Design, Modeling, and Capacity Planning for Micro-Solar Power Sensor Networks

Jay Taneja, Jaein Jeong, David Culler
Computer Science Division, UC Berkeley
Berkeley, CA 94720, USA
{taneja,jaein,culler}@cs.berkeley.edu

Abstract

This paper describes a systematic approach to building micro-solar power subsystems for wireless sensor network nodes. Our approach composes models of the basic pieces - solar panels, regulators, energy storage elements, and application loads - to appropriately select and size the components. We demonstrate our approach in the context of a microclimate monitoring project through the design of the node, micro-solar subsystem, and network, which is deployed in a challenging, deep forest setting. We evaluate our deployment by analyzing the effects of the range of solar profiles experienced across the network.

1 Introduction

The purpose of this paper is to provide a framework for the design of micro-solar subsystems in wireless sensor networks. Its motivation is simple; we were designing a microclimate network for studies of hydrological cycles in forest watersheds and needed a systematic means of engineering, sizing, and analyzing the power subsystem. Many tools and calculators are available for macro-solar installations in residential and commercial applications, but only anecdotal, point designs are represented in the sensor network literature for *in situ* micro-solar power. The basic components are obvious and well documented [13] – solar panels, regulators, and batteries – but the selection, sizing, and composition of the components is not. The problem is rather different from the macro-solar setting because of the very small power transfers involved – microwatts to milliwatts rather than kilowatts to megawatts. Micro-solar operates at very different efficiencies and every bit of power conditioning or monitoring impacts the overall performance. We do not have the luxury of putting the panels on a convenient rooftop with ample exposure, it needs to be where the measurements are to be taken, regardless of how shaded that may be. At the same time, new degrees of design freedom are presented by the tiny magnitude of the energy requirements.

As a preliminary framework, we begin by formulating a general model of micro-solar systems that is sufficient for constructing a capacity planning "calculator" to guide the sizing of the various elements. We then ground the study in a concrete design developed for the HydroWatch application. It is a well-engineered climate monitoring node and network with a flexible power subsystem that can support various specific design points and provides visibility into the solar performance in real application settings. Putting the model and empirical vehicle together, we study the design choices in each element of the solar subsystem to arrive at a deployment candidate. We then utilize this to collect detailed empirical data from the on-going deployment to drive what is expected to be an iterative refinement cycle.

2 Micro-Solar Planning Model

There have been several micro-solar power designs in the literature. [6, 8, 9, 17, 19, 24] We aim to generalize the design space using the basic micro-solar model as illustrated in Figure 1. Ultimately, the demand side is determined by the power requirements of the wireless sensor node and its associated protocols. It has been well established that this load is bimodal [16, 18] with standby current in the neighborhood of 10 uA and active current in the neighborhood of 10 mA. Thus, the duty cycle determines the average power requirement, P_{mote}, as a weighted sum of these two elements that are separated by three orders of magnitude. For example, a 1% duty cycle places the load in the neighborhood of 110 uA, or .33 mW at 3 volts.

The supply side is dictated by the incident solar energy, which is a function of the latitude, day of the year, panel orientation, and angle of inclination. Rules of thumb for various locations are widely available. To obtain greater insight into the trade-offs, we incorporated the basic astronomical

978-0-7695-3157-1/08 $25.00 © 2008 IEEE
DOI 10.1109/IPSN.2008.67

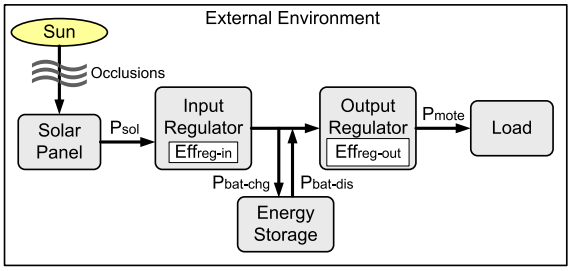

P_{sol}	Power generated from the solar panel
$P_{bat-chg}$	Power input to charge the energy storage
$P_{bat-dis}$	Power discharged from the energy storage
P_{mote}	Power consumed by the load
$P_{shunted}$	Power being shunted when in excess
Eff_{reg-in}	Power efficiency of the input regulator
$\text{Eff}_{reg-out}$	Power efficiency of the output regulator
Eff_{bat}	Charge-discharge efficiency of the energy storage

Figure 1. Micro-solar system architecture and related parameters.

