

ULTRA-LOW-POWER DESIGN

The Roadmap to Disappearing Electronics and Ambient Intelligence

Successful deployment of wireless sensor and actuator networks (WSNs) in numbers large enough to provide true ambient intelligence [1] requires the confluence of progress in several disciplines, including distributed computing, networking, wireless communications, and, most importantly, ultra-low-power design.

J. Rabaey, J. Ammer, B. Otis, F. Burghardt, Y. H. Chee, N. Pletcher, M. Sheets, and H. Qin

The latter is an essential cornerstone to the ultimate success of this innovative and paradigm-shifting concept for a number of reasons. Foremost, physical access to nodes in a deployed network may be difficult, and the sheer number of nodes makes changing batteries or manual refueling unrealistic. Hence, nodes should be energy self-contained for the lifetime of the application. In addition, node size is dominated by the energy storage and generation modules. Improved energy scavenging and lower power consumption result in smaller nodes, opening the door for a wider range of applications. Finally, energy storage and generation is responsible for a sizable fraction of the cost of a node. Ubiquitous deployment of large networks requires node costs substantially below what is available today.

In a 2000 *IEEE Computer Magazine* article [2], the authors defined the concept of the PicoNode and outlined what it would take to create a wireless sensing, communication, and computation node that would consume less than 100 μW average, be smaller than 1 cm^3 , and cost less than US\$1. It was established that achieving these goals would require a coordinated power-optimization effort across every level of the design hierarchy (application, network, media access, physical layer, computation versus communication) and every single component of the node. In this article, we describe how such an integrated approach has indeed made it possible to produce a PicoNode that meets the original goals (Figure 1). The resulting node combines innovative technologies, such as radio-frequency microelectromechanical systems (RF-MEMS) with ultra-low-power RF and digital integrated circuit (IC) design, and employs aggressive energy-scavenging and packaging techniques. For these technological advances to come to their full fruition, they must be complemented by novel opportunistic networking and wireless protocol schemes that virtually eliminate standby power while still providing robustness.

