Energy-Efficient Communication for *Ad-Hoc* Wireless Sensor Networks

Rex Min and Anantha Chandrakasan
Massachusetts Institute of Technology, Cambridge, MA
{rmin, anantha}@mtl.mit.edu

Abstract

The energy dissipated by communication is a key concern in the development of networks of hundreds to thousands of distributed wireless microsensors. To evaluate the dissipation of communication energy in this unique application domain, energy models based on actual microsensor hardware are incorporated into a simulation tool designed expressly for high-density, energy-conscious wireless networks. Assessing and leveraging the energy implications of microsensor hardware and applications is crucial to achieving energy-efficient microsensor netowrk communication.*

1. Introduction

A distributed network of thousands of collaborating microsensors promise a maintenance-free, fault-tolerant platform for gathering rich, multi-dimensional observations of the environment [1][2][3]. Microsensor networks are a specialized class of *ad hoc* networks with several distinguishing characteristics: high node density, low date rate, and an unprecedented attention to energy consumption. Unlike laptop or palmtop devices, microsensor nodes will be expected to operate from 5-10 years from an amount of energy equivalent to an "AA" cell [4], requiring innovative design methodologies to eliminate energy inefficiencies that would have been overlooked in the past. An area potentially ripe with inefficiencies is microsensor communication. Building an energy-efficient protocol stack for microsensors requires a thorough investigation of the interactions among the sensor application, network protocol, MAC layer, and radio. Energy consumption characteristics that are unique to this domain of wireless systems must be addressed and exploited for maximally energy-efficient communication.

Communication protocols for traditional ad-hoc networks generally employ multi-hop routing to ameliorate the high (r^2 to r^4) path losses incurred by radio transmission. Two general routing methodologies—source routing and distance vector approaches—are analogous to their counterparts in wired networks. Source routing specifies complete, hop-by-hop paths for each packet, while distance vector protocols maintain only next-hop information to each destination. These protocols are typically intended for wireless IP applications rather than microsensor networks. An overview of many *ad hoc* wireless routing protocols is available in [5].

Two newer protocols have been specifically designed for energy-constrained sensor networks. Directed diffusion [6] relies on local interactions among nodes to create efficient paths for data flow. No global routing state is kept anywhere in the system; rather, each node chooses its own source(s) from which to receive data, leading to reasonably efficient data propagation at a global level. LEACH (low-energy adaptive clustering hierarchy) [7] forms rotating clusters of adjacent nodes, within which nodes transmit to a single cluster head that bears the burden of a long-distance transmission. Clustering explicitly encourages data aggregation to reduce further the transmission burden on the network.

With the increasing interest in battery-powered wireless systems, energy consumption has become a primary metric of wireless communication protocols, alongside traditional metrics such as throughput and fault-tolerance. However, little prior work has characterized the energy consumption of wireless network protocols with realistic hardware models and behavioral characteristics of microsensors. Transmission energy is often modeled in μJ/bit, a model that fails to consider hardware and protocol overheads. Physical and MAC layer models have almost universally adopted energy consumption and performance characteristics from 802.11b, whose high power consumption and complexity are unsuitable for wireless microsensors.

Reducing the energy of communication in wireless microsensors demands that each aspect of communication, such as the protocol and MAC layers, is tailored to the application at hand. In this paper we present one such approach, using observations about the nature of microsen-

---

*This research is sponsored by the Defense Advanced Research Project Agency (DARPA) Power Aware Computing/Communication Program and the Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0551. Rex Min is funded by an NDSEG Fellowship.

0-7803-7147-X/01/$10.00©2001 IEEE

sor communication, accurate energy models of actual hardware, and a custom simulation tool built from the ground-up for high-density, energy-constrained wireless network simulation.

2. Microsensor Communication

2.1 Implications of Multi-Hop Communication

A wireless node in an *ad hoc* network traditionally seeks its nearest neighbors as candidates for *next-hop* transmission. However, with the high node densities that enable the high robustness and resolution of microsensor networks, we must reconsider the paradigm of routing through nearest neighbors—if not multi-hop itself.

A common model for the power consumed by multi-hop transmission is

$$P_{MH}(h, d, c_{tx}) = c_{tx}\left[(h-1)P_{rxElec} + h\left(P_{txRad}\left(\frac{d}{h}\right) + P_{txElec}\right)\right]$$

where h is the number of hops, d the total transmission distance, c_{tx} the transmit duty cycle, P_{rxElec} and P_{txElec} the power required by the receive and transmit electronics, and $P_{txRad}(x)$ the power required to overcome path loss attenuation over a distance x. Figure 1 evaluates this model for varying d and h using $P_{rxElec} = P_{txElec} = 175\text{mW}$ and r^3 path loss.