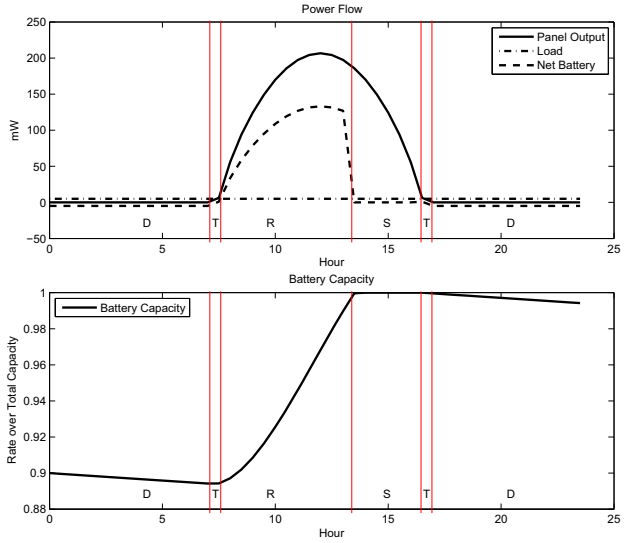

Discharge	$P_{sol} = 0, P_{bat-chg} = 0, P_{bat-dis} > 0, P_{mote} = \text{const}$
	$P_{mote} = P_{bat-dis} \cdot \text{Eff}_{reg-out}$
Transition	$P_{sol} > 0, P_{bat-chg} = 0, P_{bat-dis} > 0, P_{mote} = \text{const}$
	$P_{mote} = (P_{sol} \cdot \text{Eff}_{reg-in} + P_{bat-dis}) \cdot \text{Eff}_{reg-out}$
Recharge	$P_{sol} > 0, P_{bat-chg} > 0, P_{bat-dis} = 0, P_{mote} = \text{const}$
	$P_{sol} \cdot \text{Eff}_{reg-in} = P_{bat-chg} + P_{mote}/\text{Eff}_{reg-out}$
Saturation	$P_{sol} > 0, P_{bat-chg} = 0, P_{bat-dis} = 0, P_{mote} = \text{const}$
	$P_{sol} \cdot \text{Eff}_{reg-in} = P_{shunted} + P_{mote}/\text{Eff}_{reg-out}$

Figure 2. Energy flow and daily phases in our micro-solar model.

calculations directly in the computational model, [7].

The portion of incident solar energy that is available at the panel is determined by a variety of environmental factors. The absorption by the atmosphere is well understood, and we all recognize the spectrum of weather factors, clouds, fog, and so on. In addition, any particular point of installation will have various obstructions and shadows. This critical attenuation factor can only be characterized empirically. Experience with many deployments in different settings can provide statistical models. Care in choosing sites can potentially improve the expected availability. As a rough starting point in this study, we used a guideline that a half hour of sunlight per day should be sufficient to sustain operation. Below, we re-examine this planning guideline in light of specific model parameters and experience in the forest. (It proved to be very optimistic.)

The panel transforms available incident solar radiation to electrical power. A given panel is characterized by its IV curve and, in particular, three points: open-circuit voltage (V_{oc}), short-circuit current (I_{sc}), and its maximum power point (MPP). Internally, these are determined by the serial and parallel composition of the solar cells and the total area of the panel. Increasing temperature depresses the IV curve somewhat, reducing the power output. For the large, expensive panels used in macro-solar installations these factors are accurately characterized in data sheets and well validated. For the small, inexpensive panels used in micro-solar applications, empirical characterization is often required. More importantly, the operating point of the IV curve is determined by the load experienced at the panel, which is determined by the input regulator or the storage facility and downstream load in the absence of a regulator. For most panels, the IV curve is nearly flat for voltages less than that of the MPP, so power increases nearly linearly with V.

