Wireless networked digital devices: A new paradigm for computing and communication

by T. G. Zimmerman

The proliferation of mobile computing devices including laptops, personal digital assistants (PDAs), and wearable computers has created a demand for wireless personal area networks (PANs). PANs allow proximal devices to share information and resources. The mobile nature of these devices places unique requirements on PANs, such as low power consumption, frequent make-and-break connections, resource discovery and utilization, and international regulations. This paper examines wireless technologies appropriate for PANs and reviews promising research in resource discovery and service utilization.

In his review of previous ten-year predictions, Jean-Paul Jacob[1] observes that the explosion and effect of communication technology and infrastructure has always been under-appreciated. Few predicted the importance and ubiquity of cellular phones and the Internet. Cellular phone technology provides mobile voice communication, but mobile workers also need immediate access to digital information. The Internet provides a wealth of digital information, but the interfaces and data links are typically not portable.

The Newton**, introduced by Apple Computer, Inc., in 1993, launched the concept of a pen-based pocket computer: an electronic version of a day planner, address book, and memo pad. However, the device could not fit in a pocket and was slow and complex. The pocket organizers in the early 1990s provided the needed functions, speed, and size, but their small keyboard and inability to exchange data with personal computers (PCs) made them undesirable. 3Com Corporation's Palm Pilot** succeeded by providing a pocket-sized form factor, simple functionality, fast response, and a PC interface to popular organization software.

The fusion of cell phones and personal digital assistants (PDAs) can deliver the communication and information required by mobile users. The Nokia 9000I Communicator** was one of the first PDA/phone hybrids. However, it suffers some of the same problems as the early Apple Newton, and it may not be necessary or desirable to have a cell phone and PDA in one package. Cell phones can be considerably smaller than PDAs, which makes them convenient to carry around but difficult to use for PDA functions because of limited display and input area.

It seems probable that cell phones will adopt more PDA features and PDAs will acquire wireless connectivity. The ability to carry around small, wirelessly networked digital devices introduces a new paradigm in computer usage. The vision of ubiquitous computing, described by Mark Weiser,[2] is realized in three scales: "tabs" (about the size of an alphanumeric pager), "pads" (laptop-sized screens with pen input), and interactive whiteboards. Pads and tabs are wirelessly networked in picocells (about 6 cubic meters), allowing interaction to occur anywhere within a picocelled environment.

There are situations where carrying and holding a computer is not practical, for example when working on a jet engine or performing surgery. In both applications the worker's hands are occupied with critical tasks, yet access to computer-mediated information is needed—in the first case for blueprints and in the second for computed axial tomography (CAT) scan images. These workers would benefit from a computer they could wear. Wearable computers distribute computer components on the body for reasons of ergonomics, weight, and size. The wearable components include head-mounted displays, microphones, earphones, processors, and mass storage. Wiring these components together can create a cumbersome and restrictive web of equipment. Wearable computers would benefit from a wireless personal area network (PAN).

This paper first examines the technology used for wireless communication. Putting a radio in a digital device provides physical connectivity; however, to make the device useful in a larger context a networking infrastructure is required. The infrastructure allows devices to share data, applications, and resources such as printers, mass storage, and computation power. Defining a radio standard is a tractable problem as demonstrated by the solutions presented in this paper. Designing a network infrastructure is much more complex. The second half of the paper describes several research projects that try to address components of the networking infrastructure.

Finally there are the questions that go beyond the scope of this paper, yet will have the greatest effect on the direction, capabilities, and future of this paradigm. Will these networking strategies be incompatible, like the various cellular phone systems in the United States, or will there be a standard upon which manufacturers and developers agree, like the GSM (global system for mobile communication) cellular phones in Europe? Communication demands compatibility, which is challenging in a heterogeneous marketplace. Yet by establishing and implementing compatible systems, manufacturers can offer more powerful and useful devices to their customers. Since these are, after all, digital devices living in a programmed digital world, compatibility and interoperation are possible.

Wireless communication

Personal area networks connect mobile devices carried by users to other mobile and stationary devices. Their communicating range is scaled to the spatial interaction volume of humans (e.g., the reach of a hand or the audible distance of a voice), typically under ten meters. This distance coincides well with the limited battery capacity of small portable devices, and allows communication bandwidth (picocells) to be reused spatially. Wireless PAN communication can occur over optical, magnetic, electric, and electromagnetic channels. Table 1 presents these methods, including typical performance values.