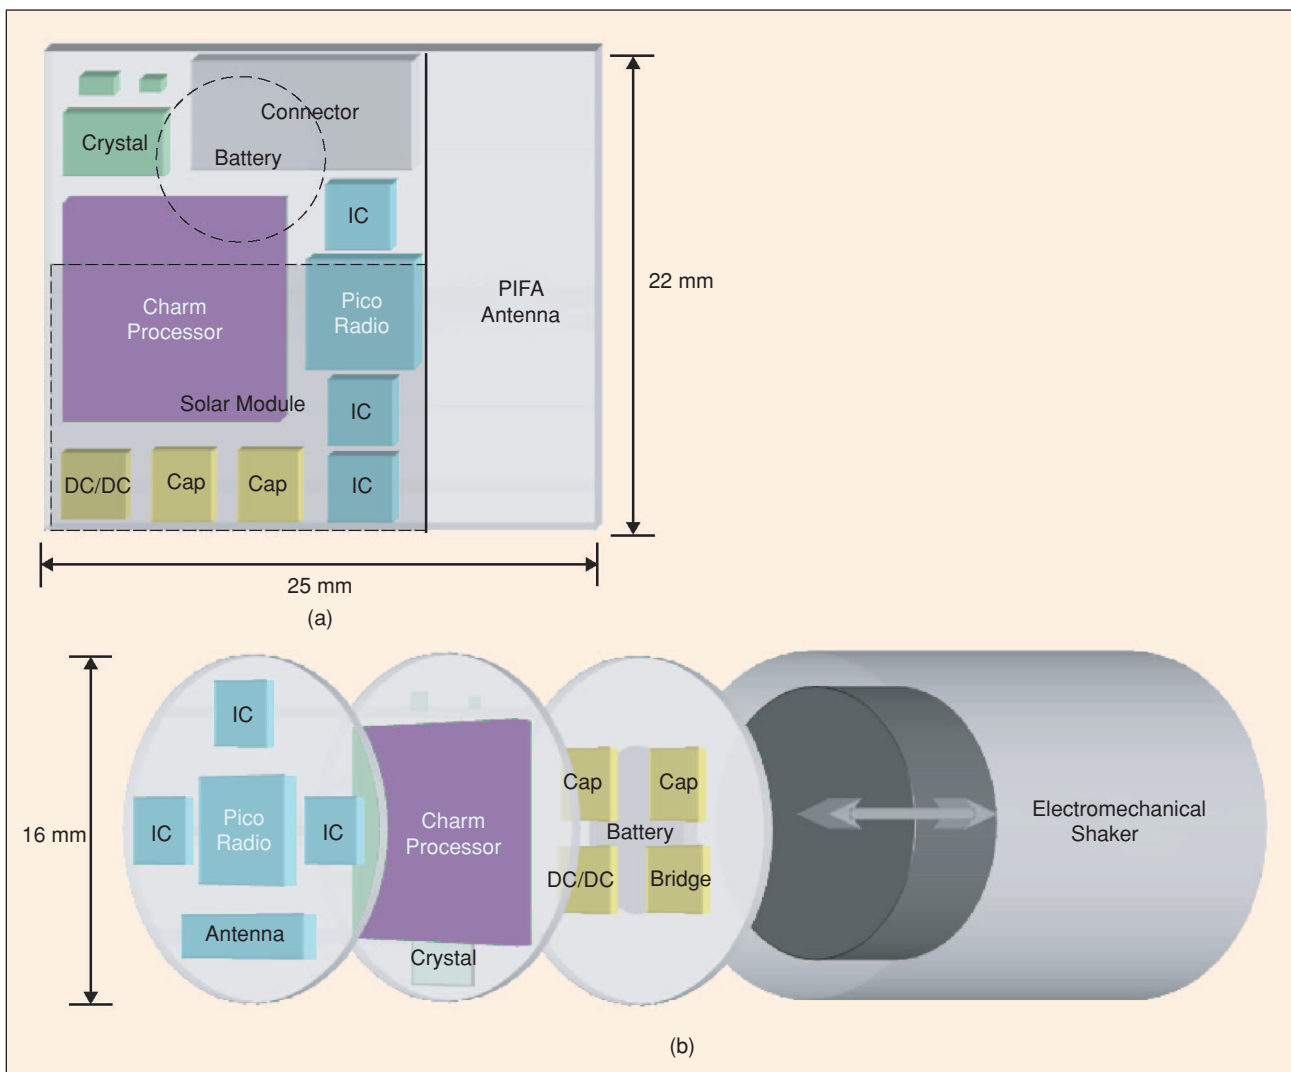
Providing a comprehensive overview of all of these techniques within the scope of a couple of pages is obviously impossible. Instead, we chose to highlight some of the elements that make a true difference in the creation of the ultra-low-power PicoNode.

POWER TRAIN

One of the main design goals is to power the node completely through the harvesting of energy from the environment. Because power available from scavenging is severely limited, efficiency in the power train is paramount. To increase the applicability, a combination of energy-scavenging methods tailored to the deployment environment can be used. Two methods we have studied in detail are light and vibration energy scavenging. Solar cells produce $10\text{--}100\ \mu\text{W}/\text{cm}^2$ in typical office environments, while prototypes of vibration scavengers have been able to supply $200\ \mu\text{W}/\text{cm}^3$ for ideally matched conditions. For applications with large displacements, a meso-scale electromechanical shaker that produces $20\ \mu\text{W}$ at $4\ \text{V}$ was developed.

Efficient storage of the harvested energy is equally important. Ceramic capacitors are commercially available with high values in small packages. Thin-film batteries offer other opportunities, but are still in the development phase. The storage device is charged with the rectified output of the energy scavenger, so if a battery is used the operating voltages of the battery technology must be matched with the energy-scavenging output voltage and impedance (about $4.5\ \text{V}$ and $3.7\ \text{k}\Omega$ for the most recent vibration-harvesting prototypes).

