HomeRF™ and SWAP: Wireless Networking for the Connected Home

Kevin J. Negus[a]
kevin_negus@hp.com

John Waters[b]
john_waters@hp.com

Jean Tourrilhes[b]
jean_tourrilhes@hp.com

Chris Romans[b]
chris_romans@hp.com

Jim Lansford[c]
jim.lansford@intel.com

Stephen Hui[d]
shui@microsoft.com

[a] Hewlett-Packard Company, Newark, CA, USA
[b] Hewlett-Packard Company, Bristol, UK
[c] Intel Corporation, Portland, OR, USA
[d] Microsoft Corporation, Redmond, WA, USA

The Shared Wireless Aaccess Protocol (SWAP) specification for wireless voice and data networking within the home will enable a new class of mobile consumer devices that draw from the power and content of the Internet and the Home PC. If cable modems and xDSL represent the "last mile" access to the home, then HomeRFTM 's mission with SWAP could be called the "very last 150 feet" within and around the home. HomeRFTM has the broad backing of the major corporate stakeholders for networking within the home and is optimized specifically for the cost/performance point needed for consumers. The technology leverages the existing PC-industry infrastructure around the Internet, TCP/IP and Ethernet and adds a standard way to connect to the Public Switched Telephone Network (PSTN) for voice telephony. First products should appear in late 1999 and future versions with enhanced features and/or higher data rates should follow by 1-2 years.

I. Introduction

Two major factors are presenting a real opportunity at last for data networking within the home. The first is the explosive growth and usage of the Internet. The Internet clearly has the potential to revolutionize the delivery of information and entertainment to the home. The second factor is the emergence of sub-$1000 powerful home PCs. With these low-cost devices the barrier to getting on the Internet and discovering the utility of the PC is low enough to reach the vast majority of middle-income households.

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, TV, videos, FM radio, 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 home PC industry formed the Home RF Working Group (or "HomeRF™") in early 1997. The key goal of the group was and remains to enable interoperable wireless voice and data networking within the home at consumer price points. HomeRF™ started by pooling market research from the member companies to produce a Market Requirements Document. This document guided the technical proposals within the group and with tremendous cooperation from major stakeholders in the RF communications industry and the nascent wireless LAN community, the Shared Wireless Access Protocol (or "SWAP") was created. HomeRF™ is now in the process of bringing the SWAP specification to its final released form. This is proceeding quickly due to the native support within SWAP for TCP/IP networking and Internet access, and for voice telephony via PSTN and voice-over-IP. SWAP achieved support for these important network stacks easily by re-using major sections of proven RF protocols and then simplifying them where appropriate for home usage.

Today the HomeRF™ organization consists of approximately 50 members representing the bulk of the PC, telecommunications and consumer electronics industries. General information on the organization is available at *www.homerf.org*. The specification described in this paper started at a Rev 0.1 from a proposal made in late 1997. As of this writing, the specification has been revised considerably to enhance performance and lower cost and is now at the 0.7 level. The final specification will reach the Rev 0.9 level by the end of 1998 and no further changes will be considered after that point. HomeRF™ will formally adopt the specification as Rev 1.0 in Q2 99 after complete working silicon is demonstrated with full air compliance to the Rev 0.9 specification.

II. Vision and Applications

HomeRF™ sees SWAP as one of several connectivity options for the home of the future. The relationship of SWAP with other connection options is shown in Figure 1. In this scenario, the main Home PC is linked somehow to an Internet gateway that might be a 56K, xDSL, or cable modem. This

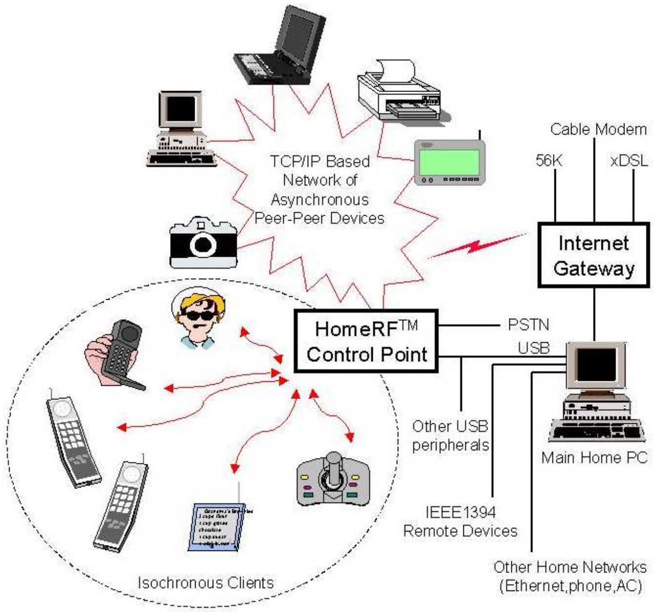

Figure 1: SWAP vision for Home Networking

link may be a simple cable, a wired network connection or even a SWAP network connection. This main Home PC would likely have 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, USB would be the bus of choice for many peripherals that do not need to be mobile or remote from the Home PC. For video applications such as connecting camcorders, IEEE1394 is the expected choice and there are no viable RF alternatives at consumer price points at this time (though Hewlett-Packard's proposed Very Fast Infrared (VFIR) protocol at 16 Mb/s is a reasonable "no-cable" choice). HomeRF™ also expects that other networking choices will also be viable for sharing resources amongst multiple PCs as shown in Figure 1. These options include conventional 10/100 BaseT Ethernet, home phone line Ethernet and AC power line networking. The last option is particularly well suited for many Home Automation scenarios where very low data rates are acceptable.

