Design Considerations for Next Generation Wireless Power-Aware Microsensor Nodes

David D. Wentzloff, Benton H. Calhoun, Rex Min, Alice Wang, Nathan Ickes, Anantha P. Chandrakasan
Microsystems Technology Laboratories
Massachusetts Institute of Technology
Cambridge, MA 02139 USA
email: {ddw, bcalhoun, rmin, aliwang, nickes, anantha}@mtl.mit.edu

Abstract

In order to break the 100 μW average power barrier of a wireless microsensor node, aggressive design methodologies need to be developed. Dynamic voltage scaling should be more aggressive, reaching subthreshold operation, and knobs should be available for adapting hardware bit-precision and latency. Since the nodes operate in a sleep state most of the time, standby leakage currents must be reduced and the power supply voltage regulated to a near-optimum value. This paper presents insight and simulation/experimental results addressing some of the challenges of designing next generation wireless microsensor nodes.

1. Introduction

A power-aware wireless microsensor network is a collection of a few nodes to several hundred nodes, each one operating to minimize the total energy consumed by the network. A power-aware node is defined by its graceful scalability of the performance of individual blocks, with the energy consumption varying accordingly [1]. The overall challenge of next generation nodes is the further reduction of energy through aggressive optimization across all layers of abstraction.

Applications for wireless sensor networks range from military, such as target tracking, to consumer electronics and industrial equipment, such as home lighting or distributed sensing. The ZigBee Alliance, a consortium formed in 2002, is developing an industry standard for wireless networking of remote monitoring, control, and sensory nodes [2]. This communication standard will be tailored for low power, low data rate, secure wireless communication in the US and European ISM (Industrial, Scientific, and Medical) bands.

First generation microsensor nodes were built using mostly commercial parts that were limited in their power-aware capabilities [3][4]. A summary of the first generation μAMPS-1 (Adaptive, Multi-Domain, Power-Aware Sensors) node is presented in Section 2 of this paper. This node demonstrated many power-aware knobs, however next generation nodes must show a substantial improvement in energy consumption. Some of the key enablers and challenges of low energy design are discussed in Section 0.

Our next generation design goal for average power is 100 μW because below this point it becomes possible for nodes to harvest their energy solely from the environment. Various schemes have been proposed to eliminate the need for batteries in a portable digital system by converting ambient energy in the environment into electrical energy which is stored and utilized by the node. The most familiar sources of ambient energy include solar power, thermal gradients, RF, and mechanical vibration [5]. Advances in MEMS technology have enabled the construction of a self-powered capacitive resonator which delivers 10 μW of power with the appropriate regulation circuitry [6].

2. First Generation μAMPS Node

The MIT μAMPS project is developing power-aware hardware, protocol, and algorithmic building blocks for microsensor networks. The microsensor research community needs a flexible hardware substrate in which the sensors, signal processing algorithms, and network protocols can be easily changed. While the eventual goal of μAMPS is a one-chip solution that incorporates sensing, processing, and radio communication, an intermediate goal is designing a hardware substrate for developing and demonstrating microsensor applications. The μAMPS-1 sensor node provides a versatile, small, and power-aware hardware substrate for distributed microsensor networks.

Proceedings of the 17th International Conference on VLSI Design (VLSID'04)
1063-9667/04 $ 20.00 © 2004 IEEE

2.1. µAMPS-1 Node Architecture

The µAMPS-1 microsensor node uses commercial, off-the-shelf (COTS) components for rapid construction. A COTS design sacrifices the power savings of integration for rapid design time and programmability.

Power is aggressively managed by designing the hardware to anticipate the application requirements of the microsensor domain. To address the low duty cycle of microsensor nodes, we minimize the energy dissipated by an idle node through fine-grained subsystem shutdown. To address the unpredictable variability in operating conditions of an active node, we ensure that energy consumption can be scaled gracefully with performance for both radio range and computation. Figure 1 illustrates the power scaling controls implemented in the µAMPS-1 node.