The input regulator conditions the output of the panel to meet the operational constraints of the particular battery, including voltage limits, current limits, and charge duration. Whereas macro-solar inverters operate in the neighborhood of 95% efficiency, in the sub-watt range, regulator efficiencies of 70-80% and below are more typical. The product of such low efficiencies translates into a significant overall supply:demand ratio.

A wide range of battery organizations and chemistries are available for storing charge, as well as supercapacitors. They have differing operating voltages, charge algorithms, and complexities. From a system design perspective, it is desirable for the power subsystem to be able to charge a fully discharged battery without software in the loop, so that when placed in sunlight the device is guaranteed to eventually become active.

The portion of energy transferred into the battery during the day and discharged during the night incurs an additional transfer efficiency, Eff_{bat}, about 66% for NiMH chemistries. The capacity of the battery determines the potential lifetime in darkness, but also how much energy can be harvested while the sun shines, as discussed below.

The output regulator matches the battery characteristics to the requirements of the mote. It too is characterized by its efficiency, $\text{Eff}_{reg-out}$, and in particular its efficiency at two very different operating points: 10s of microwatts most

of the time and 10s of milliwatts during short active periods. For a typical bimodal P_{mote}, effective efficiency of 50% or less is expected. This determines the load experienced by the supply and storage components of the power subsystem.

In general, the daily power cycle has five phases, as illustrated in Figure 2. From sundown to sun up, the battery *discharges*, supplying the device load. As the panel is initially illuminated, a *transition period* occurs during which the battery provides only a portion of the device load. With sufficient illumination, the panel supports the entire load and delivers charge into the battery. If this *recharge period* is sustained sufficiently long, the battery becomes fully charged and the system operates in *saturation*, shunting power. Eventually, a dusk transition occurs similar to dawn. The efficiency coefficients dictate the net change in battery capacity over the daily cycle, given the starting capacity, supply power, and demand power. Our sizing guideline was to assume that the recharge period would need to be no more than half an hour, possibly distributed throughout the day. Saturation merely preserves capacity. Of course, a series of overcast days may result in a progressive drop in battery capacity, which would then increase the recharge duration when the weather clears. In the micro-solar setting, given the ratio of mote load and typical battery capacities, it is even reasonable to consider design points that absorb entire seasonal variations in weather patterns.

Just the back-of-the-envelope calculation of solar availability during 2% of operation (i.e., a half hour of radiation during the day) and a 3:1 supply/demand ratio from the product of efficiencies $\text{Eff}_{reg-out} \cdot \text{Eff}_{bat}$ suggest that the solar panel needs to be sized at 150 times the average demand. This makes every aspect of the micro-solar subsystem design critical and motivates the detailed design and analysis in the remainder of the paper.

3 Node and Network Design

The HydroWatch Project [2] aims to collect widespread, high-frequency, and automated observations of the life cycle of water as it progresses through a forest ecosystem. To gather this data, we aimed to deploy a robust network of low-maintenance sensor nodes that could collect scientifically-relevant data indefinitely while withstanding a challenging wet forest environment. This section details the network architecture, node mechanical design, and micropower solar subsystem present on each node.

3.1 Network Architecture

The sensor network architecture follows the canonical habitat monitoring form described in [21], but is somewhat of a second-generation sensor network utilizing several commercially available elements. The sensor node is

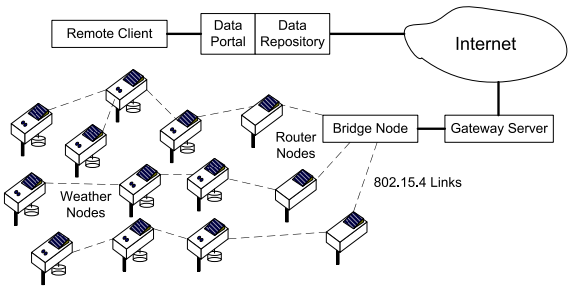

Figure 3. System architecture for the HydroWatch microclimate network.