The SWAP networking vision is also apparent in Figure 1. The HomeRF™ technology supports both isochronous clients that are slaves to the main Home PC and an asynchronous network of peer devices that is effectively a wireless Ethernet. In most cases the system starts with a *Control Point* usually connected to the main Home PC via USB. This Control Point (CP) is not absolutely necessary for devices in the asynchronous network of Figure 1, but even in that case it offers some interesting power saving options for ultra-portable devices as will be described in the MAC overview section. The isochronous clients such as cordless telephones, wireless headsets, or remote I/O devices to the home PC (a consumer Personal Information Manager or PIM) are always bound to the Control Point which assigns them guaranteed bandwidth for bounded latency communication. The asynchronous peers can also communicate to the main Home PC as with any other peer device. Note that data transfer on the asynchronous network between any two peers is directly to each other as opposed to routed through the Control Point. Consider now three major applications for the HomeRF™ technology.

First example is PC-enhanced cordless telephony. Today there are no standards-based digital cordless telephones for consumer use in the USA where interoperability of multiple vendors is enabled. HomeRF™ defines a new standard for interoperable digital cordless telephones both in the USA and globally. Furthermore, the SWAP 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, caller ID information could be sent to a PC application to look up the caller's name and then route the call to an individual handset (rather than number) and display the caller's name on any given handset. For outbound calls, the PC could interpret a spoken destination name (i.e. "Call Mom") through voice recognition and then based on date/time determine the likely number for the person and route the call using the lowest cost approach (which might be 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 3 quarts of milk to my shopping lists") or control home automation features (i.e. "Turn the temperature up 3 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 multi-million unit volume today.

A second interesting 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 radio network connection. Such a device could be either an extension of the Home PC (like an X-terminal) 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 back-up storage device or an Internet connection. Another possibility for this is multi-player gaming. Clearly these resource-sharing applications have received considerable attention from other home wiring-based alternatives to networking. It is important to note that the market for HomeRF™ is not strictly multi-PC homes. Any home with a modern Home PC or an Internet gateway is a candidate for compelling mobile devices enabled by the SWAP specification.

III. Network Topology

The SWAP architecture is a unique combination of a managed network that provides isochronous services such as interactive voice, and an ad-hoc peer to peer network that provides traditional data networking. The protocol has been optimized to provide the kinds of services that are most needed from untethered devices in the home. Three kinds of devices can be in a SWAP network:

1. a Connection Point (CP), which acts as the gateway between the personal computer, the PSTN, and SWAP-compatible devices,
2. Voice devices (Isochronous data devices, also called I-nodes), and
3. Asynchronous data devices also called A-nodes.

The Control Point is usually connected to the main home PC, typically via USB. It may also have a connection to the PSTN. It is capable of performing data transfers to and from other data devices using an asynchronous, contention-based protocol. The control point manages the network to provide priority access to the radio medium for isochronous devices.

Thus, the SWAP protocol is a hybrid in several ways; it is client- server between the Control Point and voice devices, but is peer to peer between data devices. The interactive voice transactions are circuit switched, time division multiple access but the asynchronous transactions are packet switched, carrier sense multiple access. It is precisely this richness that gives SWAP the capability to be broadly used 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. There are several usage scenarios that the SWAP protocol was designed to support; these focus on the exchange of voice and data by portable devices within the home environment. Typical usage that the network was designed to support include:

1) A busy family

- Three children play interactive PC games with each other from their rooms. They can talk to each through their PCs.
- Mom is in the kitchen listening to her phone messages recorded by the PC answering machine application.
- Dad is in the garage completing an oil change. Using a handheld pad, he enters the data into his car maintenance log which he maintains on his desktop PC.

