



PicoRadio Supports Ad Hoc Ultra-Low Power Wireless Networking

One of the most compelling challenges of the next decade is the “last-meter” problem—extending the expanding data network into end-user data-collection and monitoring devices. PicoRadio supports the assembly of an ad hoc wireless network of self-contained mesoscale, low-cost, low-energy sensor and monitor nodes.

*Jan M.
Rabaey*

*M. Josie
Ammer*

*Julio L. da
Silva Jr.*

Danny Patel

Shad Roundy

University of
California,
Berkeley

Technology advances have made it conceivable to build and deploy dense wireless networks of heterogeneous nodes collecting and disseminating wide ranges of environmental data. An inspired reader can easily imagine a multiplicity of scenarios in which these sensor and actuator networks might excel. The mind-boggling opportunities emerging from this technology indeed give rise to new definitions of distributed computing and user interface.

Crucial to the success of these ubiquitous networks is the availability of small, lightweight, low-cost network elements, which we call PicoNodes. These nodes must be smaller than one cubic centimeter, weigh less than 100 grams, and cost substantially less than one dollar. Even more important, the nodes must use ultra-low power to eliminate frequent battery replacement. We envision a power-dissipation level below 100 microwatts, as this would enable self-powered nodes using energy extracted from the environment—an approach called energy-scavenging or harvesting.¹

POWER DISSIPATION TODAY

To put power dissipation into perspective, we can compare it with the state-of-the-art commercial devices available today. One of the closest matches is the Bluetooth transceiver, an emerging standard for short-range wireless communications. While meeting the volume requirement, Bluetooth radios cost more than 10 dollars and consume more than 100 milliwatts. Although Bluetooth's price point and power con-

sumption will inevitably drop with technology scaling, these modifications would still not address the orders-of-magnitude reductions required for sensor network applications.

To reach these aggressive power dissipation levels, we must limit the effective range of each PicoNode to a couple of meters at most. Extending the reachable data range requires a scalable network infrastructure that allows distant nodes to communicate with each other. A self-configuring ad hoc networking approach is key to the deployment of such a network with many hundreds of nodes.

Reducing the PicoNode's energy dissipation to this level is our focus here. The secret lies in a meticulous concern for energy reduction throughout all layers of the system design process. The largest opportunity lies in the protocol stack where a trade-off between communication and computation, as well as elimination of overhead, can lead to a many orders-of-magnitude energy reduction. Other opportunities lie in the adoption and introduction of novel self-optimizing radio architectures and opportunities for energy scavenging.

PICORADIO APPLICATIONS

Applications of such sensor and monitoring networks include environmental control in office buildings; robot control and guidance in automatic manufacturing environments; warehouse inventory; integrated patient monitoring, diagnostics, and drug administration in hospitals; interactive toys; the smart home providing security, identification, and personalization; and interactive museums.

Building control

As one example of an application for PicoRadio networks, consider the management of environmental control systems in large office buildings. Any person who has spent a significant amount of time in such an environment is acutely aware of its problems: The temperature or the airflow is never right, and there is too little or too much light. A distributed building monitor and control approach might go a long way in addressing these problems—for example, by creating local microclimates adapting to an occupant's preferences through distributed air ducts—and might vastly improve the living conditions for the building's population. At the same time, such an approach can dramatically reduce the energy budget needed to manage the environment. First-order estimations indicate that such technology could reduce source energy consumption by two-quadrillion BTUs (British Thermal Units) in the US alone. This translates to \$55 billion per year, and 35 million metric tons of reduced carbon emissions.

Wiring the huge number of sensor and actuator nodes needed to deploy such a system is impractical and uneconomical. The cost of installing wiring for a single sensor in a commercial building averages \$200 in addition to the cost of the sensor. For low-cost devices such as temperature sensors, the cost of the wiring may be as much as 90 percent of the installed cost. In these cases, eliminating the cost of wire by using a wireless connection could reduce the installed cost per sensor by an order of magnitude and enable the deployment of ubiquitous sensor networks in contrast to the currently used sensor-starved solutions. We can even envision a future in which the sensor nodes are prebuilt into construction materials such as ceiling and floor tiles. To realize this vision, the communication/sensor nodes must be completely self-contained for the lifetime of the building.