built around the TelosB-compatible Tmote Sky [3], and is described in detail below. The mote software, which provides periodic data acquisition, thresholding, power management, remote command processing, and health monitoring, is a modified Primer Pack/IP [1] based on TinyOS 2.0. The patch network is an implementation of IPv6 using 6LoWPAN over IEEE 802.15.4 radios [14]. It utilizes a packet-based form of low-power listening [15] to minimize idle listening. Data collection is implemented as UDP packets with the routing layer using hop-by-hop retransmissions and dynamic rerouting in a redundant mesh (up to three potential parents) to provide path reliability on lossy links. It utilizes Trickle-based [12] route updates for topology maintenance. Source-based IPv6 routing is used to communicate directly to specific nodes and dissemination is performed as a series of IPv6 link-local broadcasts.

In the initial HydroWatch deployment at the Angelo Reserve in Northern California, the sensor patch contains 19 nodes over a 220m x 260m area stretching across a deep ravine formed by Elder Creek, up the deeply forested north slope of the watershed area, and bending to the east to a particular tall stand of Douglas Fir trees.

The transit network between the base-station and the patch is implemented using the same node and link technology as the patch, so there is no specific gateway node in the patch. To provide redundancy in the transit network, multiple micro-solar router nodes cover a 120m stretch from a shed housing the gateway across an old apple orchard. These nodes are just patch nodes without the environmental sensors. The network depth is 5 hops or greater. The IEEE 802.15.4 bridge node attached to the base-station uses a high-gain (19 dBi) parabolic antenna pointed out through a window in the shed. The back haul network is a WiFi-based IP network with repeaters on peaks and tree tops to reach a T1 line.

The base-station is a Linux-class gateway server that provides a web services frontend, a PostgreSQL database for information storage and retrieval, and a web-based man-

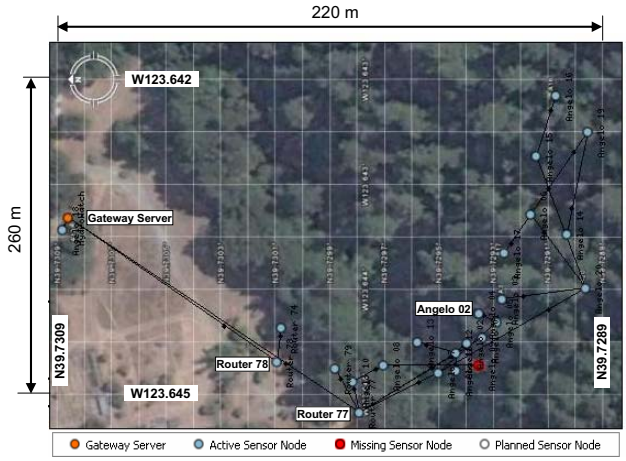

Figure 4. Snapshot of the HydroWatch forest watershed deployment.

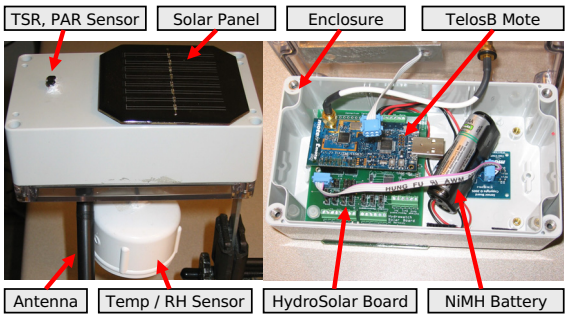

Figure 5. HydroWatch weather node.

agement console. It is also an IP router, permitting end-to-end connectivity to the patch nodes. The server facilitates such tasks as monitoring overall network health remotely, diagnosing misreporting or missing nodes, and checking the quality of links a node has to its neighbors, a function which proved critically important during the deployment phase.

