


WiseNET: An Ultralow-Power Wireless Sensor Network Solution



The WiseNET platform uses a codesign approach that combines a dedicated duty-cycled radio with WiseMAC, a low-power media access control protocol, and a complex system-on-chip sensor node to exploit the intimate relationship between MAC-layer performance and radio transceiver parameters.

Christian C. Enz

Amre El-Hoiydi

Jean-Dominique Decotignie

Vincent Peiris

Swiss Center for Electronics and Microtechnology

A wireless sensor network consists of many energy-autonomous microsensors distributed throughout an area of interest. Each node monitors its local environment, locally processing and storing the collected data so that other nodes can use it. Network nodes share this information via a wireless link. Using data fusion, specific features of interest to the end user can be extracted from the information that several nodes collect while a multihop communication scheme propagates this information to a base station node.¹⁻⁴

Since these networks often are deployed in regions that are difficult to access, the nodes should not require maintenance. They must be energetically autonomous, using batteries that do not need to be replaced or recharged. In many application scenarios, the targeted node lifetime typically ranges from two to five years, imposing drastic constraints on power consumption. With a single 1.5-V AA alkaline battery, the average power consumption ranges from 100 to 10 microwatts, for a node lifetime ranging between two and seven years. Given that today's commercially available radio transceivers typically consume several tens of milliwatts, keeping the transceiver constantly active is clearly impossible. Maintaining the required power con-

sumption requires having the nodes sleep most of the time. This can be achieved by introducing duty cycling on the order of 0.1 percent to 1 percent, while keeping a low sleep-mode current no larger than the battery leakage current.

Reducing power consumption requires optimization across all layers, from the physical layer, channel coding, and media access control layer up through the routing, transport, and application layers. The MAC layer plays the most crucial role in the communication protocols' overall energy efficiency, especially for networks with low-duty-cycling radios.

To optimize power consumption, the Swiss Center for Electronics and Microtechnology has developed WiseNET, an ultralow-power platform for the implementation of wireless sensor networks that achieves low-power operation through a careful codesign approach. WiseNET combines a dedicated duty-cycled radio with WiseMAC, a low-power MAC protocol designed for low-duty-cycle wireless sensor networks.⁵ The WiseNET solution consumes about 100 times less power than comparable solutions available today.

WISEMAC PROTOCOL

At the outset, we analyzed WiseMAC's operation to identify the most important radio parameters

A sensor node can be integrated into a single system on chip to minimize power consumption and reduce the cost.

affecting power consumption. Given that available radio transceivers did not meet expectations and could not be modified, we designed a dedicated WiseNET radio transceiver that optimized these parameters for WiseMAC. This codesign approach—which optimizes overall power consumption by exploiting the intimate relationship between MAC layer performance and the radio transceiver parameters—is the heart of the WiseNET solution.

Typically, a microsensor node performs several functions, including

- sensing environmental physical parameters,
- processing the raw data locally to extract characteristic features of interest,
- storing this information momentarily, and
- using a wireless link to transmit the information to its neighbors.

The node must also operate as a relay for implementing multihop communication by receiving the data coming from one or several of its neighbors and then processing it before routing it to the next neighbor toward the destination.

To perform these functions, a sensor node—which includes many subsystems—can be integrated into a single system on chip to minimize power consumption and reduce the cost. Therefore, in addition to the radio optimized for the WiseMAC protocol, we developed a complete sensor node SoC that includes most required functions on-chip.

WIRELESS NETWORK ARCHITECTURES

Developers use wireless networks in a spectrum of applications that lie between two extremes: the infrastructure mode and the ad hoc mode. In the infrastructure mode, mobile nodes communicate through *base stations*, special nodes that link together through a conventional network.

When a mobile node wants to communicate, it first registers at a base station in the direct communication range. When a node wants to send a packet to another node, it sends it to the base station. If the destination is registered to that base station, the base station transmits the packet directly to the destination node. Otherwise, the base station forwards the packet through the infrastructure to the base station where the destination node has registered. Typical examples of this infrastructure mode include cellular telephony, paging systems, and wireless LANs that use IEEE 802.11.