Interactive environments

Consider another scenario: A science museum for children—for example, the San Francisco Exploratorium shown in Figure 1—that presents a collection of exhibits featuring a combination of data measurements and cause-and-effect experiments. Making the museum exciting requires creating a close interaction with the visitors controlling the exhibits and providing feedback on the experiments. In an even more aggressive scenario, the children can be active participants in the experiments. Keeping the exhibits flexible and easily modifiable is hence desirable. The availability of cheap and easily deployable wireless sensor, monitor, and actuator networks could create a true revolution in how these museums operate. In addition to the sensing/control/monitoring functions, ad hoc wireless networks could also provide paging, intercom, and localization functionality.

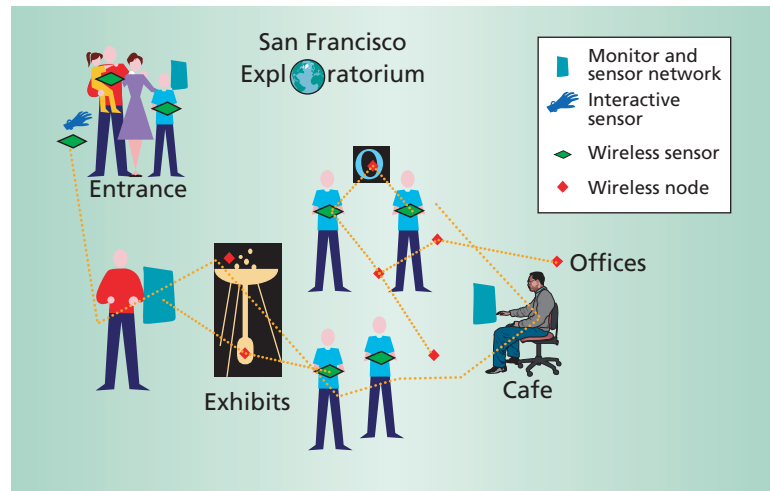


Figure 1. PicoRadio network facilitating interactive museum exhibits.

In short, these networks make it possible to take apart functions that traditionally were localized in a single point and distribute them over a much wider space—hence leading to potentially more optimal systems.

ULTRA-LOW ENERGY PICORADIO NETWORKS

The scenarios expose both the challenges and opportunities that PicoRadio networks offer in terms of energy efficiency. A number of prime properties are worth identifying:

- Sensor data rates are quite low, typically less than one hertz.
- Sensor nodes don't need to be awake all the time; in fact, a single node's activity duty cycle is typically less than 1 percent.
- Sensing data without knowing the sensor's location is meaningless. Localization should therefore be considered an implicit feature of the sensor network. This greatly simplifies the network discovery and maintenance effort and leads to substantial energy savings. For example, the sensor network can prune requests for information and direct them to the region of interest.
- Sensor networks require different addressing techniques than traditional data networks. Data requests are typically in the style of "Give me the temperature readings in room 30," compared to "Set up a connection between nodes A and B." The content- and localization-based addressing concepts make the overall network discovery and management a lot simpler.

Based on these specifications and properties, we can develop energy-efficient network, transport, media-access, and physical layer protocols. These in turn set the constraints and requirements for the hardware

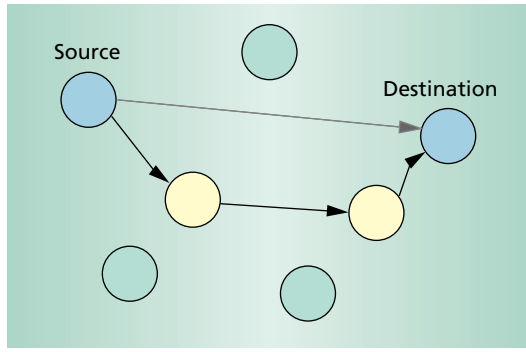


Figure 2. With multihop networks, using several short intermediate hops to send a bit is more energy-efficient than using one longer hop.

architecture and components of the transceiver nodes, including radio frequency (RF), baseband, and protocol processors. A number of innovations at the protocol stack level will make the intended energy reductions possible.