A final and crucial power train challenge is to convert the voltage of the storage node into the various supply needs of the system components, which diverge quite dramatically. The current prototype node requires $2.1\ \text{V}$ (digital), $1.0\ \text{V}$ (sensors), and $0.5\ \text{V}$ (RF). Previous node prototypes using the best available off-the-shelf converters saw power train losses of up to 50%. These concerns are overcome by using a combination of strategies: 1) Switched capacitor charge pumps are relatively efficient at low current levels, can adapt to varying current needs by changing the clock frequency, and require relatively few external components. 2) Since these circuits



1. Various incarnations of PicoNode: (a) flat aspect ratio with solar cells and (b) stacked boards with electromechanical shaker as energy source.

have relatively high output impedance, the amount of current they can provide while maintaining regulation is limited. To overcome the current requirements for the digital and RF components, which often occur in bursts, large-capacity bypass capacitors must be added at each supply voltage.

THE WIRELESS TRANSCEIVER— THE PHYSICAL LEVEL

The physical layer provides a wireless communication link with the range (10 m) and reliability (bit-error rate of at least 10^{-4}) required by typical WSN applications, supporting packets of up to 512 b. In addition, efficient media access also requires the availability of a carrier sense function [3].

Transceiver

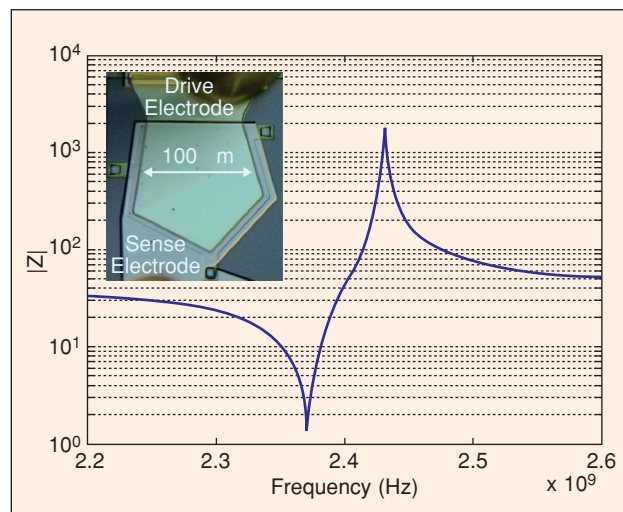
One of the great challenges of implementing ultra-small sensor nodes is the wireless link. In general, the link must be bidirectional, which rules out passive radio frequency ID (RFID) tag receivers. In contrast to traditional wireless networks, it is the receiver which dominates the overall power consumption [3]. Data traffic in sensor networks is fairly rare, meaning that sensor nodes typically spend substantial time listening to an empty channel. At the same time, the receiver should provide sufficient sensitivity, limiting the required transmitter power and the associated interference. In addition, the transceiver must be physically small. A traditional implementation with off-chip components and a quartz resonator is prohibitively large and expensive, so an entirely thin-film, batch-fabricated solution is necessary.

In recent years, MEMS technology has become an increasingly mature technology. These devices, fabricated directly on top of silicon wafers, enable high quality factor resonances ($>1,000$) and tightly controlled resonant frequencies. Measured data from such a device [a 2.4-GHz bulk-acoustic wave (BAW) resonator] is shown in Figure 2.


The sharp peaks in this plot clearly show the high Q nature of the resonator. These devices have been used for steep band-

pass filters and duplexers. Here, we utilize the high Q and stable oscillation to create ultra-low-power oscillators, which combined with innovative architectures are an essential component of any narrow-band radio. By performing the filtering and RF frequency stabilization with a passive mechanical structure instead of active circuitry, substantial power savings are realized. In addition, MEMS components make it possible to integrate the entire transceiver into a few cubic millimeters. Wafer thinning and flip-chip assembly can reduce this volume further. Ultimately, this technology can be fabricated on the same wafer as submicron CMOS circuitry, creating a truly micro-scale transceiver implementation.

