Abstract

New short-range wireless communication technologies would enable environment-aware, mobile, personal area networks. These new technologies will serve as enablers for ubiquitous, low-cost, low-complexity, small-sized information appliances. These appliances will serve as interaction tools between humans and computer-driven services and applications existing in either the close or distant vicinity of humans. In this article the new application paradigms these new technologies will enable are explored. Furthermore, an experimental wireless personal access networking platform called WiSAP, developed to research these new technologies and applications paradigms, is presented. Finally, some of the experiences gained from WiSAP while designing a consumer-oriented portable wireless communication system suitable for wireless mobile personal access networks are also presented.

WiSAP: A Wireless Personal Access Network for Handheld Computing Devices

CHATSCHIK BISDIKIAN, PRAVIN BHAGWAT, BRIAN P. GAUCHER, FRANK. J. JANNIELLO, MAHMOUD NAGHSHINEH AND PUNEET PANDOH, IBM THOMAS J. WATSON RESEARCH CENTER

IBRAHIM KORPEOGLU, UNIVERSITY OF MARYLAND

Emerging technologies for short-range personal communication networks promise continuous end-user connectivity to a plethora of data services from portable and/or wearable computer-enhanced devices. Integrated low-cost, low-power wireless modules will become commodity items. It is expected that these modules will find their way into all sorts of electronic devices, including personal digital assistants (PDAs), laptops, cellular phones, information terminals, data communication equipment, printers, and home and office appliances, thus facilitating the development of a unified connectivity solution among them.

With such universal connectivity becoming available, printing, Web surfing, accessing online services, e-mail, and corporate application access from anywhere (in an office building, at an airport, or in a hotel conference room) will become common. Mobile workers will be able to connect to their home offices at will, and urgent notifications will find their recipient at any (appropriate) place and time.

In this article any device, such as those mentioned above, capable of information processing will be called an *information processing appliance* (IPA). A device will be classified as an IPA even if its primary function is not traditional data processing. For example, a cellular phone's primary function is to carry real-time voice conversations, but it is considered an IPA because of its capability to process and store data (e.g., phone books), as well as its potential to serve as a bridge to the Internet over a cellular network, in order to enhance its primary functionality.

Ubiquitous connectivity has the potential to foster a revolution of sorts, spearheaded by the new communications-based application paradigms it will enable. The paramount enabler for these new paradigms will be the fact that connectivity itself will no longer be an (end-user) issue. Communications will be practically hidden from the end user and, in many situations, will not even require conscious intervention of the user. According to these new applications paradigms, end-user IPAs will serve as interaction tools that will assist end users not only in contacting their businesses but in interacting with their surroundings in their everyday life (e.g., at home, at the supermarket, at the exhibit hall). In this article end-user IPAs such as PDAs, cellular phones, pagers, and smart-cards will be referred to as *personal computing and communication assistants* (PCCAs).

Wireless connectivity and seamless mobility are not new concepts. Wide-area cellular networks already enable connectivity to cellular service subscribers from (practically) everywhere. Moreover, the cornucopia of wireless technologies — satellites, wireless WANs and LANs, personal communications services (PCS), paging, infrared, and so on — and standards — Global System for Mobile Communications (GSM), Cellular Digital Packet Data (CDPD), IEEE 802.11 and variations, Infrared Data Association (IrDA), High-Performance LAN (HIPERLAN), wireless asynchronous transfer mode (ATM), and so on — are manifestations of the high industry interest in keeping end users on the move "connected" to the world around the clock. With the exception of the voice-oriented cellular telephony technologies, most data-centric wireless communication solutions have not lived up to their expectations for universal connectivity yet, due to cost concerns, power requirements, restricted bandwidth, device size and weight, lack of sufficient system integration, absence of compelling applications, and other factors.

The new breed of emerging wireless radio technologies are capable of providing connectivity over short distances (about 10 m, and possibly up to 100 m) at up to a few megabits per second in link speeds [1, 2]. These technologies take advantage of the Industrial, Scientific, Medical (ISM) unlicensed frequency bands, especially in the 2.4 GHz band and, possibly, in the 5.0 GHz band in the not too distant future. Since the ISM bands are unlicensed, RF-based wireless connectivity solutions can be freely developed[1] without the need to obtain

[1] *FCC part 15 rules for unlicensed radio transmissions impose transmission power limitations typically at about 100 mW (20 dBm). Also, while no modulation restrictions exist when the transmission power is below 1 mW (0 dBm), at higher power levels spread-spectrum-based modulation is required.*

 1070-9916/98/$10.00 © 1998 IEEE

exclusive usage rights (usually at a high cost) of these frequency bands from regulatory authorities. The latter usually implies a fee-based service provision, which, although sometimes necessary, imposes a rather unappealing feature for ubiquitous use of PCCAs.

Due to power constraints, ISM-based solutions need to be tuned and optimized for the small coverage they provide. Thus, applications based on ISM band connectivity are used for communicating between two or more devices in proximity to each other, for example, to support groupware applications on an ad hoc wireless LAN in a conference room, or to communicate between a PDA and a desktop computer. ISM-based networks could also be used as access networks to wireline networks. The latter case enables access to services available to wireline backbone networks, Internet access, access to corporate applications, resource sharing (print and file servers), and so on.

Due to the capability to allow individuals access to services available through a backbone network, the ISM-based networks of interest in this article will be referred to as *personal access networks* (PAcNets). PAcNets parallel traditional last-mile access networks, sometimes referred to as "the ramps to the information highway," such as public switched telephone network (PSTN) subscriber loops, cable TV (CaTV), and satellite distribution networks used to bring, say, the Internet to buildings. A PAcNet goes a step further by bringing the Internet to an individual.