Protocol support

The three main layers we concentrate on are the physical, media access control (MAC), and network layers. Communication between two nodes requires creating a physical link between two radios. The physical layer handles the communication across this physical link, which involves modulating and coding the data so that the intended receiver can optimally decode it in the presence of channel nonidealities and interference.

Next, because many radios have to share the same interconnect medium (the aether), messages can interfere with each other, and access to the medium needs to be coordinated. The MAC layer provides this service.

When radios that are not within physical range of each other need to communicate, the network layer determines the path for a packet to take through other nodes that forward packets on their behalf. This forwarding of packets is often referred to as multihop networking.

Multihop networks

A main challenge in the design of an energy-efficient wireless network is that sending a bit of information through free space directly from node A to node B incurs an energy cost E_t , which is a strong function of the distance d between the nodes. More precisely, $E_t = \beta \times d^\gamma$, with $\gamma > 1$ as the path-loss exponent (a factor that depends on the RF environment, and is generally between 2 and 4 for indoor environments). β is a proportionality constant describing the overhead per bit. Given this greater than linear relationship between energy and distance, using several short inter-

mediate hops to send a bit is more energy-efficient than using one longer hop (as shown in Figure 2). For example, assuming $\gamma = 4$, which is a common case in indoor environments, and $\beta = 0.2$ femtojoules/meter⁴, one hop over 50 meters requires 1.25 nanojoules per bit, whereas five hops of 10 meters require only 5×2 picojoules per bit. The multihop approach in this example reduces transmission energy by a factor of 125. This situation is somewhat analogous to the problem of sending a bit over a wire on a chip, where the introduction of intermediate repeaters can help to increase the performance and energy efficiency.

In its simplest form, multihop network energy analysis argues for an infinite number of hops over the smallest possible distance. In reality, however, the number of intermediate hops is limited by the number of nodes between A and B. Moreover, we must include not only the energy radiated through the antenna, but also the energy dissipated in the radio for receiving the bit and readying the bit for retransmission. (Given the relative costs of transmission and processing, we can compute an optimal number of hops.)

This leads to some interesting observations:

- Technology scaling will gradually reduce the cost of processing, with transmission cost remaining constant. Thus, shorter hops will become more favorable over time.
- Computation cost is not a constant either. Using compression techniques, we can reduce the number of transmitted bits, thus reducing the cost of transmission at the expense of more computation. This only makes sense if the communication cost dominates, as with long-distance connections.

This *communication-computation trade-off* is one of the core ideas behind the low-energy networks we propose. The optimal trade-off has to be determined adaptively, based on data properties, node densities, and environmental circumstances. This dynamic nature has a profound impact on the hardware composition and architecture of the network nodes.

Energy trade-offs in network protocols

Establishing multihop networks seems to be the ideal way of transmitting a bit in an energy-efficient fashion. Yet some major caution is necessary. Nodes cannot know a priori the optimal route to other nodes because this path changes as nodes move, enter, or leave the network. Therefore, the network protocol coordinates the discovery and tracking of routes in the network. This discovery and tracking consumes energy because it requires communication between nodes. With the low data rates and the relatively fast dynamics of some nodes, the network discovery and maintenance overhead may well dominate the energy

consumed for data transmission itself. This is actually the case for a large number of the ad hoc networking protocols currently in vogue. The general ways to do this tracking and discovery are proactive and reactive routing.