Electric field. Zimmerman[3] and Gershenfeld have demonstrated a PAN that uses the body as a conduit for information. Electronic devices placed on and near the body modulate an electric field that induces small currents throughout the body. Data connections are established by touch or close proximity (within 2 meters). Communication is through capacitive (near-field) coupling, and negligible energy is radiated, allowing global unlicensed operation. The body, however, also acts as a shield. Placing a hand on the device blocks any electric field from leaving or reaching the device. In practice the best location for an electric-field-based PAN device is a shoe. This location provides low impedance paths from the device to the body and to the ground, enabling digital communication links to be easily established through touch and handshakes. The combination of power extraction by walking, demonstrated by Kymissis and Nivi,[4] and the existence of 3.6 centimeter disk drives,[5] suggest an interesting shoe-based computer for the future.

Magnetic field. Richley[6] has implemented a PAN for Weiser's ubiquitous computers using magnetic fields. The quadratic attenuation of magnetic field with distance allows the creation of picocells, 6 cubic meters per base station, allowing spatial allocation of bandwidth. The system uses a 5 MHz (megahertz) carrier to achieve 240 Kbps (kilobits per second) and consumes only 10 milliwatts power.[7] Magnetic field is unimpeded by bodies, so human shielding is not an issue.

Infrared. Infrared (IR) is used in television remote controls as a low-cost method of moderate-range, low-data-rate wireless data communication. Laptop

Table 1 Typical performance values for PAN wireless communication

Channel Medium	Carrier Frequency (MHz)	Maximum Bit Rate (Kbps)	Range (Meters)	Advantages	Problems	Application
Electric field	0.1–1	20	2	Capacitive coupling limits range, data transmitted by touch, international use	Signal blocked by body, grounded objects	Identification
Magnetic field	5	250	6	Signal goes through body, international use	Antenna size	Picocells data messaging
RF UHF	300–500	40	30	Low cost, efficient for distance	Local regulations	Access control, telecommand
RF microwave ISM	2400	1000	10–100	International, fast	Cost, power	Data networking
Optical infrared (low data rate)	Light	10	10	Simple, very inexpensive, no regulations	Directional	Home remote control
Optical infrared (high data rate)	Light	1000	1	Simple, inexpensive, no regulations	Directional or high power	Infrared Data Association (IrDA) for laptops

computers use faster IR transceivers[8] to provide data exchanges in excess of 1 Mbps (megabit per second) at 1 meter. Kahn[9] presents analysis showing data rates as high as 100 Mbps are possible by using infrared laser diodes located in the walls and ceilings of office buildings. Although the transmitters require a large amount of power (300 megawatts per cubic meter), the receivers can operate at a fraction of the transmit power. The biggest problem with IR communication is signal blockage by opaque objects. The directional property of IR can be beneficial when intentional secure-data exchanges are required, for example a PDA may be used as an electronic wallet to authorize a purchase at a checkout stand. However, the unencumbered view required by IR is not practical for most applications, particularly for devices worn or carried on the body.

Radio frequency. Wireless radio frequency (RF) technology is well suited for PANs. The short-range (approximately 10 meters) requirements of PAN allow the use of low-power high-bandwidth (up to 1 Mbps) data connections created by proximity. The greatest limitations of radio PANs are international emission regulations and standardization of the physical and data-link layers. Once standards are established, volume production and improvements in design and process will allow the cost and power of PAN radios to drop.

In the late 1970s the Hewlett-Packard Company began experimenting with direct-sequence spread-spectrum transmission for wireless interterminal networking and petitioned the Federal Communications Commission (FCC) to release some spectra. In 1985, after four years of study, the FCC released the ISM (industrial scientific medical) bands.[10] The ISM bands (particularly 2.4 GHz [gigahertz]) have the advantage of unlicensed worldwide support (with some restrictions in Europe and Asia), high bandwidth capability, maturing technology, and emerging standards. UHF (ultrahigh frequency, 300–450 MHz) radios deserve consideration for some country-specific PAN applications due to their simplicity, low cost, and low power consumption.