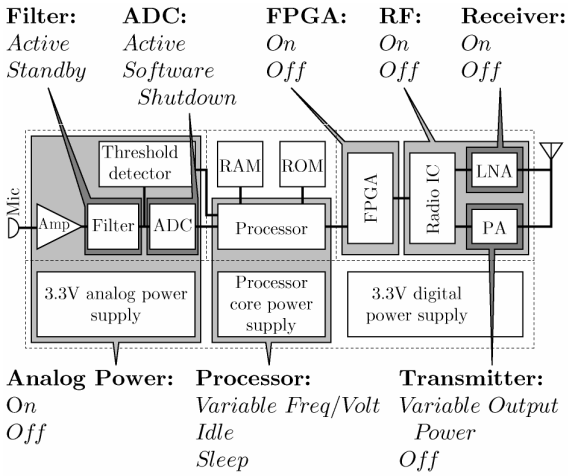

Figure 1. Power-scaling controls on the µAMPS-1 node. Most node components can be shut down. Many components incorporate more complex power controls that allow gradual power scaling.

A µAMPS-1 node consists of a stack of three or four printed circuit boards. The top board contains the radio, including the RF circuitry and the FPGA used for digital coding and decoding. The second board contains an Intel StrongARM processor and associated RAM and flash ROM. Also on the processor board are an acoustic sensor (microphone, amplifier, filter, and analog-to-digital converter) and a collection of dc/dc power converters that service the entire node. The optional third board in the stack is an additional sensor module to replace the acoustic sensor on the processor board. The µAMPS-1 node can be easily adapted to different applications by designing an appropriate sensor board. The bottom board in the stack contains the power source: four AAA batteries.

The default acoustic sensor for the µAMPS-1 node occupies a corner of the processor board and consists of an electret microphone, variable-gain amplifier, analog-to-digital converter with anti-aliasing prefilter, and a threshold detector. The entirely analog threshold detector enables the node to operate in an ultra-low power mode where only the microphone and a few op-amps are powered. If the microphone detects a sufficiently loud noise, the threshold detector is tripped, and a signal is sent to wake up the processor.

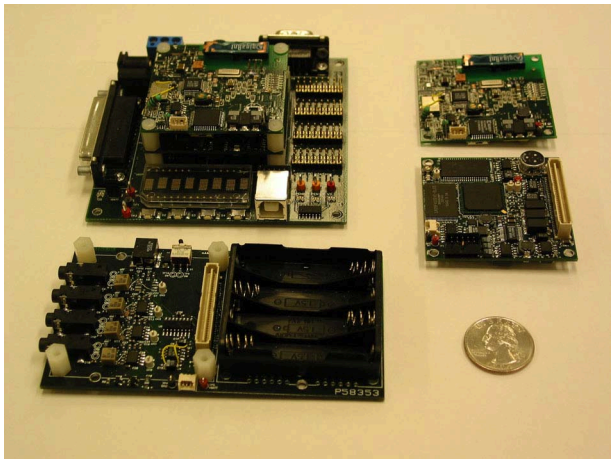

Figure 2. Variations on the µAMPS-1 node. (Upper right) Radio board. (Lower right) Processor board. (Bottom) Optional, enhanced sensor and battery adapter board for four-channel acoustic sensing. (Upper left) µAMPS base station, consisting of processor and radio boards stacked atop a PC interface board with serial and USB connectivity.

The µAMPS-1 processor block consists of a StrongARM microprocessor along with low-power static RAM and flash ROM. The StrongARM processor was chosen because of its high performance/power ratio (235 MIPS and <400 mW at 206 MHz [7]) and its built-in variable frequency (59-206 MHz) core clock generator. An adjustable clock frequency enables dynamic voltage scaling, in which processor voltage and clock frequency are adjusted together to provide an energy-latency trade off. Reducing the voltage applied to the processor core to the lowest level possible to support the current operating frequency increases power savings at low clock frequencies by up to 60% [8] by reducing both switching and leakage currents. Voltage is dynamically adjusted through a digitally programmable dc/dc converter built into the processor board.