In the ad hoc mode, there is no base station infrastructure. If the destination is in range of the source node, that node sends the packet to the destination

node. If the destination is not in range, the source node sends the packet to an intermediate node, which forwards the packet to other nodes until the packet reaches its destination or fulfills some other termination criterion. With this architecture, successive hops transport a packet from the source to the destination over several nodes—hence the term *multihop transmission*. PicoRadio is one example of such a network architecture.²

Both architectures assume that there is a way to find the route a packet must follow from its source to its destination. Defining this route is the subject of intensive research. Here, we assume that nodes are static enough for this problem to be considered a secondary issue.

The infrastructure-based architecture is popular for several reasons, particularly its relative simplicity. Base stations do not have power restrictions and enjoy a better spectrum usage because they allow frequency planning. On the other hand, advocates of ad hoc networks cite their higher versatility and potentially lower power consumption. Since no planning and no infrastructure are required, ad hoc networks can be deployed quickly and in remote areas.

Exploiting multihop communication can reduce the transmission range for each hop. Theoretically, as the attenuation increases at least quadratically with the distance, it is more beneficial from the power-consumption viewpoint to use two hops of length L than a single hop of length $2L$. While true in principle, with current transceiver efficiency, thermal power dissipation dominates radio-frequency-radiated power, which significantly diminishes the power advantage of multihop transmission.

Other hybrid solutions lie somewhere in between these two extreme architectures. For example, an infrastructure-assisted architecture would mix both approaches. A mobile node beyond a base station's range could use other nodes in the range to relay the packets to the base station. In this architecture, a packet may or may not go through the infrastructure depending on the location of the source and destination nodes. Similarly, two or more separate infrastructure-based networks can exchange packets through a sequence of hops on mobile nodes, thus forming a single network.

A building sensor system is an example of such a hybrid architecture. Rooms can be equipped with wireless sensors for temperature, moisture, gas, and fire detection. These sensors relay the information toward a number of base stations located in the building, which link to one another through a wired network. The information can thus be made

Without proper design, communication will increase network power consumption significantly.

available to control rooms and, possibly, other remote locations.

To offer maximum flexibility for different kinds of applications, the WiseNET project sought to achieve low-power operation for all three network operation modes.

REDUCING POWER CONSUMPTION

To reduce power consumption, nodes remain sleeping until they need to undertake a specific task. At some defined time, a sensor node will wake up and perform a measurement. An external event also can trigger this wake-up. The node can then decide to communicate the gathered information to a neighbor and send it a message. Unfortunately, the neighbor might be sleeping to save energy. The node must thus keep sending the information until the neighbor awakens and acknowledges receipt of the information.

If a node needs information from a neighbor, it can transmit a request until it receives a response. Alternatively, the requesting node can stay awake and wait until the neighbor decides to send the information spontaneously.

These two simplistic examples show that, without proper design, communication will increase network power consumption significantly because listening and emitting are power-intensive activities. The challenge, then, is to minimize energy consumption by reducing and eventually even eliminating the energy waste caused by

- idle listening—a node waits, listening until another node emits a message;
- overemitting—a node sends a message and the destination node is not ready to receive;
- overhearing—a node listens for a message that is sent to another node and for which it is not the destination; or
- collisions—two nodes transmit at the same time with the consequence that they must retransmit the messages later.

Reducing wasted energy guided development of WiseNET's MAC design.

HARDWARE AND SOFTWARE CODESIGN

Achieving the WiseNET project team's ultimate goal—to reach the lowest possible power consumption for hybrid networks—requires proper hardware and software codesign. Often, communication protocols are designed for existing integrated circuits, forcing developers to take into account the special features and limitations of those circuits.

On our project, we decided from the start to design the radio and protocol concurrently. A few important parameters of the radio clearly impact the higher protocol layers, including

- power consumption in receive and transmit mode;
- wake-up time—the time necessary to switch from idle mode to receive or transmit mode; the node cannot receive while power consumption is approximately the same as in receive mode;
- bit and frame synchronization time at reception because the frame synchronization pattern has an impact on both the overhead—its length—and the frame error rate when the system detects frame starts erroneously;
- the presence of an effective *receive signal strength indicator*, which can be used to improve the reliability of frame detection; the RSSI also reduces idle listening when used as a silence detector;
- some way to filter incoming packets to reduce overhearing;
- the time necessary to switch from receive to transmit mode or vice versa, during which no communication is possible while power consumption is at least that of the frequency synthesizer;
- receiver sensitivity and maximum transmit power, which affect the required number of hops;
- the capacity to adjust transmit power and receiver sensitivity to reduce the frame error rate and collision probability; and
- power consumption in sleep mode with a running, accurate clock that, depending on the application, can consume the major share of the current during the node's life.

