Design and Power Management of Energy Harvesting Embedded Systems

Vijay Raghunathan
NEC Labs America
Princeton, NJ 08540
vijay@nec-labs.com

Pai H. Chou
University of California
Irvine, CA 92697
phchou@uci.edu

ABSTRACT

Harvesting energy from the environment is a desirable and increasingly important capability in several emerging applications of embedded systems such as sensor networks, biomedical implants, *etc*. While energy harvesting has the potential to enable near-perpetual system operation, designing an efficient energy harvesting system that actually realizes this potential requires an in-depth understanding of several complex tradeoffs. These tradeoffs arise due to the interaction of numerous factors such as the characteristics of the harvesting transducers, chemistry and capacity of the batteries used (if any), power supply requirements and power management features of the embedded system, application behavior, *etc*. This paper surveys the various issues and tradeoffs involved in designing and operating energy harvesting embedded systems. System design techniques are described that target high conversion and storage efficiency by extracting the most energy from the environment and making it maximally available for consumption. Harvesting aware power management techniques are also described, which reconcile the very different spatio-temporal characteristics of energy availability and energy usage within a system and across a network.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems

General Terms

Management, Measurement, Performance, Design

Keywords

Energy harvesting, power management, wireless sensors, solar power

1. INTRODUCTION

Energy harvesting embedded systems (EHES) have received growing attention in recent years. Miniaturization and wireless or logging capabilities open up brand new applications by enabling a complete system to be mounted on or implanted inside many more objects than ever before. Although some applications such as automobiles provide their own infrastructure for supplying power, many other targets such as trees being monitored at a remote location do not readily supply electrical power. Batteries with limited capacities will eventually drain long before the service life of the system. Although radioisotopes and other radioactive materials can supply steady power for decades and are currently used in exit signs and other applications, their radioactive nature poses additional user concerns that prevent them from wider adoption [1]. As a result, for the great majority of embedded applications, designers will have to choose from more conventional energy sources that will scale well with economy.

Energy harvesting itself is not new, but what is new is how to build efficient energy harvesting capabilities into modern embedded systems while satisfying all their constraints. For instance, windmills and hydroelectric generators have been in use for a long time, and solar panels have been powering satellites and space stations for decades. It is possible to miniaturize these power sources and use them to power embedded systems. However, straightforward implementations will usually result in low efficiency.

Efficiency can be divided into several parts: *conversion efficiency* from one form of energy to another (*e.g.*, from light to electricity), *transfer efficiency* from the source to the supply, *buffering efficiency* once it has been harvested, and *consumption efficiency* in terms of the amount of useful work given the harvestable energy. Although research work has been proposed to optimize the efficiency at each of these levels, it is crucial to consider these techniques together in the context of the entire system, or else the gain at one level may come at the price of efficiency loss or high overhead elsewhere, rendering the system much less efficient than before.

1.1 System classification

One way to classify EHESs is based on the energy source of their application in relation to the system. They can be divided into *environmentally embedded* and *wearable*. Environmentally embedded systems are those that are embedded into an environment, which can be a building, a habitat, a greenhouse, or many other environments. In some such environments, abundant energy is available for harvesting, including sunlight [2, 3, 4, 5], wind [6], etc [7], A special case of the environment is a wearable [8, 9] or implantable system, where the subject is a person or an animal. In such systems, the source of energy can be from the subject itself, possibly in addition to the environment in which the subject operates.

In other environments, if the device is buried underground or inside walls, then very little energy can be harvested. In these cases, energy harvesting is difficult, but *wireless energy transfer* can be performed. For instance, inductive charging can be used to receive energy from electromagnetic emission [10]. Although related, this

ISLPED'06, October 4–6, 2006, Tegernsee, Germany.

paper will consider wireless energy transfer to be a separate topic as it requires additional coordination with the energy source.