This methodology was used to design a low-power super-regenerative transceiver [5], as shown in Figure 3. This design used two high-frequency 1-mm² BAW resonators alongside a 2-mm², 0.13- μ m CMOS chip. The receiver consumes less than 400 μ W when active and supports data rates up to 50 kb/s. A three-wire serial peripheral interface (SPI) bus controls all



2. Measured impedance of single 2.4-GHz BAW resonator. Inset photograph shows the 100 μ m \times 100 μ m active resonator area [4].



3. CMOS/BAW resonator transceiver implementation. The CMOS chip is flanked by two BAW resonators, which can be wire-bonded directly or assembled by chip-chip flip-chip technology.

receiver functionality, including variable data rate, RF chain bias current, baseband gain, and baseband bandwidth. A sensitivity of approximately -100 dBm allows for a very low transmitted power level. The design techniques, demonstrated for low-power sensor networks, can be applied with similar benefits in the high-performance design space as well.

While minimizing the receiver dissipation is essential, the transmitter should not be ignored. For a communication range of 10 m, the required transmit power is less than 1 mW. At such low transmit power levels, the energy efficiency of the overall transmitter is not determined by the power amplifier only but also by all the supporting circuitry such as synthesizers, modulators, and antenna matching. Accordingly, a simple modulation scheme, such as on-off keying (OOK), is employed to simplify the transmitter topology. Further improvements in efficiency can be accomplished through architectural and circuit innovation. For instance, converting the power amplifier into a power oscillator reduces the load on the previous stages, as the oscillator is self-driven. Frequency stability and accuracy is obtained by injection-locking to a BAW reference oscillator. Such an injection-locked transmitter achieves an overall efficiency of 32% while delivering 1 mW to the antenna [6]. Even higher efficiencies (up to 48%) can be obtained by codesigning the power amplifier and the antenna [7]. These numbers should be compared to the efficiencies obtained for industrial designs (for the 802.15.4 standard), which range between 3–5%.

Baseband Processing

Synchronization consists of estimating channel parameters so that symbols can be detected. The term “channel” is loosely defined to include impairments in the transmit and receive front-ends, as well as the traditional wireless channel (over the air). WSNs use packet-based communication where the synchronization parameters need to be re-estimated for every packet due to the large separation between consecutive packets. Packet lengths for typical applications range from 30 b for control packets to 200 b for typical data packets with a maxi-

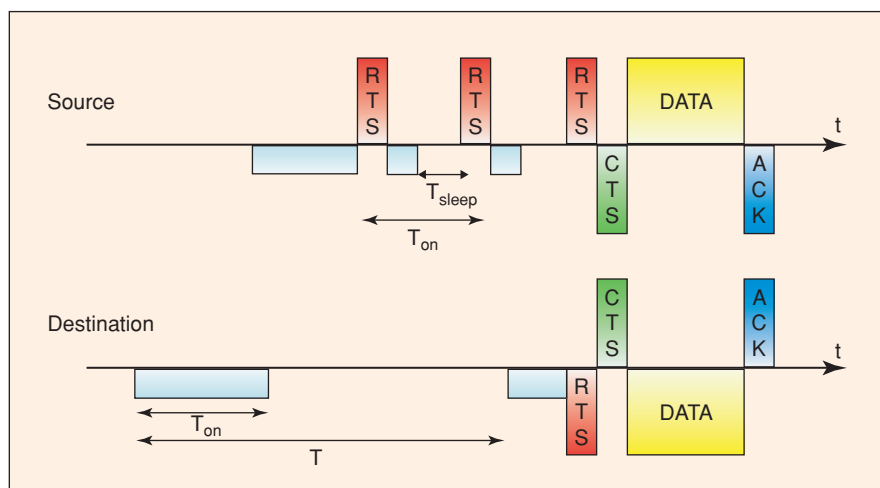
mum of 500 b [3]. The small packet lengths require short synchronization preambles. For instance, the 128-b preamble of the 802.11b protocol would impose over 300% overhead for control packets.

