Design Considerations for Solar Energy Harvesting Wireless Embedded Systems

Vijay Raghunathan, Aman Kansal, Jason Hsu, Jonathan Friedman, and Mani Srivastava
Networked and Embedded Systems Lab (NESL)
Department of Electrical Engineering
University of California, Los Angeles, CA 90095
Email: {vijay, kansal, jasonh, jf, mbs}@ee.ucla.edu

***Abstract*— Sustainable operation of battery powered wireless embedded systems (such as sensor nodes) is a key challenge, and considerable research effort has been devoted to energy optimization of such systems. Environmental energy harvesting, in particular solar based, has emerged as a viable technique to supplement battery supplies. However, designing an efficient solar harvesting system to realize the potential benefits of energy harvesting requires an in-depth understanding of several factors. For example, solar energy supply is highly time varying and may not always be sufficient to power the embedded system. Harvesting components, such as solar panels, and energy storage elements, such as batteries or ultracapacitors, have different voltage-current characteristics, which must be matched to each other as well as the energy requirements of the system to maximize harvesting efficiency. Further, battery non-idealities, such as self-discharge and round trip efficiency, directly affect energy usage and storage decisions. The ability of the system to modulate its power consumption by selectively deactivating its sub-components also impacts the overall power management architecture. This paper describes key issues and tradeoffs which arise in the design of a solar energy harvesting, wireless embedded system and presents the design, implementation, and performance evaluation of Heliomote, our prototype that addresses several of these issues. Experimental results demonstrate that Heliomote, which behaves as a plug-in to the Berkeley/Crossbow motes and autonomously manages energy harvesting and storage, enables near-perpetual, harvesting aware operation of the sensor node.**

I. Introduction

The application space for wireless sensor networks is dominated by the longevity constraint, since the cost of physically deploying the sensor nodes often outweighs the cost of the nodes themselves. Energy is the limiting factor in achieving extreme (months to years) systemwide lifetime. Fortunately, a promising technique to forestall this network energy crisis is emerging – environmental energy harvesting. Exploiting energy sources ubiquitous to the operating space of the sensor nodes raises the possibility of infinite lifetime. Achieving this (through harvesting aware design) represents a new frontier in the natural progression of energy optimization techniques, which started from low power design [1], evolved into power aware design [2], and recently, battery aware design [3].

Table I shows the power generation potential of several energy harvesting modalities [4]. While a wide variety of harvesting modalities are now feasible, solar energy harvesting through photo-voltaic conversion provides the highest power density, which makes it the modality of choice to power an embedded system that consumes several mW using a reasonably small harvesting module. However, the design of a solar energy harvesting module involves complex tradeoffs due to the interaction of several factors such as the characteristics of the solar cells, chemistry and capacity of the batteries used (if any), power supply requirements and power management features of the embedded system, application behavior, *etc*. It is, therefore, essential to thoroughly understand and judiciously exploit these factors in order to maximize the energy efficiency of a solar harvesting module.

TABLE I
POWER DENSITIES OF HARVESTING TECHNOLOGIES

Harvesting technology	Power density
Solar cells (outdoors at noon)	$15mW/cm^3$
Piezoelectric (shoe inserts)	$330\mu W/cm^3$
Vibration (small microwave oven)	$116\mu W/cm^3$
Thermoelectric (10^oC gradient)	$40\mu W/cm^3$
Acoustic noise (100dB)	$960nW/cm^3$

A. Paper contributions

This paper makes the following contributions. First, we identify the various considerations and tradeoffs that are involved in the design of a solar energy harvesting module, and describe their impact on harvesting efficiency. We illustrate how these considerations differ from conventional battery based systems. Second, we discuss the desired features of such a solar harvesting module, and the services it should provide to the rest of the system to enable harvesting aware power management. We also illustrate how such harvesting aware operation can further improve system lifetime compared to state-of-the-art battery aware power management. Finally, we present the design, implementation, and performance evaluation of Heliomote, our plug-and-play solar energy harvesting module for the Berkeley/Crossbow motes. Heliomote autonomously performs energy harvesting, storage, and power routing, and enables harvesting aware operation by providing instantaneous solar and battery-state information through a simple one wire interface. Our experimental results demonstrate the feasibility of self-sustained near-perpetual operation of outdoor sensor networks using solar energy harvesting.