Industry leaders have recognized the great potential of these new technologies and the new services they enable, so they are working frantically to develop the basic technology building blocks, specifications, and protocols that will facilitate deployment and wide acceptance of these technologies. The IEEE 801.11 standards group has formed a wireless personal area network (W-PAN) subgroup to study the requirements of the personal area network market and develop related communications standards [2]. The Bluetooth group is developing specifications for low-cost small-footprint RF-based connectivity solutions targeting businesses and mobile workers [1]. The HomeRF group is developing specifications for RF-based connectivity solutions targeting the home networking market [3]. When the next millennium arrives, some of these efforts are expected to bear fruit, and the first computing devices incorporating these communication technologies may be available in the market.

A viable end-to-end solution does not rely on the existence of any single technology enabler (e.g., the appropriate radio module, a proper battery) or even of several enablers in isolation. Plenty of work needs to be done selecting and combining all the appropriate enablers in a manner that is technically feasible and satisfies the market requirements of an end product. Unfortunately, such solutions do not exist in the market yet.

Our Wireless Services Access Platform (WiSAP) project aims to realize our vision for pervasive wireless connectivity of PCCAs via low-cost low-power integrated solutions. WiSAP is an umbrella project for developing end-user systems and solutions for realizing our vision by researching, selecting, and combining the proper technology enablers for this task. This vision and our description of WiSAP will be presented in the following sections of this article.

The rest of the article is organized as follows. In the following section we motivate and formulate our vision of universal connectivity. We then present the various connectivity usage models we envision. We go on to summarize WiSAP, an experimental wireless development platform for researching the creation of end-to-end solutions enabled by universal connectivity. We discuss some of the experiences gained while working with WiSAP, and conclude with a summary of our work.

A Personal Communications Vision

Market forces and trends have the power to render simple gadgets highly successful, or products with great intentions and aspirations irrelevant. Properly designed products that address the right end-user problem, at the right time, in the right form, and at the right price are more successful than products that are technologically superior but do not address the proper problem, are ahead of their times, or are not in the right form. A proper interpretation of the market forces is a necessary element in the development of highly successful and widely used products.

One very strong current market trend, as exemplified by its strong user interest and demand, is Internet access. Another is the need to stay continuously reachable, as is evident by the inroads that pagers and cellular phones have made in everyday life. Yet another market trend is manifested by the ever increasing proliferation of PDAs and PDA-type wearable PCs used in assisting users with personal information management (PIM), which include things like to-do lists, calendars, and phone books. PIM data are usually kept in sync with similar data residing on desktop computers.

A common denominator of these trends is a communications-enhanced lifestyle. Using these market trends as a product design guideline will create end-user devices, services, and solutions that can further enrich this common denominator. Such products should be easy to use, low-cost and low-weight, and be able to, say, access the Internet for remote updating of PIM information, Web access and e-mail on the move, instant notification and group paging, and so on.

In general, abstracting from the above market trends, we envision wireless technologies enabling information access and communications among intelligent devices and network-based services. These devices will have diverse capabilities in terms of usage models, power consumption, and user interfaces, and will support multiresolution applications with a wide range of content/media presentation capabilities. The end-to-end transport will be based on an adaptive networking framework with a collection of collaborating network elements. The low cost of the envisioned wireless technologies would translate to a universal presence of these technologies in all sorts of personal IPAs from the simplest to the most advanced, enabling connectivity transparent to the end user.

We envision that a heterogeneous collection of IPAs will be able to communicate with each other, each able to provide a set of distinct functions on its own. When these devices are found in close proximity to each other, the combination of their distinct functions will enable advanced usage scenarios that will further enrich an end user's experience interacting with his/her personal IPA and ultimately increase his/her productivity.

As an example, consider a PDA and a cellular phone. Each has a very well defined, useful, and distinct functionality. If these devices can communicate directly with each other in close proximity, these two devices can coordinate their functions to provide enhanced services to their owner which are different from the primary services provided by each of them. The PDA could connect to the cellular phone, and instruct it to dial up an Internet service provider (ISP) and establish a wide-area connection between the PDA and the ISP. This enhanced service provision occurs transparently to an end user. The PDA could be programmed to search for the cellular phone or any other device in its vicinity, such as a data access point on the ceiling of a conference room, that can provide Internet connectivity. If the two devices communicate via an RF link, the cellular phone in this example need not even be in sight of the PDA. The cellular phone may as well be hiding inside its owner's pocket or briefcase. This example

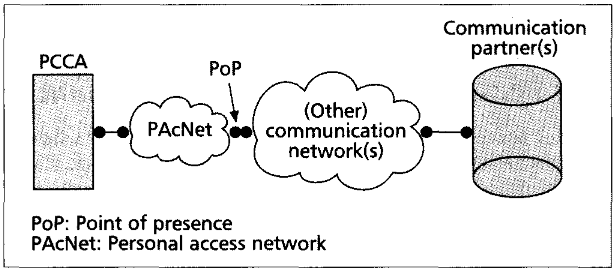

■ **Figure 1.** *A PCCA communication model.*

represents an instance of the *hidden and unconscious personal computing and communications* paradigm which is enabled by the short-range RF-based type of connectivity considered here. This paradigm is at the heart of our vision.

Our vision further encompasses a device- and application-aware communications infrastructure able to adjust information content to the constraints imposed by the underlying networks and end-user IPAs. For example, considering the previous example, wireless WAN bandwidth constraints and PDA display limitations could be considered when accessing a Web site so that contents of the Web site are transmitted efficiently over the wireless WAN and presented in a manner consistent with the PDA display limitations.

The potential of device-aware enhanced services resulting from combining the functions of several IPAs together is a powerful paradigm. This paradigm will free an end user from the mundane chores of connecting, configuring, maintaining, and in general worrying about unnecessary technical issues. These will be totally hidden from the end user, who will only be concerned with engaging with the services of interest.