The µAMPS-1 radio consists of a digital baseband processor (implemented on an FPGA) and an RF front-end. The digital component is responsible for encoding, decoding, and error detection/correction. It also controls the timing of the transmitter and receiver according to the TDMA scheme employed by the network protocol. The core of the RF circuitry is an LMX3162 2.4 GHz radio transceiver chip from National Semiconductor. A Maxim MAX2242 2.4 GHz variable power amplifier provides 0-

20dBm of transmit power. Transmit power can be set to one of six different levels by the processor. This is an important power-awareness feature that can be controlled at higher layers of the protocol stack. For example, the MAC protocol can change the radio range to include only the optimal number of nodes.

2.2. µAMPS-1 Node Performance

Figure 3 illustrates the effectiveness of the µAMPS-1 node's power-management controls while running an acoustic target-tracking application. The figure shows the instantaneous power consumption of a µAMPS-1 node as it collects data samples from the microphone, performs a line-of-bearing calculation on the collected data, and relays the results of this calculation to other nearby nodes. While collecting data, the processor alternates rapidly between active and idle modes, because no computation is needed in between data samples. When the computationally-intensive beamforming algorithm is run, the processor remains in active mode continuously. Power consumption rises sharply when the radio transmitter is enabled, and then falls rapidly as the entire node is put to sleep.

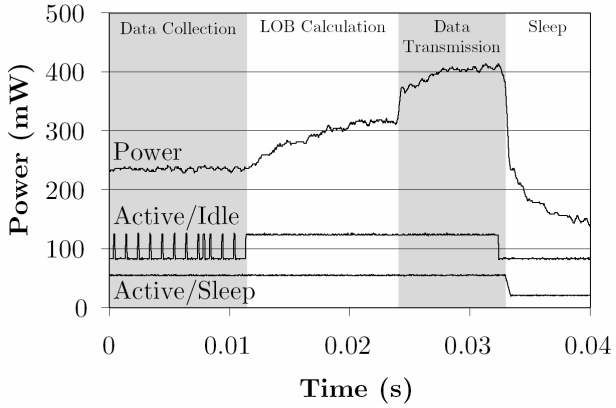

Figure 3. µAMPS-1 node power consumption during the execution of a beamforming application. The upper trace indicates node power consumption. The middle trace is high when the processor is active, and low when the processor is in idle mode. The lower trace goes low when the processor enters sleep mode. The node's power consumption continued to fall beyond the right edge of the graph, stabilizing at approximately 3.5 *mW*.

3. Challenges for Next Generation Nodes

The µAMPS-1 nodes demonstrated how the overall energy consumption of a node can be carefully managed by making use of the available power knobs. However, in order to meet the next generation design goal of 100 µW, the node must be much more aggressive in terms of rationing energy. Techniques such as dynamic voltage scaling must be extended into the subthreshold region. Because the node will spend most of its time in a sleep state, large savings can be made by reducing leakage currents and lowering V_{DD} in this state. The DSP must be adaptable to changing requirements for latency and bit precision. Since wireless communication is extremely costly to the energy budget, a power amplifier that scales gracefully over a wide dynamic range is required for optimizing radio links.

3.1. Circuit Level Power Reduction

3.1.1. Subthreshold Operation. Subthreshold circuits scale V_{DD} below V_T where subthreshold leakage currents (dis)charge load capacitances, limiting performance but giving orders of magnitude energy savings over nominal V_{DD} operation. Figure 4 shows the simulated performance (diagonal lines) and energy dissipation (thin curves) of an 0.18-µm CMOS 16-bit ripple carry adder. The dashed line indicates the (V_{DD}, V_T) points that minimize energy consumption for a given frequency. The curves show that subthreshold CMOS circuits are viable for some low performance wireless and medical applications [9]. Subthreshold circuits are highly sensitive to process variations and operating temperatures. Fine-grained threshold voltage control can counteract the effects of process and environmental variation using adaptive back biasing [10] or leakage controlled feedback circuits [11].

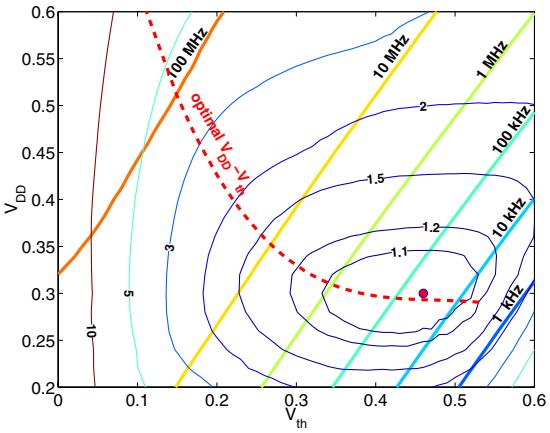

Figure 4. Average energy dissipation and performance for 16-bit adder operating in subthreshold.

