

UCLA

Papers

Title

Energy Harvesting Aware Power Management

Permalink

<https://escholarship.org/uc/item/1qx6t77n>

Authors

Kansal, Aman
Srivastava, Mani B

Publication Date

2005-05-05

Peer reviewed

Chapter 9

Energy Harvesting Aware Power

Management

9.1 Introduction

The true autonomy of wireless sensor networks depends on their reliable operation for extended times without human intervention. Energy supply is a critical factor in this design. Wireless and ad-hoc deployment, which is essential in some scenarios and cost-effective in others, precludes the use of a wired energy infrastructure. The sensor nodes are thus forced to operate on limited battery reserves and low power design is an important design consideration [1, 2].

Unlike human carried devices such as hand-helds or cell-phones which can be returned to charging docks periodically, sensor node batteries are limited in supply. They cannot be replaced in the large numbers of nodes as the embedded nature of deployment makes it hard to access each individual node. A limited amount of energy supply, however, is not sufficient to ensure uninterrupted operation for the several year long lifetimes typically expected from an embedded deployment. The small node size puts constraints on the

maximum battery size. Batteries already dominate the node volume in prototype sensor nodes. The energy density for common battery technologies varies in the range of 1200 J/cu.cm. (Alkaline) to 3780 J/cu.cm (Zinc-air). At such energy density, assuming a sensor node operating at 1mW (average consumption after power management) and assuming the full battery capacity can be utilized, a year-long operation requires a battery-size of the order of 10cu.cm which is rather large. Thus, batteries alone cannot be expected to reliably supply a sensor network deployment for several years.

A viable alternative then is to endow the sensor nodes with appropriate energy harvesting technologies such as solar, vibrational [3], wind/water flow, thermal gradient scavenging [4 - 6], electromagnetic direct conversion [7] and others [8 - 11]. These sources can supplement or even entirely replace the battery energy supply. A fundamental difference between environmental energy and battery supply is that the environmental energy can be scavenged for as long as desired and if efficiently utilized, can enable a system to last eternally (until its hardware is outdated).

However, the introduction of harvesting components into sensor nodes requires design changes spanning the hardware of the node, the node level power management, and network wide energy scheduling. Consider for instance, that the power consumption at a node is to be matched to the environmental energy available to it. Then, the node must be equipped with additional hardware to measure the environmental energy input, rather than just the residual battery measurement. Further, power management decisions in a network differ when harvested energy is available. For instance, consider a solar energy harvesting network. At a particular instance, the network may have two alternative data routes available to satisfy the immediate data transfer requirement. However, the harvested energy available at nodes along the two routes may be different, say due to the presence of a shadow on part of the network. In such a case, the network requires a harvesting aware power management strategy, along with the environmental energy input measurement, which allows it to choose the route passing through nodes outside the shadow.

The remainder of this chapter addresses the resulting issues in harvesting aware system design, including both the hardware and power management software.

9.2 Harvesting Technologies

We define a *harvesting node* to be any sensor node with at least one form of environmental energy harvesting as part of its power supply. Typically, such a device will also have an energy storage mechanism, such as a battery or an ultracapacitor, to allow energy harvesting and consumption to occur without total synchrony. However, the storage device is not essential in all scenarios, and harvesting nodes such as described in [10] use the energy generated from the press of a button immediately to transmit a packet and are inactive otherwise.

A network of such devices will be referred to as a harvesting network. The design of the network involves further considerations than just the individual nodes, since the nodes in such a network may be heterogeneous and the environmental energy available at each node may be different.

We consider the harvesting node first. Figure 9.1 shows the new modules that are part of a harvesting node in addition to the usual sensor node components. The various blocks are discussed in the subsequent sections, including an example implementation. While most of the blocks shown are implemented as hardware circuits, the block labeled Harvesting Aware Power Management is best implemented as a set of algorithms on the sensor node processor itself. A design component not shown in the figure but which influences the hardware and software design is knowledge about the application behavior with respect to energy consumption and this knowledge should be exploited to the extent available.

Figure 9.1: Block diagram of a harvesting node.