There are several ways IPAs could connect with each other and other service providers. In the next section we present some of the most commonly expected connectivity models involving personal IPAs (i.e., PCCAs).

Connectivity Models

The ever increasing number of PCCAs such as laptops, PDAs, and palmtops can communicate with other devices using the general connectivity model shown in Fig. 1.

As Fig. 1 shows, the PCCA communicates over a communications infrastructure with a *communications partner*. The communication partner could be another PCCA, a server device, or, more generally, a service provider. The (other) communication network(s) cloud in Fig. 1, representing a PSTN, an intranet, a LAN, or some other form of network, may or may not be present in all cases, but if one is present a personal access network (PAcNet) is also present that connects the PCCA with a point of presence (PoP) on one of these other communication networks. The PAcNet could be a simple point-to-point RF link or a small shared medium network.

With this in mind, we identify the following alternative connectivity models. Each of these alternatives may enable its own particular set of applications, while other applications may be viable over several of these alternatives.

The Direct Connectivity Model

In this alternative, communicating devices are in close proximity and communicate directly with each other. In this case, the (other) communications network(s) cloud in Fig. 1 does not exist, and the PCCA and its communication partner are directly connected on the same PAcNet. Depending on the nature of the access network, the following connectivity models are possible.

Dedicated Point-to-Point Link – In this case only two devices communicate. A symmetric communication protocol is usually used with each communicating device having similar communications capabilities.

Shared Medium Link – In this case multiple devices can be connected on the PAcNet. There may be a master communications device playing a central role in establishing and managing the communications between the various communicating devices (a master/slave relation), or each device is at par with any other with respect to communications (a peer-to-peer relation). In either case, whenever needed, applications could view the underlying communication infrastructure as a dedicated point-to-point link.

The Indirect Connectivity Model

With this alternative, communicating devices may be located beyond the coverage area of a single radio, and one or more additional communication networks are needed to connect the PCCAs with their communication partner(s). The PAcNet serves as a "pure" access network with the responsibility of connecting the PCCA with a PoP of the backbone collection of networks to which the service provider is ultimately connected.

The backbone could be a collection of privately owned enterprise intranets (LANs, WANs, etc.), the public Internet, public carriers (PSTN, public cellular networks, ISPs, etc.), or any combination. From an end-user point of view this should be immaterial, even though throughput and latency constraints imposed by these networks may affect the kinds of applications enabled by the various connectivity models. Even though the (other) networks cloud in Fig. 1 may comprise multiple networks, we distinguish the following general cases of connectivity; note that while we still consider an RF-based PAcNet, no such requirement is imposed on the (other) communication networks.

LAN-Based – In this case we consider all cases where the server applications are accessed via private (enterprise or home) networks and PCCAs connect to these servers via privately owned and managed networks.

WAN-Based – In this case we consider all cases where the server applications are accessed via at least one public network provider by leasing service offerings from third-party operators (e.g., accessing a wireless cellular data network like CDPD).

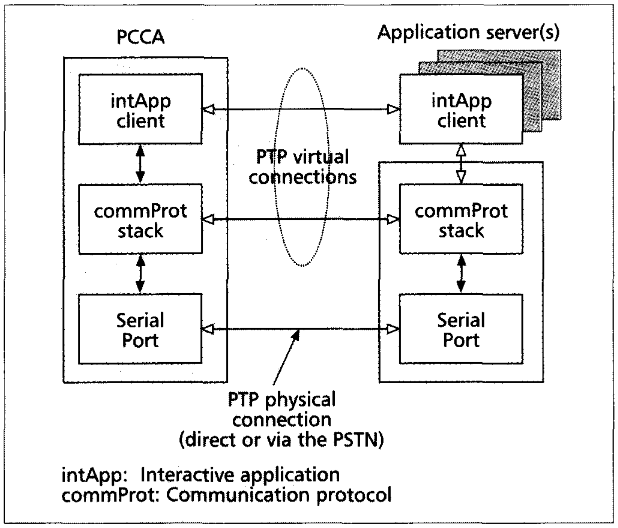

■ **Figure 2.** *A typical connectivity scenario for PCCAs.*

Due to the high concentration of people needing to access information resources and to be reachable while "on premises," it is our belief that one of the first instantiations of our vision would be in "restricted" areas, such as corporate office and academic buildings, hotel conference rooms, cafeterias, airport lobbies, and hospitals. Within such areas one can leverage the already existing dense backbone network infrastructure to enable two-way communications to people's PCCAs. At the periphery of the backbone network, data access points could be added to serve as a dense set of PoPs to the backbone network and all the services available through it. Leveraging the existing backbone infrastructure and the possibility of gradual deployment of this enhanced connectivity service lowers both the cost and risk of deploying the service. Both of these features are very attractive for any kind of new technology deployment.

The model of connectivity that enables the low-cost, low-risk first instantiation of our vision is the LAN-based indirect connectivity. WiSAP comprises an end-to-end prototype solution for connecting PCCAs to corporate intranets and home networks using the LAN-based indirect connectivity model.