II. Related Work

Energy efficient design techniques have been studied for sensor networks [5] at all levels from hardware design [6] to protocols for medium access control [7], routing [8], data gathering [9], topology management [10], [11], *etc*. Tools and techniques for energy and battery life estimation have also been proposed [12].

Environmental energy harvesting has been considered for improving the sustainable lifetimes of wearable computers [13], [14], sensor networks [15]–[17], *etc*. Numerous harvesting modalities have been successfully demonstrated including solar, vibrational, biochemical, and motion based [18]–[21], and several others are currently being developed [22]. While harvesting technology provides the ability to extract energy from the environment, it must be efficiently integrated into an embedded system to translate that harvested energy into increased application performance and system lifetime. A solar harvesting augmented high-end embedded system was described in [23], in which a switch matrix was used to power individual system

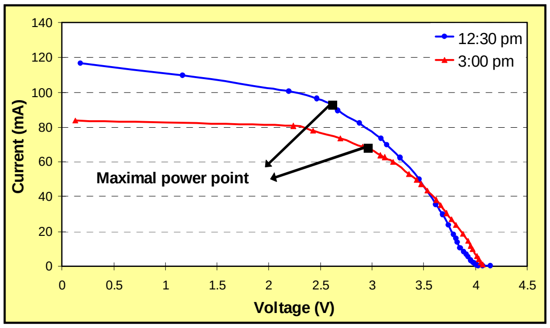

Fig. 1. Measured V-I characteristics of the Solar World 4-4.0-100 solar panel.

components from either the solar panel or the battery. Harvesting aware protocols have also been proposed for data routing [24], and distributed performance scaling [25], [26].

III. Solar Harvesting Module Design

This section presents a description of the various components of a generic solar energy harvesting module, design considerations that impact its efficiency, and the associated tradeoffs. Since the power consumption characteristics of the target embedded system heavily influence the various design decisions, for the remainder of this paper, we tailor our discussion towards "mote class" embedded systems, which consume on the order of few tens of mW.

A. Solar cell characteristics

Solar cells have vastly differing characteristics from batteries. The V-I characteristics of the 4-4.0-100 solar panel (formed by a series/parallel combination of solar cells) from Solar World Inc. are shown in Figure 1. The characterization was performed on Nov. 28, 2004, with a panel that measured 3.75" x 2.5". Solar panels are characterized by two parameters, the open circuit voltage (V_{oc}) and the short circuit current (I_{sc}). These form the x- and y- intercepts of the V-I curve, respectively. Several observations can be made from the figure. First, it is clear that a solar panel behaves as a voltage limited current source (as opposed to a battery which is a voltage source). Second, there exists an optimal operating point at which the power extracted from the panel is maximized. Finally, as the amount of incident solar radiation decreases (increases), the value of I_{sc} also decreases (increases). However, V_{oc} remains almost constant.

Due to its current source-like behavior, it is difficult to power the target system directly from the solar panel, since the supply voltage would depend on the time varying load impedance. Hence, an energy storage element, such as a battery, is used to store the energy harvested by the panel and provide a stable voltage to the system.

B. Energy storage technologies

Perhaps the most complex (and crucial) design decision involves the energy storage mechanism. The two choices available for energy storage are batteries and electrochemical double layer capacitors, also known as ultracapacitors. Batteries are a relatively mature technology and have a higher energy density (more capacity for a given volume/weight) than ultracapacitors, but ultracapacitors have a higher power density than batteries and have traditionally been used to handle short duration power surges. Recently, such capacitors have been explored for energy storage, since they are more efficient than batteries and offer higher lifetime in terms of charge-discharge cycles. However, they involve leakage (intrinsic and due to parasitic paths in the external circuitry), which precludes their use for long term energy storage. While it is also possible to implement a tiered energy storage mechanism using an ultracapacitor and a battery, the tradeoff is a decrease in harvesting circuit efficiency due to the increased overhead of energy storage management.