The key distinguishing characteristic of the harvesting node appears at the top left in the block diagram - the harvesting device. It is based on one of the harvesting technologies relevant to the deployment environment and outputs energy. The key harvesting technologies of interest to embedded systems are discussed below.

9.2.1 Solar Energy

Solar or other light energy can be converted to electric power using solar cells. The magnitude of energy generated varies from approximately 15mW/sq.cm in noon-time sunlight to 10 μ W/sq.cm in indoor incandescent lighting. The energy output depends on the material used. Crystalline materials such as silicon and gallium arsenide have moderate absorption efficiency and high conversion efficiency (15-30%) while thin film materials such as cadmium telluride have high absorption efficiency and lower conversion efficiency (\leq 10%). The choice of material also depends on its spectral response and the light source of interest.

For the purposes of circuit design, the solar cell may be modelled as a voltage source with an internal resistance. The output voltage is fairly constant in the useful operating range and the supply current varies with light intensity. A single solar cell output is 0.6V but panels with series of such cells can generate any required voltage for the circuit.

9.2.2 Vibration Energy

Vibrations are available in many environments of interest including commercial buildings, parking structures, aircrafts, trains, industrial facilities and even residential buildings. Preliminary analysis and experiments presented in [3] show that 300 μ W/cm³ is available in such environments. The sources of vibrations which may be heavy machinery, home appliances, HVAC vents, movement of people or vehicles, and other movements vary a great deal in their acceleration characteristics and the frequency spectra [3]. Methods to convert this energy to electricity can be classified into electromagnetic, electrostatic and piezoelectric [12 - 16].

Electromagnetic conversion uses vibration to move a conductor in a magnetic field. Existing prototypes [15, 12] generate very low voltage output to be usable. Electrostatic conversion uses vibration energy to move the conductors of a charged capacitor. The disadvantage of this approach is that a separate voltage source is required to charge the capacitor. An advantage however is that the output voltage is in the usable range of two to several volts. Piezoelectric conversion uses materials which when mechanically deformed generate an electric potential. The piezoelectric method combines the advantages of electromagnetic and

electrostatic conversion but are difficult to implement at micro-scale. With the current technology, they have the greatest available energy density among the three methods.

9.2.3 Other sources

Wind or water flow can be converted to energy. While macro-scale generators based on these flows are widely used, compact technologies to extract such energy are lacking. A sensor networking application of wind energy is also for locomotion, such as used in NASA's Tumbleweeds [17]. These are inflatable spheres which can roll along the deployment surface using wind energy, and are aimed at Martian and polar exploration.

Thermoelectric generation using Seebeck effect (flow of current in a loop made from two wires of certain metals when a temperature difference is applied to the wire junctions) and other methods have been demonstrated to yield $10\mu\text{W}/\text{cm}^2$ to $40\mu\text{W}/\text{cm}^2$ using a 5-10 degree Celcius temperature gradient [18, 6, 4].

Pressure variations, such as the pressure of a fluid or gas in an enclosed space changing with time of day, can also be used to generate energy. In the Atmos clock for example, invented in 1928 by Jean-Leon Reutter, a mixture of gas and liquid enclosed in a sealed capsule expands as the temperature rises and contracts as it falls, moving the capsule back and forth providing sufficient motion to run the clock. Methods to convert such limited pressure variation or motion into electricity are not readily available, however.

The exact choice of the harvesting technology depends several factors, including on the achieved energy density from a particular technology, sensor node form factor, and most importantly, the availability in the deployment scenario. For instance, a sensor network deployed for environmental monitoring or precision agriculture in outdoor settings may have ready access to sunlight and could hence use solar energy, while a sensor network operating indoors for industrial applications such as machine health monitoring may have plenty of vibrations available for harvesting energy.

9.3 Descriptions of the Components of a Harvesting Node

Figure 9.1 alluded to various new components that are required when a sensor node is converted into a harvesting node. This section discusses the reasons why these blocks are needed and their design requirements.