Regulations. The number of wireless consumer products in use, such as pagers, cellular phones, and wireless headphones, is dramatically increasing, yet the electromagnetic spectrum (EMS) remains a finite resource. Each country regulates its own EMS. Most bands are licensed to limit interference. The licensee charges a fee to customers for using their EMS ("air time"). Licensed spectra are used for long-distance communication (>100 meters), for example, cellular phones and pagers. PAN systems require local communication (approximately 10 meters) and users do not want to pay a fee to have their PDA talk to their desktop computer. Coordination is essential so that digital devices do not jam one another. The coordination is provided by protocol standards.

Manufacturers of RF wireless digital devices and individuals who use their devices while traveling must consider international regulations. Table 2 offers a summary of international unlicensed EMS regula-

Table 2 Unlicensed international EMS regulations

RF Frequency (MHz)	Transmit Power	Range (Meters)	Comment/Usage	Country
13.56	High	10	Very narrow bandwidth, ISM band, industrial plasma welding, "contactless" smart cards	World
303.825	Low	2	Very low power, car door alarm, control	Japan, Korea, USA
303.825	Medium	20	Car door alarm, control	USA, Australia
315.0	Medium	20	Car alarms, garage doors	USA, Canada, Italy
433.92	Medium	20	Car alarms, garage doors, telemetry	Europe
868.0–870.0	High	10–30	ISM band, data networks, telemetry	Europe
916.5	High	10–30	ISM band, high power, telemetry, data networks	USA
2400.0	High	10–100	ISM band, microwave oven and RF lighting, data networks, telemetry	World (Asia and Europe restrict bandwidth)

tions.[11] Although many of the frequencies listed are actually frequency bands, a discrete center frequency is often listed, based on the commercial availability of frequency setting components.

Radio fundamentals. The bit rate of a communication channel is a statement about the channel capacity, which has a theoretical limit defined by the Hartley-Shannon law[12]

$$C = Blog(1 + S/N)$$

where C is channel capacity in bits per second, B is bandwidth, S is signal, and N is noise. To achieve a high channel capacity for a fixed bandwidth, the signal-to-noise ratio must be maximized. The signal-to-noise ratio increases with transmitter signal strength and receiver sensitivity, but decreases with distance. Transmit power is limited by EMS regulations and, in the case of portable devices, battery life.

The signal strength of an RF transmitter falls off with the square of the distance in "free space" (e.g., an open area), but can be as high as the fourth power of the distance for dense office spaces. Table 3 shows some experimental data of RF propagation range for a low power (1 milliwatt) 916 MHz radio operating at 19.2 Kbps with a −92 decibel sensitivity receiver.[13]

Table 3 RF signal propagation of a 1 milliwatt 916 MHz transmitter in various environments

Environment	Range (Meters)
Dense cubical office space	13
Open retail space	31
Large open area	62
Free space (outside)	175

Data are transmitted over a radio by modulating the radio signal (carrier). There are two basic types of carrier modulation: linear and nonlinear. Linear modulation translates the spectrum of the data up to the carrier frequency. The transmission bandwidth is narrow, never exceeding twice the bandwidth of the message, and the signal-to-noise performance can only improve by increasing the transmit power. On-off keying (OOK) is the simplest form of linear modulation. The carrier oscillator is turned on to represent a "1" and off to represent a "0." Amplitude modulation (AM) varies the amplitude of an oscillator, which is faster (higher data rate) than waiting for an oscillator to power up, but consumes more power than OOK.

Nonlinear modulation uses a nonlinear process to modulate the carrier, producing a broadband transmission much greater than the data bandwidth. The signal-to-noise performance of nonlinear modulation increases with the bandwidth of the transmission, trading bandwidth for power. Frequency shift keying (FSK) and spread spectrum (SS) are two common forms of nonlinear modulation. Nonlinear modulation is less susceptible to noise and other interference since the contribution of any one frequency represents a small portion of the total signal, as demonstrated by the lower noise floor (less "background hiss") of FM (frequency modulation) vs AM radio broadcasts.