Four types of rechargeable batteries are commonly used: Nickel Cadmium (NiCd), Nickel Metal Hydride (NiMH), Lithium based (Li+[1]), and Sealed Lead Acid (SLA). Of these, SLA and NiCD batteries are less used because the former has a relatively low energy density, and the latter suffers from temporary capacity loss caused by shallow discharge cycles, termed as the *memory effect*. The choice between NiMH and Li+ batteries involves several tradeoffs. Li+ batteries are more efficient than NiMH, have a longer cycle lifetime, and involve a lower rate of self-discharge. However, they are more expensive, even after accounting for their increased cycle life. Li+ batteries also require a significantly more complicated charging circuit. Further, charging Li+ batteries at very low rates is often not possible due to charge acceptance issues, and they are known to degrade if subjected to deep discharge cycles. An additional consideration is battery aging due to charge-discharge cycles. For example, NiMH batteries (when subjected to repeated 100% discharge) yield a lifetime of about 500 cycles, at which point the battery will deliver around 80% of its rated capacity. This does not mean the battery cannot be used further, rather that it will have only 80% of the capacity of a new battery. The residual capacity is significantly higher if the battery is only subjected to shallow discharge cycles. At the rate of one discharge cycle per day, the battery will last for several years before its capacity becomes zero.

Several other battery related factors, which are usually insignificant for conventional mobile devices, also play a role due to the nature of the target system/application. First, the battery non-ideality termed as rate capacity effect is non-existent since the system's current draw (few tens of mA) is an order of magnitude less than the rated current of most current day batteries. Second, since solar energy harvesting requires outdoor system deployment, the operating temperature of the batteries will vary, leading to changes in battery characteristics. For example, battery self-discharge rate approximately doubles with every 10^{o}C increase in ambient temperature. Finally, since the amount of current generated by the solar panel is limited, the round trip efficiency of the battery becomes crucial. This is in contrast to devices that are recharged using wall chargers, which provide an unlimited supply of charging current. As we will show in Section IV, scheduling strategies that eliminate the round trip to the battery by utilizing solar energy directly significantly improve overall harvesting efficiency.

Thus, the choice of battery chemistry for a harvesting system depends upon its power usage, recharging current, and the specific point on the cost-efficiency tradeoff curve that a designer chooses.

C. Harvesting circuit design

The core of the harvesting module is the harvesting circuit, which draws power from the solar panels, manages energy storage, and routes power to the target system. The most important consideration in the design of this circuit is to maximize efficiency. There are several aspects to this.

As mentioned in Section III.A, solar panels have an optimal operating point that yields maximal power output. The harvesting circuit should ensure operation at (or near) this maximal power point, which is done by clamping the output terminals of the solar panel to a fixed voltage. Since the maximal power point changes slightly with

[1]We group Lithium ion and Lithium polymer batteries together since their characteristics are similar.

the time of day (*i.e.*, as the incident radiation changes), a maximal power point tracker (MPPT) circuit can be used to continuously track and operate at the optimal point. However, commercially available MPPT ICs are designed for high power applications such as solar based water heaters, and hence consume a lot of power, precluding their use in a low power, solar harvesting, embedded system.

A DC-DC converter is often used to provide a constant supply voltage to the embedded system. The choice of DC-DC converter depends on the operating voltage range of the particular battery used, as well as the supply voltage required by the target system. If the required supply voltage falls within the voltage range of the battery, a boost-buck converter is required, since the battery voltage will have to be increased or decreased depending on the state of the battery. However, if the supply voltage falls outside the battery's voltage range, either a boost converter or a buck converter is sufficient, which significantly improves power supply efficiency.

Although specialized battery charging ICs are available, they are designed to regulate charging at significantly higher currents (*e.g.*, wall chargers) than the few tens of mA provided by a small solar panel and are inefficient (if still operable) at such low currents.

Finally, since the primary goal is to efficiently harvest and utilize every precious mW of power provided by solar panel, it is desirable to make the harvesting circuit as application- and system- specific as possible. For example, building a harvesting circuit that works with several solar panels and charges NiCd, NiMH, and Li+ batteries is a poor design decision because of the efficiency loss that inevitably accompanies such a general purpose solution.