1.2 Mechanisms for Energy Harvesting

The form of energy that can be harvested include mechanical, thermal, photovoltaic, electromagnetic, biological, and chemical. Mechanical energy is possibly the most prevalent and are found in the forms of wind [6], limb movement [8], strain [11], ambient vibration [12, 13, 14, 15, 16], car wheel rotation, and many more. Heat differential can also be used to generate electricity [17]. The most well known source is light. Each given form of energy can be harvested by a different class of generator that performs conversion to electricity. To the system, the key differences are the output power level (current, voltage), AC vs. DC, the dynamic range, and the impedance model. For instance, windmills [6], magnetic coil generators [8], piezoelectric generators [11, 12, 13, 14, 15, 16], and magnetic induction [10] output AC power, whereas thermal [17] and photovoltaic [2, 4, 3, 5] power sources output DC power. Since most digital systems run on DC power, the default option is to rectify the current. An emerging alternative is to design self-timed circuits that will run directly on rectified AC power with minimal conversion loss [18]. Before such technologies become available in all components, it might not be possible to adopt them in more conventional designs. Even for DC power sources, it is often necessary to convert the DC power to a different voltage. In both cases, AC-DC or DC-DC conversion will incur additional loss.

The goal of this paper is to survey techniques that can be readily adopted in conventional systems that harvest energy. We divide the discussion into system design issues and power management issues. The former covers system functionality that must be designed-in in order to support efficient energy harvesting. The latter covers power management policies that specifically optimize for energy harvesting.

2. SYSTEM DESIGN ISSUES

The issues at the system level can be divided into voltage and current of the supply, form of energy storage, maximum power point tracking, and use of multiple power sources. It is important to consider the cost and operating range (V, I) associated with each level when trying to improve its efficiency, or else it can be counterproductive.

2.1 Voltage and Current

The very first consideration in energy harvesting is the voltage. Without high enough voltage, it is difficult or impossible to either power the system directly or to charge an energy storage device. Ideally, the circuit would adapt its power consumption and performance by tracking the available power without additional conversion. Asynchronous circuits can do this, but most commercially available components use synchronous designs. Moreover, systems such as wireless sensor nodes contain analog and RF components, which are sensitive to noise in the power supply. Therefore, it is necessary to perform voltage conversion to a known, controllable level, and then consume regulated power.

Voltage regulators are used to bridge the gap between the supply and the consumer. Linear regulators output clean, stable power, which are required by analog/RF components. However, they have lower conversion efficiency, incur a voltage drop, and dissipates more heat. On the other hand, switching regulators have much higher efficiency and are commonly used for digital subsystems, which are more immune to noise.

Switching regulators are further divided into buck, boost, and buck-boost regulators. Buck regulators perform voltage step-down conversion and are efficient, but the input voltage must be higher than the output or else it does not work properly. Boost regulators perform voltage step-up conversion but are less efficient. Buck-boost regulators, as the name implies, can work as a buck or a boost depending on if the input voltage is higher or lower than the output.

In the case of power supply design, there is almost always an energy storage device such as a battery or a supercapacitor. It acts as a consumer to the harvesting device when charging, and acts as a supply to the system it is powering. Thus, regulators may need to be inserted on both ends. The overall conversion efficiency is no longer just a ratio of output / input power; it must also take the *operating range* of all stages into account. Specifically, even though a buck regulator has higher *power efficiency*, it does not necessarily have a higher *energy efficiency* when regulating solar output to charge a battery. This is because by performing step-down conversion, it stops charging the battery when the solar output drops below the battery voltage [16]. A similar issue occurs in EHESs that use a voltage comparator and prioritize drawing current from the energy harvesting source when its voltage is higher than the battery or capacitor. When it is lower, the power below the useful threshold level is simply discarded [3, 2], causing the system to draw energy from the storage. In both cases, the power conversion efficiency may be higher while within the operating range, but the overall energy conversion efficiency may be lower. This is known as the *(internal) power fragmentation* problem.

A few solutions are possible. One simple way is to use a boost regulator to raise the voltage above the threshold, which then makes the power usable again. For instance, the MIT shoes harvest from a high-impedance piezoelectric generator, which outputs AC. To power 5V digital devices, the authors propose a circuit that includes a full-wave rectifier bridge, uses a boost (step-up) regulator to raise the voltage to charge the capacitor to up to 12.6V, above which it turns on the 5V regulator. However, boost regulators are usually much less efficient than buck (step-down) regulators or DC-DC converters. The large gap between the input and output voltage may be another source of inefficiency. An interesting problem is to choose an alternative capacitor configuration (*e.g.*, parallel vs. series composition, difference capacitance values) that will decrease this voltage gap for higher regulator efficiency.

