection-oriented interface, not the connectionless interface traditionally used. In addition to the basic miniport functions, I-node drivers provide call management functions such as those required for setting up and terminating voice calls. Call control of SWAP-CA I-node adapters is managed by the *telephone application programming interface* (TAPI) [24]. TAPI is a simple, generic set of objects, interfaces, and methods for establishing connections between devices; TAPI communicates with NDIS via a TAPI service provider also called a *TAPI proxy*. TAPI applications set up, control, and tear-down calls on SWAP-CA devices via the TAPI proxy. The various software modules are illustrated in Fig. 12, once again the shaded region are part of the most recent version of the operating system.

Some applications may wish to stream voice conversations between SWAP-CA adapters and another adapter within the PC in real-time. An example scenario is that of a SWAP-CA I-node (a cordless phone) user communicating with a SWAP-CA connection point, which is part of the PC. This connection is forwarded on to another adapter in the PC, possibly a modem attached to a phone line or a sound card attached to speakers and a microphone. To implement this application, the voice data is streamed between the SWAP-CA adapters via the *DirectShow Streaming* module

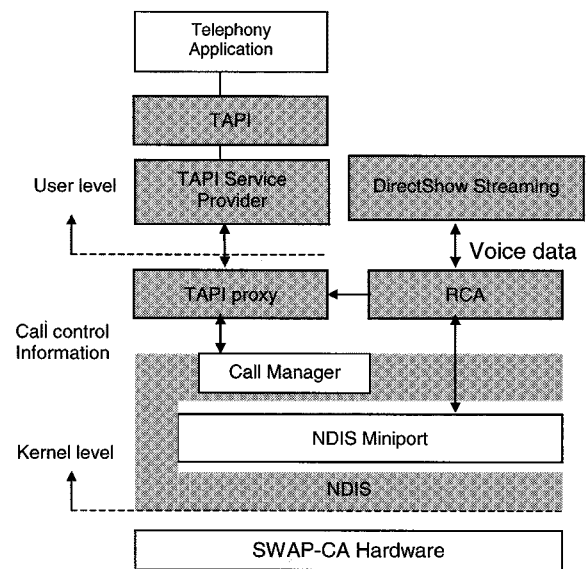


Fig. 12. SWAP-CA for isochronous service driver architecture.

[24]. The DirectX architecture was designed to accommodate real-time traffic within a PC. A DirectShow filter is plumbed from the data source (in this case, the SWAP-CA adapter) to the data sink (the modem or sound card). Voice streams coming in via NDIS are redirected to the *raw channel access* (RCA) filter, which in turn sends them into DirectShow. The RCA filter is part of the operating system so connection-oriented miniports are automatically voice stream-enabled.

VII. POSITIONING WITH OTHER TECHNOLOGIES

Comparisons are generally always controversial, but we still present Table 5, which compares the features of various similar home networking technologies. In our opinion, several of these the technologies included are more complementary than competitive.

The HomeRF and HomePNA [3] technologies are very synergistic for home electronics manufacturers because they share much networking infrastructure even though the physical media are quite different. In both cases, there are simplified “single-cell” networks where voice to the PSTN and data to the Internet can be combined simultaneously. For in-room (or in-car) point-point or point-multipoint connectivity, the proposed Bluetooth protocol [27] and the industry-proven IrDA standards (with over 60 million units shipped) [8] are most appropriate. Amongst these two technologies, Bluetooth offers greater physical convenience in its usage model as it is not line of sight and can pass through minor obstructions, but the IrDA standards are very hard to beat by any radio technology in terms of their data rate, cost, or physical size.

The few competitive wireless networking options that are either available today or will be available in the near future have been designed primarily for the enterprise market. For example, the IEEE 802.11 [14] and HIPERLAN [23] are effectively wireless Ethernet technologies developed for the enterprise network. Both support multiple cell handoffs