The *recharging circuit* receives part of the energy output of the harvesting device and stores it. This block is needed since the current and voltage characteristics of the harvesting device output may not be directly suitable for charging the storage device. For example, a rechargeable battery must be charged a voltage equal to or higher than its output, and if the harvesting device outputs current at a lower voltage than this, a voltage conversion circuit may have to be used.

The recharging circuit also depends on the energy storage mechanism, and in case of batteries, on the battery chemistry used. For a battery charged from a solar cell, the simplest implementation is to directly connect a solar panel to the battery through a diode to prevent the battery being drained through the solar cell. However, this circuit would yield low longevity for the battery as the battery could be repeatedly damaged from overcharge (continued charging of battery even when it is charged to its full capacity) and undercharge (continued discharging after the battery voltage has dropped below a battery chemistry specific threshold). In general, to ensure battery longevity, a recharging circuit is required to produce an acceptable charging profile and to disconnect the battery at its under-charge limit.

The *consumption arbiter* uses a combination of the energy from the harvesting source and the storage device to supply the power requirements of the harvesting node. There are three reasons which necessitate this arbiter.

First, the harvested energy may not always be sufficient to power the load and hence a mechanism is needed to share the load between the harvested and stored energy. Storage efficiency is always less than 1. The arbiter must ensure that the environmental energy is supplied to the load directly, and only the deficit is drawn from the storage. When the energy available from the harvesting device exceeds the consumption, the excess energy is routed to the recharging circuit, which may store it if the storage is not already full.

Second, for a given environmental stimulus, the energy output of the harvesting device depends on its own internal resistance and the characteristics of the load presented to it. The consumption arbiter may also modify the load characteristics appearing at the output of the harvesting device to maximize the scavenged energy.

Third, the storage mechanism used may not supply the energy at the voltage required by the sensor node and the arbiter regulates the supply voltage.

The arbiter may also be designed to prioritize storage of energy in certain cases such as when the battery has drained to a very low level. In such a case, if the harvested energy level is just sufficient to power the sensor node but no excess is available to charge the battery, then the arbiter may shut off the node and charge the battery first, in order to ensure system availability in emergency situations.

The *energy tracker* monitors the energy available from the harvesting device and also the current state of the energy store. Such data may be used by the harvesting aware power management algorithms for learning the energy environment. The instantaneous battery status and environmental energy availability along with the long term availability behavior is used by several network-wide scheduling algorithms to distribute workload as discussed later in sections on power management.

The *sub-module power switching block* is typical to any device which provides for shutting down parts of its circuit for power management. This is very crucial in sensor networks where sensors and minimal direct memory access peripherals may be kept active for much longer durations than the energy intensive communication, processing and larger memory modules. In a harvesting node, this block is useful for turning on or off various components as the amount of environmental availability varies. This switching functionality is controlled by the harvesting aware power management algorithms.

A prototype harvesting node, implemented at the Networked and Embedded Systems Laboratory, University of California Los Angeles, is shown in Figure 9.2. This node uses a pair of solar cells connected in parallel as the harvesting device. The storage is a pair of AA sized NiMH re-chargeable batteries. The recharging circuit provides the charging profile required for NiMH batteries. The arbiter supplies the solar

cell output energy to the load while the excess or deficit flows through the battery. The heliomote arbiter also regulates the voltage output of the NiMH battery pack which varies between 2.2 and 2.8V to a stable 3V required by the sensor node. The energy tracker module monitors the total energy charged into the battery, the instantaneous battery current and voltage. The measured data is provided to the sensor node processor via a one-wire interface. The detailed circuit schematics are available in [19].

Figure 9.2: Heliomote.

9.4 Harvesting Aware Power Management

Power management strategies differ significantly when the power source changes from a fixed battery supply to a harvesting device. There are two reasons which lead to these differences. First, the environmental energy source is highly variable. Unlike the stored supply which is simply characterized by the amount of residual energy of the battery and is reliably available, the environmental energy requires a more sophisticated characterization. Second, the environmental energy is not a limited resource and has the potential to be used eternally. Below, we present a model to characterize the production and consumption of environmental energy which can be used to design power management algorithms for harvesting nodes.