D. Energy measurement

To enable harvesting aware power management decisions, the harvesting module should have energy measurement capabilities. Low power battery monitor ICs can be used to provide this feature. The target system should be able to query the harvesting module for the instantaneous power being provided by the solar panels, and the battery terminal voltage. In addition, a harvesting module can also learn the solar power availability pattern, and build and train a power macro-model that provides information about future power arrival. As shown in Section IV, this information can be used by the system to intelligently schedule the execution of its application workload.

IV. Harvesting Aware Power Management

Another aspect of harvesting system design is to exploit the harvesting capability through the use of appropriate power management strategies. This is especially important in a distributed harvesting system, such as a sensor network, where each node may have different environmental harvesting opportunity and hence, instead of just minimizing the total energy consumption, it becomes necessary to adapt the power management scheme to account for these spatio-temporal variations. This section illustrates the benefits of such harvesting aware power management decisions.

Observe that environmental energy availability at a node is not characterized by the residual battery level. In fact, energy may not even be an appropriate metric to use since energy availability is virtually unlimited along the temporal axis; power seems to be a more appropriate metric instead. As a result, state-of-the-art residual battery based power management policies are insufficient, as illustrated by the following example.

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:

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

As is evident from the above description, 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 also exposed another crucial aspect of harvesting aware power management. To make better routing choices, available solar power at more than just one node was needed. Further, information about future energy availability was required (*e.g.*, at noon, $\mathcal{H}$ used information about energy arrival in the afternoon to choose node B for routing for the remainder of the day). Such information about the future, or more precisely, a prediction of the same, can be obtained by developing parametrized macro-models for the power source. Power models could range from simple ones that model the average solar power availability over a long duration to more advanced ones that predict the complete solar power profile. These models can be combined with battery state and system power usage information to drive a distributed, harvesting aware, workload scheduling framework, as illustrated in Figure 2. We are currently developing such a framework to demonstrate the impact of harvesting aware decisions on network level performance and longevity.

V. Heliomote Design and Evaluation

This section presents the design of our solar energy harvesting module, Heliomote (derived from the Greek root *Helios*, which refers to the sun), that augments the Berkeley/Crossbow motes with energy harvesting capability. The experience gained over the course of building Heliomote forms the basis for the design considerations and recommendations described in the preceding sections.

A. Overview

Our driver application is an ecosystem sensing one in support of biological science objectives. System prototypes are already being

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

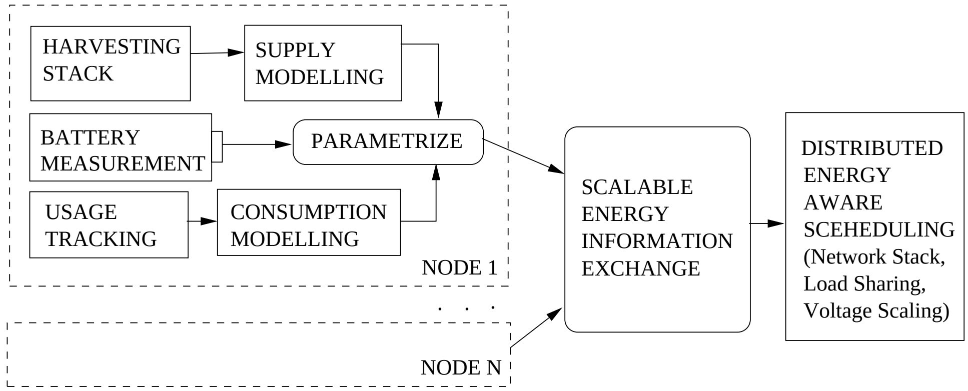

Fig. 2. Coordinated energy harvesting framework for a distributed system.

Fig. 3. Heliomote, a solar harvesting sensor node.

developed for this application [27] and the known deployment scenario enables appropriate understanding of the energy environment. The system is deployed outdoors and has access to solar radiation for most of the year.