In a direct-sequence spread spectrum system the data to be transmitted are multiplied by a digital pseudo-

Table 4 Representative UHF radio designs

Bit Rate (Kbps)	Transmitter, Modulation	Receiver Filter, Demodulation	Application	Cost per Transmission Received
1.2	LC, OOK	LC, super regenerative	Command and control	$ 2–$ 3
19.2	SAW, OOK	SAW, ASH/baseband	Telemetry	$ 5–$10
40	SAW, FSK	SAW, super heterodyne	Network, telemetry	$10–$15

noise (PN) sequence, spreading the spectrum of the data. The receiver retrieves the data by multiplying the incoming signal by the identical PN sequence. In frequency-hopping spread spectrum (the more common of the two), the carrier frequency is shifted (hopped) in a PN sequence. With both methods, the transmitter and receiver PN sequences must be synchronized in order for the receiver to extract the data from the sequence. Synchronization is the greatest challenge of a spread-spectrum communication system. A common receiver synchronization strategy is to acquire synchronization when a transmitter or receiver is first turned on, then continuously track synchronization while the devices are communicating.

UHF radios. UHF radios are suitable for providing low bit-rate (<20 Kbps), low-cost PANs. Each country has its own RF regulations, and unfortunately there is no unlicensed UHF signal frequency band that is allowed internationally. The frequency band around 433.92 MHz has the greatest acceptance, but is not international. Some applications under development in our lab at the IBM Almaden Research Center use this frequency to implement remote asset tracking, personal activity detectors, and other low-bandwidth telemetry. Garage door openers, car alarms, and wireless doorbells are common commercial applications of low-cost UHF radios.

UHF radio design offers the engineer many choices. The radio can be optimized for bit rate, cost, power, sensitivity, and size. Table 4 shows some possible constructions of UHF radios. The cost estimates are highly variable based on volume, design, vendor, and technology and are only listed to give a relative comparison.

Transmitter. A radio transmitter requires a frequency reference to establish the carrier frequency. There are several choices, of varying cost and accuracy. The lowest-cost UHF oscillator consists of a capacitor (C), inductor (L), and amplifier (typically a single transistor). The LC oscillators run at UHF frequencies (300–450 MHz) but are subject to component and environment variation (they lose tuning in proximity to conductors and dielectrics). They are typically used for garage door openers and car alarm transmitters. Surface acoustic wave (SAW) devices can resonate at UHF frequencies and provide a good frequency reference. They are used in higher quality car alarms and telemetry. Crystals cannot operate at UHF frequencies, but their overtones can be used or they can provide a reference for digital frequency synthesizers (e.g., phased locked loops).

Super regenerative receiver. A radio receiver must amplify a tiny signal induced on an antenna and convert it back to data, while rejecting all other frequencies and signals. The lowest cost, yet surprisingly sensitive, receiver is a super regenerative receiver. The super regenerative receiver is based on the principle that an oscillator circuit needs some electrical noise to start the oscillation. The noise basically provides a "push." The oscillator will start faster if the applied stimulus (the push) is at the oscillator's resonant frequency. By turning an oscillator on and off repeatedly, the start-up time, hence the presence of the resonance frequency, can be measured as the average oscillator current. This is the principle of a super regenerative receiver and can be implemented with a few dollars' worth of parts (a few transistors, coils, capacitors, and resistors). One by-product of super regenerative receivers is the generation and radiation of carrier frequency (electromagnetic pollution) by the receiver. Super regenerative receivers are only suitable for low data-rate OOK or AM modulation. They are primarily used for car alarms, wireless doorbells, garage door openers and other low-cost, low-bit-rate receivers.

SAW/ASH receiver. The ASH receiver (named after its inventor Darrell Ash) demonstrates a clever use of SAW devices, as a filter in the frequency and time domain.[14] The ASH receiver uses one SAW device to select (filter) incoming RF energy, rejecting all other RF energy. The second SAW device is used to help stabilize RF amplification. A receiver must amplify the desired frequency band up to 60 decibels (a power gain of one million). An amplifier with this much gain would run the risk of self oscillating due to cross coupling of the output to the input. Ash's solution is to amplify the signal in two stages, each with modest gain, and to have only one amplifier on at a time. The signal is applied to the first amplifier and the output is launched across a SAW device. As the signal travels across the SAW device (which acts as a 0.5 microsecond delay line), the first amplifier is turned off and the second amplifier is turned on. The signal reaches the second amplifier where it is further amplified. The result is extremely high gain (similar to a super heterodyne receiver) with stable performance. The time multiplexing scheme used by Ash, however, limits the receiver to OOK and AM detection.