Proactive routing. In proactive routing, the network layer periodically updates routes, and hence always has an up-to-date picture of the optimal routes. A proactive network finds the routes between many nodes at once in an efficient manner. Thus, it consumes less energy than finding each particular route separately. When it needs to transmit a packet of real data, a proactive network knows the route and it sends the data with little extraneous network activity. In a sense, periodic updates generate a fixed amount of traffic, but any specific packet requires less network overhead.

Reactive routing. Reactive routing discovers routes only when the network needs them. In the reactive scheme, the network generally does not maintain routes until it uses them. With this method, periodic updates do not generate a fixed amount of traffic, but there is network overhead for each specific data packet or stream.

To communicate infrequently with a small number of nodes, there is no advantage to maintaining infrequently used routes, so a reactive approach is preferable. However, if the data rate is high and the network communicates with a large number of nodes, proactive routing is more desirable. Of course, we can use hybrid methods to optimize the network for a specific application.

Hybrid solutions

In fact, our application offers a major opportunity for reducing overhead. As mentioned earlier, sensor networks are best served by content- and localization-based addressing schemes—in which data in the network is accessed not through an absolute address, but through a query for information of a certain type in a certain location (similar to database queries). This approach avoids the setup and maintenance of extensive routing tables, relying instead on the broadcast propagation of queries, pruned by information content and geographical data. Research along these lines has yielded some very promising protocols, which may help to reduce the network overhead to a small fraction of the overall cost. Directed diffusion routing,² geographical routing,³ and swarm-intelligence⁴ are just a few of the techniques to watch.

Energy trade-offs at the MAC Layer

The MAC layer affects the energy efficiency in a number of ways.

First, MAC-layer power management can minimize the standby power of the network—that is, the power

consumed by a radio when it is not transmitting or receiving. Standby power is typically much lower than transmit or receive power. Yet, with the low data rates of the sensor nodes, radios have a small active duty cycle, and standby power easily dominates the overall power dissipation. We can virtually eliminate standby power by putting the radio in sleep-mode when nonactive, powering down all but a few functions. This poses the interesting problem of establishing a coherent network in which most of the nodes are solidly asleep most of the time.

Second, careful control of access to the aether reduces the number of wasted (re)transmissions corrupted by interference from neighboring nodes in the network. We can accomplish this by carefully separating potentially conflicting transmissions in time and frequency/code space.

To reach our ultra-low energy target, we need a MAC protocol that lets radios sleep most of the time and yet lets them awaken precisely when they need to transmit or receive data. (Most networking protocols assume that network nodes are “in listening or receive” mode when nonactive. This is a costly assumption because the computational energy for receiving a bit supercedes the cost of transmitting one—ignoring small amounts of radiated energy for short-distance communications). Unfortunately, current radio technology does not easily allow a radio to be awakened upon request. Hence, a radio must wake up periodically, see if anyone wants to talk to it, and, if not, go back to sleep. A mechanism that allows the radio to be awakened precisely when there is data for it could reduce the awake-time and hence the overall node energy consumption.

Our specific application domain once again comes to the rescue. Most “content-addressed” datagrams require broadcasting to all the neighboring nodes. Waking all neighboring nodes at once is easier than waking one specific neighbor. The reactive broadcast channel, thus established, can also advertise information regarding specific communication times and channels for point-to-point connections, hence reducing interference. Our estimates indicate that this MAC protocol can help reduce overall energy consumption by at least a factor of 50 (assuming radio duty cycles below 1 percent).

PICONODE IMPLEMENTATION

Implementing network optimizations requires a platform that fulfills the demanding low-power requirements yet has enough flexibility to enable the dynamic reconfiguration and adaptability the network requires. We are conceiving an architecture that attempts to satisfy these challenging requirements.

The media access layer can help reduce overall energy consumption by at least a factor of 50.