Heliomote, shown in Figure 3, consists of an off-the-shelf sensor node (Mica2 motes), and a custom circuit board for solar harvesting. To ensure ease of use, Heliomote was designed to be a plug-and-play enhancement to the motes. Users can simply turn off the mote, plug in the Heliomote harvesting board (shown in Figure 4) with its solar panels and batteries, and switch it on. Sensor boards can be stacked on to the Heliomote, and the sensor connector is pin compatible with the MTS300/310 sensor boards from Crossbow. Advanced users can also desolder and remove the mote's original battery pack to reduce weight/volume. The Heliomote module can easily be adapted for use with other sensor nodes in the same power class by replacing the mote-compatible 51 pin connector. Our hardware design, including the schematics and layout files, is available for download at [28].

B. Design choices and description

To increase efficiency, our design choices were guided by the goal of minimizing energy wastage in the harvesting module. We used the 4-4.0-100 solar panels from Solar World Inc., which have a rated V_{oc} of 4.0V and I_{sc} of 100mA. The maximal power point of this panel lies at 3.0V and varies slightly depending on the time of day, as evident from its V-I characteristics shown in Figure 1.

The solar panel is connected to a battery whose terminal voltage determines the panel's operating point along its V-I curve. We ensure operation at the maximal power point while avoiding the use (and overhead) of an MPPT circuit through our choice of battery. Using two NiMH batteries to operate the Heliomote, the battery voltage varies between 2.2V and 2.8V, which, together with a diode used to prevent reverse current flow into the solar panel, ensures that the voltage across the solar panel terminals remains close to optimal. In addition, by avoiding the use of a Li+ battery, our charging circuit is considerably simplified, leading to increased efficiency. The AA size also retains compatibility with the mote form factor and enables reuse of the mote's original battery case, thereby reducing cost.

To avoid problems such as decreased radio range caused by decreased battery voltage, we use a step up DC-DC converter to provide a constant 3V supply voltage to the mote. The board also features power routing switches that provide overcharge and undercharge protection for the batteries. Finally, the Heliomote has an on-board power monitor chip that measures the instantaneous current provided by the solar cell, the battery terminal voltage, and also provides current aggregates over a specified time interval.

C. Software interface

The Heliomote board offers a one-wire interface to provide information regarding instantaneous solar power availability, battery terminal voltage, and accumulated current. We have written custom one-wire device drivers for the Mica2 mote, as well as the associated nesC components that provide an API, shown in Table II, for TinyOS applications to easily retrieve this information. To retain flexibility, the Heliomote design does not feature inbuilt power macro-modeling for the solar source. At the cost of a small overhead, the mote can build and characterize a simple power model using the snapshot information provided by the Heliomote board.

TABLE II

NesC interface for querying the Heliomote

```
async command result_t getData();
async event result_t dataReady(uint16_t volt,
 int16_t current, uint32_t rtc, uint8_t acc);
```

D. PCB design considerations

We briefly overview a few PCB design decisions that help improve the Heliomote's reliability. Note that these are not specific to solar energy harvesting, but represent good design practice in general. The ground plane is placed low in the stack, which helps decouple radiated electromagnetic interference (EMI) and prevent it from affecting other circuits and the mote below. Since the majority of the EMI is sourced from the switching supply, boosting the pulse width modulation

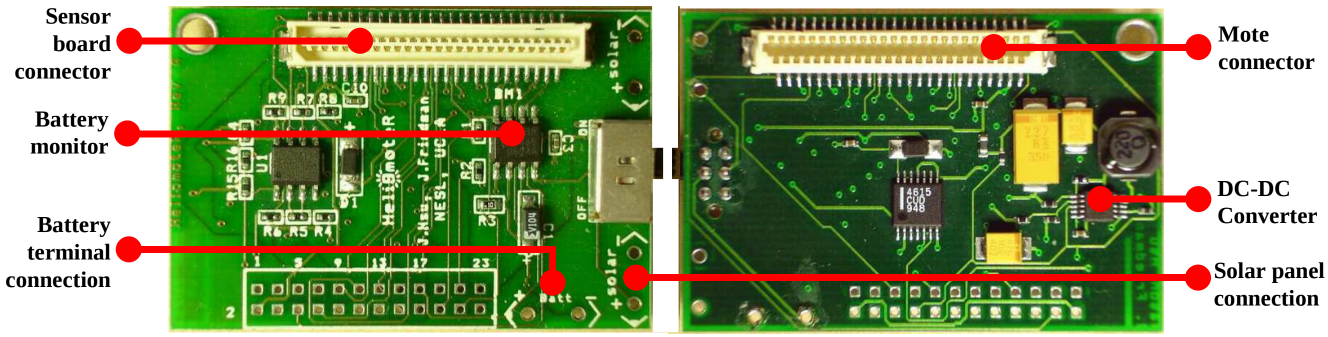

Fig. 4. Top and bottom sides of the Heliomote board.