Super heterodyne receiver. The most common type of radio receiver is a super heterodyne receiver. The receiver achieves high gain without feedback by amplifying at two different frequencies. The first stage of amplification occurs at the carrier frequency (e.g., 400 MHz). The amplified signal is mixed to an intermediate frequency (typically 10.7 MHz), filtered, amplified, and applied to a detector that recovers the data. The super heterodyne receiver is used in audio radios as well as pagers and other telemetry equipment, and allows many types of modulation including OOK, AM, and FSK (the most popular due to noise rejection).

Spread spectrum receiver. Spread spectrum receivers consist of an analog front end that amplifies the signal from the antenna and a digital processing unit that gathers and recovers the data. The amplified RF is either converted down to a lower intermediate frequency or directly digitized for digital signal processing. Spread spectrum is used in wireless LANs and some cordless phones.

2.4 GHz radio. The most promising radios for widespread PAN deployment are in the 2.4 GHz ISM band, owing to the international availability and bandwidth of this frequency band. Spread spectrum is typically employed to reduce interference and utilize the bandwidth. Wireless LAN products implementing the Institute of Electrical and Electronics Engineers (IEEE) 802.11 standard wireless LAN protocol are commercially available, with PCMCIA (Personal Computer Memory Card International Association) cards (to network laptop computers throughout a building, for example). However IEEE 802.11 products consume too much power and have excessive range for many PAN devices. An IEEE special interest group (WPAN)[15] has formed to modify the IEEE 802.11 standard for PAN application. Two other groups developing radio specifications operating at 2.4 GHz are HomeRF[16] and Bluetooth.[17] The HomeRF Working Group, a consortium of several major consumer electronic and computer companies, is developing a specification for wireless communications in the home: a means to interconnect PCs, peripherals, and remote displays. The Bluetooth Special Interest Group, an industry group of several major cell phone and computer companies, is developing a global specification for wireless technology. The Bluetooth technology was initiated as a means to connect cell phones to laptop computers.

Each of the radio technologies is very similar, yet incompatible. All use the 2.400–2.4835 GHz frequency band and frequency-hopping spread spectrum (FHSS), and all support data rates of at least 1 Mbps. The main differences among them are the hop rate and power (hence distances) as shown in Table 5.[18]

It would appear from this table that the needs of IEEE 802.11 WPAN could be served by the HomeRF specification. The incompatible HomeRF and Bluetooth specifications fragment wireless operation into home and business markets. It would be advantageous for both markets to adopt one standard. Customers would benefit from a continuity of connection that is independent of environment, and manufacturers would benefit from higher volumes and lower cost for radio components.

Networks

Networks allow computers to share information and resources. Connecting computers together is no easy task. For computers to communicate, they must agree on a physical means of connectivity and a communication format. This allows data to flow from one machine to another, but what makes data truly useful is context, typically an application. Having reviewed the physical forms of wireless communication, we now turn to a discussion of data formats and methods for making data available to applications.

ISO networks. The ISO (International Organization for Standardization) network standard (ISO 7498, released in 1984)[19] provides an explicit means to incorporate devices into existing computer systems.

Table 5 2.4 GHz PAN radios

Parameter	IEEE 802.11 WPAN	HomeRF	Bluetooth
Distance	Not determined	50 meters	10 meters
Hop rate	2.5 hertz	50 hertz	1600 hertz
Transmit power	<1 watt (US) 100 milliwatts (Europe and Japan)	100 milliwatts (North America)	1 milliwatt

The model defines the functions of multiple layers to allow each layer to be developed independent of the others and to assure that they will operate together. The first layer is the physical layer, the electrical specification to transfer unstructured bits across a communication medium. The second layer establishes a data link, providing reliable data transmission through synchronization and framing techniques. The third layer provides networking to upper layers independent of the specific physical and data link techniques used. The last layer, the application layer, provides services to the user, such as file transfers.

We have discussed the physical layer, examining the technology and evolving standards of the wireless communication channel. As noted, the 2.4 GHz radios aptly address the lower layers, establishing a reliable information link and managing media access control (MAC). We now turn our attention to the upper layers, some of which are unique to the mobile wireless nature of PAN devices.