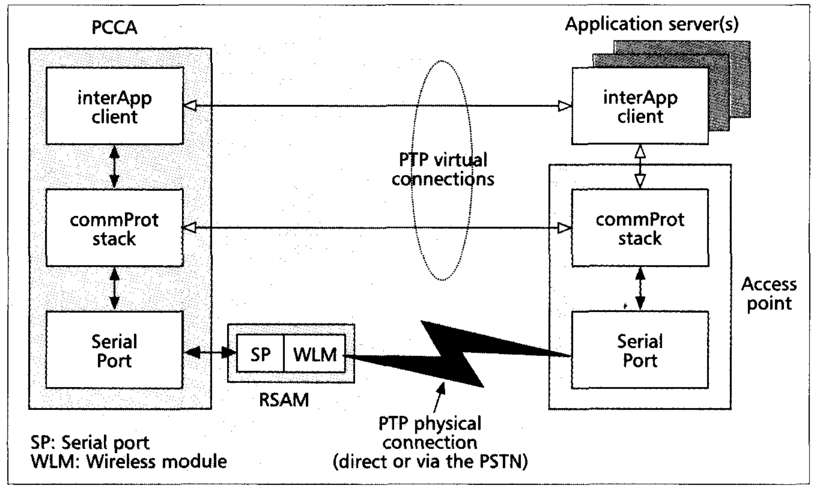

■ **Figure 3.** *A wireless-enhanced scenario for PCCAs.*

WiSAP: A Wireless Personal Access Network Solution for PCCAs

Our vision for continuous PCCA connectivity cannot be adequately realized using wireless solutions available in the market yet. For example, an 802.11-based solution runs contrary to the requirement for low cost and low power. A typical 802.11 PCMCIA card costs about US$500, which is more than what the most popular PDAs cost, while the power saving features of the protocol are hardly used.

For creating a total wireless solution addressing the envisioned networking and application frameworks within the system usability requirements, a different set of design "compromises" must be made. One of our goals is to bridge the gap between the wireless solutions currently existing in the market and our vision.

WiSAP is an umbrella name for a series of ongoing projects for developing end-to-end networking solutions enabled by the (ubiquitous) existence of low-cost RF-based connectivity modules for PCCAs. WiSAP serves as a development platform for creating, implementing, experimenting with, and testing combinations of the various technology enablers that will allow the realization of our vision, described earlier. WiSAP is not merely an exercise in building a wireless network, or a proof of concept for a wireless-based application. WiSAP is a research project on PAcNets, where usability requirements for realizing a pervasively wireless connected system via low-cost low-power integrated devices are as important as the technologies researched.

The project comprises various phases, including the development of medium access control (MAC) protocols for the wireless PAcNet, the development of backbone network processes and protocols for supporting PCCA mobility, and service discovery. Next, we provide a short overview of the WiSAP wireless platform.

WiSAP is targeted at enhancing PCCAs with RF-based wireless connectivity, while preserving their primary role as simple, low-cost, highly portable computing companions. WiSAP will provide PCCA users with seamless wireless access to network-based services (e.g., corporate applications, data store, Web content). WiSAP will be scalable and secure, and support conditional access without requiring changes to the network backbone infrastructure and, more important, without the need to modify the PCCAs or the applications running on them.

The basic WiSAP approach is summarized in Figs. 2 and 3. The overwhelming majority of PCCAs are equipped with an RS-232 port for serial communication. Figure 2 shows a typical communications scenario for such a device. The device is connected to a communications partner either directly (using "null-modem" wiring) or via a modem. In the latter case, the PCCA uses the connectivity services provided by the modem to connect to application servers located remotely to the PCCA.

WiSAP leverages the serial port connectivity of PCCAs by introducing an attachment to the PCCA called the *remote station access module* (RSAM). The RSAM attaches externally to the PCCA via its serial port[2] and allows it to communicate with a WiSAP *access point* over a shared RF link (Fig. 3). The access point, which is capable of communicating with multiple PCCAs via their RSAMs, serves as a multiplexing and demultiplexing device, bridging traffic between the wireline corporate intranet backbone and the PCCAs on the wireless access network. The access point also serves as the traffic coordinator for traffic originating from or destined to a PCCA. With respect to Fig. 1, the access point plays the role of the PoP entity for the wireline intranet.

Both the RSAM and the access point contain a wireless module (WLM). The WLM currently in use in WiSAP is based on IBM's Cordless Computer Connection (CCC), which is a cordless modem solution. The WLM radio is an FCC part 15 compliant radio operating at the 900 MHz ISM band; actually, it uses any of 10 different frequency shift keying (FSK) modulated channels in the 902–928 MHz range; channel selection is based on noise conditions on these channels. The transmission power is 1 mW (0 dBm) with about 30 m coverage in open space. The power consumption of the radio is roughly 180 mA at 4.5 V. The link transmission speed is 115 kb/s.

The overall WiSAP network consists of collections of IPAs with specific functionalities and responsibilities. By far, most of the devices are PCCAs located at the periphery of the backbone

[2] *The specific physical attachment currently available attaches the RSAM to the serial port of the PalmPilot/WorkPad family of PDAs. However, the RSAM itself assumes nothing about the actual PDA, and changing the physical attachment could allow the RSAM to connect to practically any device with a serial port.*

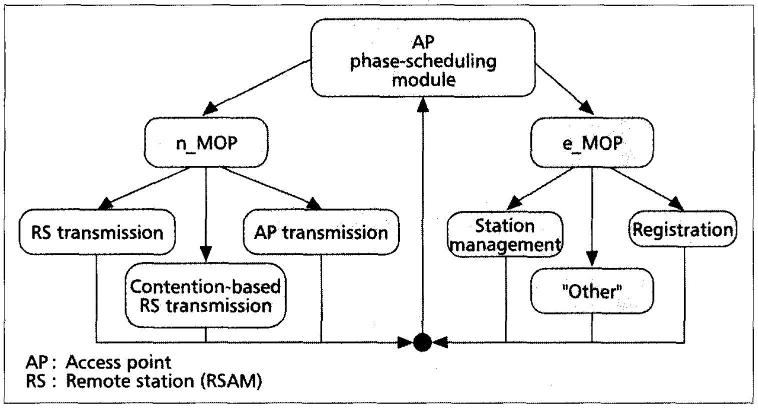

■ **Figure 4.** *The transmission decision process in an access point.*

network. Also, at the periphery of the network are the application servers that provide services to the PCCAs. The rest of the devices serve as facilitators for the transport of information between the PCCAs and the application servers. The access points are located at the edges of the backbone network. Within the backbone network there are *network support processes*, which are entities representing the various intranet and application support servers, such as DHCP and DNS servers for IP network support, mobility support processes, and transcoding servers for adjusting the information content sent by application servers to the capabilities of PCCAs.