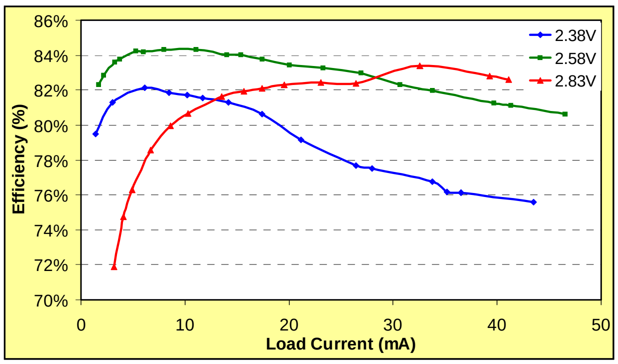

Fig. 5. Efficiency of the Heliomote board as a function of load current at different supply voltages.

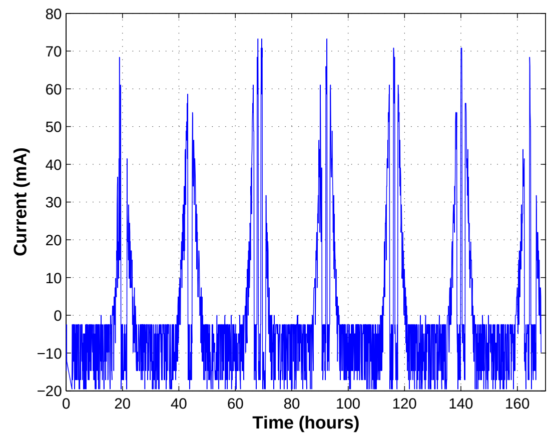

Fig. 6. Plot of battery charging current vs. time

(PWM) frequency of the DC-DC converter to a value well above the operating frequency of the mote enabled us to use smaller filter components, suffer smaller peak currents, and make the filtering more effective since the noise frequencies fall deeper into the filters' stop-bands. The benefits of the increased PWM frequency also apply to the mote itself, making its local decoupling filters more effective against radiated and conducted noise. In addition, special care is taken to place the EMI sourcing components far from the sensitive RF circuits (radio) of the mote and a tuned second order filter topology is present in the mote's power supply path. While we have not encountered any noise related problems during our testing/operation of the Heliomote so far, we are also currently performing a detailed noise analysis to quantitatively characterize this aspect.

E. Performance evaluation

We ran several experiments to evaluate the performance of Heliomote. Our first experiment was to measure the efficiency of the circuit board itself. For this experiment, we disconnected the solar panel, and used a potentiometer as the load to vary the load current. Figure 5 plots the efficiency as a function of the load current, for three different battery voltages. Although the DC-DC converter by itself is specified to be over 90% efficient, power losses in other parts of the board, including the diode, filters, overcharge and undercharge protection circuitry, and the battery monitor IC, decrease the overall efficiency of the Heliomote board to between 80% and 84% for the typical operating range of the motes, as shown in the figure.

Next, we investigated the impact of energy harvesting capability on the residual battery capacity of the motes. A Mica2 mote, augmented with a Heliomote board, was placed outdoors for a week and operated at a 20% duty cycle, which is much higher than the 1%-5% envisioned for most sensor network applications. The mote periodically queried the Heliomote module for solar and battery related information, and transmitted the data to a nearby base station for logging. Figure 6 plots the battery charging current as a function of time. A negative value of current implies that the battery is being discharged. The periodic nature of the curve follows from the periodic (*i.e.*, daily) pattern of solar energy availability. Figure 7 shows the evolution of the battery terminal voltage over time. As can be seen, the battery receives sufficient current from the solar panel each day to remain almost fully charged. In fact, overcharge protection was often activated to re-route solar power away from the battery, as observed by the sudden drops in battery current in Figure 6 when the battery is almost fully charged (*e.g.*, at t = 44 hours). Note that while Figure 6 is indicative of the state of the solar panel, Figure 7 indicates battery state. Although they are not completely independent, it is, in general, not possible to obtain one from the other. Finally, Figure 8 shows the current accumulator reading (initialized to 1000mAH) over the duration of the experiment. The reading increases each day[3], which implies that overall, the residual battery capacity increases. Together, these curves indicate the feasibility of near perpetual operation of the sensor node using our solar energy harvesting module.