2.2 Maximum Power Point Tracking

Maximum power point tracking (MPPT) refers to drawing power from an energy harvesting source at a level that maximizes the power output. For DC sources such as solar panels, the maximum power point (MPP) is a voltage-current combination that maximize the power output under a given sunlight condition, and a given temperature[1]. The power is maximized when the supply and the load are impedance matched [4]. For AC sources such as piezoelectric, on the other hand, the MPP is actually related to the resonant frequency of the device in addition to the magnitude of the physical oscillation. Although MPPT is not strictly required for energy harvesting to work, the efficiency loss can be tremendous that 65% to 90% of the available power may simply be discarded.

2.2.1 Measurement Method for MPPT

MPPT requires that the input intensity be known, so that the MPP in terms of voltage and current can be determined. The input intensity can be determined by measurement either *before* or *after* conversion to electricity.

For instance, the MPP for solar panels is determined mainly by the light intensity and secondarily by temperature. One may perform direct measurement of sunlight before conversion by includ-

---

[1] The MPP is only weakly dependent on temperature.

ing a light sensor and possibly a temperature sensor to provide the readings, so that the MPP can be looked up or computed. While straightforward, the light sensor covers a much smaller area than the solar panel and might not yield a representative reading, if the dust or shadow on the panel does not cover the light sensor in the same proportion. MPPT applies to other energy sources as well, including windmills and even fuel cells. For a windmill, a rotation speed sensor can be used.

An alternative is to sense the input intensity after conversion. This means measuring the voltage and current from the solar panel. One may measure either the open-circuit voltage or short-circuit current. Both require the load to be temporarily disconnected from the supply during its measurement. This is possible if there is another energy storage device such as either a battery or a capacitor to continue powering the system while the measurement is being taken. This method can better track the entire area of exposure, although discrete sampling assumes the input level does not alter rapidly.

2.2.2 MPPT for AC

For AC generators that rely on vibration, the MPP depends on not only the amplitude but also frequency of the vibration [8]. Ottman *et al.* [16] show that the power is maximized if the rectifier voltage is maintained at 1/2 open circuit voltage, and the duty cycle at approximately the square-root:

$$V_{rect} = \frac{I_P}{2\omega C_P} \tag{1}$$

$$D_{opt} \approx \sqrt{\frac{4\omega L C_p f_s}{\pi}} \tag{2}$$

Applying their tracking method to control a DC-DC converter has been shown to result in 400% improvement in efficiency. However, the controller is implemented with a DSP which may consume nontrivial power, making this difficult to apply to μW-level energy harvesting.

2.2.3 Software vs. Hardware MPPT Controller

The control for MPPT can be implemented either in hardware or software. A hardware implementation usually means a special designed circuit. It entails taking the output of a sensor, usually before conversion to electricity, and control a DC-DC converter or a programmable regulator, all without additional computation. On the other hand, a software implementation entails sampling the voltage level, usually after conversion to electricity, either performing a table look-up or running a DSP algorithm, and then controlling the power circuitry accordingly. The former is usually called *autonomous* in that MPPT can be performed with low overhead and be part of the power subsystem in a modular way, without the involvement from the microcontroller or DSP. The latter requires the MPPT to be scheduled as part of the software, and it is acceptable if it is scheduled as a low duty cycle task. Also, software-based schemes are more suitable for higher power energy harvesting systems, since the minimum power requirement for processors are usually significantly higher. Other downsides include more complex software, the use of precious I/O pins for control, and inability to operate if the DSP or MCU itself is powered down. Hardware implementations are almost mandatory for very low-power MPPT, or if the input power level can change rapidly. However, most hardware implementations are very simple in order to keep the overhead low, but they tend to track the MPP with a hysteresis band [6].

The MPPT overhead must be considered for the entire system, or else MMPT may be counterproductive. Software MPPT usually is counterproductive for microwatt-level energy harvesting. For example, MPPT enables one piezoelectric generator to improve its power output by 400%, from 16.43mW to 70.42mW after 18.87mW efficiency loss due to the DC-DC converter. However, the TMS320C31 DSP that runs the MPPT algorithm consumes 80mA of current, which is already higher than the whole energy harvesting source. The voltage at 67.3V_{oc} is also very high for harvesting 30.66mW [19]. Therefore, it can actually be a loss in this case, unless the overhead can be amortized over a large enough array of these harvesting devices.