3.1.2. Leakage Power Reduction. Process scaling at deep submicron levels increases subthreshold leakage, gate leakage, gate-induced drain leakage, and reverse biased diode leakage [12][13]. Multi-threshold CMOS (MTCMOS) is a popular technique for reducing standby leakage power by severing a circuit from the power rails with high V_T sleep devices [14]. Most MTCMOS designs use large sleep devices at the block level, but local sleep devices allow circuit partitioning into local sleep regions. Any unused circuit regions can enter sleep mode while

surrounding circuits remain active. This approach only provides savings if sneak leakage currents are prevented. Figure 5 shows how sneak leakage can occur in a data dependent way across hierarchical levels. The leakage path in the figure disappears if A=1, so non-comprehensive simulation might not detect it. A careful design methodology can prevent sneak paths. A fabricated 0.13-µm, dual V_T testchip employs such a methodology for a low-power FPGA architecture with over 8X measured standby current reduction [15]. The local sleep regions reduce active chip leakage by up to 2.2X (measured) for some configurations.

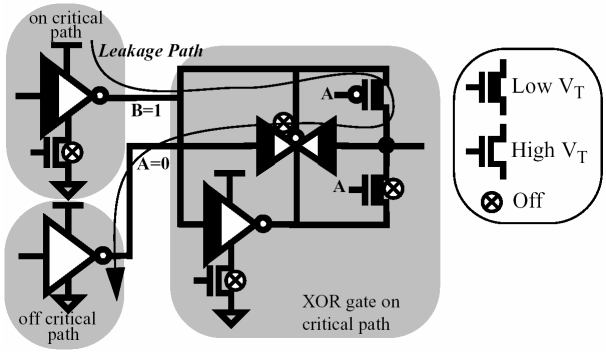

Figure 5. Sneak leakage paths in an MTCMOS circuit.

3.1.3. Standby Voltage Scaling for Leakage Reduction. Lowering V_{DD} during standby mode will reduce power by decreasing V_{DD} along with subthreshold current and gate leakage [16]. In most instances, V_{DD} scaling has a limit because circuit state must be preserved. One open-loop approach pinches in the rail voltages during standby using diode stacks and a power gating MOSFET, reducing V_{DD} by about 40% [17]. The open-loop approach cannot capitalize on the available savings. Figure 6 shows the measured dependence of current and power savings on V_{DD} for a 0.13-µm, MTCMOS testchip in sleep mode [18]. Since power savings increase rapidly for V_{DD} below about 200 mV, reducing V_{DD} to near the point where state is lost gives the best savings.

The critical path flip-flop (FF) on the testchip holds its state during sleep mode [19] down to 95 mV (measured). A bank of "canary" FFs sized to fail consistently at higher supply voltages than the core FFs show how close the critical path FFs are to failure under different environmental conditions, such as varying temperature. Closed-loop control of the standby voltage supply based on feedback from the canary flip-flops can lower the V_{DD} aggressively close to the minimum value without causing the critical path FFs to fail. This closed loop approach offers significant power savings even over an optimal open-loop approach without loss of state [18].

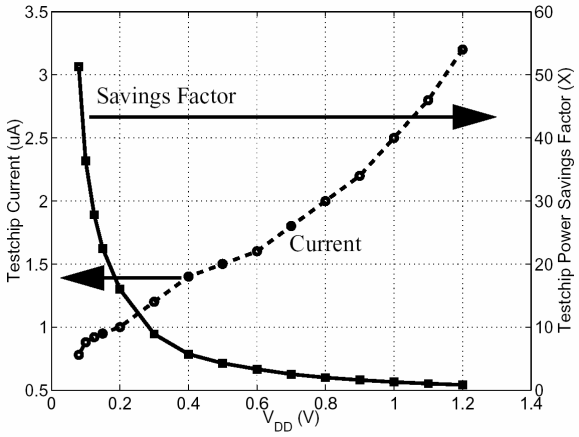

Figure 6. Measured testchip current and power savings versus V_{DD}.