VI. Conclusions and Future Work

Environmental energy harvesting has recently emerged as a viable option to supplement battery supplies in energy constrained embed-

[3]For our choice of sense resistor, the current accumulator register can hold a maximum value of 1250mAH, which is why the plot clips off at that value.

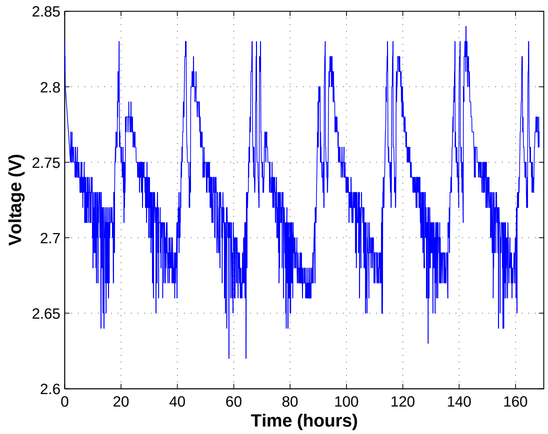

Fig. 7. Plot of battery terminal voltage vs. time

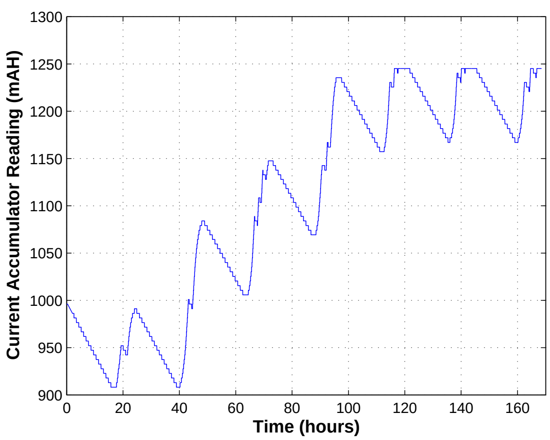

Fig. 8. Plot of the current accumulator reading vs. time

ded systems. However, designing an efficient solar harvesting system involves an understanding of several factors. This paper systematically analyzed the various components, design choices, and tradeoffs involved in the design of a solar energy harvesting module and their impact on the its efficiency. We illustrated how harvesting aware power management improves energy usage compared to battery aware approaches. We presented the design and performance evaluation of Heliomote, our efficient plug-and-play solar harvesting module for the Berkeley/Crossbow motes. Our experimental results indicate the feasibility of near perpetual operation of harvesting aware, outdoor sensor networks.

Acknowledgements

This paper is based in part on research funded through the DARPA PAC/C program, the ONR AINS program, the National Science Foundation, and the Center for Embedded Networked Sensing. We would also like to thank Dunny Potter for his contributions to the Heliomote board.

References

[1] A. P. Chandrakasan and R. W. Brodersen, *Low Power CMOS Digital Design*. Kluwer Academic Publishers, MA, 1996.

[2] M. Pedram and J. Rabaey (editors), *Power Aware Design Methodologies*. Kluwer Academic Publishers, MA, 2002.

[3] K. Lahiri, A. Raghunathan, and S. Dey, "Battery-driven system design: a new frontier in low power design", in *Proc. IEEE International Conference on VLSI Design*, pp. 261–267, 2002.

[4] S. Roundy, *Energy Scavenging for Wireless Sensor Nodes with a Focus on Vibration to Electricity Conversion*. Ph. D. Dissertation, Dept. of EECS, UC Berkeley, May 2003.

[5] V. Raghunathan, C. Schurgers, S. Park, and M. Srivastava, "Energy aware wireless microsensor networks", *IEEE Signal Processing Magazine*, vol. 19, iss. 2, pp. 40–50, March 2002.