2) A montage of voice applications

- Mom is using her personal handset to record a message for the kids to listen to when they return from school
- Dad is asking for stock quotes from Internet and getting vocalized response through TTS (text to speech) engine
- Dad checks the temperature in the garage and turns on a space heater via voice commands
- Uncle Ed listens to a soccer match broadcast via audio over the web to his wireless headset on the front porch

3) A montage of display applications

- Mom pulls up a recipe from the PC and adds oregano to the shopping list
- Dad updates the family financial portfolio

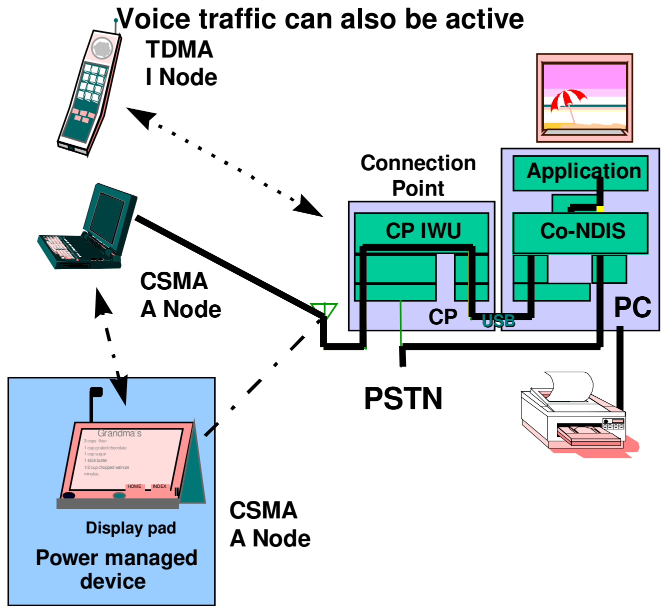

Figure 2: SWAP network topology flexibility

- Junior plays Tetris
- Daughter reads the latest on-line issue of Teen magazine

Figure 2 shows an example of a typical SWAP network consisting of two A-nodes, one I-node, and a Connection Point. 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 has a unique ability among networking protocols to mix intense, high demand packet traffic with infrequent command and control traffic and with high quality voice traffic. The personal computer is an integral part of the SWAP system although peer-peer data networking is available even when the PC is inoperative.

IV. Software Architecture for the PC

SWAP A-nodes will interface with Microsoft Windows operating systems through the existing Network Driver Interface Specification (NDIS) [6]. NDIS is a library of driver functions common to all networking device drivers. With NDIS, network adapter manufacturers need only write a miniport driver which provides functionality specific to their hardware. They can then immediately use their devices with higher-level protocols knowledgeable about their media type. For example, the standard IP protocol stack, which is Ethernet knowledgeable, will work over any new device that declares itself as and behaves like a member of the Ethernet media type.

SWAP A-node miniport drivers will declare themselves as members of the Ethernet media type. Because of this, SWAP A-nodes will be indistinguishable from regular Ethernet adapters to higher level protocols, and any application that works over Ethernet will be immediately enabled over SWAP as well.

Call control of SWAP I-node adapters and CPs will be done through the Telephone Aapplication Programming Interface

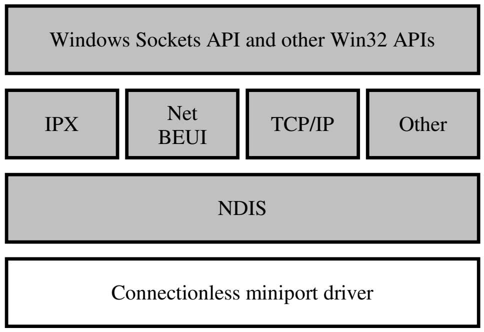

Figure 3: SWAP A-node driver architecture

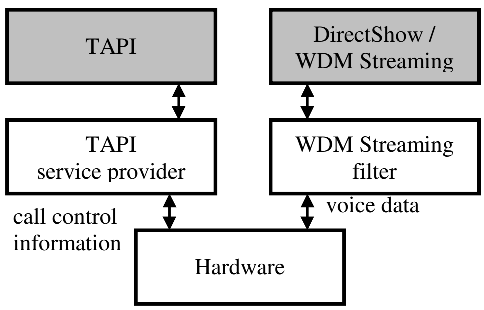

Figure 4: Windows 98 streaming and call control architecture

(TAPI). TAPI is a simple, generic set of objects, interfaces, and methods for establishing connections between devices. TAPI communicates with hardware via TAPI service providers. In Windows 98, Windows 95, and Windows NT 4.0, SWAP I-node manufacturers will provide TAPI service providers that interface directly with their hardware. Call control of I-nodes and CPs in Windows NT 5.0 will also be provided by TAPI; the details of the interface are currently being defined.

Some device designers may wish to stream voice conversations between the SWAP adapter and another adapter within the PC in real time. An example scenario would be that of a voice conversation between the SWAP adapter (i.e. a user with a SWAP handset communicating with the SWAP adapter) and another adapter in the PC (such as a modem attached to a phone line or a sound card attached to speakers and a microphone).