2.3 Power Defragmentation

One problem with harvesting energy from environmental sources is the wide dynamic range of power. Even with MPPT, the available power may be so low that it is below the useful threshold. To solve this problem, one may wish to harvest energy from multiple sources. For instance, one might include both a solar panel and a wind generator in a heterogeneous power harvesting system. This will harvest more energy than relying just on one single source. However, it is difficult to “add” power by simply composing heterogeneous sources such as a battery and a windmill in series or in parallel. Depending on the target system to be powered, each source might not be sufficiently powerful to power the entire system. In this case, the power must be discarded, even if it is available at a non-trivial level. This is known as the *(external) power fragmentation* problem.

To address the power fragmentation problem, power matching switches have been proposed [20]. The idea is to divide up the system into subsystems that can be powered separately. The supply-to-power are related by a one-to-many mapping. This way, instead of discarding the power when it is below the threshold of the system, now it is possible to continue utilizing power until the subsystem threshold, which is much lower.

2.4 Energy Storage Devices

Many energy harvesting embedded systems need energy storage because they need to continue operation even when there is no energy to harvest (*e.g.*, at night for a solar-powered system). By default, rechargeable batteries are used for longer term storage (several days to weeks). Batteries have the problem of non-ideal effects including aging and rate capacity effects.

More recently, alternative energy storage technologies have become available, including supercapacitors and fuel cells. Supercapacitors, also called ultracapacitors, are high-value (*i.e.*, from mF to hundreds of farads), portable sized (10s of cubic-cm) capacitors. They are commonly used for buffering transient energy (from minutes to hours). For instance, electric and gasoline-electric hybrid vehicles use supercapacitors to store energy in regenerative brakes, and they are starting to be used in EHESs as well [4, 5, 21, 22, 10, 6, 3]. Supercapacitors do not have the aging and rate-capacity problems. They do have limited energy capacity and higher leakage, but these are less of a problem if energy is regularly replenished. Another often cited problem is the linear discharge curve and inaccessible charge below the target operating range, but this can be addressed easily with the buck-boost regulator described in Section 2.1.

Even though capacitor charging and MPPT are active areas of research, the combination of MPPT using supercapacitors as an energy storage device poses new challenges. A capacitor of hundreds of farads that is attached to the power rails appears as a short circuit on a cold start and throughout much of the charge cycle. The capacitor will still charge but at a very inefficient level. One attempt to solve this problem is to insert a voltage regulator between the supply and the capacitor. However, this will not work because the regulator as a feedback control circuitry see the low output voltage and attempts to drive up the voltage, which also results in effectively a short circuit. Other converters either require AC [23] or incur high overhead [24]. Instead of a feedback mechanism, a feed-forward PFM (pulse frequency modulation) regulator [4] has been proposed to address these problems.

3. POWER MANAGEMENT ISSUES

Complementary to the problem of designing a harvesting device that efficiently extracts, stores, and transfers power to the load is the issue of adapting the embedded system's power management policy to be aware of the operating conditions and state of the energy harvesting device, a process termed as *harvesting aware power management*. This section illustrates how harvesting aware power management improves upon conventional battery-based power management and surveys recent work in designing such harvesting aware power management methodologies.

As mentioned in Section 2, there are several non-idealities (*e.g.*, power fragmentation, inefficiency of energy storage elements, *etc.*) that manifest in energy harvesting systems. We use the following example to illustrate the potential benefits of harvesting aware power management.