As already mentioned, the WiSAP project aims to develop wireless solutions where ubiquitous connectivity weighs heavily on the design choices. Some of the solutions developed and experiences gained so far working on this project will be discussed next.

Discussion: Design Lessons Learned

The first and foremost lesson learned during our involvement in WiSAP is also the most obvious one. To provide value to end users, one needs to carefully study their current needs and develop products that reflect these needs. While one could build a state-of-the-art system with advanced solutions regarding each and every element of the system (e.g., battery size and lifetime, advanced MAC with QoS provisions, range coverage, support for mobile devices), the system will be branded useless immediately if it does not satisfy the basic user requirements: low cost, right form factor, good performance, low maintainability, transparent system operation, and so on.

Design decisions need to be made that bring the final product in harmony with end users' requirements. Some of these designs, the reasons behind them, and additional experiences gained and observations made during our involvement will be presented next.

Cost and Complexity

With several PDA-type devices costing around US$300–400, and with prices of these devices dropping steadily,[3] there is a natural limit on the price people would be willing to pay for any additional options available for such devices. The cost of materials of any module supporting wireless connectivity should be no more than a few tens of dollars if one wants to manufacture such a module and market it successfully.

To keep the cost low, all the elements that make up the overall end-to-end solution need to reflect this requirement. The WiSAP system has been designed for low-cost wireless connectivity. This is reflected in both the system architecture selected and the MAC protocol designed.

System Design – The use of a simple plug-and-play RSAM and access point combination has achieved the following. It provides enhanced connectivity services and at the same time protects end-user investment in hardware and software already purchased. The latter implies that an end user does not need to purchase a new PCCA, and he/she does not need to install special software to achieve the enhanced services. Also, any application that uses the serial port (e.g., connecting to the Internet using a PPP connection over a modem) can transparently operate in the new enhanced environment. It is the exclusive responsibility of RSAM to provide the enhanced services.

Furthermore, any MAC protocol operations running in the RSAM for sharing the RF link (i.e., the PAcNet) between multiple RSAMs and a access point are kept at a minimum. This is achieved by off-loading most of the MAC intelligence to the access point, whose capabilities can be shared among several RSAMs. Reducing MAC complexity at the RSAM permits us to design low-cost RSAMs with a small footprint.

The key operational elements of the MAC protocol are highlighted below.

The MAC Protocol – The RSAM receives data in bytestreams from the PCCAs via its serial port (Fig. 3). The RSAM forms data packets out of these bytestreams and then transmits them to the access point over the wireless PAcNet via a shared medium MAC protocol.

The data packet transport between the RSAMs and the access point is performed under the control of the very simple wireless network MAC (vsWINMAC) protocol. The protocol is a simple polling-based protocol with an advanced scheduling module sitting at the access point. The access point is responsible for any transmission decisions, while the RSAM simply waits to be polled by the access point. Decisions as to when to poll an RSAM, or when to send a message to it are made by the access point phase-scheduling module (AP_PSM) (Fig. 4).

Even though this protocol is very simple, it is capable of supporting quality of service (QoS) and accommodating dynamic changes in traffic loads. The AP_PSM can be replaced with more advanced or simpler versions of it as needed. Hence, based on the desired complexity of the access point and AP_PSM, transmission decisions could be made according to traffic demand to/from particular RSAMs, the requirement to maintain a minimum bandwidth to/from particular RSAMs, the requirement of periodic transmissions to/from particular RSAMs, or other requirements.

The output of the AP_PSM module is a decision of what type of transmission should happen next. It could be either a normal mode of operation (n_MOP) or an exceptional mode of operation (e_MOP) transmission. The former kind of transmission includes regular user payload transmissions from either the access point or RSAMs, which is accomplished by sending a poll from the access point to a single RSAM or a group of RSAMs. The latter includes station and link management-related transmissions, such as registration and handoff. While the AP_PSM could comprise anything from simple round-robin to highly sophisticated scheduling algorithms, the AP_PSM intelligence and associated complexity is hidden from the RSAM, which simply needs to know how to respond to the type of transmission it receives from the access point.

[3] *As this article was prepared, news reports were announcing the introduction of a sub-$100 PDA.*

Given the simplicity of vsWINMAC within an RSAM, this portion of vsWINMAC can be implemented on the PCCA's own processor. In this case, the RSAM will consist only of the radio front that will transmit the byte-stream received through the serial port (WLM in Fig. 3). Such a solution will further simplify the PCCA and hence lower its cost.

Higher-Layer Network Support Services – The MAC protocol for coordinating RSAMs and access point transmissions in the PAcNet serves as a transport mechanism for higher-layer network protocols such as PPP and IP. To simplify the system design of the access point, we decided to "push" the termination of PPP, or another higher-layer protocol, away from the access point and into network support processes located inside the network. Thus, the access point serves as a layer 2/3 bridge between the wireless and wireline media for traffic from/to the RSAMs with which it is associated. Network support processes can share their protocol processing resources for traffic arriving from or destined to not only multiple RSAMs, but multiple access points as well.

Power Consumption

Since a PCCA is designed to be portable, any add-on feature needs to be power-cautious. Using the cellular phone power consumption model, something to which most mobile workers are accustomed and most probably will use as a comparison criterion, the wireless module should be able to stay powered at an absolute minimum for at least a day in standby, with two to three hours of active operation.