The initial physical deployment of this architecture is illustrated in Figure 4. The base-station appears at the left (west) side of the image. The router nodes form a sparse stretch reaching southeast. A fairly rich interconnection is provided among the several nodes up the watershed. Node sites were chosen to achieve both vertical (up the hill) and horizontal (across the hill) profiles to enable an initial understanding of the microclimates in the watershed.

3.2 Engineering the Node

While the mote platform (microcontroller, radio, and flash), system software, and networking are fairly common across many applications, the sensor suite, power subsystem, and mechanical design of the node tend to be application-specific and highly inter-related.

The sensor suite for this microclimate monitoring application is essentially that developed for tracking weather fronts in Redwoods [22] and available natively on the TelosB platform – total solar radiation (TSR), photosynthetically active radiation (PAR), temperature, and relative humidity. However, to provide high-quality data, the sensors must be exposed properly to the environment while the rest of the electronics are protected. We used the TelosB-compatible Tmote Sky mote [16] with an attached SMA connector for an external antenna and no on-board sensors. We connected the mote to external sensorboards using custom cables with IDC connectors, providing a degree of freedom to determine sensor orientation on the node.

The two photodiodes used to measure incident PAR (Hamamatsu S1087) and TSR (Hamamatsu S1087-01) are tiny discs connected to two long leg-like contacts. Previous experience [22] recommends that these sensors should be elevated from their surroundings to avoid collecting water on the sensing surface and to obtain unobstructed indications of solar illumination.

A Sensirion SHT15 sensor provides relative humidity (RH) and temperature, factory calibrated to exhibit a maximum +/- 2% RH and +/- 0.3 degrees Celsius error. To accurately measure humidity requires the sensor be exposed to naturally-aspirated air flow, whereas to measure temperature it should be shaded and decoupled from large thermal masses and sources of self-heating. To meet these demands, the SHT15 was placed within a 2-in PVC cap and suspended from the underside of the node. To prevent electrical malfunction, the hole in the top of the PVC cap for the interface wire was sealed and conformal coating was applied to the sensorboard. Though we recognize that a hanging design may be prone to connection disturbances caused by curious wildlife, the accuracy of sensor data was worth the risk of a small number of unavailable nodes.

The RF environment was expected to pose a critical challenge for our network, due to interactions of foliage and water vapor with 2.4 GHz radio connectivity. Evident from past deployments [8, 11, 20, 23, 24] and more specifically forest deployments [5, 22], natural environments are capable of eliciting tremendous swings in link quality causing wireless networks to fail drastically and unexpectedly. To address this issue, we provisioned our nodes with a 7 dBi omnidirectional antenna with a flexible orientation. Nodes were attached to the top of 3 ft and 4 ft metal fence posts by a simple angle bracket.

Overall, the mechanical design of the node sought to provide accurate sensor readings over a long duration. To protect the internal electronic components from environmental damage, we limited the number of node features that required holes in the enclosure; every opening in the en-

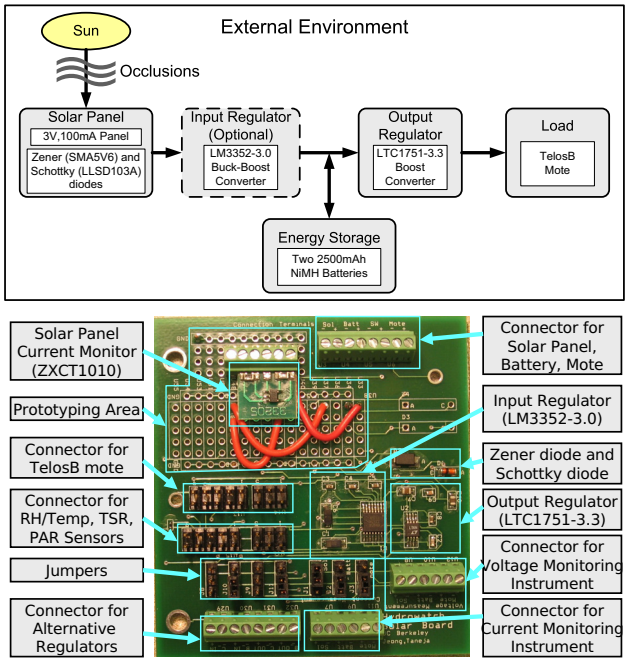

Figure 6. HydroSolar micro-solar power subsystem.

closure is a potential opportunity for a leak. Additionally, packets of indicating desiccant were placed inside each enclosure box, preventing water vapor from damaging interior electronic components.