The special context of sensor networks required making some basic choices. First, we decided not to minimize the transmit power, but rather to choose a reasonably high constant value—such as the maximum allowed transmit power in the chosen bands, which is 10 dBm for the European 434-MHz ISM band. Because nodes usually transmit rarely, the transmit energy is not the most important parameter to be optimized. Second, we concentrated on reducing both the energy consumption and wake-up time in receive mode. Since WiseNET uses a form of preamble sampling at the MAC layer, waking up quickly has a clear impact on energy consumption.

For these reasons, we carefully elaborated the RSSI information. For a third level of optimization, we added robust bit synchronization and packet filtering based on a programmable pattern. Both these measures reduce frame losses, while packet filtering additionally reduces overhearing. In the first MAC simulations, we found that the capacity to adjust transmit power and receiver sensitivity was less important; therefore, we chose not to implement this feature because doing so would require additional circuitry.

PREAMBLE SAMPLING

WiseMAC is a single-channel contention protocol based on nonpersistent *carrier sense multiple access*. Combining nonpersistent CSMA with preamble sampling mitigates idle listening. The preamble sampling technique consists of regularly sampling the medium to check for activity. In this context, sampling means periodically measuring the received signal strength.

All nodes in a network sample the medium with the same constant period, but their relative sampling schedule offsets are independent. If a node finds the medium busy, it continues to listen until it receives a data packet or the medium becomes idle again. At the transmitter, a wake-up preamble of size equal to the sampling period is transmitted ahead of every data packet to ensure that the receiver will be awake when data transmission begins. This technique enables low power consumption when traffic is low, as is usually the case in sensor networks. It also provides the lowest possible power consumption in the absence of traffic and for a given wake-up latency using a conventional receiver.

The main disadvantage of the fixed-length preamble protocol is its high power consumption overhead, both in transmit and receive, due to the wake-up preamble. Also, in an ad hoc network, both the intended destination and all other nodes overhearing the transmission pay the cost of reception.

WiseMAC introduces a novel scheme to reduce the length of this costly wake-up preamble. As Figure 1 shows, the scheme learns the sampling schedule of direct neighbors and exploits this knowledge to use a minimized wake-up preamble.

The nodes learn or refresh their neighbor's sampling schedule during every data exchange by piggybacking into the acknowledgment messages the remaining time until the next sampling instant. Every node keeps an updated table of sampling time offsets for all its usual destinations. Since a node has only a few direct destinations, it can manage such a table even with limited memory resources.

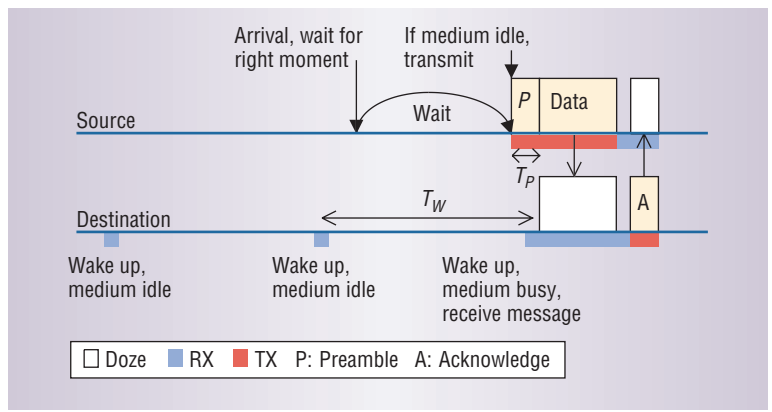


Figure 1. WiseMAC preamble minimization. A low-power, media access control protocol, WiseMAC uses a scheme that learns the sampling schedule of direct neighbors and exploits this knowledge to minimize the wake-up preamble length.