Battery efficiency increases at a very slow rate, while the cost of the most efficient is not in line with a low-cost design objective. Thus, it is imperative that the system be architected in a power-consumption-aware manner (e.g., by selecting a low-power microprocessor). However, low power consumption of components in itself is not always sufficient, and more drastic power conserving measures will be needed as well.

To address this problem, vsWINMAC is designed to facilitate power conservation. Special station management messages can be exchanged between a PCCA and the access point that can place an RSAM in a powered-down sleep mode. During the sleep period of an RSAM, the access point queues any traffic destined to that RSAM. Either the access point or the RSAM may initiate the message exchange for placing the RSAM in sleep mode. When an RSAM goes to sleep, the AP_PSM is informed of this event and suspends transmissions to/from the RSAM for the duration of the sleep period. The RSAM may wake up periodically to either check whether the access point has traffic for it — the AP_PSM will schedule message transmissions to the RSAM during the periods when an RSAM wakes up — or notify the access point that it has traffic to send and will suspend its sleep mode. For an informative treatise on low-power technologies and solutions see [4].

Low-Cost Radio Technologies

Even though the radio used currently in WiSAP is relatively low-cost, there is great interest in leveraging new technologies that will bring the integration level, and hence the cost and power consumption, of the radio modules even lower.

As the transmission frequency moves up (i.e., 2.4 and 5x GHz), the manufacturing options are very limited. For low-volume applications, a discrete radio design using pseudomorphic high electron mobility transistor (PHEMT) devices for low-noise amplifiers (LNAs), and straight GaAs for power and Si for the lower IF sections, using separate discrete filters, and so on may be a reasonable approach. But for the high volumes wireless enhanced PCCAs will command, the volumes appear to drive integration, cost, and power targets that are unreachable with standard technologies.

Radio module designs using SiGe technology to achieve higher integration, lower cost, and reduced power are very promising. SiGe offers the ability to fully integrate the radio down to the IF and then do analog-to-digital (A/D) and D/A conversion directly. With this accomplished, the baseband processing can be done to demodulate the signal either on the same chip or off chip, depending on the aggressiveness and needs of the application. There are trade-offs to be made in terms of the overall system performance, such as the noise figure: whereas there is no beating GaAs PHEMT technology in low noise, in virtually every other area SiGe offers a designer more options and greater flexibility.

Currently, there are SiGe-based filters, mixers, LNAs, PAs, and so on available on the market. If one can design a system within the bounds of a given technology, taking into account all performance parameters, and thereby integrate these components onto a single piece of silicon, doing away with costly items such as surface acoustic wave filters (SAW), there can be a significant advantage.

Radio Integration – The degree of hardware integration is a key contributor in the overall system cost. Properly selected radio module architectures could permit higher or lower degrees of integration using a particular technology. Given what the cell phone manufacturers have achieved, one would like to capitalize on their efforts and build a radio similar to those built for their digital systems. If one takes this approach, the result could be the architecture shown in Fig. 5, a fairly complex and versatile radio, although a somewhat costly design. The performance of this radio generally exceeds all other standard architectures, but if it cannot be built in time or at an attractive cost it does little good.

If the requirement is cost-driven and time is critical, and if one can live with slightly degraded system performance such as a slightly higher noise figure, in-band image, and a little bit more phase noise, there are a number of less complex architectures with lower technical integration risks and fewer components. An example of such an architecture is shown in Fig. 6.

The architecture in Fig. 6 is a single conversion design which can essentially accommodate all the necessary elements onto a single chip. There are fewer phase-locked loops (PLLs), and less on-chip filtering, single-stage mixing, direct modulation, and so forth. It uses standard technology and RF components that currently exist in many libraries in several technologies. This architecture will end up with reduced operating range (distance), less adjacent channel interference (ACI), and therefore a slightly higher risk of jamming. However, these "drawbacks" can be factored into the overall network architecture and made acceptable. For example, in WiSAP multiple access points can be placed to cover a large area. Using the low-coverage but low-cost design of Fig. 6 for the multiple RSAMs and to the few access points supporting these RSAMs in large areas (e.g., an auditorium) still keeps overall system cost low. With a simpler radio design, as in Fig. 6, higher integration levels are possible (e.g., using SiGe technology). Alternative radio architectures can be found in [5].

Regarding power consumption, the final PA is a critical design choice for the system since its DC efficiency will directly affect the operational battery life of the system. If the coverage range requires additional output power (i.e., greater than 0 dBm and more like 20 dBm), one needs to consider the efficiency achievable in a particular technology. Off-the-shelf amplifiers can be as low as 5 percent efficient, and the best class A or AB amplifiers can be near 50 percent efficient. Based on this alone, it will cost the system at least 60 mA at 3.3 V to get 20