In Windows NT 5.0 and Windows 98, the Windows Driver Model (WDM) Streaming / DirectShow architecture enables real-time data streaming. A DirectShow filter graph is plumbed from the data source (in this case, the SWAP adapter) to the data sink (the modem or sound card). In Windows 98, I-node manufacturers who wish to enable voice streaming can write WDM Steaming filters to send data from their hardware into the streaming architecture. Windows NT 5.0 will likely use a NDIS interface for I-nodes; Microsoft provides a filter with the operating system (the Raw Channel Access filter) that links NDIS and the streaming architecture. In contrast, Windows 95 and Windows NT 4.0 do not include the WDM Streaming / DirectShow architecture, so manufacturers who wish to implement voice streaming in these operating systems will have to provide their own streaming implementations.

Table IV summarizes SWAP driver architectures for recent

Operating system	**A-node**	**I-node / CP**
Windows NT 5.0	connectionless miniport	currently being defined
Windows 98	connectionless miniport	manufacturer-supplied TAPI service provider and WDM Streaming filter
Windows NT 4.0	connectionless miniport	manufacturer-supplied TAPI service provider and streaming architecture
Windows 95	connectionless miniport	manufacturer-supplied TAPI service provider and streaming architecture

Table 1: SWAP driver architectures for recent Microsoft operating systems

Microsoft operating systems.

V. Medium Access Control Protocol - Overview

The SWAP Medium Access Control (MAC) protocol has been optimized for the home environment. It is designed to carry both voice and data traffic while interoperating with the Public Switched Telephone Network (PSTN) using a subset of the Digital European Cordless Telephone (DECT) standard [1, 2]. The MAC is designed for use with a frequency hopping radio and includes a Time Division Multiple Access (TDMA) service to support the delivery of isochronous data (e.g. interactive voice), and a Collision Sense Multiple Access with Collision Avoidance (CSMA/CA) service, derived from the IEEE 802.11 standard [4], to support the delivery of asynchronous data. The SWAP MAC provides the following features:

- Support for voice and data by using both TDMA and CSMA/CA access mechanisms
- Support for 6 high quality voice connections with 32kb/s ADPCM [3]
- High data throughput of 2 Mb/s
- Data security - None/Medium/High levels of encryption
- Power management for both Isochronous and Asynchronous nodes
- 24-bit Network ID

The MAC protocol uses a superframe as illustrated in Figure 5, which incorporates two contention-free periods (CFPs) and a contention period. The start of the superframe is the point at which a station begins to hop to a new channel and ends immediately before the station starts to hop to the next channel. The duration of the superframe is fixed and is the same as the dwell or hop period. The access mechanism used during each CFP is TDMA, whilst the access mechanism used during the contention period is CSMA/CA.

Each of the contention free periods is 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 Connection Point to a node (downlink) and the second is used to transmit voice data from a node to the Connection Point (uplink). In a managed network a Beacon is transmitted immediately after the hop. This Beacon is used to maintain network synchronization, control the format of the superframe and manage when each node should transmit and receive data.

The CFP at the end of the superframe (CFP2) is used for the initial transmission of the voice data, whilst CFP at the start of the superframe (CFP1) is used for the optional retransmission of any data which was not received or was incorrectly received in the previous dwell. The dwell period is fixed at 20ms to provide acceptable performance with respect to latency. The length of the dwell period also means that each voice data message contains 20 ms of ADPCM data (640 bits), equivalent to an extended DECT B-field and 56 bits of control data, equivalent to the DECT A-field plus some additional addressing information . With a 20 ms superframe the MAC is able to provide upto 6 voice connections, and a CFP that is long enough to enable up to 4 retransmissions.

The CFP2 in which initial transmission occurs and the CFP1 in which retransmission occurs are separated by a frequency hop, giving frequency and time diversity, which is particularly important given the potentially noisy environment in which the protocol operates. At the end of CFP1 in the superframe there is a space reserved for a service slot that is used by voice nodes to request connections from the Connection Point. Each voice data packet transmitted by a node includes in the packet header a piggyback acknowledgement of the last voice data message received by the node. i.e. in the uplink packet, the voice node acknowledges the downlink packet sent by the Connection Point. This system allows the Connection Point to determine prior to a hop which voice data transmissions were lost, to determine the retransmissions required, and advertise these retransmissions in the Beacon at the start of the next superframe; each voice data packet can only be retransmitted once.

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 and microwave oven interference. Peak effective user throughputs of over 1.2 Mb/s are possible under lightly loaded conditions in the 2 Mb/s 4FSK mode. Furthermore, data transfer rates exceeding 300 kb/s for distant nodes or 600 kb/s for nearby nodes can occur even with 4 voice calls active simultaneously.

The CSMA/CA mechanism is similar to Ethernet (802.3), enabling easy integration with an existing TCP/IP protocol stack within a host platform; the main difference with Ethernet is the slotted contention mechanism and the addition of MAC level acknowledgement of unicast packets. Figure VIII illustrates how the medium is accessed during the contention period.