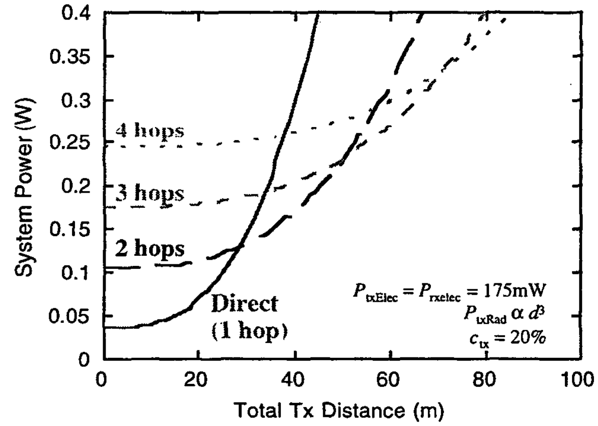

Figure 1: Evaluation of the ideal model discussed above. As the transmit distance increases, increasing the number of hops improves energy-efficiency.

The seemingly even spacing of crossover points between n-hop and $(n+1)$-hop routing is no coincidence: the existence of a single minimum-energy hop distance, a fact that is proved in [8], yields this periodic behavior.

As the distance of the required transmission increases, it becomes advantageous to increase the number of hops. However, it is clear that there is a large range of distances for which direct transmission is more energy-efficient than multi-hop transmission. For the above model and parameters, this corresponds to about $d < 30$ meters.

Our model above has assumed ideal multi-hop communication with no overhead whatsoever. Our energy characterization must account for protocol and MAC overhead, sub-optimal node spacing, and the fact that radio receivers, not being omniscient to packet arrivals, must occasionally poll for packets. These sources of overhead introduce the additional power $P_{overhead}$ into multi-hop routing. For thousand-node microsensor networks, these overheads can and will be substantial. Figure 2 plots the energy of direct transmission *versus* a multi-hop routing protocol for which

$$P_{MHActual}(h, d, c_{tx}) = P_{overhead} + P_{MH}(h, d, c_{tx})$$

This is the communication energy as a function of both transmission distance and the transmit duty cycle c_{tx}, using an arbitrary choice of $P_{overhead} = 2P_{rxElec}$ for the sake of example.

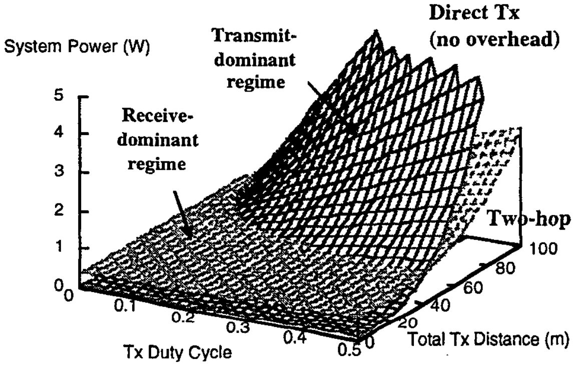

Figure 2: When receive and protocol overhead are taken into account, the efficacy of multi-hop routing becomes an increasing function of not only how far data is transmitted, but how much as well. Multi-hop has no advantage until transmit power dominates receive and protocol overhead.

We now have a more complete picture for the regime in which multi-hop routing is beneficial. In the *receive-dominant regime*, the hardware and protocol overhead of receiving packets outweighs the energy savings of shorter radio transmissions. In the *transmit-dominant regime*, both the transmission distance and the number of bits transmitted are sufficiently large that multi-hop routing is beneficial. As intuition would suggest, both the transmission distance *and quantity* of transmitted bits determine the breakpoint between the two regimes, a result concealed by the simpler model. If transmit duty cycles are sufficiently small, then protocol-free direct transmission will be more energy-efficient even at very large distances. With overhead accounted for, the transmission distance at which multi-hop transmission becomes advantageous over direct transmission becomes *much greater* than our initial $d = 30$ m.