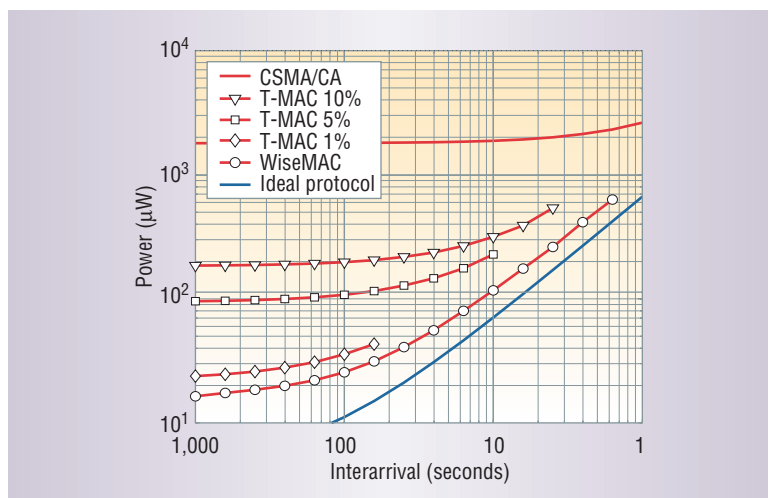


Figure 2. Lattice multihop topology simulation. The results show average power consumption as a function of the network's node traffic. The T-MAC protocol can provide either low-power consumption in low-traffic conditions or high throughput.

The duration of the wake-up preamble must cover the potential drift between the source clock and the destination clock. This drift is proportional to the time since the last acknowledgment was received. Thus, the wake-up preamble's required duration is given by $T_p = \min(4\theta T_c, T_w)$, where θ is the frequency tolerance of the time-based quartz, T_w is the sampling period, and T_c the interval between communications. WiseMAC schedules a transmission so that the middle of the wake-up preamble coincides with the destination's expected sampling time. Using a randomized wake-up preamble mitigates the systematic collisions that this synchronization can introduce.

Because the preamble's length is proportional to the interval between packets, WiseMAC's overhead adapts to the traffic. WiseMAC can thus provide both ultralow average power consumption in low-traffic conditions and high energy efficiency in high-traffic conditions. Figure 2 shows simulation results of a lattice multihop topology.

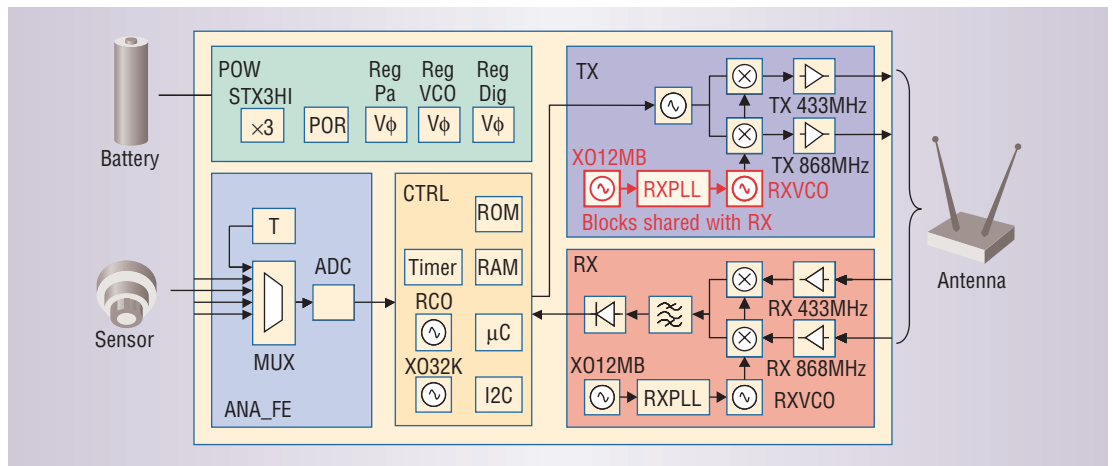


Figure 3. Generic WiseNET SoC building blocks. In addition to the ultralow-power dual-band radio transceiver (Tx and Rx), the architecture includes a sensor interface with a signal conditioner and two analog-to-digital converters (ANA FE), a digital control unit based on a Cool-RISC microcontroller (μC) with on-chip low-leakage memory, several time-basis and digital interfaces, and a power management block (POW).