The CSMA/CA access procedure is designed to provide fair access to the medium to all nodes by using a contention window and backoff counter as shown in Figure VIII. 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 it's 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 also 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 or transmission failure the size of the collision window is increased from an initial value of 16 exponentially up to a maximum 128 to avoid congestion.

The primary function of the beacon signal is to enable all nodes to synchronize to the hopping pattern of the network. The beacon transmitted by the Connection Point is also used to manage the network during the contention free periods. The Connection Point Beacon (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. The CPB can also provide Power Management for Isochronous and Asynchronous nodes to maximize the battery life of portable devices.

For the case when no Connection Point is present, the data nodes can create an ad-hoc network in which control of the network is distributed between all the nodes.

V.A. Power Management

The procedure for power management of isochronous nodes is straightforward. In this process, during an active connection (e.g a voice call) the isochronous nodes power-on, initially only for the duration of the CPB, to receive slot assignment information. They then power-down until their assigned slots are due. When not in an active connection state, isochronous nodes need only power-up every N dwells, where N is chosen by the system designer according to the application being supported, and as a compromise between power-saving and speed of response to a new connection.

The Connection Point provides power management services for asynchronous nodes. Figure VIII illustrates the process for sending broadcast messages to power-saving asynchronous nodes (PS-nodes) which is as follows: The Connection Point maintains a countdown to the next dwell when PS-nodes should wake-up which is broadcast in the CPB. In step 1, PS-node 3 powers-up and receives the 'dwells-to- broadcast' counter. The maximum value of the 'dwells-to-broadcast' counter is system design parameter that allows the designer to tradeoff latency, CP buffer size and broadcast reliability against battery-life. Node 1 transmits a broadcast message, which is received and stored by the CP. Node 2 transmits a broadcast message, which is received and stored by the CP. The PS-node wakes up when its dwell counter decrements to 0 (Step 4). The CP transmits the buffered broadcast messages during the dwell (Step 5). The PS-node receives the broadcast messages (Step 6) and then goes back to standby mode.

Figure VIII illustrates the process for sending unicast messages to power-saving asynchronous nodes. The PS-node is controlled by a 'wake-up' flag transmitted in the CPB. This flag is activated by a request from the sending node. Thus in this example the PS-node wakes-up and listens for a wake-up flag in the CPB (Step 1). Since the flag is not set, it powers down. In step 2, node 1 transmits a request to the CP to wake-up the PS-node. In step 3 the CP asserts the wake-up flag in the CPB. Some time later (dependant on the designer's tradeoff of power-saving versus latency) the PS-node wakes-up, and receives the wake-up flag (step 4). The PS-node and node 1 transfer data using the normal CSMA/CA access method (step 5). The PS-node powers-down a time after the final message exchange (step 6).

VI. PHY and Components

The PHY specification for SWAP was largely adapted from the IEEE 802.11FH standard with significant modifications to reduce cost while maintaining more than adequate performance for home usage scenarios. The PHY specification takes tremendous advantage of the very nature of frequency hopping in the 2.4 GHz band. That is, significant interference sources are easier to hop away from or momentarily defer to rather than attempt to filter them out. Some of the key SWAP physical layer specifications include:

- Transmit power – +16 to +20 dBm (or nominally just under 100 mW)
- Receiver sensitivity in 2FSK or 1 Mb/s mode – -76 dBm (or 4 dB less stringent than IEEE802.11FH)
- Optional low transmit power mode – 0 to +4 dBm (for portable devices with limited peak current capability)
- Hopping time – 300 us (to allow conventional synthesizers to be used)
- Transceiver turnaround time – 25 us (toughest SWAP PHY spec though easier than 802.11)
- Adjacent and alternate channel filtering – no requirement (substantial relaxation from 802.11)

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). As with IEEE 802.11FH, a 4FSK or 2 Mb/s mode is available. However, in SWAP the 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 inter-symbol interference due to filter group delay variations in the passband. And third, the SWAP packet headers for 4FSK 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 2 Mb/s data rate is really available with SWAP and adds virtually no cost to the 1 Mb/s solution.

Although the hopping time is easy to meet, the transceiver turnaround time creates challenge for many conventional RF transceiver architectures and components. This low transceiver turnaround limit is essential for SWAP 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 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 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 Figure 9. For many of the digital devices envisioned by HomeRF™, the digital MAC baseband portion of the component solution can be integrated into a large ASIC already in the device. At $\tilde{3}$0K gates for the SWAP data core, this is extremely low cost in the sub-0.25 um 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 IC as shown because of the specific technical requirements on filtering and modulation chosen by HomeRF™. 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% that 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.

With such high levels of integration and optimized front-end, the RF modem section overall cost in multi-million unit volume should be well below $10 (similar to the situation today with DECT) 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 $100's 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 HomeRF™ should not be a serious barrier.