Table 5
Comparison of SWAP-CA With Other Connectivity Options

Properties	HomeRF [2]	Bluetooth [27]	IEEE 802.11 FH [14]	HomePNA [3]	IrDA [8]	HIPERLAN-I [23]
Operational Spectrum	2.404 – 2.478 GHz	2.402 – 2.480 GHz	2.40 – 2.4835 GHz	Phone Line	Infrared, 850 nm	51.15-5.3 GHz (and 17.1 – 17.3 GHz)
Physical layer	FHSS, 50 hops/sec, 2-FSK & 4-FSK	FHSS, 1600 hops/sec, GFSK	FHSS, 2-FSK, 4-FSK, ...	?	Optical	Differential GMSK
Channel Access	CSMA/CA + TDMA	Master-slave, Polling	CSMA/CA with RTS/CTS	?	Polling	EY-NPMA
Peak Raw Data Rate	0.8 and 1.6 Mbps	0.721 Mbps	1 and 2 Mbps	10 Mbps	16 Mbps (VFIR)	23.5 Mbps
Range	< 50 m	< 10 m	< 50 m	Phone jack	~ 1 m	< 30 m
Standby & Peak current	< 1 mA & ~ 300 mA	< 1 mA & ~ 60 mA	~ 10 mA & ~ 400 mA	?	< 10 uA, & ~300 mA	> 2A
Data Traffic	via TCP/IP	via PPP	via TCP/IP	via TCP/IP	via PPP	via TCP/IP
Voice Traffic	via IP & PSTN	via IP & Cellular	via IP	via IP & PSTN	via IP	via IP
Error Robustness	CRC / ARQ Type 1	1/3 rate FEC,, 2/3 rate FEC and ARQ Type 1	CRC / ARQ Type II	?	?	BCH (31,26)
Mobility support	- NA -	- NA -	- Yes -	- NA -	- NA -	- Not specified -
Energy Conservation	Directory based	Yes	Directory based	- No -	- No -	Directory based
Guaranteed latency	< 20 msec. for voice		None	None	- not specified -	< 10 msec for voice
Speech Coding	32 Kbps ADPCM	64 Kbps with CVSD/logPC M	- NA -	-not specified-	- not specified -	32 Kbps ADPCM
Security	56-bit Shared-key encryption	Stream cipher algorithm	64-bit shared key encryption	?	?	?
Communications Topology	Peer-to-peer, MS-to-BS	Peer-to-Peer, Master-slave	Peer-to-peer, MS-to-BS	Broadcast	Master-slave	Peer-to-peer, multi-hop, MS-to-BS
Price Point (estimate)	Medium	Medium	Medium/High	Medium /Low	Low	High

and roaming so entire campuses may be covered. These technologies permit users to trade data-rate with cost and power consumption. Note that while HIPERLAN is legal in Europe only, similar technologies are likely to develop in the U.S., where the unlicensed national information infrastructure (U-NII bands) located at 5.15–5.35 GHz and 5.725–5.825 GHz has recently been created [26].

Note that in Table 2, the standby current refers to the average current draw for the transceiver portion of portable devices while retaining full network availability for the given technology.

VIII. FUTURE OF HOMERF

The HomeRF organization is discussing a variety of future derivatives from the initial SWAP-CA specification. One possible derivative is simply to increase the data rate within the existing the 2.4-GHz band while retaining full

backward compatibility with the initial specification. The group is presently considering options in this regard that would scale SWAP-CA to 10 Mb/s in the 2.4-GHz band. In addition, HomeRF is also developing two major new market requirements documents. The first is SWAP-MM (for multimedia), which is looking at true video applications within the home enabled by wireless networking. This work will likely proceed to a formal technical proposal for the 5-GHz band. It is unclear at this point whether the SWAP-MM specification can ever be as near global as the 2.4-GHz SWAP-CA. As with the SWAP-CA case, achieving consumer price points with a SWAP-MM solution will be critical. The other direction HomeRF is considering is an ultra-low-cost version called SWAP-lite that could be developed to be interoperable with future SWAP-CA devices while achieving much lower price and power consumption points. Keyboards, mice, joysticks, remote controls, toys, etc., are the products that might use SWAP-lite. Clearly, such

a system overlaps in capability with infrared technology, which sets a tough measure to compete with in terms of price/performance.