9.4.1 Harvesting Theory

Since the harvesting source is highly variable, an important design consideration is to provide a reliable system performance. Another objective is to decide if the environmental energy alone is sufficient for the system to operate at the desired level of performance. If so, the system can operate indefinitely. Otherwise, the designer may want to determine the required harvesting additions. The goal of developing a harvesting theory is to enable analysis of a wide variety of harvesting technologies with respect to system performance. We begin with a source characterization.

Definition 1. $(\rho, \sigma_1, \sigma_2)$ – source : Suppose $E(t)$ is a continuous and bounded function of a continuously varying parameter t . $E(t)$ is said to be a $(\rho, \sigma_1, \sigma_2)$ – source if and only if for any finite real number T , it satisfies:

$$\int_T E(t)dt \geq \rho T - \sigma_1 \quad (9.1)$$

$$\int_T E(t)dt \leq \rho T + \sigma_2 \quad (9.2)$$

This definition uses only three parameters keeping it analytically tractable but still allowing it to model a wide variety of variations in energy sources. $E(t)$ models the power output at time t . The model captures the asymptotic rate of availability, which is the maximum power at which the system can operate. Since we are modelling physical energy sources, the restrictions placed on the function $E(t)$ are justified. It may be noted that the unit of ρ is power, e.g. Watts and the unit of σ_1 and σ_2 is energy, e.g. Joules.

The consumption itself may not be constant and the following definition can characterize most consumption profiles.

Definition 2. (ρ', σ) – consumer : A device is said to be a (ρ', σ) consumer if its power consumption, $E_c(t)$, satisfies the constraint

$$\int_T E_c(t)dt \leq \rho' T + \sigma \quad (9.3)$$

for any value of T .

With this definition, the following theorem specifies the minimum performance that can be guaranteed. The achievable performance may in fact be higher, such as if an application requires node operation only when the environmental energy is available.

Theorem 1. *Sustainable Performance at Eternity (Variable Consumption Profile):* If a (ρ', σ) -consumer device is powered by a $(\rho, \sigma_1, \sigma_2)$ -source, has an energy storage capacity of $\sigma + \sigma_1 + \sigma_2$, and $\rho' < \rho$, then the device can operate forever.

The proof is available in [20]. The following example demonstrates an immediate application of the above theorem.

Table 9.1: Solar cell parameters in experimental environment

Parameter	Value	Units
ρ	23.6	mW
σ_1	1.4639×10^6	J
σ_2	1.8566×10^6	J

Example 1. Figure 9.3 shows the power flowing into the battery observed by a test heliomote, when placed in sunlight, in the month of January in the Northern hemisphere. Negative values represent battery drain.

Figure 9.3: Solar energy based charging power recorded for 9 days.

Measuring the battery current instead of the current out of the solar cell ensures that only the actual solar power available, and not any power lost due to circuit inefficiencies, is considered. Since the negative portions of the waveform dominate, the batteries suffer a net discharge. For the waveform plotted in Figure 9.3, the source characterization parameter values are given in Table 9.1.

With these values the solar cell is a $(\rho, \sigma_1, \sigma_2)$ -source in the test environment. We now need a (ρ', σ) classification of the consumer node. The sensor node used in the heliomote has a sleep mode power drawn $P_{sleep} \leq 3\text{mW}$ and the maximum current drawn $P_{max} = 100\text{mW}$. Thus, $(\rho' = 100\text{mW}, \sigma = 0)$ is a valid classification. However, to achieve $\rho' = \rho$, we can set the node to sleep for 78.7% of the time. If the minimum wake up duration is 2s, this leads to $\sigma = 153\text{J}$. Theorem 2 implies that a rechargeable battery with capacity $= \sigma + \sigma_1 + \sigma_2 = 3.32 \times 10^6\text{J}$, or 922.43mAh, is required for sustaining the node indefinitely from solar energy. This helps the designer to choose a AAA sized NiMH battery. Using a larger battery does not help improve long term performance.