3.2. Energy Scalable Computing

Energy efficient digital signal processors (DSPs) are becoming increasingly important in wireless sensor networks, where tens to thousands of battery-operated microsensors are deployed remotely and used to relay sensing data to the end-user. Given the constantly changing environments of portable devices and the extreme constraints on battery lifetimes, system level energy-aware design considerations should be taken into account. Energy-aware design is in contrast to low power design, which targets the worst case scenario and may not be globally optimal for systems with varying conditions. A new metric for design is to maximize energy-awareness [20]. The energy-awareness of a system can be increased by adding hardware to cover functionality over many scenarios of interest and by tuning the hardware such that the system is energy-efficient over a range of scenarios.

The Fast Fourier Transform (FFT) is an algorithm that is widely used in sensor and wireless applications for frequency domain beamforming, source tracking, harmonic line association, and classification. An energy-aware FFT will be able to adapt energy consumption as energy resources of the system diminish or as performance requirements change. Therefore it is advantageous to design the FFT with energy scalability hooks such as variable memory size and variable bit precision so that it can be used for a variety of scenarios. Our design focuses on a real-valued FFT (RVFFT) that can scale between 128-512-point FFT lengths and can operate at both 8 and 16-bit precision computation.

One example of an energy-scalable bit precision datapath design is showcased by the Baugh Wooley (BW) multiplier for two's complement arithmetic. In the RVFFT, there are four BW multipliers in a complex multiplication. A non-scalable design would optimize the worst case scenario by building a single 16-bit BW multiplier. A scalable BW multiplier design is shown in Figure

7. When 16-bit multiplication is needed, the entire multiplier is used. However, if only 8-bit multiplication is needed, the 8-bit feedthrough logic is enabled and the *Y* inputs are gated. Therefore only the adders in the lower left are driven, thus reusing hardware from the 16-bit multiplier [21].

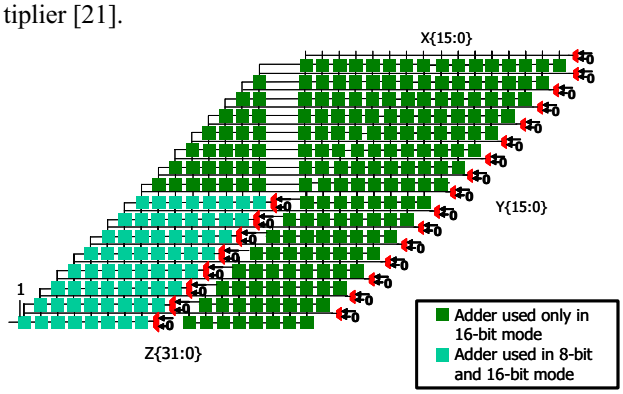

Figure 7. 8-bit and 16-bit scalable Baugh Wooley Multiplier using the reuse of point solution method. In the reuse method the 8-bit multiplier is reused for the 16-bit multiplication, thereby adding scalability without a large area penalty.

Design parameters and energy simulations of the different butterfly datapath architectures are shown in Table 1. The scalable architecture proves to be globally more energy efficient than a non-scalable solution. For the 8-bit butterfly datapath, the data and input gating to the unused hardware leads to a 66% energy dissipation reduction over a non-scalable design. However, for 16-bit computations, the energy dissipation of the scalable designs increases over the non-scalable design due to overhead logic. Nevertheless, scalable design is globally more energy-efficient than the non-scalable datapath if lower bit precision results are required.

Table 1. Comparing RVFFT datapath implementations

	Non-scalable	Scalable
Area	0.13 mm^2	0.14 mm^2
Transistor Count	47k	50k
Energy (16-bit, 512pt.)	331.9 nJ	341.2 nJ
Energy (8-bit, 512pt.)	287.9 nJ	96.4 nJ

3.3. Low Energy Communication Techniques

The energy required for radio communication scales with distance as d^2 to d^4. Since this path loss of radio transmission scales with distance in a super-linear fashion, communication energy may be reduced by dividing a long-distance transmission into several shorter ones. Intermediate nodes between a data source and destination can serve as relays that receive and rebroadcast data. This concept, known as multihop communication [22], is analogous to the use of buffers over a long, on-chip interconnect.