Data formats. The Web provides a working model of how to exchange text and graphics through a network. Documents are sent using the markup language HTML (HyperText Markup Language). Although HTML tells a browser basic layout information (e.g., paragraph, font size), it does not provide a means to convey data structures. The recently introduced Extensible Markup Language (XML)[20] addresses this problem by providing a data format for structured document interchange. Unlike HTML, which is a fixed format, XML is a language that lets a designer create a markup language appropriate to the information and application.

WAP (Wireless Access Protocol)[21] is a wireless protocol specification under development. It is designed for small devices, particularly cellular phones. WAP will define a microbrowser, scripting, telephone functionality, some content formats (e.g., business cards and calendars), transport, and security. WAP has its own markup language (WML) and is intended to be implemented in less than 10 kilobits. It will most likely first appear on GSM phones in Europe.

Transport systems and service discovery. The Web was designed for desktop computer users to view and link to multimedia (text, graphics, and sound) files. WAP will allow cellular phones to display text and simple graphics and accept user selections. More ambitious goals are for wireless networked digital devices to use a shared office printer, capture lecture notes off a "smart" whiteboard, and receive a menu, place an order, and pay for a meal. These tasks require a means to discover and use services. The following are several innovative research projects that attempt to address network service problems.

Jini. Jini**[22] is a research project started in 1994 by Bill Joy and Jim Waldo of Sun Microsystems, Inc. Jini allows devices to create "spontaneous networks" when plugged in to each other. "Plugging in" can include two devices coming into wireless proximity. Jini is a networking infrastructure running on top of Java** to allow devices running Java Virtual Machines (JVMs) to announce and share services across a network. Jini is Java code (about 48 kilobyte core) consisting of class library forms and conventions. A goal of the project is to eliminate device configuration and driver installation. When a device needs a service (for example, a digital camera needs a printer), the device looks up the relevant services on the network.

JavaSpaces. JavaSpaces[23] is an event-driven system written in Java using remote method invocation (RMI) to allow "buying" and "selling" requests to be fulfilled. RMI is a set of application programming interfaces and a model for remote object communication to allow distributed computing. Java RMI includes means to negotiate use of resources, to commit to a set of operations, and to notify objects of state changes. Some anticipated applications of JavaSpaces include reservation and trading services.

Hive. Hive[24] is a research project at the Massachusetts Institute of Technology (MIT) Media Laboratory to allow the construction and operation of distributed systems through networked computers. Hive provides a structure for communication and control of devices and Java applications. It uses RMI and object serialization to call and move Java objects running on other JVMs.

TSpaces.* TSpaces[25] is a research project at IBM's Almaden Research Center led by Toby Lehman. TSpaces is written in Java and provides group communication, databases, URL-based file transfer, and event notification services. TSpaces provides a foundation for client applications that can be downloaded by proximity. By providing a common platform of data exchange, TSpaces can be used as a universal print, e-mail, and pager service.

GinJo. GinJo[26] is a research project at IBM's Tokyo Research Laboratory. GinJo (geographical interaction network for jumping objects) is an *ad hoc* wireless network where messages and code can hop from one wireless device to another when the devices meet on the street, in a building, or in vehicles. GinJo will allow advertisements to beam down from billboards, handing out electronic coupons for restaurants or free parking for events.

Conclusion

Portable computing devices with wireless short-range links are seen as a new paradigm for computing and communication. The availability of low-cost, low-power, short-range international 2.4 GHz digital radios will provide the required communication technology for PAN. Convergence of the competing radio standards would allow a seamless connection among the digital wireless devices in our homes, offices, and environment. Spontaneous networks and service discovery and delivery are vital to the usefulness of PAN devices. A browser can point to any Web page on the Internet—a PAN device should be able to likewise "plug in" to innumerable services anywhere in the world.

Acknowledgments

The author would like to thank the USER Group of the IBM Almaden Research Center and Neil Gershenfeld and the Physics and Media Group of the MIT Media Lab for their collaboration and support of PAN work.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Apple Computer, Inc., 3Com Corporation, Nokia, or Sun Microsystems, Inc.