3.3 Micro-Power Subsystem

The core of the node design is a flexible power subsystem board that ties together a solar panel, an optional input regulator, a battery, and a switching output regulator – as shown in Figure 6. It provides measurement points for a number of electrical parameters that can be connected to the mote ADCs, sampled, and recorded along with the environmental measurements. In our configuration, these monitoring features produce time-series logs of solar panel voltage, solar panel current, and battery voltage, in addition to the logs of sensor data from the application and link/neighbor data. All of these are collected and stored by the gateway server, enabling deeper analysis of the performance of the node and network under varying solar conditions. The solar board also provides the mechanical structure that attaches the mote to the enclosure.

The HydroSolar board was designed to permit the study of a variety of power subsystem options. The solar panel and the battery are attached through screw terminals. Headers and mounting holes permit direct attachment of motes of the TelosB form factor, but a mote of any other type can be attached to the board through screw terminals. Additionally, the board has a prototyping area which can be used to change power subsystem configuration. In fact, we were able to change any of the circuit elements originally used in the board schematic by simply changing jumper settings and populating the prototyping area. We used this flexibility to evaluate candidate parts for each component and quantify their contribution to the efficiency of the entire system.

4 Micro-Solar Sizing and Selection

While Section 2 outlines the components of micro-solar power subsystems, this section provides the rationale and key criteria for selecting specific components as seen through the lens of our experience designing the HydroSolar board. We begin with an analysis of application load – this directly impacts the selection of the other components in the design. The components ultimately selected for the HydroSolar micro-solar board are summarized in Table 1.

Table 1. Characteristics of the components used for the HydroSolar board.

(a) Solar Panel (Silicon Solar #16530)	
V_{oc} , I_{sc}	4.23V, 111.16mA
MPP	276.0mW at 3.11V
I-V curve	Shown in Figure 7
Dimension	2.3in x 2.3in
Material, Efficiency	Polycrystalline silicon, 13%
(b) Input Regulator (LM3352-3.0: Optional)	
Manufacturer-provided efficiency	65%-83% (I_{out} = 5mA-100mA, V_{out} = 3.0V, V_{in} = 2.5V-3V)
Measured efficiency	54.71%-65.40% (I_{solar} = 0mA-100mA, V_{out} = 3.0V)
(c) Energy Storage	
Configuration	Two AA NiMH batteries in series
Voltage	2.4V nominal, 2.6V-3.0V at charge
Capacity	2 × 1.2V × 2500mAh = 6000mWh
(d) Output Regulator (LTC1751-3.3)	
Manufacturer-provided efficiency	55%-60% (I_{out}=0.1mA-20mA, V_{in}=2.75V, V_{out}=3.3V)
Measured efficiency	49.69%-52.15% (I_{out}=3mA-6mA, V_{in}=2.55V-2.71V, V_{out}=3.3V)
(e) Load	
Mote platform	Tmote Sky / TelosB mote
Vcc	2.1V - 3.6V, 2.7V - 3.6V with flash
Average current	App.-Dependent; 0.53mA for ours
Maximum current	23mA with MCU on, radio RX

4.1 Load

To get a notion of the power requirements of a node, we empirically measured the load created by our application.

Table 2. Estimated lifetime of a node using each energy storage element without recharging.