Figure 8 illustrates multihop communication to a base station across a distance d using h hops. The power consumed by multihop may be modeled as

$$P(h,d) = h[\alpha + \beta(d/h)^n]$$

where α is the sum of the distance-independent components of communication energy, such as the receiver, bias currents, and startup time, and β is the sum of the distance-dependent terms such as power amplifier and antenna losses, and n is the path loss exponent.

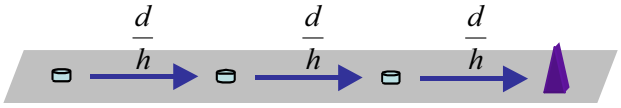

Figure 8. Multihop routing with *h* hops reduces the distance of each transmission by *h*.

The introduction of relay nodes is a balancing act between reduced $\beta(d/h)^n$ and increased α. Hops that are too short lead to excessive distance-independent overhead. Hops that are too long lead to excessive path loss. Between these extremes is an optimum transmission distance called the *characteristic distance*, d_{char} [23]. The characteristic distance, which determines the optimal number of hops, depends only on the energy consumption of the hardware and the path loss coefficient.

In order to demonstrate the tradeoff between multihop routing and direct transmission between two nodes, a prototype radio was built using commercial off-the-shelf components. The radio consists of a Chipcon CC1010 transceiver and a Maxim MAX2235 power amplifier (PA). The CC1010 is equipped with an internal PA capable of transmitting up to 4 dBm, and an integrated 8051 core microcontroller for handling basic protocol processing. The radio can be configured to use only the CC1010's PA, or to use the variable gain MAX2235 external PA, as shown in the block diagram in Figure 9. The maximum combined output power of the CC1010 and external PA together was measured at 27.7 dBm. The gain of the external PA can reduce its output power to below 4 dBm, however at these levels it is much more energy efficient to disable the external PA and use the CC1010 alone.

An experiment was conducted in an outdoor, line-of-sight environment to test the link distance of the radio. The antennas were held roughly 1.5 m above ground. Two radios were communicating at 915 MHz, using binary frequency shift keying (FSK). The power levels of the CC1010 PA and the external PA were varied via digital settings, and for each power setting the distance at which packets began to drop was recorded. The link was tested for the CC1010 PA alone, and for the CC1010 with external PA. RF and total DC power levels were later measured in a lab for each of the digital settings tested.

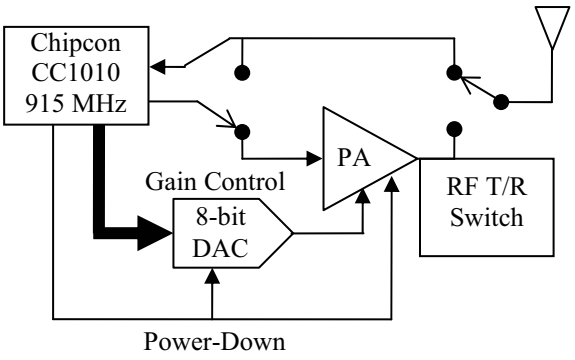

Figure 9. Block diagram of the transceiver and power amplifier.

Figure 10 characterizes the channel for the experimental scenario by plotting the required output power as a function of the distance covered by the transmission. With both axes plotted in a logarithmic scale, the data are amenable to a linear fit representing a path loss of approximately $n = 2.9$. As the radio range exceeded the dimensions of the test site at maximum power, the maximum-power point in Figure 10 was extrapolated for $n = 2.9$.