The PicoNode architecture illustrated in Figure 3 aims to provide both flexibility and low energy. This architecture is composed of four modules:

- an embedded processor subsystem for application- and protocol-stack layers, which require more flexibility but have low computational complexity at relatively low update rates;
- configurable processing modules for the protocol's more speed-intensive layers;
- a parameterized and configurable digital physical layer; and
- a simple direct-down conversion RF front end.

A flexible interconnection scheme that is optimized for low-power operation connects the modules.

This architecture is inspired by observations regarding the evolution of technology and new trends in energy-efficient architectures. We also have addressed

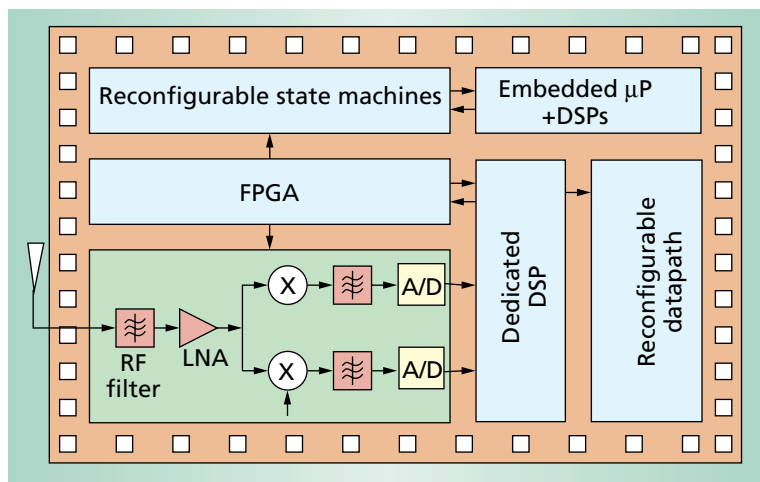


Figure 3. PicoNode conceptual architecture.

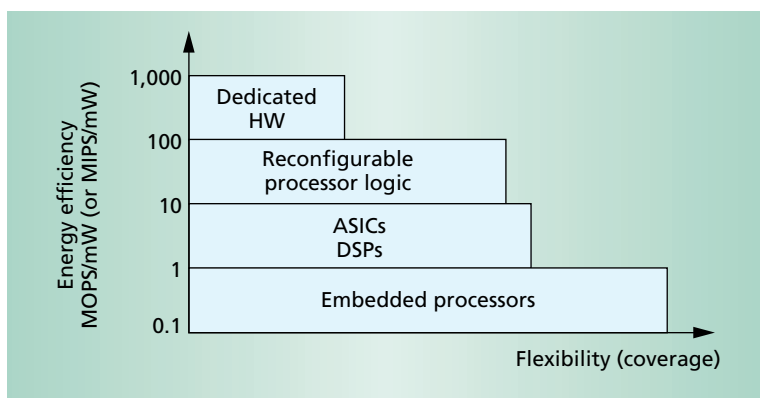


Figure 4. Flexibility versus energy trade-off in implementation (for a 0.25-micron CMOS technology).

the partitioning of operations between the processor and configurable processing modules and the digital physical layer and RF front end.

Digital back end

Some of our previous efforts in the area of reconfigurable computing⁵ have demonstrated that a dynamic matching between application and architecture leads to spectacular energy savings for signal-processing applications while maintaining implementation flexibility.

The goal here is to establish that a similar scenario holds for protocol- and network-oriented applications. As Figure 4 shows, the key insight is that processor implementation is three orders of magnitude more expensive (in terms of energy consumption) than implementation on dedicated hardware. Systems designers can trade flexibility and programmability for energy consumption in a continuum from hardware to software. Designers must match system elements individually to implementation platforms in order to reach the optimum in the energy-flexibility, space. In addition, the overall architecture must weave together these heterogeneous blocks with a flexible, yet energy-efficient interconnect scheme.