EXAMPLE 1. *Consider the task of routing data in a simple sensor network where two route options exist from the data source to the sink, one of which uses node A and the other node B. Nodes A and B receive the same amount of solar energy per day, E_s, but due to obstacles such as trees, node A receives all of its energy in the morning, whereas node B receives all of its energy in the afternoon. Both nodes begin with the same residual battery energy, E_b, and the battery round trip efficiency is η. A node uses energy E_r for one hour of routing activity, and the daily workload consists of an hour of routing activity in the morning and another hour in the afternoon. We compare two routing schemes, $\mathcal{H}$, which explicitly uses information about the solar energy availability pattern, and $\mathcal{B}$, which operates based on residual battery levels alone and is representative of state-of-the-art power aware routing schemes. On the first morning, $\mathcal{H}$ chooses node A to route data (since it knows that node A receives solar energy in the morning) while $\mathcal{B}$ may pick either node, as each has the same battery level. Say that it chooses node B. At noon, in the system running $\mathcal{H}$, node A has energy[2] $E_b + (E_s - E_r)\eta$, and B has energy E_b. In the system running $\mathcal{B}$, node A has $E_b + E_s\eta$ and B has $E_b - E_r$. In the afternoon, algorithm $\mathcal{H}$ will choose node B (since it is aware that node B receives solar energy in the afternoon), and the residual battery energy at the end of the day is $E_b + (E_s - E_r)\eta$ at A and $E_b + (E_s - E_r)\eta$ at B. Algorithm $\mathcal{B}$ will instead choose node A due to its higher battery level, resulting in battery levels of $E_b + E_s\eta - E_r$ at A and $E_b + E_s\eta - E_r$ at B. At each node, the nodes following algorithm $\mathcal{H}$ have a higher energy, ΔE, given by:*

$$\begin{aligned}\Delta E &= E_b + (E_s - E_r)\eta - (E_b + E_s\eta - E_r)\\ &= E_r(1-\eta)\end{aligned} \tag{3}$$

Note that, at the end of the day, both nodes in each system have equal energy. Hence the process may repeat on the next day, increasing the energy gap between $\mathcal{H}$ and $\mathcal{B}$.

The above example shows that modifying the power management policy to be harvesting aware can provide improved energy efficiency compared to a conventional system-level power management scheme that operates without any knowledge of the spatio-temporal characteristics of the environmental energy source.

3.1 Energy Neutrality in Harvesting Systems

A significant difference between battery-powered systems and energy harvesting systems from the perspective of power management is that conventional energy optimization metrics might not be suitable in an energy harvesting scenario. For instance, a commonly used objective in battery-powered sensor networks is to maximize network lifetime under a total energy constraint. Clearly, this changes if energy harvesting is allowed since the amount of energy available itself depends on the time duration for which the system operates. Instead, a more relevant design objective might be to operate in an *energy neutral* mode, consuming only as much energy as harvested. Such a mode of operation raises the possibility of indefinitely long lifetime, limited only by the hardware longevity. Note that reducing the power consumption below the level needed for energy neutrality will not increase system lifetime any further.

Achieving energy neutrality in a harvesting system depends on several factors such as the average power generated by the harvesting device, the capacity of the energy storage device, *etc.*, which are influenced by design choices made by the system architect. In order to help designers of harvesting systems make the right design decisions, a systematic framework is needed that can capture and analyze the various energy-neutrality related requirements quantitatively. In [25], the authors develop such a framework to better understand the energy neutral mode of operation.

3.1.1 Modeling environmental energy sources

The first step in analyzing a harvesting system's energy neutrality is to analytically model the power generated by the harvesting transducer. However, this is a non-trivial task due to the inherent temporal variability present in environmental energy sources. In [25], the authors introduce the following model for characterizing time-varying environmental energy sources.

DEFINITION 3.1. *A non-negative, continuous, and bounded function $P(t)$ is said to be a $(\rho, \sigma_1, \sigma_2)$ function if and only if, for any positive, finite real numbers τ and T, the following is satisfied:*

$$\rho T - \sigma_2 \quad \leq \quad \int_{\tau}^{\tau+T} P(t)dt \quad \leq \quad \rho T + \sigma_1 \tag{4}$$

For instance, let $P_s(t)$ denote the power output by the transducer at time t. Intuitively, if $P_s(t)$ is a $(\rho_1, \sigma_1, \sigma_2)$ function, then the average rate at which energy is available from the transducer is ρ_1, and the burstiness caused by temporal variations is bounded by σ_1 and σ_2. Note that this model can also be used to model power consumers. The power consumption profile of the load, $P_c(t)$, may be modeled as a (ρ_2, σ_3) function, where the parameters ρ_2 and σ_3 are used in a constraint similar to the upper bound inequality of Equation (4) to place a upper limit on the power consumption, while no constraint is placed on the minimum power consumption.