A tradeoff of packet length, data rate, clock accuracy, and modulation scheme leads to simplified synchronization requirements. Implementation schemes across the digital/analog boundary were explored in [8]. It was determined that the digital algorithm always results in lower system power because of its shorter preamble. The chosen digital synchronization scheme uses an analog-to-digital converter (ADC) that oversamples the signal from the RF front-end by a factor of ten. Amplitude threshold and timing are estimated using maximum-likelihood feed-forward data-aided estimators requiring a preamble length of 18 symbols. During data reception, symbols are matched filtered and sliced using the estimated threshold. The total implementation loss is 1.3 dB, thereby achieving better than $1e-4$ bit error rate (BER) at 13 dB input signal-to-noise ratio (SNR). The power consumption of the synchronization system (including the ADC) is under $320 \mu\text{W}$ (when on).

THE DATA LINK LAYER

The data link protocol is responsible for transmitting packets between neighbors. It performs dynamic node ID assignment, maintains a table of neighbors, and manages media access control (MAC). Nodes that are in direct radio contact with each other are said to be in the same neighborhood. Node IDs must be unique in any neighborhood in which the node resides as well as any neighborhood containing its neighbors (one-hop neighbors). In a dense network, neighborhoods intersect, and with dynamic fluctuations in the channel conditions (as is typical in RF wireless), membership in a neighborhood changes over time. Essentially, membership in a neighborhood implies that there is “acceptable” communication between all nodes in the neighborhood. In certain situations, such as a lost packet or hidden nodes, the assumption of local ID uniqueness may be violated. A detection mechanism is used to identify and correct this situation.

One of the most straightforward methods for reducing the overall power consumption of the wireless link is to maximize the amount of time the transceiver spends in sleep mode. One approach to accomplish this is through duty cycling; that is, each node turns on its receiver periodically to monitor the channel for transmissions. Two different styles of such a cycled receiver [called transmitter or receiver-initiated cycled receiver scheme (TICER and RICER)] were explored for the PicoNode. An alternative option is to add an ultra-low-power carrier-sense module—called a wake-up receiver—to the node, the task of which is to awaken the main transceiver when it detects activity on the channel [5].



4. TICER helps reduce the standby power of the node when no active transmissions are happening.

The Cycled Receiver

To explain the idea of the cycled receiver, consider the following scheme called TICER. When a sensor node has no data to transmit, it wakes up to monitor the channel with period T and goes back to sleep after a wake-up duration T_{on} , as shown in Figure 4. As soon as the node needs to transmit data, either generated earlier from the upper layers of the protocol stack or forwarded by another node, it wakes up and monitors the channel for the duration of T_{on} . If the channel is free, it starts transmitting requests-to-send (RTSs) to the destination node repeatedly and it monitors the channel for T_l seconds and then goes to sleep for T_{sleep} seconds after each RTS transmission. The destination node, upon waking up according to its regular wake-up schedule, acquires and receives the RTS and responds with a clear-to-send (CTS) signal to the source node. The source and destination nodes then follow the CTS-DATA-ACK handshake. In an inverse scheme (RICER), one can imagine a destination node waking up periodically and transmitting a CTS requesting for data to come its way.

The cycled-receiver synchronization technique leads to substantial power reduction, virtually eliminating the channel monitoring power. How much can be gained depends upon the observed traffic and the required latency.