The key utility of the above theory is that it enables characterizing the energy availability using a small number of parameters which can be used for power management as discussed next.

9.4.2 Scheduling Algorithms

In battery powered systems, the power management strategies are designed merely to minimize the energy consumption. Several power scaling strategies to minimize battery consumption have been studied [21 - 24]. Many of these can be used to minimize battery drain in harvesting nodes also. However, these strategies are not sufficient for efficiently utilizing the harvesting source. In harvesting systems, the problem space expands. The aim is not always to minimize consumption but adapt the consumption to the available environmental energy within the performance constraints. For instance, the tasks may be rescheduled when possible to operate directly off the environmental energy. This does not reduce the energy consumed in executing the task, but reduces the wastage due to battery inefficiency. Research has recently begun in the area of harvesting aware power management [25, 26].

Various power scaling mechanisms are typically available to control the power consumption in embedded systems, such as dynamic voltage scaling (DVS), transmit power control, and duty cycling the node. For each mechanism, reducing power affects the performance. For DVS and duty cycling for instance, using a lower energy mode increases the latency of executing a task. The algorithms below are illustrated using duty cycling as an example power scaling mechanism since it is available in most processors.

Theorem 1 gives the asymptotic rate of energy consumption available to a harvesting node. As mentioned, a fundamental difference in harvesting systems compared to battery powered systems is that if the system performance is adapted to the environmental availability, the system can last eternally. To adapt the performance in this manner, the system only needs to estimate the parameters ρ , σ , σ_1 , and σ_2 for its deployment scenario. This can be achieved if the environmental behavior over a learning phase is assumed to be representative of the long term behavior, which is valid for many sources. For instance, solar energy follows a diurnal and annual cycle. Winds have known repetitive patterns. The sustainable performance depends directly on ρ and hence we wish to estimate ρ to within an error margin Δ . The power parameter ρ can be estimated by averaging the energy obtained from the energy source over time. Let the average at time t , $\rho(t)$

be calculated as:

$$\rho(t) = \frac{\int_0^t E(t)dt}{t} \quad (9.4)$$

Since the device will sample $E(t)$ at discrete times, the above integral will be evaluated as a discrete summation. Apart from the running estimate of average, also store the most recent local minima and maxima observed in $\rho(t)$. When the difference between the maxima and the minima reaches Δ we assume that estimation of ρ is complete. The $E(t)$ waveform can also be used to estimate σ_1 and σ_2 [20].

Once the source has been characterized, the device can adjust its performance to operate at the available rate ρ . Suppose active mode power is P_{max} and sleep mode power is P_{sleep} . Then the duty cycle x satisfies:

$$\rho = xP_{max} + (1 - x)P_{sleep} \quad (9.5)$$

neglecting the energy consumption of switching between modes. The value of x determined from this equation can be used to decide the sleep duration if the minimum time spent in active mode, T_{min} is known. Thus, performance can be adjusted at a single node to match the environmental availability.

9.5 Distributed Harvesting

In a distributed system with several harvesting nodes, the harvesting aware power management problem would be to distribute the workload among the nodes in such a way that the overall performance of the system is maximized. Finding the optimal solution requires complete knowledge of what energy resources are available at every node and the complete set of tasks to be performed by the system. Many of these tasks, such as routing a data packet from one node to another involve energy consumption at several intermediate nodes, making the scheduling decisions coupled among nodes. Also, in a distributed system, scalability concerns dictate that all the information at every node not be communicated to a central node for scheduling decisions. Moreover, in many distributed systems such as sensor networks, communication itself is the major energy consumer and distributed decisions are the assumed norm. Distributed power management solutions are hence used, even though they may not be globally optimal.

Distributed power management is a nascent area of research. Methods have been proposed for distributed battery powered systems [27 - 31]. These methods attempt to reduce and balance the energy usage across multiple nodes to ensure system-wide sustainability rather than minimizing the total energy consumption which could cause certain nodes to deplete faster and make the overall system useless. In harvesting networks on the other hand, instead of balancing the energy consumption among nodes, the objective is to share the workload in proportion to the harvesting opportunity available to each node. The distribution of workload may be highly application-dependent and such methods typically require an understanding of the application characteristics. One such method which addresses a typical usage scenario is discussed below.