3.1.2 Analyzing Energy Neutrality Requirements

The harvesting theory as presented in [25] considers two non-idealities associated with energy storage, namely the round trip efficiency and self-discharge. Round trip efficiency is related to the energy loss that occurs when energy is stored into an energy storage element such as a battery and retrieved later. Self-discharge is the energy loss due to leakage paths in the energy storage element. Based on the energy production, consumption, and energy storage models discussed above, the condition for energy neutrality, the

---

[2] Since the energy for routing is supplied from the solar panel and only the remainder is stored in the battery. It is assumed that $E_s \geq E_r$, although a similar reasoning can be followed when $E_s < E_r$.

following theorem provides a sufficient condition for guaranteeing energy neutrality of the system [25]

THEOREM 3.2. *Consider a harvesting system in which the energy production profile is characterized by a* $(\rho_1, \sigma_1, \sigma_2)$ *function, the load is characterized by a* (ρ_2, σ_3) *function, and the energy buffer is characterized by parameters* η *for storage efficiency, and* ρ_{leak} *for leakage power. The following conditions are sufficient for the system to achieve energy neutrality:*

$$\rho_2 \leq \eta\rho_1 - \rho_{leak} \tag{5}$$

$$B_0 \geq \eta\sigma_2 + \sigma_3 \tag{6}$$

$$B \geq B_0 \tag{7}$$

where B *denotes the capacity of the energy buffer and* B_0 *is the initial energy stored in the buffer.*

This theorem has two important design implications. First, it characterizes the sustainable performance level that may be supported in energy neutral mode. This is significant for system design both in hardware and software. At the hardware level, if the sustainable power consumption supported is too low, changes may be made to increase the harvested energy (*e.g.*, by using a larger solar panel). In software, this level will help determine the appropriate power scaling required based on the relationship between energy consumption and system performance. Second, it specifies the minimum capacity of the energy storage element required to achieve energy neutrality for given burstiness bounds on energy production and consumption. This size can be directly used while designing the system. Using a higher capacity battery will not yield an increase in sustainable performance or lifetime.

3.2 Node Level Power Management

While the harvesting theory presented above enables us to design the embedded system in order to guarantee energy neutrality, equally crucial are techniques that adapt the performance and power consumption of the embedded system at runtime in response to the spatial and temporal variations in harvested energy. The goal of these techniques is to maximize system performance while not violating the energy neutrality requirement.

In [26], the authors present an algorithm for harvesting-aware duty cycling of wireless sensor nodes. The authors choose to use duty-cycling between active and low power modes for the purpose of performance/power scaling since most sensor nodes provide at least one low power mode in which the power consumption is negligible. More sophisticated performance/power scaling methods, such as dynamic voltage scaling, may be used when available. Their algorithm consists of three steps, namely (a) learning the harvested energy profile at run-time, (b) adapting the power consumption level to match the harvested energy, and (c) fine tuning the power scaling algorithm to account for battery non-idealities. The last step is important because it helps minimize the energy loss due to battery inefficiency.

A significant challenge with harvesting-aware power management is that determining the optimal duty cycle for a node at a given point in time requires information about the harvested energy availability in the future. The authors overcome this by learning the daily energy generation profile for the harvesting device and use this information to predict the energy availability for the near future. The authors argue that, on a typical day, the energy generation is expected to be similar to the energy generation at the same time on previous days. Based on this, they use an Exponentially Weighted Moving-Average (EWMA) filter based prediction model. The method is designed to exploit the diurnal cycle in solar energy but at the same time adapt to weather and seasonal variations. The predicted energy generated for each 30 minute slot is calculated as a weighted average of the energy received in the same time slot during previous days. The weights are exponential, resulting in decaying weights for older data.

Using a user defined system utility function that is a function of the duty cycle, the authors then solve a lightweight optimization problem that determines duty cycles for each time slot in order to maximize the utility of the system over the course of the entire day. As a post processing step, they account for any errors in the predicted energy profile by adjusting the duty cycle as the day progresses using an approach similar to slack redistribution techniques during power-aware task scheduling [27]. The authors demonstrate that their algorithm utilizes environmental energy more efficiently (up to 58%) compared to duty cycling techniques that are not harvesting aware.

3.3 Network Level Power Management

The technique presented above adapts the performance at a single node in response to the temporal variations in harvested energy. Next, we consider a distributed network of harvesting enabled nodes and discuss how the network as a whole can be power managed to address the spatial variations in harvesting opportunity across the different nodes. The authors of [7] addressed this problem in the context of a data routing application presented a harvesting aware routing scheme which attempts to align the allocation of routing energy consumption with the harvested energy profiles across the network.