Carrier Sensing for a Reactive Radio Architecture

The cycled-receiver scheme, while effective, leads to latency in the network and some power inefficiency, as the receiver must be periodically active to listen for packets even if no data is being transmitted. An alternative approach is continuous monitoring of the channel using a separate physical receiver, called a carrier-sense receiver (Figure 5). This receiver operates in tandem with the main data radio, listening for data transmissions and only waking up the data receiver when packets are being transmitted. We call this architecture a reactive radio, where the main data receiver remains in standby mode and only reacts when data is sent to the node [2].

Because the carrier-sense receiver is continuously monitoring the channel, network latency is significantly reduced. However, to save power over the cycled-receiver approach, the carrier-sense receiver must meet a challenging power budget of less than $100 \mu\text{W}$. The choice of circuit topologies that may operate with very low supply voltages, even below the nominal V_{dd} of the CMOS technology, is one potential power-saving technique. In addition, current consumption may be reduced through subthreshold transistor biasing, which is becoming viable in deeply scaled CMOS due to the increased speed of the devices. Recent results have shown that power dissipation levels of around $50 \mu\text{W}$ for a wake-up receiver are indeed obtainable.

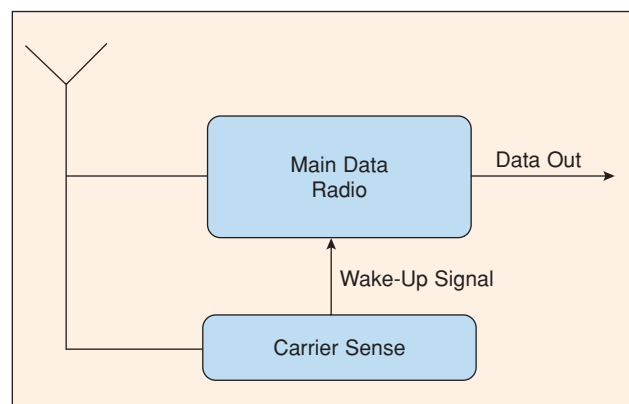
THE NETWORK LAYER

At the network level, PicoNodes communicate with each other in an ad hoc manner with no central point of contention or failure. The network protocol is responsible for global end-to-end routing—that is, from requestor to sensor node and vice-versa. A network path can be many “hops”

long. The protocol uses a form of the directed diffusion algorithm [9]. A requesting node sends a broadcast interest packet directed toward a three-dimensional (3-D) box-shaped region defined by the locations of its opposite corners, and the request is forwarded in that direction by other nodes along the way. As the request propagates outward through forwarding, each node collects data on the most immediate source of that request (a neighbor). Since request propagation depends on directed broadcast forwarding, a node may see multiple requests from the same source arriving on different paths (from different neighbors). Included in this data are an energy metric as well as the location of the node that originally requested information. The node records the energy and source address for each neighbor ID. When responses in the form of data packets are returned from the nodes within the target region, the network protocol knows which return paths are available and chooses a path based on energy and a probability algorithm that ensures no path is overused. Modifications of this algorithm were developed to improve the reliability and energy efficiency. Nodes in the network can appear and disappear randomly, and a good network layer should be able to manage these changes without a glitch.

LOCATIONING

The network routing algorithm works only because all nodes know their relative position in space. In a self-configuring sensor network, most nodes are deployed without any presumed position. Only a small number of anchor nodes are given a priori information about their positions with respect to a global coordinate system. The rest of the nodes then calculate their positions using the position and distance to these anchors. The locationing subsystem performs 3-D triangulation using at least four anchor node reference points. If more anchor nodes are present in the network, the error in the calculation is reduced. During triangulation, internodal distances are expressed in terms of the reference and unknown position coordinates. These equations are then cast into an overdetermined linear system, which is solved using a least-squares algorithm. The hop count between an anchor and the node is



5. Radio architecture utilizing a carrier sense radio. This radio continuously monitors the channel, waking up the data radio only when a neighboring node is initiating communication.

used as a substitute for the real Euclidean distance, as initially proposed by Savarese [10]. The sensor network is assumed mostly static, so only slow position changes are expected and tracked. The location update rate is programmable and usually on the order of minutes.