One of the applications for sensor networks is to monitor a deployment scene for specific events and report the occurrence of such events with low latency. The sensing transducers consume minimal energy and are kept active, while the power intensive modules- processor and radio, are woken up when an event is triggered by a transducer. The event must now be communicated with low latency to a central base station. The power management problem is to provide an energy efficient communication topology for the network, which adapts the energy consumption at different nodes to their individual harvesting opportunities. The scheme below achieves this in a completely distributed fashion. Also, it can report the expected performance to the base with minimal communication overhead.

9.5.1 Communication Protocol

The communication protocol followed is depicted in Figure 9.4. When a node has data to transmit to its next hop neighbor along the data path, the node transmits a BEACON packet and listens for response for a period T_{ack} . It repeats this process until an ACK is received. T_{ack} is the time required by an active node to send an ACK after receiving a BEACON packet. Suppose the time required to transmit a BEACON packet is T_{beacon} . Every node in the network wakes up for a duration $T_{min} = 2T_{beacon} + T_{ack}$ to listen for any BEACON messages from nodes attempting to send data to it. After every awake period, it sleeps for a

duration

$$T_{sleep} = \frac{1-x}{x} T_{min} \quad (9.6)$$

where x is the sustainable duty cycle at this node, determined from (5). Whenever a node receives a BEACON packet, it transmits an ACK and stays awake until it has received the relevant data and forwarded it to the next hop. With this arrangement, the worst case delay in receiving an ACK when repeatedly sending a BEACON is $T_{sleep} + T_{min}$. The energy spent in transmitting the event data is negligible compared to the energy spent in periodically entering active mode to listen for potential data when events are infrequent.

Figure 9.4: Communication with sleep cycle.

The route discovery from all nodes to the base is based on the formation of a data gathering tree, adopted from [32]. The base node transmits an INIT packet. All nodes which receive this packet, treat the sender of the INIT packet as their parent and send an ACK with random back-off delay. They now know the route to the base, which is one hop away. These nodes assign themselves depth = 1. They retransmit the INIT message with their own ID. All nodes which have not already assigned themselves a depth acknowledge this packet. These nodes now know their next hop on the route to the base and assign themselves depth = 2. The process continues. When a node transmits an INIT message but does not receive any responses for a timeout duration, it assumes that there are no nodes deeper than itself along the path that passes through it. Such nodes are denoted leaf nodes.

9.5.2 Distributed Power Management

The network starts at an arbitrary performance level. Once the estimation of ρ, σ_1 and σ_2 parameters is completed, using the same methods as discussed in the previous section, the nodes will begin transition to the sustainable performance level. Each node sets its response latency, L , equal to $T_{min} + T_{sleep}$. If a node is not receiving any energy from the environment, it may have to select a preset value of L for itself. When a leaf node has estimated its L , it sends a LATENCY message to its parent containing its estimate of L . A

node which is not a leaf, waits for the LATENCY packet from all its children. Let L_i denote the latency value heard from i -th child and let it have N_c children. When it has collected all LATENCY packets from its children and has estimated its own L , it sets

$$L' = \max_{i \in [1, \dots, N_c]} \{L_i\} + L \quad (9.7)$$

It sends this cumulative value of L' to its parent. L' represents the total worst case data transfer latency along the path through this node from the worst leaf. The process repeats. The base then chooses L_{max} equal to the largest L' among the ones it receives from its children. This L_{max} is the maximum latency of data transfer in the network along the worst path. Each node switches to the duty cycle based on its estimated ρ after having sent the LATENCY message. Thus the nodes gradually transition to their individual performance levels, with no centralized coordination.

The in-network processing of the transmitted L values at each parent reduces the data sent to the parent nodes and the amount of data sent to the base node is thus proportional to the number of depth 1 nodes and not to the total number of nodes in the network. This ensures scalability while at the same time returning a network wide performance metric to the base. If the achieved performance is not acceptable, additional resources may have to be added.