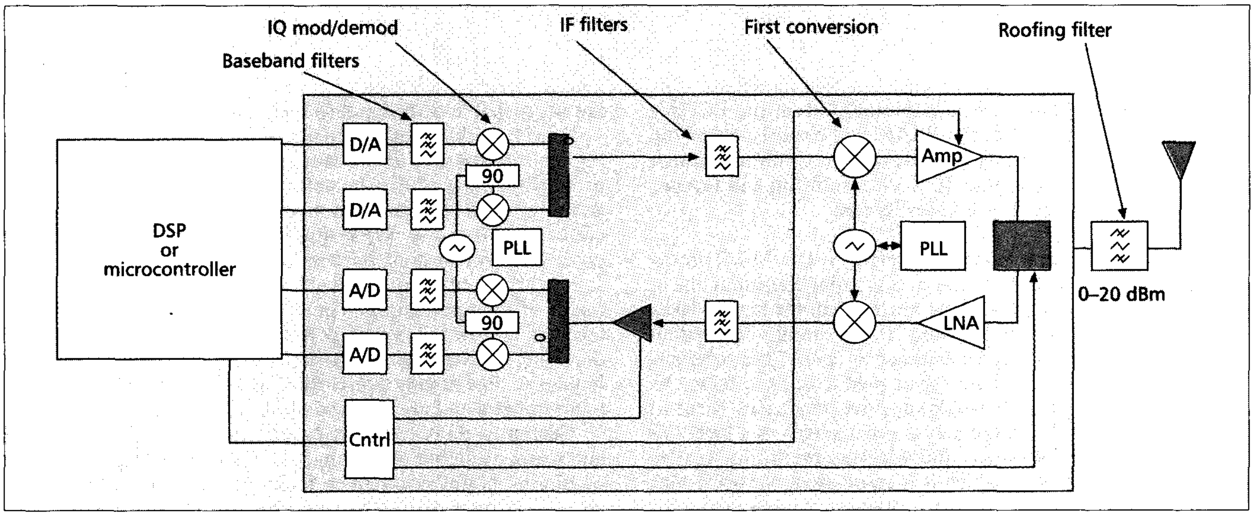

■ **Figure 5.** *Dual-conversion-based architecture.*

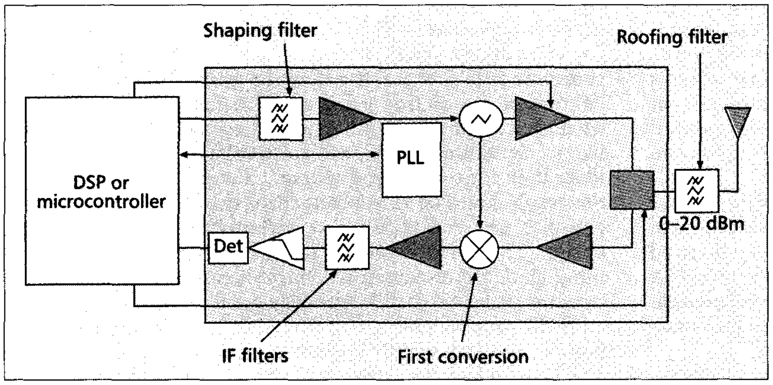

■ **Figure 6.** *Low IF architecture with analog detection.*

dBm out to the antenna (which is another design topic left out here)! Since this has obvious implications on battery life, one may need a solid justification to require spending this much DC power for higher coverage. When compared to roughly 1 mA at 3.3 V to get 0 dBm transmission power, one should begin to get a sense of the problem. A typical battery may have roughly 600 mA/hr of energy available, and a single 20 dBm power amplifier requires about 10 percent of that all by itself. In this and many other instances, to keep battery life long, integration high, and cost low, it is preferable to have shorter-range units and a few more access points. As technology improves and efficiency increases this will change, but for the near term this seems to be the proper approach.

Other Issues

To further enhance the unconscious and hidden computing paradigm, the following issues have also been identified as potential areas of study. Solutions to some of these issues may already exist in the literature; nevertheless, such solutions need to be revisited and new ones developed to once again be in harmony with the simple and low-cost PCCAs.

Enabling Connectivity for the Unconscious and Hidden Computing Paradigm — A wireless personal access data network could be highly diversified; multitude of different end devices with different capabilities may participate in it. Due to the requirement of no end-user intervention in setting up and configuring such a network, devices participating in the network need to follow a rich yet lightweight series of negotiations to self-configure and enable communications between them.

A wireless PAcNet is fundamentally different than conventional data networks, either wireline or wireless. In the latter networks exists a well-defined set of services provided by a well-defined set of network nodes. Services and nodes are closely monitored and maintained by appropriately trained personnel, and any change in these services or nodes is usually well advertised. While services (Internet access, printing, data sync) to be accessed by a PCCA over a PAcNet can be well defined (most of the time), their availability cannot be guaranteed at all times. Services may disappear abruptly as a worker steps through the corridors of an office building, while other services may come into existence temporarily. The fluidity of services is caused by the fact that as people move around, they come closer or go further, consciously or unconsciously, to/from a variety of other IPAs, each of which may provide a different set of services, again consciously or unconsciously, to PCCAs in close proximity to them. We may say that PAcNets "suffer" from a high mobility-to-coverage ratio (MCR), meaning that the area covered by a PAcNet is too small compared with the speed at which the nodes in the PAcNet may move. Hence, service providers may be accessible only for a very limited timed.

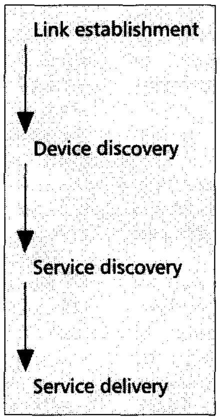

■ **Figure 7.** *The enabling procedures.*

Therefore, to support a rich and dynamically changeable set of end-to-end services between devices on a PAcNet, (at least) the following set of procedures and supporting protocols must be established (Fig. 7). Moreover, the following procedures should occur transparently to the end user, whose main interest is whether a particular service can be provided to him/her.

Link Establishment – A PCCA, upon a request for remote device connectivity, will first try to establish link connectivity with other devices that share a common air protocol. Using the services of MAC- and PHY-level protocols, the PCCA will "search" for available networks and register its presence within these networks.

Device Discovery – Due to the hidden and unconscious connectivity capabilities of the envisioned networks, quite often devices registered in a network during the link establishment phase above may not be visible, or their existence may be altogether unknown, even to the user of a PCCA. Thus, following the establishment of a link(s) with other devices, the devices need to identify each other so that proper device drivers related to each device can be loaded to allow communication between them. For example, devices can identify themselves as cellular phones, printers, PDAs, fax machines, pagers, toasters, or whatever.

Service Discovery – Each device can provide support for a number of services over the wireless network (e.g., printing services, Internet access services, fax services). Based on the required application, a device will establish an end-to-end connection with only those devices (if any) that support the desired services and applications.