This observations holds two noteworthy implications for microsensor networks. First, as most microsensor networks utilize a small mean distance between nodes, nearest neighbors are often the wrong candidates for energy-efficient next-hops. Second, large classes of applications exist for which the *entire network diameter* is in the receive-dominant regime. For these classes of networks, such as those completely enclosed within a room, machine, or small lawn, transmission is most efficient with *no multi-hop protocol at all.* In these situations, it is increasingly important to focus on the energy dissipation characteristics of the hardware.

2.2 Application-Driven Routing

Presuming that the network is operating in a transmit-dominated regime, we next consider techniques for reducing $P_{overhead}$ for multi-hop wireless microsensor communication. Exploiting the microsensor network's architecture and application-specific characteristics allows for a great deal of energy-conscious optimization in the protocol stack.

Reduction of radio receive energy is our primary concern. A radio receiver that is on and idle consumes a substantial amount of power—often as much as transmission. Given that many nodes will likely be in the receiving range of any node's transmission, it is desirable to shut down the radio receiver in the majority of idle nodes. Unfortunately, most wireless protocol and MAC layers in use today utilize unique addresses to route packets to specific destinations, with the expectation that these destinations are actively listening for packets. With radio shutdown, this assumption no longer holds, and routing tables—whether they contain source routes or next-hop information—suddenly become very unstable.

We observe, at the application level, that communication in a microsensor network is decidedly one-way, from the observer nodes to a base station. There are many data sources and relays, but few actual sinks. As the individual relays have no need for the data they are relaying, the entire notion of *addressing* a packet to a specific relay node is unnecessary. Our only concern is that packets move progressively closer to a base station.

We thus propose that microsensor nodes not utilize explicit addresses at the protocol or MAC level, but rather a metric of their approximate distance to the nearest base station. This metric could be propagated across the network by flooding, with the base stations initially broadcasting a zero metric and each node adding a constant factor to the smallest value heard.

A node that with a packet destined for the base station simply broadcasts its packet with its current distance metric. Nodes that receive the packet compare their own distance metric to that of the packet. To minimize the number of hops, the receiving node that is closest to the base station and farthest from the originating node would relay the packet onward. Such behavior could be implemented, for instance, with a delay timer proportional to the difference between the packet's and relay node's distance metrics, or simply the RSSI. The node with the lowest delay would forward the packet, and the others, hearing the forward, would drop their respective copies.

Address-free forwarding allows any active node—rather than one that is specifically addressed—to relay a packet. This enables flexible, protocol-independent radio receiver shutdown.

3. Case Studies

We now consider the energy consumption of simulated microsensor networks for both the transmit-dominant and receive-dominant regimes.

3.1 System Models and Simulation Framework

The MIT μAMPS project (Adaptive Multidomain Power-Aware Sensors) [9] is developing the enabling technologies for energy-efficient microsensor networks. The μAMPS-1 node developed within this project is the basis for our hardware energy consumption models. The μAMPS-1 node consists of sensing, processing, and radio subsystems. The sensing system consists of an acoustic sensor and low-power A/D converter. Data processing, as well as some network functions, are carried out by a StrongARM SA-1110 microprocessor. The radio transmits and receives at 1 Mbps at half-duplex in the 2.4 GHz range. Figure 3 illustrates the hardware of μAMPS-1.

The measured energy consumption and performance of μAMPS-1 form the basis of our hardware model. Key parameters are listed in Table 1. As the state of the art in 2.4 GHz radios has advanced somewhat since the introduction of μAMPS-1, we incorporate some parameters of modern Bluetooth radios as well [10]. As node antennas would essentially be at ground level model, we model the radio path loss with an empirical r^3 rather than the conventional two-ray ground wave propagation model.

All simulations are on a custom network simulator designed specifically for high-density wireless networks. The simulator, written in Java, provides a rich framework for protocol and MAC implementation, hardware description and energy characterization. We have chosen to deploy this custom tool as an alternative to making substantial low-level changes to existing tools to support our behavioral and energy models.

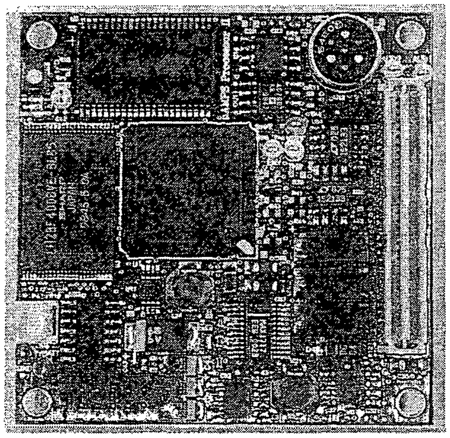

Figure 3: The µAMPS-1 sensor/processor board is 55mm on a side. A radio board (not pictured) is stacked with this board to form a complete node.