The above method provided algorithms to adapt the power consumption on a network wide basis to the local energy availability at the various network components. Further research is underway to exploit more detailed information about the energy availability profile, including its temporal variations and spatial patterns to optimize performance in given energy environments.

9.6 Conclusions

Energy harvesting provides a viable alternative to create long lasting wireless sensor networks, using harvested energy to supplement or replace stored battery reserves. Harvesting is not only useful but essential in some usage scenarios such as space exploration and extra-terrestrial sensor network deployments where

the operating lifetime required is very long and battery replacements are next to impossible. The unmanned vehicles exploring the surface of Mars use solar energy harvesting for all their energy needs. Different types of harvesting opportunities exist in most deployment scenarios and this form of energy should be actively exploited for enhanced system performance. The use of harvested energy leads to several research issues which are different from conventional power management based on only stored battery energy. Not only is the design of the hardware different to account for the harvesting source, but workload scheduling now depends on the nature of the environmental source. The scheduling is also different in case of distributed systems with multiple harvesting nodes. Harvesting theoretic techniques can be used to determine the sustainable performance levels and additional resources required for achieving requisite performance. These techniques also facilitate the design of distributed scheduling methods. An ideal system would utilize its starting battery resources and harvested energy to achieve the maximum possible application throughput with the available resources. Such power management for harvesting networks is an open problem with a rich set of research challenges.

Bibliography

- [1] Raghunathan, V . , C. Schurgers, S. Park, and M. Srivastava, “Energy aware wireless microsensor networks,” *IEEE Signal Processing Magazine*, Vol. 19, No. 2, 2002, pp. 40–50.
- [2] Min, R., M. Bhardwaj, S. Cho, A. Sinha, E. Shih, A. Wang, and A. Chandrakasan, “An architecture for a power-aware distributed microsensor node,” in *IEEE Workshop on Signal Processing Systems (SiPS '00)*, October 2000.
- [3] Roundy, S . , P. K. Wright, and J. M. Rabaey, *Energy Scavenging for Wireless Sensor Networks with Special Focus on Vibrations*. Kluwer Academic, 2004.
- [4] Stordeur , M . , and I. Stark, “Low power thermoelectric generator- self sufficient energy supply for microsystems,” in *16th International Conference on Thermoelectrics*, pp. 575–577, 1997.
- [5] Weber, W . , “Ambient intelligence: industrial research on a visionary concept,” in *Proceedings of the 2003 international symposium on Low power electronics and design*, pp. 247–251, ACM Press, 2003.
- [6] Pescovitz, D . , “The power of small tech.” *Smalltimes*, Vol. 2, No. 1, 2002.
- [7] “Darpa energy harvesting projects.” <http://www.darpa.mil/dso/trans/energy/projects.html>.
- [8] Kymisis, J . , C. Kendall, J. Paradiso, and N. Gershenfeld, “Parasitic power harvesting in shoes,” in *Second IEEE International Conference on Wearable Computing, (ISWC)*, pp. 132–139, IEEE Computer Society Press, October 1998.