VII. Positioning with Other Technologies

Table 2 gives an overview of the HomeRF™ SWAP attributes along with comparative estimates for other "similar" technologies. Comparisons are always controversial. In this case, these technologies are really complementary more than competitive. HIPERLAN [7] and IEEE 802.11 [4] are effectively wireless Ethernet technologies for the enterprise network. Both support multiple "cell" hand-offs and roaming for coverage of entire campuses and together permit users or Information Technology (IT) managers to trade data rate with cost and power consumption. Note, currently HIPERLAN is legal in Europe only, however similar technologies are possible in the recently created National Information Infrastructure (NII) bands at 5 GHz in the US. The HomeRF™ and Home Phoneline Networking Alliance (HomePNA) [8] technologies are very synergistic for home electronics manufacturers because they share so much networking infrastructure in common even though the physical media are quite different. In both cases, these are simplified "single-cell" networks where voice to the PSTN and data to the Internet can be combined simultaneously. Finally, for in-room (or in-car) point-point or point-multi-point connectivity, the proposed Bluetooth protocol [9] and the industry-proven IrDA standards (over 60 million units shipped) are most appropriate. Amongst these two technologies, Bluetooth offers far greater physical convenience in its usage model as it is not line of sight and can pass through minor obstructions. The IrDA standards are very hard to beat by any radio technology in terms of their data rate, cost or physical size. 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 HomeRF™ Derivatives

The HomeRF™ organization is already discussing a variety of future derivatives for the initial SWAP 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 to 6 or even 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 multi-media) 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 case. As with the SWAP 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 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] ETS 300 175-1. DECT Common Interface, Part 1: Overview. Second Edition, September 1996.

[2] ETS 300 175-3. DECT Common Interface, Part 3: Medium Access Control Layer. Second Edition, September 1996.

[3] CCITT Recommendation G.726. 40, 32, 24, 16 kbit/s Adaptive Differential Pulse Code Modulation (ADPCM). Geneva 1990.

[4] IEEE Std. 802-11.1997. IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification. Approved 26 June 1997.

[5] C. Romans, J. Tourrilhes, "A Medium Access Protocol for Wireless LANs which supports Isochronous and Asynchronous Traffic," 9th International Symposium on Personal, Indoor and Mobile Radio Communications, September 8-11, 1998. Boston.

[6] Microsoft Corporation, Microsoft Windows NT 5.0/WDM 1.1 Driver Development Kit, 1998. *Please note that the Windows NT 5.0 Driver Development Kit is not itself part of the SWAP specification, and that this article is based upon the draft DDK; changes to the DDK may affect the SWAP PC architecture.*

[7] ETSI, "High Performance Radio Local Area Network (HIPERLAN)," Draft standard ETS 300 652, March 1996

[8] HomePNA, "Home Phoneline networking Alliance," http://www.homepna.org

[9] J. Haartsen, W. Allen, J. Inouye, O. J. Joeressen, and M. Naghshineh "Bluetooth: Vision, Goals, and Architecture," ACM Mobile Computing and Communications Review, vol. 2, no. 4 (this issue)

Biographies

Kevin J. Negus received the B.A.Sc., M.A.Sc, and Ph.D. degrees all in Mechanical Engineering from the University of Waterloo in Canada. Since 1988 he has worked at the Hewlett-Packard Company (formerly Avantek, Inc.) on silicon bipolar device modeling, RFIC design and RF system analysis for most digital cellular, digital cordless and wireless data systems in use today. Formerly the silicon and gallium arsenide design department manager for HP, he now holds the position of R&D Pathfinder responsible for new business development at HP for wireless data applications. Dr. Negus has published over 40 papers and holds 6 US patents.

John Waters received a BSc in Electrical and Electronic Engineering from the University of Bath in the UK. In 1976 he joined the GEC Hirst Research Centre where he conducted research on various microwave components and sub-systems for digital radio relay systems. In 1988 he joined Hewlett-Packard, and for the last 7 years has been working in HP Labs in the field of mobile communications. He has worked with a variety of communications technologies, managing projects in the areas of digital cellular and paging systems, cordless telephony and wireless LANs. His current interests include the development of wireless systems for the home, and for the mobile professional.

Jean Tourrilhes received an "Ingenieur" diploma from the Ecole Nationale Superieure des Telecomunications in Paris in 1994. Since 1996 he has worked on wireless LANs, radio technology and MAC protocols at Hewlett-Packard Laboratories Bristol, England. He is one of the main designers of the SWAP specification.