Energy aware routing protocols in sensor networks typically use battery energy based routing cost metrics [28]. The objective is to choose data routes appropriately such that the routing load is distributed uniformly across the network to leverage the total battery resource for maximizing lifetime and preventing individual nodes from being over-used and hence, running out of battery quickly. In an energy harvesting network, battery awareness is not sufficient to select the best routes, as we saw in Example 1. In [7], the authors propose a harvesting enhanced cost metric for data routing. The authors assume that each node can learn the expected rate of energy harvesting, ρ_i, at node i. They propose an enhanced routing cost metric that considers both the harvesting potential of a node as well as its residual battery level (B_i). They define an energy potential, E_i, at node i as follows:

$$E_i = w \cdot \rho_i \; + \; (1 - w) \cdot B_i \tag{8}$$

where w is a weight parameter, $0 \leq w \leq 1$ and is typically set close to 1. A low value of w may be relevant in networks powered predominantly by batteries with a small harvesting opportunity. The authors use the inverse of the energy potential at a node as the communication cost for all wireless links into that node. In other words, in the directed graph representing each sensor node as a vertex, v_i and each wireless link between two nodes i and j as an edge e_{ij}, they associate the following cost for each edge into node i:

$$c_{ki}(e_{ki}) = \frac{1}{E_i} \quad \forall k \in \{k | e_{ki} \in E_{comm}\} \tag{9}$$

where E_{comm} represents the set of edges across which radio communication is feasible for the deployed network topology and the radio hardware used. Thus, the authors derive a graph representation of the wireless network with each link represented as a weighted directed edge. They use a distributed Bellman-Ford algorithm to compute minimum cost routes between a given source and destination. They show that such a harvesting-aware data routing scheme significantly increases the energy scalability of the network.

4. CONCLUSIONS

Energy harvesting in embedded systems is representing a fruitful area of research as made possible by the convergence of low-power designs, miniaturization, and advances in materials and mechanical devices. The power consumption has been reduced to the same level as the harvesting devices are capable of outputting. The potentially perpetual operation of these EHESs are just starting to enable a brand new class of applications. However, it seems unlikely that existing systems can automatically operate efficiently by just adding an energy harvesting module. We believe the entire system must be optimized in a holistic way from the design of the architecture to power management at the application and networking levels.

5. REFERENCES

[1] Amit Lal, Rajesh Duggirala, and Hui Li. Pervasive power: A radioisotope-powered piezoelectric generator. *IEEE Pervasive Computing*, 4(1):53–61, January–March 2005.

[2] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava. Design considerations for solar energy harvesting wireless embedded systems. In *Proc. 4th Int. Conf. on Information Processing in Sensor Networks*, pages 457–462, Apr. 2005.

[3] X. Jiang, J. Polastre, and D. Culler. Perpetual environmentally powered sensor networks. In *Proc. 4th Int. Conf. on Information Processing in Sensor Networks*, pages 463–468, Apr. 2005.

[4] Farhan Simjee and Pai H. Chou. Everlast: Long-life, supercapacitor-operated wireless sensor node. In *Proc. ISLPED*, October 2006.

[5] Henri Dubois-Ferrière, Laurent Fabre, Roger Meier, and Pierre Metrailler. TinyNode: a comprehensive platform for wireless sensor network applications. In *IPSN '06: Proceedings of the fifth international conference on Information processing in sensor networks*, pages 358–365, New York, NY, USA, 2006. ACM Press.

[6] Chulsung Park and Pai H. Chou. AmbiMax: Efficient, autonomous energy harvesting system for multiple-supply wireless sensor nodes. In *to appear, Proc. Third Annual IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks (SECON)*, 2006.

[7] Aman Kansal and Mani B. Srivastava. An environmental energy harvesting framework for sensor networks. *Proceedings of the 2003 international symposium on Low power electronics and design*, pages 481–486, 2003.

[8] John Kymissis, Clyde Kendall, Joseph A. Paradiso, and Neil Gershenfeld. Parasitic power harvesting in shoes. In *Proc. of the Second IEEE International Symposium on Wearable Computing ISWC*, pages 132–139, 1998.