TABLE 1 Simulation Parameters

Parameter	Value
Sensing Subsystem	40 mW
Digital Electronics	110 mW
Sensing Energy	1.75µJ/bit
Radio Electronics (Tx and Rx)	175 mW
Radio Startup Time	466 µs
Radio Data Rate	1 Mbit/s

3.2 Reducing Protocol Overhead in the Transmit-Dominant Regime

We have implemented a multi-hop routing protocol patterned on our development of address-free relay in Section 2.2. CSMA is employed to reduce collisions.

The network under simulation consists of a uniformly random distribution of nodes with an average density of one node per 200 m^2, over an area 50 meters wide and a length that is varied from 200 to 800 meters. The largest region simulated thus contains 2000 nodes. Transmitting nodes relay with a transmit power sufficient to cover nearly 7000 m^2, meaning that the network is heavily over-provisioned from a multi-hop perspective. A stimulus that produces observation packets from nearby nodes, and a base station that receives these observations, are placed at opposite ends of the simulation region. The sum of the 840-byte packets generated by the observers occupy about 5% of the available radio bandwidth. Figure 4 illustrates this simulation scenario.

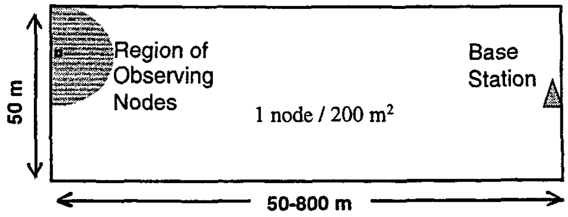

Figure 4: Simulation scenario for multi-hop routing.

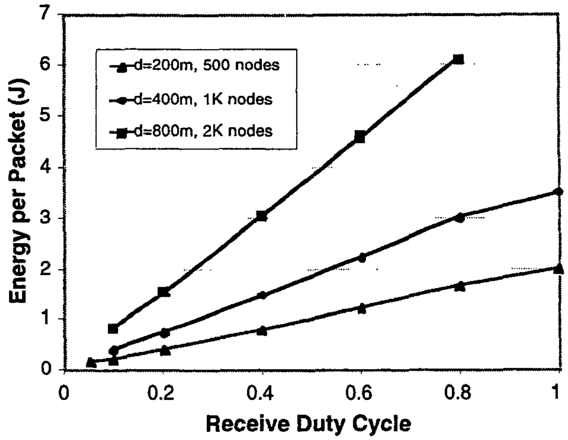

Figure 5: Reducing the duty cycle of the nodes' receivers provides a nearly linear reduction in the cost of communication. The energies plotted are the *total system* energy per 840-byte packet received at the base station.

Figure 5 illustrates the impact of applying a duty cycle of suggest that most of the energy overhead arises from the receive electronics; the total system energy per packet received at the base station is nearly proportional to the nodes' receive duty cycle.

3.3 Impact of Hardware in the Receive-Dominant Regime

We next consider the receive-dominated regime, where transmission distances and sizes are too small to justify the overhead of *any* multi-hop protocol. In this scenario we simulate 100 nodes of uniform random distribution within a 10m x 10m space, representative of nodes placed densely in a room, on a lawn, or within a machine. The base station is placed at a corner of the network. Each node generates a single packet every second of 2 to 128 samples (3 to 192 bytes). Packets are transmitted directly to the base station, and since the nodes need not act as relays, the nodes' radio receivers are always off.

When packets are extremely short, the energy required for radio startup exceeds the energy of the actual transmission. Figure 6 illustrates this concern. While the radio

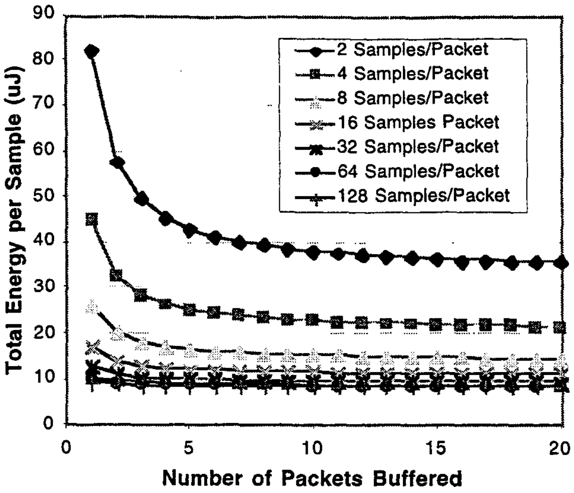

Figure 6: The system energy of transmitting short packets is dominated by the startup energy of the radio. Buffering multiple short packets into a single transmission trades latency for system energy savings.