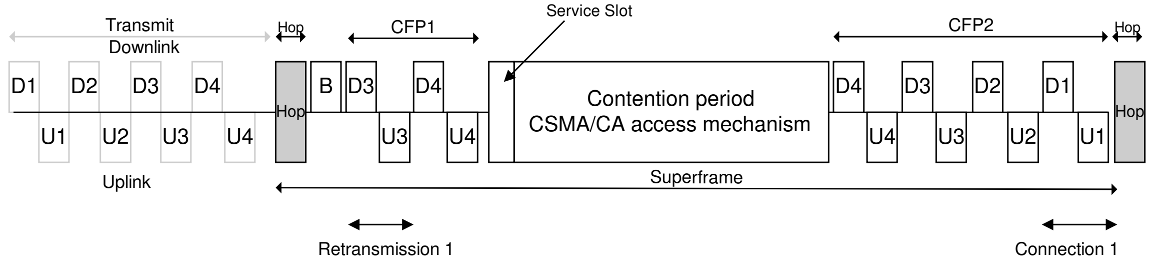

B - Beacon
Dn - Downlink slot
Un - Uplink Slot
CFP1 - Contains two slots per connection for data that requires retransmission
CFP2 - Contains two slots per connection, one for downlink data and the other for uplink data

Figure 5: SWAP frame description

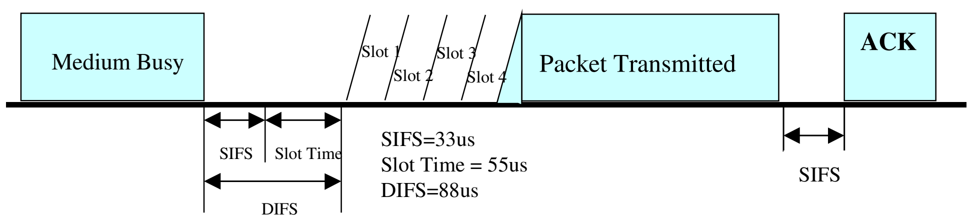

- *Sender selects a slot (backoff counter) and then decrements the counter whilst the medium is clear*
- *Medium must be free for a DIFS period before the backoff counter is decremented.*
- *Example shows transmission of a packet in slot 5.*