Improved energy scavenging and lower power consumption result in smaller nodes, opening the door for a wider range of applications.


new packet or the expiration of a timer. After an event is processed, the node is otherwise idle. This article uses algorithms, circuits, and architectures to exploit this event-driven burst behavior to minimize power consumption in idle mode and maximize the time spent idling.

PROTOCOL STACK IMPLEMENTATION

Background

A common observation of sensor nodes is that they have a low average activity factor. This work further notes that this activity occurs in bursts after specific events, such as the reception of a

- ◆ Standby power reduction tackles the large fraction of power wasted by leakage currents in the digital processors.
- ◆ Energy scavenging slowly recharges the energy store for a burst data transmission.
- ◆ Network routing perturbs as few nodes as possible and can choose different routes to maximize network survivability.
- ◆ Reactive radio only activates the node when a packet is present.




6. Die photograph of Charm protocol processor.

Protocol Processor ASIC

The Charm ASIC implements the complete digital protocol stack for a PicoNode [11]. The chip contains a synthesized 8051-compatible embedded microcontroller with 64 kB of program/data random access memory (RAM), two 1-kB packet queues, a custom data-link layer, a neighborhood management subsystem, the digital portion of a custom baseband, a location computation subsystem, and one custom and several standard external interfaces. The chip is implemented in a 0.13- μ m, triple-well process, and the 2.7×2.7 mm² die photo is shown in Figure 6.

The chip components are divided into eight primary subsystems called power domains (PDs). The mode of each PD can be separately set to active, standby, or sleep, enabling the power manager scheduler to place each subsystem in the lowest power mode that still meets the performance requirements. In general, the PDs follow the natural layers of the protocol stack, which allows most layers to sleep while others



7. Actual PicoNode realizations. (a) Flat form factor. (b) PicoCube, a modular node implementation with cubic form factor.

are processing packets. Additional PDs, such as packet queues and a neighborhood information table, are used when multiple subsystems need access to the same data at different times. This decouples the generation and consumption of data and allows the accessing subsystems to remain in a lower power mode for a longer time.

Within each PD, the Charm chip uses power-rail gating to implement a sleep mode that significantly reduces leakage power consumption. In this method, the voltage on the power rail is lowered, which reduces the leakage current. To simplify the restoration to active mode, the PDs are designed to retain their state during sleep mode so the standby supply voltage can only be reduced to the data retention voltage (DRV), below which the state is lost. Due to the effect of process variations and temperature, the DRV is expected to range between 60–390 mV for different die. Measurements, taken on a 4-kB SRAM at 100 mV above the minimal DRV, show that even a conservative standby voltage can reduce leakage 85% in the SRAM components [12]. An on-chip voltage regulator produces the standby voltage from the nominal 1-V external supply.

Another significant power-saving feature is a low-power clocking scheme. At the architectural level, a centralized system timewheel enables blocks to schedule a future wake-up event and sleep until it occurs. The clock tree itself is heavily gated, both at the leaf (register) level and trunk (subsystem) level, which reduces switching activity for unused (sleeping) portions of the clock tree. Lastly, the custom crystal oscillator circuit generates a suitable clock with $<10 \mu\text{W}$ power consumption, which is over 100 times better than general-purpose off-the-shelf oscillators. This savings is achieved at the expense of a guaranteed absolute center frequency, since successful communication depends more on the relative frequency offset between nodes.

NODE INTEGRATION

Bringing together the diverse components that make up the PicoNode in a small volume and at low cost requires aggressive and innovative packaging strategies. Using a number of diverse assembly approaches, we have shown that nodes of about 1 cm^3 can indeed be realized in different form factors, as illustrated in Figure 7.