REFERENCES

- [1] The HomeRF Working Group, , <http://www.homerf.org>.
- [2] HomeRF Technical Committee, "Shared wireless access protocol, cordless access (SWAP-CA) specification,," Revision 1.0 (January 1999); Revision 1.1 (May 1999); Revision 1.2 (October 1999).
- [3] Home Phoneline Networking Alliance (HomePNA), , <http://www.homepna.org/>.
- [4] American National Standards Institute, ANSI/IEEE, . New York, 1985, Std. 802.3-1985.
- [5] "Carrier sense multiple access with collision detection,," IEEE, New York, IEEE 802.3, 1985.
- [6] Universal Serial Bus, <http://www.usb.org/>.
- [7] IEEE 1394—High Performance Serial Bus Bridges Working Group, <http://www.ieee.org/groups/1394/1>.
- [8] Infrared Data Association, <http://www.irda.org/>.
- [9] D. Radford, "Spread spectrum data leap through AC power wiring,," *IEEE Spectrum*, pp. 49–55, June 1994.
- [10] "DECT common interface, Part 1: Overview,," ETS 300 175-1, 2nd ed., Sept. 1996.
- [11] "DECT common interface, Part 3: Medium access control layer,," ETS 300 175-3, 2nd ed., Sept. 1996.
- [12] "DECT common interface, Part 7: Security features,," ETS 300 175-7, 2nd ed., Sept. 1996.
- [13] "CCITT recommendation G.726 40, 32, 24, 16 kbit/s adaptive differential pulse code modulation (ADPCM),," Geneva, Switzerland, 1990.
- [14] *IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification*, IEEE Std. 802-11.1997, June 26, 1997.
- [15] "OpenAir™ specification, wireless LAN interoperability forum,," <http://www.wlif.org>, 1998.
- [16] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression,," *IEEE Trans. Inform. Theory*, vol. 23, no. 3, pp. 337–343, 1977.
- [17] D. A. Huffman, "A method for the construction of minimum redundancy codes,," *Proc. Institute of Radio Engineers*, vol. 40, no. 9, pp. 1098–1101, Sept. 1952.
- [18] P. Deutsch, "DEFLATE compressed data format specification version 1.3,," in *Internet Engineering Task Force: RFC 1951*, May 1996.
- [19] S. W. Richard, *TCP/IP Illustrated Volume 1*. Reading, MA: Addison-Wesley, 1994.
- [20] R. A. Powers, "Batteries for low power electronics,," *Proc. IEEE*, vol. 83, pp. 687–693, Apr. 1995.
- [21] E. P. Harris, S. W. Depp, E. Pence, S. Kirkpatrick, A. Sri-Jayntha, and R. R. Troutman, "Technology directions for portable computers,," *Proc. IEEE*, vol. 83, pp. 636–658, Apr. 1995.
- [22] H. Taub and D. L. Schilling, *Principles of Communication Systems*, 2nd ed. New York: McGraw-Hill, 1986.
- [23] "Radio equipment and systems (RES): High performance radio local area network (HIPERLAN); function specification,," ETSI RES 10, Jan. 1995.
- [24] Microsoft Corporation, , "Microsoft Windows NT 5.0/WDM 1.1 Driver Development Kit,," 1998.
- [25] The ATM Forum, <http://www.atmforum.org>.
- [26] FCC, , "Amendment of the commission's rules to provide for operation of unlicensed NII devices in the 5 GHz frequency range,," ET Docket No. 96-102, Jan. 9, 1997.
- [27] J. Haartsen and S. Mattisson, "BLUETOOTH: A new radio interface for ubiquitous connectivity,," *Proc. IEEE*, vol. 88, Oct. 2000.



Jim Lansford received the B.S. degree in electrical engineering, with highest honors, from Auburn University, Auburn, AL, in 1980, the M.S. degree in electrical engineering from the Georgia Institute of Technology, Atlanta, in 1982, and the Ph.D. degree in electrical engineering from Oklahoma State University, Stillwater, in 1988.

He is currently Vice President of Business Development at Mobilian Corporation, Hillsboro, OR, where he is promoting Mobilian's multistandard "True Connectivity" technology. He was previously a Wireless System Architect with Intel Corporation, Santa Clara, CA, where he was responsible for the research and development of wireless consumer technologies for a variety of computing, home control, and entertainment applications. In this capacity, he served as the industrial liaison from Intel to the Berkeley Wireless Research Center, Berkeley, CA. He was also Chairman of the Technical Committee for the HomeRF Industry Working Group, a wireless technology industry consortium of more than 80 companies. Prior to joining Intel, he co-founded a company called Momentum Microsystems, which developed WLAN products. He was with the Georgia Tech Research Institute as a principal investigator in various radar and communications signal processing studies, which included the development of advanced digital signal processing hardware. Earlier, he served as a group leader at Harris Corporation, where he supervised a team of engineers in developing parallel signal processing algorithms and hardware for intelligence applications. He taught electrical engineering at Oklahoma State University from 1985 to 1987, and served as an Assistant Professor at the University of Colorado, Colorado Springs, from 1990 to 1995. He has authored more than 20 articles in the fields of stochastic processes, encoding, digital signal processing, and wireless communications.



Paramvir Bahl (Senior Member, IEEE) received the Ph.D. degree in computer systems engineering from the University of Massachusetts, Amherst.

He is currently with Microsoft Corporation, Redmond, WA, where he heads the wireless networking research efforts, investigating problems related to wireless Internet access, location-determination systems and services, multihop multiservice *ad-hoc* sensor networks, and real-time audio-visual wireless communications. Previously, he was with Digital Equipment Corporation, where he initiated, led, and delivered several seminal multimedia projects, including the industry's first hardware and software implementations of audio/video compression and rendering algorithms. He is the Co-founder and Vice Chairman of ACM Special Interest Group in Mobility (SIGMOBILE). He is the Founder and Editor-in-Chief of the *ACM Mobile Computing and Communications Review* and serves on the editorial boards of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS and the *ACM Journal on Wireless Networking*. He has served as a guest editor for several IEEE and ACM journals on topics related to mobile multimedia communications and on several NSF panels on networking research. He served as the General Vice Chairman of ACM MobiCom '99 and has participated in the Technical Program Committees of more than 20 international conferences and workshops. He has lectured extensively and is active in standards bodies, including the HomeRF WG and the Bluetooth SIG, where he co-chairs the working group on location positioning. He has authored more than two dozen scientific papers and 26 patent applications in the areas of wireless communications, home networking, digital signal processing, and computer communications.

Dr. Bahl is a Member of ACM, and a Past President of the electrical engineering honor society Eta Kappa Nu, Zeta Pi chapter.