[9] Joseph A. Paradiso and Thad Starner. Energy scavenging for mobile and wireless electronics. *IEEE Pervasive Computing*, 4(1):18–27, January–March 2005.

[10] Olivier Chevalerias, Terence O'Donnell, Daithi Power, Neil O'Donovan, Gerald Duffy, Gary Grant, and Sean Cian O'Mathuna. Inductive telemetry of multiple sensor modules. *IEEE Pervasive Computing*, 4(1):46–52, January–March 2005.

[11] David L. Churchill, Michael J. Hamel, Christopher P. Townsend, and Steven W. Arms. Strain energy harvesting for wireless sensor networks. *Proceedings of SPIE*, 5055:319, 2003.

[12] S. Kim, W.W. Clark, and Q.M. Wang. Piezoelectric energy harvesting using diaphragm structure. *Proceedings of SPIE*, 5055:307, 2003.

[13] Henry A. Sodano, Daniel J. Inman, and Gyuhae Park. A review of power harvesting from vibration using piezoelectric materials. *The Shock and Vibration Digest*, 36(3):197–205, 2004.

[14] Sunghwan Kim. *Low power energy harvesting with piezoelectric generators*. PhD thesis, University of Pittsburgh, 2002.

[15] Shad Roundy, Eli S. Leland, Jessy Baaker, Eric Carleton, Elizabeth Reilly, Elaine Lai, Brian Otis, Jan M. Rabaey, Paul K. Wright, and V. Sundararajan. Improving power output for vibration-based energy scavengers. *IEEE Pervasive Computing*, 4(1):28–36, January–March 2005.

[16] Geffrey K. Ottman, Heath F. Hofmann, Archin C. Bhatt, and George A. Lesieutre. Adaptive piezoelectric energy harvesting circuit for wireless remote power supply. *IEEE Transactions on Power Electronics*, 17(5):669–676, 2002.

[17] Ingo Stark. Invited talk: Thermal energy harvesting with Thermo Life. In *International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06)*, pages 19–22, Los Alamitos, CA, USA, 2006. IEEE Computer Society.

[18] J. Siebert, J. Collier, and R. Amirtharajah. Self-timed circuits for harvesting AC power supplies. In *Proc. International Symposium on Low Power Electronics and Design*, 2005.

[19] Geffrey K. Ottman, Heath F. Hofmann, and George A. Lesieutre. Optimized piezoelectric energy harvesting circuit using step-down converter in discontinuous conduction mode. *IEEE Transactions on Power Electronics*, 18(2):696–703, March 2003.

[20] Chulsung Park and Pai H. Chou. Power Utility Maximization for Multi-Supply Systems by a Load-Matching Switch. In *ISLPED'04*, pages 168–173, August 2004.

[21] P. Enjeti, J.W. Howze, and L. Palma. An approach to improve battery run-time in mobile applications with supercapacitors. In *IEEE 34th Annual Power Electronics Specialists Conference*, 2003.

[22] T.A. Smith, J.P. Mars, and G.A. Turner. Using supercapacitors to improve battery performance. In *Power Electronics Specialists Conference*, 2002.

[23] H. Pollock. High efficiency, high frequency power supplies for capacitor and battery charging. In *IEE Colloquium on Power Electronics for Demanding Applications*, pages 901 – 910, Apr. 1999.

[24] R.M. Nelms and J.E. Schatz. A capacitor charging power supply utilizing a ward converter. *IEEE Trans. Ind. Electronics*, 39:421–428, Oct. 1992.

[25] Aman Kansal, Dunny Potter, and Mani B. Srivastava. Performance aware tasking for environmentally powered sensor networks. In *SIGMETRICS*, pages 223–234, 2004.

[26] J. Hsu, S. Zahedi, A. Kansal, V. Raghunathan, and M. Srivastava. Adaptive duty cycling for energy harvesting systems. In *Proc. ACMInternational Symposium on Low Power Electronics and Design*, 2006.

[27] Niraj K. Jha. Low power system scheduling and synthesis. In *ICCAD*, pages 259–263, 2001.

[28] Rahul C. Shah and Jan M. Rabaey. Energy aware routing for low energy ad hoc sensor networks. In *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, pages 350–355, 2002.