In the chosen radio range, every node has eight neighbors. Traffic flows through each node, with the interarrival time shown on the x -axis. In addition to providing power consumption below 20 microwatts in low-traffic conditions, WiseMAC can approach the power consumption of an ideal protocol in high-traffic conditions. With an interarrival time of 100 seconds, the power consumption amounts to as little as 25 microwatts—which translates into more than a five-year lifetime for a single AA alkaline battery.

For comparison, Figure 2 also shows the performance of

- T-MAC⁶—an improved version of S-MAC⁷ with different duty cycles,
- CSMA/CA—CSMA with collision avoidance, and
- an ideal protocol.

We used the WiseNET transceiver’s power consumption and timing parameters to simulate these protocols.

CSMA/CA’s power consumption is limited at low-traffic levels by the power consumed in receive mode because the transceiver is never switched off. The ideal protocol represents the lower bound that low-power MAC protocols should strive for—the minimum power consumption required to transmit the data without any overhead.

The T-MAC protocol requires choosing the duration of the listen and sleep phases. In the absence of traffic, this ratio is actually the transceiver’s duty cycle. In Figure 2, the T-MAC protocol is plotted only up to the point above which more than 5 percent of the packets drop because of congestion. T-MAC can thus provide either low-power consumption in low-traffic conditions, or high throughput.

We selected a single-channel contention protocol to ease self-configuration. WiseMAC requires no setup signaling or network-wide time synchronization. The combination of preamble sampling and wake-up preamble-length minimization provides both ultralow power consumption in low-traffic conditions and high energy efficiency in high-traffic conditions.

Although designed initially for multihop networks, WiseMAC has also proven suitable as the downlink of an infrastructure network. In low-traffic conditions, WiseMAC provides lower power consumption than the power-save scheme that IEEE 802.11 and IEEE 802.15.4 specify.⁸ WiseMAC can thus be used in a hybrid network topology, to receive data from both battery-powered nodes and energy-unconstrained base stations.

WISENET NODE ARCHITECTURE

Since a WiseNET network has many distributed microsensors, we do not envision replacing or recharging the batteries. The nodes therefore require long-term autonomy—typically more than three years—and, consequently, low average power consumption. In addition, the nodes must be tiny to fit into all kinds of spaces and, given their high number, they must also be inexpensive. Building from these basic specifications, we used a SoC approach to design the nodes as highly integrated devices that use a dedicated integrated circuit.

Figure 3 shows the basic WiseNET SoC architecture. In addition to the ultralow-power dual-band radio transceiver, the SoC also includes a sensor interface with a signal conditioner and two analog-to-digital converters, a digital control unit based on a Cool-RISC microcontroller with on-chip low-leakage memory, several time-basis and digital interfaces, and a power management block.

Although all these blocks contribute to power consumption, the most critical block is the RF transceiver.

Researchers have found that power consumption as low as one milliwatt in receive mode can be achieved with a -95 -dBm sensitivity at a 24-Kbps data rate in the 434-MHz European ISM band.^{9,10} These encouraging results led to the design of the WiseNET transceiver along roughly the same lines, but taking advantage of moving from a 0.5-micrometer to a 0.18-micrometer standard digital CMOS process.

Design objectives

In addition to incorporating the features listed in Table 1, our most important design objectives were to

- keep the power consumption within the 1-milliwatt range while in receive mode, adding the 868-MHz SRD band to extend the operation to dual band;
- achieve several years of autonomy by operating from a single 1.5-V AA alkaline battery with a supply voltage as low as 0.9 V, corresponding to the battery end-of-life voltage;
- use a 0.18-micrometer standard digital CMOS process that has no precision analog components such as capacitors and resistors or dedicated RF technology options such as substrate isolation; and
- minimize both the number of external components and the cost.

Dual-band operation, combined with the use of several channels, allows for frequency diversity to solve the difficult problem of rejecting strong nearby interferers. We chose to stay in the 434-MHz and 868-MHz bands instead of moving to the popular and globally available 2.4-GHz band mainly to limit power consumption. Indeed, about 50 percent of a receiver's power consumption relates directly to the circuits operating at or close to the RF frequency, such as the frequency synthesizer and the RF front end, which consists of a low-noise amplifier and a power amplifier.

Since the power consumption of these blocks is approximately proportional to frequency, choosing sub-GHz bands limits power consumption to the milliwatt range. In addition, operating at 2.4 GHz with the chosen 0.18-micrometer CMOS process would require a supply voltage on the order of 1.8 V, which is incompatible with the target 0.9-V minimum supply voltage.