Reconfigurable architectures—which utilize a “programming-in-space” approach⁶—have proven to be very efficient for the signal-processing component of communications processing, delivering high-performance computation at an energy cost close to custom hardware implementations and maintaining enough flexibility to adapt to the varying conditions of the system and environment. We are currently establishing that the same holds for the control-oriented component of the transceiver—the protocol stack. Experiments have shown that FPGA and configurable finite-state-machine implementations of the protocol stack are two orders of magnitude more efficient than embedded microprocessor or microcontroller solutions.

Physical layer implementation

As Figure 2 shows, analog circuitry represents only a small fraction of the PicoNode implementation. Traditionally, wireless transceivers were almost completely implemented using RF and analog circuit modules. A mostly digital approach is currently coming into vogue. This trend is inspired by the observation that digital circuitry improves exponentially with the scaling of technology, while analog circuits get linearly worse, mostly due to reduction of the supply voltage. Hence, it is better to rely on a small, noncritical analog front end and use digital back-end processing to correct for the nonidealities.

Although the data rates are on the order of one hertz, which seems trivial when compared to the state-of-the-art radios reaching rates of hundreds of mega-

Table 1. A comparison of energy sources.

Energy source	Power density	Energy density
Batteries (zinc-air)		1050 -1560 mWh/cm ³
Batteries (rechargeable lithium)		300 mWh/cm ³ (3 - 4 V)
Solar (outdoors)	15 mW/cm ² (direct sun) 0.15mW/cm ² (cloudy day)	
Solar (indoors)	0.006 mW/cm ² (standard office desk) 0.57 mW/cm ² (< 60W desk lamp)	
Vibrations	0.01 - 0.1 mW/cm ³	
Acoustic noise	3E-6 mW/cm ² at 75 Db 9.6E-4 mW/cm ² at 100 Db	
Passive human-powered systems	1.8 mW (shoe inserts)	
Nuclear reaction	80 mW/cm ³ 1E6 mWh/cm ³	

* Values are estimates taken from literature, analyses, and a few experiments.

** Values are highly dependent on the amplitude and frequency of the driving vibrations.

hertz, there are some challenges at the physical layer, mostly related to the low-energy targets and variable demand from the network.

To satisfy variable demand from the network, the PicoNode physical layer must be parameterizable. Parameters include power control modes, modulation scheme, and bit rate. The goal of this step is to define the parameters and identify the ranges for those parameters that allow limited flexibility at the physical layer without affecting energy consumption.

In order to achieve low-energy operation, the physical layer must meet two usually mutually exclusive goals: fast signal acquisition and low standby power. Many radio architectures expend a bit of energy during standby in order to achieve a fast lock once awoken or, conversely, incur a longer acquisition time in order to sleep more deeply during standby. In order to make the less than 1 percent duty cycle a reality, the physical layer must be able to quickly acquire a signal when awoken, receive the burst of data, and then immediately return to sleep. On the contrary, since the radio is asleep 99 percent of the time, the standby power of the radio must be very low in order to meet the energy budget. These goals can only be met by advancing the state of the art in low-energy physical layer design.

ENERGY SCAVENGING

Our project's Holy Grail is for the PicoNodes to be self-contained and self-powered using energy extracted from the environment (energy-scavenging). Reaching this goal requires new advances both in reducing the nodes' energy consumption and in increasing the amount of energy the nodes can extract from the environment.

Harvesting ambient energy requires compliance with two major constraints: applicability within the

environments envisioned for the PicoNodes (office buildings and homes) and the size constraint of the one-cubic-centimeter chip. Table 1 shows a comparison of energy sources based on a combination of published studies, theory, and experiments.

Although batteries can store harvested energy that can't be used immediately, a continuous source of energy is desirable. Table 1 shows that solar cells can contribute up to 15 milliwatts per square centimeter during direct sunlight hours and up to 0.15 milliwatts on cloudy days. Averaging over daylight and nighttime hours, and considering nodes in the interior of the building or embedded in ceiling tiles, shows that solar cells can just barely serve as the sole energy source for PicoNodes, and additional sources of energy would be welcome.