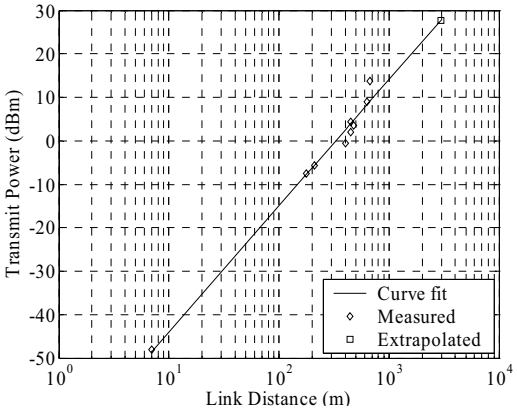

Figure 10. Channel charasteristic: radiated power (P_{rad}) versus maximum transmit distance.

Figure 11 plots the *total* energy consumed by the radio as a function of the distance covered and compares the energy-efficiency of direct transmission versus multihop transmission. The large step discontinuity within each plot signals the point at which the external PA must be turned on; the dramatic increase in power is due to the static current required by the external PA. From Figure 11, the crossover distance between h-hop and h+1-hop communication is equal to the distance at which the internal PA operates at maximum power. Beyond this point, the additional distance provided by the external PA does not sufficiently compensate for the additional, distance-independent energy overhead. Before this point, the distance-independent power consumption in the radio—even without the external PA—negates the benefits of reducing the power of the internal PA and introducing additional hops. The latter is a result common to most short-range radios (+20 dBm or less radiated) in common use today [24].

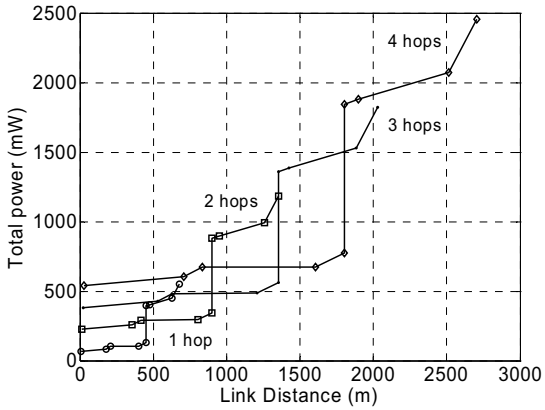

Figure 11. Multihop curves showing crossover points.

A protocol-level technique such as multihop has disadvantages. The increased protocol complexity of multihop may result in additional energy consumption. Since each data packet is received and rebroadcast by relay nodes, the end-to-end delay increases, and data throughput decreases. Moreover, reliance on multihop requires that relay nodes be in range of every transmitter, placing a minimum-density constraint on node placement. For applications with stringent latency and bandwidth requirements, or low or variable node density, multihop may not be an option.

To minimize the need for multihop routing, next-generation power amplifiers should be designed for high efficiency at a wider range of operating points, which would increase the range of distances for which direct transmission is energy-efficient. A series of power amps highly optimized for a small range of outputs is one potential solution, reflecting an "ensemble of point systems" that enable continuous, wide-ranging scalability [25].

4. Conclusion

The design and results of the first generation μAMPS-1 node, built using commercial off-the-shelf components, have been presented, demonstrating several power-aware techniques. The challenge of next generation nodes is to further reduce energy consumption by optimizing power-awareness over all layers of abstraction. This is accomplished using the methodologies for leakage reduction and subthreshold operation in both the active and sleep states, scalable hardware with minimal overhead, and a graceful scaling of radio performance. By applying these techniques, we believe the average power of a next generation node will meet the 100 μW design goal, opening up the

possibility of eliminating the battery and harvesting energy entirely from the environment.

5. Acknowledgments

Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) Power Aware Computting/Communication Program and Air Force Research Laboratory, under agreement numbers F30602-00-2-0551 and F33615-02-2-4005. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon; by Hewlett-Packard under the HP/MIT Alliance; and by Texas Instruments. Ben Calhoun is supported with an Infineon fellowship. Alice Wang is supported with an Intel Fellowship. The authors acknowledge Cypress Semiconductor for chip fabrication and Sarah Wentzloff.

6. References

[1] M. Bhardwaj, R. Min, A. Chandrakasan, "Power-Aware Systems," *34th Asilomar Conference on Signals, Systems, and Computers*, vol. 2, pp. 1695-1701, November 2000.