Table 1. WiseNET sensor network characteristics.

WiseNET radio parameters

Operating frequency	433 MHz (ISM) and 868 MHz (SRD)
Channel separation	600 kHz (primary), 200 kHz (secondary)
Propagation range	~2 km outdoors — ~10 m indoors
Data rate/modulation	<100 Kbps with FSK ($\Delta f = 25$ kHz) <4 Kbps with OOK
Power consumption in Rx mode	1.8 mW
Power consumption in Tx mode	31.5 mW
Wake-up time	800 μ s
Rx to Tx and Tx to Rx turnaround time	400 μ s

Main measured results for the Rx and Tx blocks.

Supply voltage	$V_{DD} = 0.9$ V — 1.5 V (Rx and Tx)
Sleep current	3.5 μ A
Receiver (Rx) (measured at $V_{DD} = 1$ V and at 25 °C)	
Sensitivity	-105 dBm @ BER = 10^{-3} and 25 Kbps
PLL phase noise	-120 dBc/Hz @ 600 kHz offset
Supply current	$I_{Rx} = 2$ mA
Transmitter (Tx) (measured at $V_{DD} = 1$ V and at 25 °C)	
Output power	10 dBm
Efficiency @ 10 dBm	30%
Supply current @ 10 dBm	$I_{Tx} = 24$ mA (PA-preamp)

Radio parameters

Choosing a 0.18-micrometer CMOS standard digital process allows the selected subGHz-frequency bands to trade the high-frequency capability of minimum-length transistors with lower current consumption by biasing the devices at lower current densities, even for devices working at radio frequency. We achieve this by moving the transistor operating points to the moderate and weak inversion regions.¹¹ Biasing the devices in moderate inversion also offers a good tradeoff between high-current efficiency, low-voltage operation, and reasonable linearity.

The selected radio architecture also strongly conditions power consumption. Although we sought a highly integrated solution, implementing the rejection of strong nearby interferers on-chip would have required prohibitive dynamic range and power consumption due to the low-voltage requirement. To circumvent this problem, we used an external RF filter.

The radio architecture we chose, which builds on similar work done by other researchers,¹² consists of a superheterodyne with a high intermediate frequency, followed by conversion to DC. This architecture offers the advantage of achieving a significant gain at the intermediate frequency without prohibitive power consumption, while the channel selection is performed around DC. Having sufficient gain at the intermediate frequency also reduces the signal-to-noise ratio degradation from the strong $1/f$ noise present in deep-submicron CMOS processes. This noise strongly affects the




Figure 4. The WiseNET system-on-chip sensor node. Key SoC components include the dual-band transceiver (Rx/Tx), the sensor interface with two ADCs (ANA), the power management block (POW), the control unit (μC) with an 8-bit CoolRISC processor, and the embedded low-leakage memory (RAM).

analog baseband. Other external devices include a high-Q inductor used for the LC-tank circuit to achieve the desired low-phase noise and a few capacitors for impedance matching to the antenna.

Power consumption

We also designed a special on-chip varactor to operate at the required low voltage, while still offering a sufficient tuning range and a high-quality factor that helps avoid degrading the tank's overall Q-factor.¹³ The WiseNET transceiver offers a digital RSSI as well, which the WiseMAC protocol uses for preamble sampling and carrier sense activities.

If all contributors within the receive chain are on, the total current drain makes long-duration operation on a single battery impossible. Analyzing the individual contributions reveals the current consumption to be significantly larger for the RF blocks and smaller for the baseband blocks. This disparity occurs because the current directly relates to the frequency of operation or the required bandwidth.

On the other hand, turn-on times will vary inversely with the frequency of operation because baseband blocks require more time for all the nodes in the circuit to reach a quiescent state after being turned on than RF blocks. Therefore, the system can save significant energy by waking up the lower-power baseband blocks before the power-consuming RF circuits, which wake up quickly.

Optimization parameters

The WiseNET chip offers a high degree of flexibility in controlling each block's individual operation. A wake-up sequence implementation can, for example, start with powering on the low-frequency reference clock; then the baseband path of channel filters, limiters, and RSSI; followed by the frequency synthesizer; and, eventually, the intermediate frequency amplifiers and the RF front end's low-noise amplifier and down-conversion mixers.