CONCLUSION

When we embarked on the PicoNode project, it was by no means obvious that a sub-100- μW average power energy-scavenged node would really be achievable. Careful examination of the main components of the power budget and the introduction of the wide variety of power-reduction techniques described above have ultimately established that such a node is indeed possible with today's technologies. In fact, we now believe that even lower numbers can be accomplished. Truly "disappearing" electronics are in the realm of the possible and will ultimately change the way the environment and humans interact with one another (that is, ambient intelligence). The main stumbling block is still cost.

Cheap and simple assembly as well as reduced cost of the individual components (such as energy scavengers and energy storage devices) are necessary. Many opportunities are still left to be explored and will make for exciting research in the years to come.

ACKNOWLEDGMENTS

Portions of this work are funded by the MARCO Gigascale Systems Research Center focus center, Defense Advanced Research Projects Agency, the California Energy Commission, and the member companies of the Berkeley Wireless Research Center. The authors thank ST Microelectronics for chip fabrication.

REFERENCES

- [1] F. Boekhorst, "Ambient intelligence, the next paradigm for consumer electronics: How will it affect silicon?," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2002, vol. 1, pp. 28–31.
- [2] J. Rabaey, J. Ammer, J. da Silva, D. Patel, and S. Roundy, "Picoradio supports ad-hoc ultra-low power wireless networking," *IEEE Comput. Mag.*, vol. 33, no. 7, pp. 42–47, July 2000.
- [3] En-Yi A. Lin, Jan M. Rabaey, and A. Wolisz, "Power-efficient rendezvous schemes for dense wireless sensor networks," in *Proc. Int. Conf. Communications*, Paris, France, 2004, vol. 7, pp. 3769–3776.
- [4] R. Ruby, P. Bradley, J. Larson III, Y. Oshmyansky, and D. Figueredo, "Ultra-miniature high-Q filters and duplexers using FBAR technology," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2001, pp. 120–121.
- [5] B. Otis, Y.H. Chee, and J. Rabaey, "A 400 μW Rx, 1.6 mW Tx super-regenerative transceiver for wireless sensor networks," in *Int. Solid State Circuits Conf. (ISSCC 2005) Dig. Tech. Papers*, San Francisco, Feb. 2005, pp. 396–398.
- [6] Y.H. Chee, A.M. Niknejad, and J. Rabaey, "An ultra-low power injection locked transmitter for wireless sensor networks," in *Proc. IEEE 2005 Custom Integrated Circuits Conf. (CICC 2005)*, Sept. 2005, pp. 797–800.
- [7] Y.H. Chee, A.M. Niknejad, and J. Rabaey, "A 46% efficient 0.8 dBm transmitter for wireless sensor networks," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, June 2006, pp. 52–53.
- [8] M. Josie Ammer and J. Rabaey, "Low power synchronization for wireless sensor network modems," in *Proc. IEEE Wireless Communications and Networking Conf. (WCNC 2005)*, New Orleans, LA, Mar. 2005, vol. 2, pp. 670–675.
- [9] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," *IEEE/ACM Mobicom*, 2000, pp. 56–67.
- [10] C. Savarese, K. Langendoen, and J. Rabaey, "Robust positioning algorithms for distributed ad-hoc wireless sensor networks," in *Proc. USENIX Technical Ann. Conf.*, Monterey, CA, June 2002, pp. 317–327.
- [11] M. Sheets, F. Burghardt, T. Karalar, J. Ammer, Y.H. Chee, and J. Rabaey, "A power managed protocol processor for wireless sensor networks," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, June 2006, pp. 262–263.
- [12] H. Qin, Y. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey, "Standby supply voltage minimization for deep sub-micron SRAM," *IEEE Microelectron. J.*, vol. 36, pp. 789–800, Aug. 2005.

J. Rabaey, F. Burghardt, Y.H. Chee, N. Pletcher, M. Sheets, and H. Qin are with the Berkeley Wireless Research Center, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. B. Otis and J. Ammer are with the Department of Electrical Engineering, University of Washington. E-mail: jan@eecs.berkeley.edu. CD ■