[2] C. Evans-Pughe, "Bzzzz [ZigBee wireless standard]," *IEE Review*, vol. 49, issue 3, pp. 28-31, March 2003.

[3] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Next Century Challenges: Mobile Networking for 'Smart Dust'," *Proc. Mobicom 1999*, pp. 271-278, 1999.

[4] J. Rabaey et al., "PicoRadio Supports Ad Hoc Ultra-Low Power Wireless Networking," *Computer*, pp. 42-48, July 2000.

[5] R. Amirtharajah and A. P. Chandrakasan, "Self-powered Signal Processing Using Vibration-based Power Generation," *IEEE JSSC*, pp. 687-695, May 1998.

[6] S. Meninger, J.O. Mur-Miranda, et al., "Vibration-to-Electric Energy Conversion," *IEEE Trans. on VLSI Systems*, pp. 64-76, February 2001.

[7] *Intel StrongARM SA-1110 Microprocessor: Advanced Developer's Manual*, Intel Corp., 1999. Available at http://developer.intel.com.

[8] R. Min, T. Furrer, and A. Chandrakasan, "Dynamic Voltage Scaling Techniques for Distributed Microsensor Networks," *Workshop on VLSI (WVLSI '00)*, April 2000.

[9] A. Wang et al., "Optimal Supply and Threshold Scaling for Subthreshold CMOS Circuits," *Proc. Symposium on VLSI*, pp. 5-9, 2002.

[10] S. Narendra, D. Antoniadis, V. De, "Impact of Using Adaptive Body Bias to Compensate Die-to-die Vt Variation," *ISLPED*, vol. 40, no. 10, pp. 1895-1897, Oct. 1993.

[11] H. Soeleman et al., "Robust Subthreshold Logic for Ultra-Low Power Operation," *Trans. on VLSI Systems*, Vol. 9, No. 1, February 2001.

[12] A. Keshavarzi et al., "Intrinsic Leakage in Low Power Deep Submicron CMOS ICs," *Int'l Test Conf.*, 1997.

[13] S. Borkar, "Design Challenges of Technology Scaling," *IEEE Micro*, vol. 19, no. 4, pp. 23-29, July-Aug. 1999.

[14] S. Mutoh et al., "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," *JSSC*, vol. 30, no. 8, Aug 1995.

[15] B. Calhoun et al., "Design Methodology for Fine-Grained Leakage Control in MTCMOS," *ISLPED*, 2003.

[16] R. Krishnamurthy et al., "High-performance and low-power challenges for sub-70nm microprocessor circuits," *CICC 2002*, 2002.

[17] T. Enamoto et al., "A self-controllable-voltage-level (SVL) circuit for low-power, high-speed CMOS circuits," *ESSCIRC 2002*, 2002.

[18] B. Calhoun, A. Chandrakasan, "Standby Voltage Scaling for Reduced Power," *CICC*, 2003.

[19] J. Kao, A. Chandrakasan, "MTCMOS Sequential Circuits," *Proc. ESSCIRC 2001*, 2001.

[20] M. Bhardwaj, R. Min, A. Chandrakasan, "Quantifying and Enhancing Power-Awareness of VLSI Systems," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 9, issue 6, pp. 757-772, December 2001.

[21] A. Wang, A. Chandrakasan, "Energy-Aware Architectures for a Real-Valued FFT Implementation," *IEEE Intl Symposium on Low Power Electronics and Design (ISLPED 03)*, Seoul, Italy, August 2003.

[22] L. Kleinrock, "On giant stepping in packet radio networks," *Packet Radio Temporary Note #5 PRT 136*, UCLA, Los Angeles, CA, March 1975.

[23] M. Bhardwaj, T. Garnett, A. Chandrakasan, "Upper Bounds on the Lifetime of Sensor Networks," *IEEE International Conference on Communications*, pp. 785-790, 2001.

[24] T. Basten et al., "Scaling into Ambient Intelligence," *Proc. 2003 Design, Automation, and Test in Europe*, Munich, Germany, March 2003.

[25] M. Bhardwaj, R. Min, A. Chandrakasan, "Quantifying and Enhancing Power-awareness of VLSI Systems," *IEEE Transactions on VLSI Systems*, pp 757-772, December 2001.