Harvesting energy from vibrations is promising for this application. Raised floors and dropped ceilings in most office buildings exhibit measurable vibrations (from trucks driving down nearby streets and people walking on the raised floors) that can be harnessed. Advances in MEMS devices make integrated and tiny variable capacitors a reality. These capacitors are used to make chip-scale electrostatic vibration generators that will integrate well with the other PicoNode components. Power outputs between 10-100 microwatts per cubic centimeter are plausible from vibrations in a normal office building using existing MEMS technology.

In our vision of distributed computing, thousands of tiny nodes scattered throughout the daily living environment gather, process, and communicate information in a self-organizing fashion.

The major challenge in the implementation of these wireless ad hoc sensor, monitor, and actuator networks is minimizing energy consumption.

As we show, the only way to implement an ultra-low power node is by optimizing all layers of the protocol. We present a configurable architecture that enables these opportunities to be efficiently realized in silicon. We believe that this energy-conscious system-design and implementation methodology will lead to radio nodes that are two orders of magnitude more efficient than existing solutions. *

Acknowledgments

We acknowledge the contributions of the PicoRadio Group of the Berkeley Wireless Research Center. This research is sponsored by DARPA as a part of the PAC/C program, and by the MARCO GSRC research consortium.

References

1. "The Berkeley Wireless Research Center PicoRadio Project;" http://bwrc.eecs.berkeley.edu/Research/Pico_Radio.
2. C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *Proc. ACM MobiCom 2000*; <http://netweb.usc.edu/estrin/papers/diffusion-mobicom-submitted.ps>.
3. R. Jain, R. Puri, and R. Sengupta, "Geographical Routing Using Partial Information for Wireless Ad-hoc Networks," University of California, Berkeley, Dec. 1999; <http://www-path.eecs.berkeley.edu/~rjain/geog-routing99.ps.gz>.
4. E. Bonabeau and G. Theraulaz, "Swarm Smarts," *Scientific American*, Mar. 2000, pp. 73-79.
5. "The Pleiades Reconfigurable Processor Project;" http://bwrc.eecs.berkeley.edu/Research/Configurable_Architectures.
6. H. Zhang et al, "1V Heterogeneous Reconfigurable Processor IC for Baseband Wireless Applications," *Proc. ISSCC Conf.*, IEEE Press, Piscataway, N.J., 2000, pp. 68-69.

Jan M. Rabaey is a professor and vice-chair in the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley. His interests include virtually all aspects of the design of electronic circuits and systems, with special focus on the exploration of algorithms, architectures, and implementation of signal processing systems and the related design methodologies. Current research is directed toward various aspects of portable, distributed multimedia systems, including low-power design and wireless networking.

M. Josie Ammer received a BS and an MEng in computer science and electrical engineering from MIT. Currently, she is pursuing a PhD in electrical engineering at the University of California, Berkeley. Her research interests include energy-efficient physical layer architectures and implementations.

Julio L. da Silva, Jr. is a postdoctoral student working with protocol specification, design, and implementation at the Berkeley Wireless Research Center. He received a PhD from Katholieke Universiteit Leuven. His PhD research focused on design methodologies for data-dominated applications.

Danny Patel received a BS in electrical engineering from North Carolina State University and an MS in electrical engineering from the University of California, Berkeley. His research interests include energy-efficient routing in ad hoc networks.

Shad Roundy received a BS in mechanical engineering from Brigham Young University and an MS in mechanical engineering from the University of California, Berkeley, where he is currently a PhD candidate. His research focuses on efficient conversion of vibrations to electrical energy.

Contact the authors at {jan, mjammer, julio, dpatel, shadrj}@bwrc.eecs.berkeley.edu.



The Computer Society

publishes over 150
conference proceedings
a year.

For a preview of
the latest papers in
your field, visit

computer.org/proceedings/