In addition, the RSSI can efficiently leverage WiseMAC's sampling nature by turning on only the required block in the baseband. Once it has activated the receiver chain up through the baseband channel filter, the RSSI can measure the signal strength. If it determines that the detected activity level is sufficient, the RSSI can turn on the demodulation blocks for further symbol and frame analysis and error correction. If the power level is lower than an acceptable threshold, the RSSI turns off the complete receiver, thus optimizing the power burned during both idle listening and receiving. The RSSI can use the same power sequencing techniques for the transmit path.

The second most important optimization parameter is the wake-up time. To optimize this parameter, we carefully defined the activation sequence for the different transceiver blocks. Implementing this measure mandated using proper circuit design techniques to minimize the wake-up times of the slower blocks.

This delay is not an issue for RF blocks because the larger currents and small parasitic capacitances required to achieve RF performance allow very quick settling. Given the significant energy savings it can provide, optimizing the setup delay in the intermediate frequency amplifiers, intermediate frequency to baseband mixers, and baseband blocks is worth investigating. Deep-submicron CMOS offers the clear advantage of making operation within the 100-MHz range feasible with proper biasing of the signal path transistors in weak or moderate inversion regions. In the chosen high intermediate frequency architecture, the intermediate frequency blocks can operate at frequencies that minimize setup time without degrading power consumption.

Baseband blocks exhibit an obvious limitation due to their inherent low-frequency poles. Nonetheless, techniques such as storing the bias points between channel sampling, boosting the start-up currents momentarily at wake-up, or switching the internal poles to higher frequencies for fast settling and then setting them back to the nominal frequency reduce the settling times.

The third key optimization parameter is the receive-to-transmit and transmit-to-receive switching times. During these turnaround delays, all blocks that remain on will burn power that is extraneous to the communication. Thus, we designed the WiseNET SoC architecture shown in Figure 4 to meet this particular MAC parameter. For example, the receive and transmit paths both use the same synthesizer and LO signals. The design makes this possible because it implements the receiver and transmitter according to the same high intermediate frequency superheterodyne architecture. In this way, the synthesizer remains on, working around the same operating points to yield fast turnaround times.

Using the WiseNET transceiver with the WiseMAC protocol, a relay sensor node consumes about 25 microwatts when forwarding 56-byte packets every 100 seconds. The radio transceiver used on motes¹⁴ requires 24 milliwatts in receive mode or 8 milliAmps at 3 V. Using the S-MAC or T-MAC protocols with such a transceiver, at a duty cycle of 10 percent, yields a minimum power consumption of 2.4 milliwatts—about 100 times more than that achieved with the WiseNET solution we have presented.

In a next step, we plan to study how to leverage WiseMAC for routing in more dynamic networks. Finding power-efficient solutions to achieve data consistency and action synchronization will be next on the agenda. On the hardware side, we currently are investigating new narrow-band radio architectures using high-Q bulk acoustic wave resonators. These architectures should allow even lower power consumption and higher operating frequencies. Alternatively, ultrawideband techniques using impulse radios could make the communication more robust to fading, while maintaining low power consumption. 📡

Acknowledgments

We thank the Swiss Center for Electronics and Microtechnology design team for its invaluable contribution to the success of the WiseNET project, particularly T. Melly and E. Le Roux for the information they provided. The work presented here was supported in part by the National Competence Center in Research on Mobile Information and Communication Systems, a center supported by the Swiss National Science Foundation under grant number 5005-67322.