Service Delivery – Having located an appropriate service provider, the device can now provide the requested service to its end user.

Security – Due to the ubiquitous, hidden, and unconscious nature of the connectivity envisioned, system security is of paramount importance. System security does not restrict itself to the link layer, where a third-party eavesdrops on somebody else's data exchanges. It encompasses every layer of communications, including the applications and the data that these applications manipulate. With hidden and unconscious connectivity, one may not even be aware to which devices one's PCCA is connected. Someone in the next room or on the floor above may connect to someone else's PCCA and poke around in his/hers stored e-mail, meeting schedules, and other private data.

Thus, not only must encryption be employed to avoid eavesdropping, but robust authentication procedures need to be established for connecting both trusted and nontrusted devices with each other. Data as well as services need to be flagged with different access control levels such as *private*, *secret*, or *public*. Private data, like one's phone book, will be accessible to no one but their owner. Secret data, like certain e-mail messages, will be accessible only to certain trusted devices. Public data, like electronic business cards, will be available with no constraints.

Conclusions

We present our vision of pervasive, unconscious, and hidden connectivity for personal computing devices via RF-based solutions. We then discuss WiSAP, a research project for realizing our vision for ubiquitous and wirelessly connected PCCAs while respecting market constraints with respect to device usability, cost, and power consumption. We finally discuss some of the lessons learned so far from WiSAP.

Acknowledgments

The authors would like to acknowledge the many contributions of Ephraim B. Flint and Young H. Kwark to the CCC radio front-end used in WiSAP.

References

[1] Bluetooth: http://www.bluetooth.com
[2] HomeRF: http://www.homerf.org
[3] WPAN http://grouper.ieee.org/groups/802/11/index.html
[4] Special Issue of *IEEE Pers. Commun.*, Energy Management in Personal Communications and Mobile Computing, vol. 5, no. 3, June 1998.
[5] L. E. Larson, "Radio frequency integrated circuit technology for low-power wireless communications," *IEEE Pers. Commun.*, vol. 5, no. 3, June 1998, pp. 11–19.

Biographies

CHATSCHIK BISDIKIAN [SM] received a Diploma in electrical engineering from the Polytechnic School of the Aristotle University of Thessaloniki, Greece, in 1983, and M.Sc. and Ph.D. degrees, both in EE, from the University of Connecticut, Storrs, in 1985 and 1988, respectively. Since 1989 he has been a research staff member at the IBM T. J. Watson Research Center, Yorktown Heights, New York. His current research interests are in short-range wireless communication networks, home networking, multimedia and broadband communications, HFC networks and cable modems, modeling and performance evaluation of high-speed networks, and computer system modeling and analysis.

PRAVIN BHAGWAT (pravin@watson.ibm.com) is a research staff member in the mobile networking group at the IBM Thomas J. Watson Research Center. He received a B.Tech. degree from the Indian Institute of Technology, Kanpur, in 1990, and M.S. and Ph.D. degrees in computer science from the University of Maryland, College Park, in 1992 and 1995, respectively, all in computer science. He actively serves on program committees of networking and mobile computing conferences, and has published numerous technical papers in the areas of mobile networking, routing, and application layer proxies.

BRIAN P. GAUCHER [M] performed his undergraduate work at the University of Massachusetts specializing in RF and microwave design. He received his graduate degree in electrical engineering from Northeastern University specializing in digital communication systems. From 1981 to 1983 he worked at Alpha Industries performing research and design of GaAs FET amplifiers and supercomponents. In 1983 he took a position at GTE in Needham, Massachusetts, where he worked on a variety of state-of-the-art communication and radar systems ranging from 1–60 GHz for the military. These included direct sequence and frequency hopping systems as well as frequency chirped radar applications. He was an active principal investigator throughout his career at GTE working on projects from millimeter wave components to handheld spread-spectrum transceivers. In 1993 he took a position at IBM's T. J. Watson Laboratory where he has been working on novel wireless computing methods and products. Since joining IBM's Communications Technology Department, he has helped three products to the marketplace.

FRANK J. JANNIELLO received a Bachelor's degree in electrical engineering from the University of Bridgeport and a Master's degree in computer science from Polytechnic University. He has done hardware design for the RAINBOW I and II optical networks, the 9729 wavelength-division Multiplexer, and a wireless communications system. He has been with IBM since 1990.

IBRAHIM KORPEOGLU received his B.S. degree from Bilkent University in Ankara, Turkey, in 1994, and his M.S. degree from University of Maryland-College Park in 1996, both in computer science. He is currently a Ph.D. student at University of Maryland-College Park. His main reseach interests are network protocols, mobile computing, and wireless networks. Since June 1998 he has been an intern in the Communications Systems Department of IBM T. J. Watson Research Center.

MAHMOUD NAGHSHINEH [SM] is with the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, where he currently manages the communication systems group. He joined IBM in 1988. Since then, he has worked on a variety of research and development projects dealing with design and analysis of LANs, communication protocols, fast packet-switched/broadband networks, wireless and mobile ATM, and wireless radio/infrared access broadband and LANs. His current interests are optical and wireless networking, hardware and software co-design of communication systems, and adaptive networking. He received his doctoral degree form Columbia University, New York, in 1994, his M.S. in electrical engineering and B.S. in computer engineering from Polytechnic University, New York, in 1991 and 1988, respectively, and the Vordiplom degree in electrical engineering from RWTH Aachen, Germany, in 1985.

PUNEET PANDOH received his B.S. degree in electronics and communication engineering from Delhi College of Engineering, India, in 1994. He worked in the design and implementation of network management protocols. His main areas of interest are network management and wireless network protocols. He has worked since October 1997 in the Communications Systems Department of IBM T. J. Watson Research Center.