requires 466 µs for startup, a packet transmission of 24 bytes at 1 Mbps, for instance, requires just 24 µs. By buffering packets at the local nodes and transmitting many short packets in a single transmission, we reduce the number of energy-consuming startups. By sending ten packets at 10-second intervals, for instance, we use less than half the total energy of immediate transmission. As individual packets grow larger in size, the impact of radio startup energy on total system energy is inherently reduced.

The trade-off is the increased latency of the buffered observations. Hence, for applications such as short-distance, low-rate environmental sensors, exposing the number of buffered packets as a dynamically adjustable quantity provides an effective energy-quality trade off. For instance, a node could choose to transmit observations immediately if an observed parameter fell outside a normal range, but buffer them otherwise.

4. Conclusion

The severe energy constraints on distributed microsensor networks demand the utmost attention to all aspects of energy consumption, and the use of energy models and simulation tools that are suited to the task of evaluating high-density, energy-conscious microsensor networks. Using custom tools and models appropriate to our task, we exploit the application characteristics of microsensors to reduce the overhead of communication. For multi-hop communication, the high node density of microsensor networks demands receiver shutdown, which is enabled by a distance-metric addressing scheme that takes advantage of the network's one-way communication from nodes to the base station. Nodes with sufficiently little or very short-distance transmissions are most energy-efficient with direct transmission to the base station. In this regime, we focus our attention on inefficiencies of the *hardware*, such as radio startup, rather than protocols.

While a stratified communication architecture that separates the physical, MAC, protocol, and application layers is convenient for abstraction and instruction, energy-efficiency is clearly gained by each level's awareness of the characteristics of the others. For microsensors, where energy matters most, we have demonstrated energy conservation by tailoring the MAC and protocol architectures to the specific characteristics of the radio and application domain.

Acknowledgments

The authors thank Piyada Phanaphat for sharing keen insights into the energy consumption of multi-hop routing, and Seong-Hwan Cho, Nathan Ickes, and Alice Wang for their assistance with developing the component energy and timing models for the simulation.

References

[1] B. Warneke, B. Atwood, K.S.J. Pister, "Smart dust mote forerunners," Proc. 14th IEEE Int'l Conference on MEMS, Jan. 2001, pp. 357-360.

[2] J. Rabaey et al., "PicoRadio supports ad hoc ultra-low power wireless networking," Computer, vol. 33, no. 7, July 2000, pp. 42-48.

[3] G. Asada et al., "Wireless Integrated Network Sensors: low power systems on a chip," Proc. 1998 ESSCIRC, Sept. 1998.

[4] A. Wang, S.-H. Cho, C. Sodini, A. Chandrakasan, "Energy efficient modulation and MAC for asymmetric RF microsensor systems," Proc. ISLPED 2001, August 2001, pp. 106-111.

[5] E. Royer and C-K. Toh, "A review of current routing protocols for ad hoc mobile wireless networks," IEEE Personal Communications, vol. 6, no. 2, April 1999, pp. 46-55.

[6] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed Diffusion: A scalable and robust communication paradigm for sensor networks," *Proc. MobiCOM 2000.*

[7] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-Efficient Routing Protocols for Wireless Microsensor Networks," *Proc. 33rd Hawaii Int'l Conf. on System Sciences (HICSS '00)*, Jan. 2000, pp. 3005-3014.

[8] M. Bhardwaj, T. Garnett, and A. Chandrakasan, "Upper Bounds on the Lifetime of Sensor Networks," *Proc. ICC 2001*, vol. 3, June 2001, pp. 785-790.

[9] R. Min et al., "Low-Power Wireless Sensor Networks," *Proc. 14th Int'l Conf. on VLSI Design*, Bangalore, India, Jan. 2001, pp. 205-210.

[10] F. Op't Eynde et al, "A Fully-Integrated Single-Chip SOC for Bluetooth," *Proc. ISSCC 2001*, Feb. 2001, pp. 196-197,446.