References

1. I.F. Akyildiz et al., "Wireless Sensor Networks: A Survey," *Computer Networks*, vol. 38, 2002, pp. 393-422.
2. J.M. Rabaey et al., "PicoRadio Supports Ad Hoc Low Power Wireless Networking," *Computer*, July 2000, pp. 42-48.
3. G.J. Pottie and W.J. Kaiser, "Wireless Integrated Network Sensors," *Comm. ACM*, May 2000, pp. 51-58.
4. A. Chandrakasan et al., "Power Aware Wireless Microsensor Systems," *Proc. European Solid-State Circuits Conf.*, Univ. of Bologna, 2002, pp. 47-54.
5. A. El-Hoiydi et al., "WiseMAC: An Ultra Low Power MAC Protocol for the WiseNET Wireless Sensor Network," *Proc. 1st ACM SenSys Conf.*, ACM Press, 2003, pp. 302-303.
6. T. van Dam and K. Langendoen, "An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks," *Proc. 1st ACM SenSys Conf.*, ACM Press, 2003, pp. 171-180.
7. W. Ye, J. Heidemann, and D. Estrin, "An Energy-Efficient MAC Protocol for Wireless Sensor Networks," *Proc. IEEE Infocom Conf.*, IEEE Press, 2002, pp. 1567-1576.
8. A. El-Hoiydi and J-D. Decotignie, "WiseMAC: An Ultra Low Power MAC Protocol for the Downlink of Infrastructure Wireless Sensor Networks," *Proc. IEEE Symp. Computers and Comm.*, IEEE Press, 2004, pp. 244-251.
9. A-S. Porret et al., "An Ultralow-Power UHF Transceiver Integrated in a Standard Digital CMOS Process: Architecture and Receiver," *IEEE J. Solid-State Circuits*, vol. 36, no. 3, 2001, pp. 452-466.
10. T. Melly et al., "An Ultralow-Power UHF Transceiver Integrated in a Standard Digital CMOS Process Transmitter," *IEEE J. Solid-State Circuits*, vol. 36, no. 3, 2001, pp. 467-472.
11. C. Enz, "An MOS Transistor Model for RF IC Design Valid in All Regions of Operation," *IEEE Trans. Microwave Theory and Techniques*, Jan. 2002, pp. 342-359.
12. H. Darabi and A. Abidi, "A 4.5-mW 900-MHz CMOS Receiver for Wireless Paging," *IEEE J. Solid-State Circuits*, vol. 35, no. 8, 2000, pp. 1085-1096.
13. A-S. Porret et al., "Design of High-Q Varactors for Low-Power Wireless Applications Using a Standard CMOS Process," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, 2000, pp. 337-345.
14. Crossbow, "MICA 2—Wireless Measurement System Datasheet," 2004, document part number 6020-0042-04; www.xbow.com.

Christian C. Enz is the vice president of the micro-

electronics division at the Swiss Center for Electronics and Microtechnology (CSEM). He is also a professor at the Swiss Federal Institute of Technology, Lausanne (EPFL). His research interests include low-power analog and RF IC design, wireless sensor networks, and device modeling. Enz received a PhD in electrical engineering from the EPFL. Contact him at christian.enz@csem.ch.

Jean-Dominique Decotignie is a section head at CSEM. His research interests include real-time and power-aware networks and middleware for robustness and reliability of nomadic systems. He is also a professor at the EPFL. Decotignie received a PhD in electronic engineering from the EPFL. He is Fellow of the IEEE. Contact him at decotignie@ieee.org.

Amre El-Hoiydi is an R&D engineer at CSEM. His main research interest is the design of energy-efficient communication protocols for wireless sensor networks. El-Hoiydi received an MS in electrical engineering from ETH Zurich. Contact him at amre.el-hoiydi@csem.ch.

Vincent Peiris heads the RF and Analog IC design group at CSEM. His research interests include low-power RF CMOS transceiver design for wireless sensor networks. Peiris received a PhD in electrical engineering from the EPFL. Contact him at vincent.peiris@csem.ch.

GET CERTIFIED

Apply now for the 1 September – 30 October test window.
(Deadline to apply: 15 August.)



CERTIFIED SOFTWARE DEVELOPMENT PROFESSIONAL PROGRAM

Doing Software Right

- Demonstrate your level of ability in relation to your peers
- Measure your professional knowledge and competence

Certification through the CSDP Program differentiates between you and other software developers. Although the field offers many kinds of credentials, the CSDP is the only one developed in close collaboration with software engineering professionals.

“The exam is valuable to me for two reasons:

One, it validates my knowledge in various areas of expertise within the software field, without regard to specific knowledge of tools or commercial products...

Two, my participation, along with others, in the exam and in continuing education sends a message that software development is a professional pursuit requiring advanced education and/or experience, and all the other requirements the IEEE Computer Society has established. I also believe in living by the Software Engineering code of ethics endorsed by the Computer Society. All of this will help to improve the overall quality of the products and services we provide to our customers...”

— Karen Thurston, Base Two Solutions

Visit the CSDP web site at www.computer.org/certification
or contact certification@computer.org