[6] R. Min, *et al.*, "An architecture for a power-aware distributed microsensor node", *Proc. IEEE Workshop on Signal Processing Systems*, pp. 581–590, 2000.

[7] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks", *Proc. ACM International Conference on Embedded Networked Sensor Systems*, pp. 95–107, 2004.

[8] R. C. Shah and J. M. Rabaey, "Energy aware routing for low energy ad hoc sensor networks", *Proc. IEEE Wireless Comm. and Networking Conference*, pp. 350–355, 2002.

[9] K. Kalpakis, K. Dasgupta, and P. Namjoshi, "Maximum lifetime data gathering and aggregation in wireless sensor networks", *Proc. IEEE International Conference on Networking*, pp. 685–696, 2002.

[10] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. Srivastava, "Optimizing sensor networks in the energy-density-latency design space", *IEEE Transactions on Mobile Computing*, pp. 70–80, January-March 2002.

[11] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing", *Proc. ACM International Conference on Mobile Computing and Networking*, pp. 70–84, 2001.

[12] V. Schnayder, *et al.*, "Simulating the power consumption of large-scale sensor network applications", *Proc. ACM International Conference on Embedded Networked Sensor Systems*, pp. 188-200, 2004.

[13] J. Kymisis, C. Kendall, J. Paradiso, and N. Gershenfeld, "Parasitic power harvesting in shoes", *Proc. IEEE International Conference on Wearable Computing*, pp. 132–139, 1998.

[14] T. Starner, "Human-powered wearable computing" in *IBM Systems Journal*, 35(3-4), 1996.

[15] M. Rahimi, *et al.*, "Studying the feasibility of energy harvesting in a mobile sensor network", *Proc. IEEE International Conference on Robotics and Automation*, pp. 19–24, 2003.

[16] P. Zhang, C. Sadler, S. Lyon, and M. Martonosi, "Hardware design experiences in ZebraNet", *Proc. ACM International Conference on Embedded Networked Sensor Systems*", pp. 227–238, 2004.

[17] S. Roundy, B. Otis, Y.H. Chee, J. Rabaey, P. Wright, "A 1.9GHz RF Transmit Beacon using Environmentally Scavenged Energy", *ACM International Symposium on Low Power Electronics and Design*, 2003.

[18] Panasonic solar cells technical handbook. (`http://www.solarbotics.net/library/datasheets/sunceram.pdf`)

[19] S. Meninger, *et al.*, "Vibration-to-electric energy conversion", *Proc. ACM/IEEE International Symposium on Low Power Electronics and Design*, pp. 48–53, 1999.

[20] C. Melhuish, *The Ecobot project*. (`http://www.ias.uwe.ac.uk/energy_autonomy/EcoBot_web_page.html`)

[21] J. A. Paradiso and M. Feldmeier, "A compact, wireless, self-powered pushbutton controller", *Proc. ACM International Conference on Ubiquitous Computing*, pp. 299–304, 2001.

[22] Darpa energy harvesting projects. (`http://www.darpa.mil/dso/trans/energy/projects.html`)

[23] C. Park and P. H. Chou, "Power utility maximization for multiple-supply systems by a load-matching switch", *Proc. ACM/IEEE International Symposium on Low Power Electronics and Design*, pp. 168–173, 2004.

[24] T. Voigt, H. Ritter, and J. Schiller, "Utilizing solar power in wireless sensor networks", *Proc. IEEE Conference on Local Computer Networks*, 2003.

[25] A. Kansal and M. Srivastava, "An environmental energy harvesting framework for sensor networks", *Proc. ACM International Symposium on Low Power Electronics and Design*, pp. 481–486, 2003.

[26] A. Kansal, D. Potter, and M. Srivastava, "Performance aware tasking for enrironmentally powered sensor networks", *Proc. ACM International Conference on Measurement and Modeling of Computer Systems*, pp. 223–234, 2004.

[27] James San Jacinto Mountains Reserve. (`http://www.jamesreserve.edu`)

[28] CVS repository, Networked and Embedded Systems Lab, UCLA. (`http://cvs.nesl.ucla.edu`)