Figure 6: CSMA medium access procedure

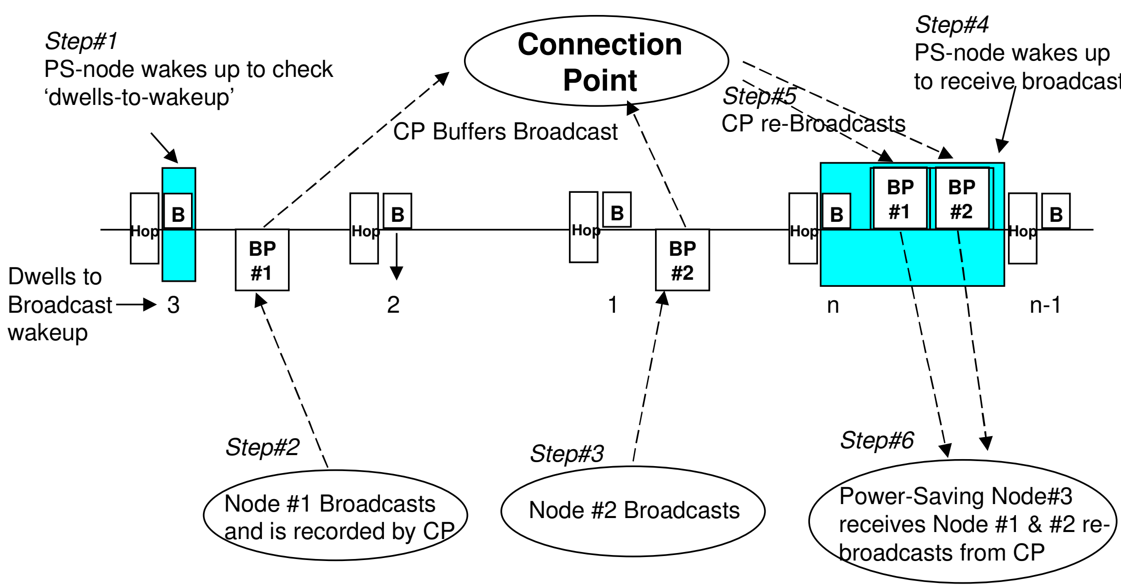

Figure 7: Power management of CSMA nodes for broadcast messages

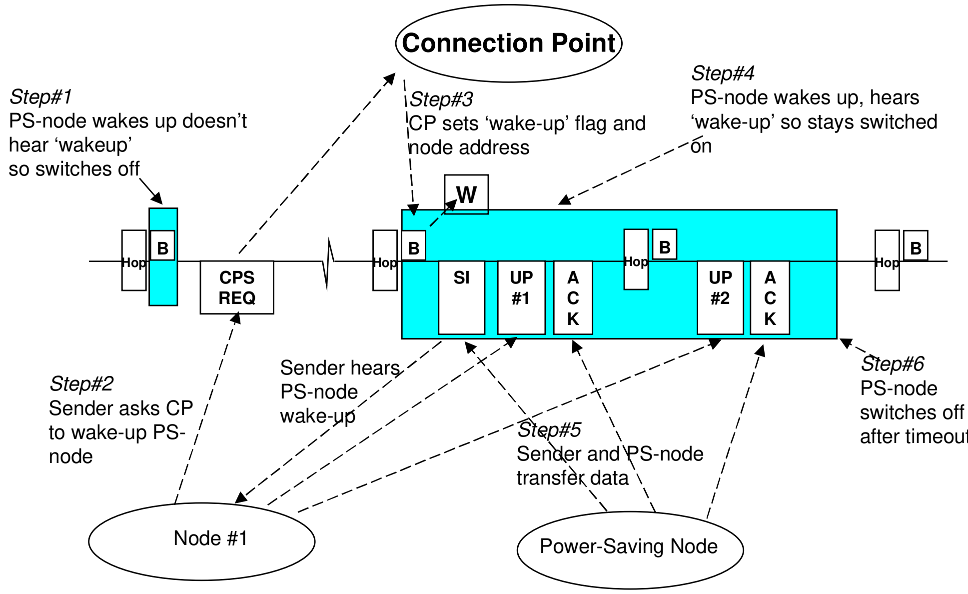

Figure 8: Power management of CSMA nodes for unicast messages

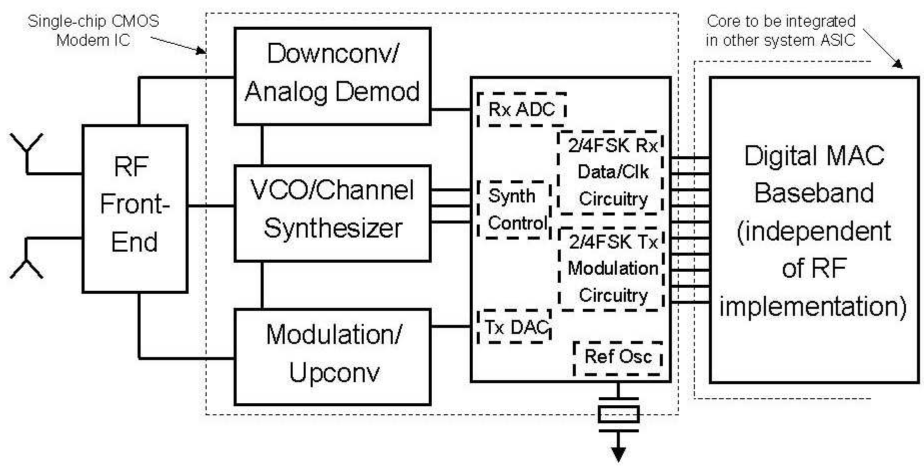

Figure 9: Partitioning of RF Modem and Digital MAC Component Sections

	Peak Data Rate	**Relative Cost**	**Data Network Support**	**Voice Network Support**	**Range in home**	**Standby & Peak Currents**
HiperLAN	23.5 Mb/s	High	TCP/IP	Via IP	$>$ 30 m	TBD, $>$2A
IEEE802.11FH	2 Mb/s	Medium/High	TCP/IP	Via IP	$>$ 50 m	$\approx$ 10 mA, $\approx$400 mA
HomeRFTM (SWAP)	2 Mb/s	Medium	TCP/IP	Via IP & PSTN	$>$ 50 m	$<$ 1 mA, $\approx$300 mA
HomePNA	1 Mb/s	Medium/Low	TCP/IP	Via IP & PSTN	All phone jacks	TBD
Bluetooth	1 Mb/s	Medium	Via PPP	Via IP & Cellular	$<$ 10 m	$<$ 1 mA, $\approx$60 mA
IrDA	16 Mb/s	Low	Via PPP	Via IP	$<$ 2 m &line of sight & aimed	$<$ 10 uA, $\approx$300 ma

Table 2: Comparison of SWAP with other connectivity options

Chris Romans is a Senior Member of Technical Staff in the Personal Systems Laboratory of Hewlett-Packard Laboratories Bristol, England, with a special interest in mobile communications. He has a degree in Physics and Computer Science and joined HP in 1992 where he has worked on cellular networks and wireless data networks including wireless LANs. Prior to joining HP he worked for several cellular network operators.

Jim Lansford is a Wireless Systems Architect in Intel's Communications Architecture Lab. He received his BSEE from Auburn University, MSEE from the Georgia Institute of Technology, and PhDEE from Oklahoma State University. Prior to joining Intel in 1996, Dr. Lansford was Chief Technical Officer of Momentum Microsystems in Colorado Springs, and was previously on the faculty of the University of Colorado at Colorado Springs. Dr. Lansford has published numerous papers in the field of communications signal processing and consulted for companies in the LEOSAT, cellular, PCS, and WLAN markets.

Stephen Hui is a Program Manager in the Windows Networking and Communications team of Microsoft's Personal and Business Systems Division. His responsibilities include collaborating with the wireless industry to build support for mobile communications into the Windows family of operating systems. Prior to joining Microsoft, he developed software for DECT PCS phones and ATM switches. He holds a BASc in Computer Engineering from the University of Waterloo in Canada.