- [9] Shenck, N. S., and J. A. Paradiso, "Energy scavenging with shoe-mounted piezoelectrics," *IEEE Micro*, Vol. 21, May-June 2001, pp. 30–42.
- [10] Paradiso, J. A., and M. Feldmeier, "A compact, wireless, self-powered pushbutton controller," in *ACM Ubicomp*, (Atlanta, GA, USA), pp. 299–304, Springer-Verlag Berlin Heidelberg, September 2001.
- [11] Starner, T., "Human-powered wearable computing," *IBM Systems Journal*, Vol. 35, No. 3-4, 1996.
- [12] Williams, C. B., C. Shearwood, M. A. Harradine, P. H. Mellor, T. S. Birch, and R. B. Yates, "Development of an electromagnetic micro-generator," *IEEE Proceedings- Circuits, Devices and Systems*, Vol. 148, December 2001, pp. 337–342.
- [13] Ottoman, G. K., H. F. Hofmann, and G. A. Lesieutre, "Optimized piezoelectric energy harvesting circuit using step-down converter in discontinuous conduction mode," *IEEE Transactions on Power Electronics*, Vol. 18, No. 2, 2003, pp. 696–703.
- [14] Glynn-Jones, P., S. P. Bebb, E. P. James, and N. M. White, "The modeling of a piezoelectric vibration powered generator for microsystems," in *Transducers 01/Euroensors XV*, June 2001.
- [15] Amitharajah, R., and A. P. Chandrakasan, "Self powered signal processing using vibration-based power generation," *IEEE Journal of Solid State Circuits*, Vol. 33, No. 5, 1998, pp. 687–695.
- [16] Shearwood, C., and R. B. Yates, "Development of an electromagnetic microgenerator," *IEE Electronics Letters*, Vol. 33, October 1997, pp. 1883–4.
- [17] "Exploring mars: Blowing in the wind." <http://www.jpl.nasa.gov/technology/features/tumbleweed.html>, 2001.
- [18] Whalen, S., M. Thompson, D. Bahr, C. Richards, and R. Richards, "Design, fabrication and testing of the p3 micro heat engine," *Sensors and Actuators*, Vol. 104, No. 3, pp. 200–208.
- [19] "HelimoteR." <http://cvs.nesl.ucla.edu/cvs/viewcvs.cgi/HeliumoteR/>. CVS Repository.

- [20] Kansal, A . , D. Potter, and M. Srivastava, "Performance aware tasking for environmentally powered sensor networks," in *ACM Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, June 2004.
- [21] Simunic, T . , H. Vikalo, P. Glynn, and G. D. Micheli, "Energy efficient design of portable wireless systems," in *ACM/IEEE International Symposium on Low Power Electronic Design*, 2000.
- [22] Min, R . , T. Furrer, and A. Chandrakasan, "Dynamic voltage scaling techniques for distributed microsensor networks," in *IEEE Computer Society Workshop on VLSI*, pp. 43 – 46, April 2000.
- [23] Schuette , M . A . , and J. R. Barr, "Embedded systems design for low energy consumption," in *Proceedings of the 1994 IEEE/ACM international conference on Computer-aided design*, pp. 534–540, IEEE Computer Society Press, 1994.
- [24] Simunic, T . , L. Benini, G. D. Micheli, and M.Hans, "Source code optimization and profiling of energy consumption in embedded systems," in *IEEE International Symposium on System Synthesis (ISSS)*, 2000.
- [25] Kansal , A . , and M. B. Srivastava, "An environmental energy harvesting framework for sensor networks," in *International symposium on Low power electronics and design*, pp. 481–486, ACM Press, 2003.
- [26] Voigt, T . , H. Ritter, and J. Schiller, "Utilizing solar power in wireless sensor networks," in *The 28th Annual IEEE Conference on Local Computer Networks (LCN)*, (Bonn/Konigswinter, Germany), October 2003.
- [27] Rong , P . , and M. Pedram, "Extending the lifetime of a network of battery powered mobile devices by remote processing: A markovian decision based approach," in *ACM Design Automation Conference (DAC)*, (Anaheim, CA, USA), June 2003.

- [28] Chang, J. H., and L. Tassiulas, "Maximum lifetime routing in wireless sensor networks," in *Advanced Telecommunications and Information Distribution Research Program (ATIRP)*, (College Park, MD, USA), March 2000.
- [29] Li, Q., J. Aslam, and D. Rus, "Online power aware routing in wireless ad-hoc networks," in *ACM SIGMOBILE*, (Rome, Italy), July 2001.
- [30] Xu, Y., J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," in *Proc. ACM Mobicom*, pp. 70–84, ACM Press, 2001.
- [31] Schurgers, C., V. Tsiatsis, S. Ganeriwal, and M. Srivastava, "Optimizing sensor networks in the energy-density-latency design space," *IEEE Transactions on Mobile Computing (TMC)*, January-March 2002, pp. 70–80.
- [32] Intanagonwiwat, C., R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *ACM Mobicom*, August 2000.