Cited references and notes

1. Personal conversation with Jean-Paul Jacob, lecturer and member of Business Development Programs, IBM Research, Almaden Research Center (November 2, 1998).
2. M. Weiser, see http://www.ubiq.com/hypertext/weiser/UbiHome.html.
3. T. Zimmerman, "Personal Area Networks (PAN): Near-Field Intra-Body Communication," master's degree thesis, MIT Media Laboratory, Cambridge, MA (September 1995).
4. J. Kymissis and B. Nivi, see http://www.media.mit.edu/ttt/groups/joep/power/index.html.
5. IBM product announcement, see http://www.research.ibm.com/news/detail/mdintro.html.
6. E. Richley and L. Butcher, *Wireless Communications Using Near Field Coupling*, U.S. Patent No. 5,437,057 (issued July 25, 1995).
7. Magnetic field PAN performance from personal e-mail from Mark Weiser (November 2, 1998).
8. Infrared Data Association (IrDA), see http://www.irda.org.
9. P. Kahn, T. H. Probert, and M. E. Chase, "Trends in Local Wireless Networks," *IEEE Communications Magazine* **33**, No. 3, 88–95 (March 1995).
10. "Radio Frequency Devices," *FCC Regulations*, Part 15, Federal Communications Commission, Washington, D.C. (1993).
11. International EMS compiled from data sheets appearing in *RFM Product Data Book*, RF Monolithics, Inc., Dallas, TX (1997).
12. A. B. Carlson, *Communication Systems*, Third Edition, McGraw-Hill Book Company, New York (1986), pp. 230, 401, 514–536, 554.
13. *ASH Transceiver Designer's Guide*, RF Monolithics, Inc., Dallas, TX (1999); available from http://www.rfm.com/products/tr_des18.pdf.
14. "A Low Power Receiver/Transmitter Set," *RFM Product Data Book*, RF Monolithics, Inc., Dallas, TX (1997); see also http://www.rfm.com.
15. IEEE 802.15 Working Group for Wireless Personal Area Networks, see http://grouper.ieee.org/groups/802/15/index.html.
16. HomeRF Consortium, see http://www.homerf.org.
17. Bluetooth Consortium, see http://www.bluetooth.com.
18. Adapted from "Wireless Initiatives Comparisons, a PHY Layer Review" by Jim Zyren and Bruce Kraemer, IEEE 802.11-98/249 (July 1998), see http://grouper.ieee.org/groups/802/11/Mailings/July/julmin.pdf.
19. The ISO Network Standard ISO 7498, see http://info.acm.org/sigcomm/standards/iso_stds/OSI_MODEL/index.html.
20. XML, see http://www.w3.org/XML/.
21. WAP Forum, see http://www.wapforum.org.
22. See http://java.sun.com/products/jini/vision.
23. JavaSpaces, see http://java.sun.com/products/javaspaces/faqs/jsfaq.html.
24. Hive, see http://hive.www.media.mit.edu/projects/hive/primer.html.
25. TSpaces, see http://www.almaden.ibm.com/cs/TSpaces/index.html.
26. GinJo, see http://www.trl.ibm.com.

General reference

T. G. Zimmerman, "Personal Area Networks (PAN): Near-Field Intra-Body Communication," *IBM Systems Journal* **35**, Nos., 3&4, 609–617 (1996). Available at http://www.research.ibm.com/journal/sj/mit/sectione/zimmerman.html.

Accepted for publication May 25, 1999.

Thomas G. Zimmerman *IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120 (electronic mail: tzim@almaden.ibm.com).* Mr. Zimmerman is a research staff member in the Computer Science Department of the IBM Almaden Research Center, working on new human-machine interface devices and paradigms. He received his B.S. degree in humanities and engineering (1980) and his M.S. degree in media science (1995) from MIT. His interest in music led him to cofound Breakaway Technologies, where he designed a voice-controlled synthesizer and later patented several musical devices to help people sing and play music. In the early 1980s he invented the DataGlove and cofounded VPL Research with Jaron Lanier. In 1993, in collaboration with Neil Gershenfeld, he developed the electric field Personal Area Network at the MIT Media Lab, work he continues at IBM. His interactive exhibits can be seen at the Exploratorium in San Francisco, the National Geographic Society in Washington, D.C., and the Great Lakes Science Center in Cleveland, Ohio.

Reprint Order No. G321-5706.