Type	Lifetime
Lead Acid (LC-R061R3P)	98.5 days (= $\frac{7800mWh}{79.2mWh/\text{day}}$)
Two NiCd (KR-1100AAU)	33.3 days (= $\frac{2\times1320mWh}{79.2mWh/\text{day}}$)
Two NiMH (NH15-2500)	75.8 days (= $\frac{2\times3000mWh}{79.2mWh/\text{day}}$)
Li-ion (UBP053048)	35.4 days (= $\frac{2800mWh}{79.2mWh/\text{day}}$)
Li-polymer (UBC433475)	42.9 days (= $\frac{3400mWh}{79.2mWh/\text{day}}$)
Supercap (BCAP0350)	3.8 days (= $\frac{304mWh}{79.2mWh/\text{day}}$)

As is typical of sensor networks for environmental data collection, nodes alternate between a low-power state roughly 99% of the time and brief higher-power active periods. The gateway server provides estimates of the duty cycle for the MCU (0.4%) and the radio (1.2%). The peak active current is 23 mA with the MCU on and the radio in RX mode, the sleep current is around 15 uA, and the RMS average current is 0.53 mA. We use our application load requirement to guide our selection of the rest of the components.

4.2 Energy Storage

Table 3 lists a number of possible rechargeable energy storage options that can be used for micro-solar power systems. We consider a number of characteristics including capacity, operating range, energy density and charging method.

Employing the measured average consumption of our application of 0.53mA at 3.3V and the efficiency of the output regulator estimated at 50%, the daily energy requirement from the energy storage element is 79.2 mWh. This energy requirement drives the storage selection process. First, we compare each type of storage based on capacity in Table 2. All options except the supercapacitor can provide energy for more than 30 days of operation without recharging, long enough to operate for a number of days in the absence of solar radiation.

For our application, even with loose physical sizing constraints, lead-acid batteries are not plausible because of low energy density. NiCd batteries have a similar footprint and charging method as NiMH batteries, but with a much smaller capacity. Additionally, NiCd chemistries are less environmentally-friendly and far more susceptible to the memory effect, which can significantly reduce battery capacity over time.

For the decision between Lithium-based chemistries and NiMH, we drew on previous experience from the Trio deployment [8]. Our desire to avoid having software in the charging loop (ultimately to allow nodes to simply charge when placed in the sun entirely independent of their software state) coupled with the complexity of integrating a hardware Li-ion charger dictated the selection of NiMH as it operates with more straightforward charging logic. This choice does present some drawbacks, however. This chemistry suffers from a self-discharge rate of 30% and an input-output efficiency of roughly 66%, both worse than for any other battery chemistry considered. The practical implication of this is that for every 3 units of energy that are input to a battery, only 2 units of energy are output. We felt this cost was overcome by the simplicity of the charging logic, though.

A 2-cell configuration would enable the potential to operate without an input regulator; this choice is further discussed in Section 4.4. For increased capacity, it would be possible to put 2-cell packs in parallel. Additionally, since the discharge curve of NiMH batteries is relatively flat, most of the discharge cycle produces a near-constant voltage.

4.3 Solar Panel

In selecting an appropriate panel for a micro-solar subsystem, the critical factors are the panel's IV curve (specifically, the MPP), its cell composition, and its physical dimensions. Care should be taken in selecting a panel that will operate near its MPP given the load it is expected to support, be it a combination of an input regulator and energy storage or energy storage alone. The cell composition – that is, how many cells are present and their serial/parallel arrangement – becomes a factor when the solar panel is partially occluded. We discuss this in greater depth in Section 5.2. Last, the physical dimensions of the panel should be compatible for the choice of enclosure.

For the HydroSolar power subsystem, we selected a 4V-100mA panel from Silicon Solar Inc. whose characteristics are summarized in Table 1(a) and whose IV and PV curves are in Figure 7. The MPP of this panel occurs at 3.11V, which makes it appropriate for charging 2 NiMH cells directly. Using the astronomical model from Section 2, we were able to vary the latitude, day of year, time of day, panel orientation, and angle of inclination to match the conditions expected in our field deployments. As a basic approximation using our rule of thumb of 30 minutes of sunlight per day, the solar energy generated by this panel at its MPP is 139 mWh, satisfying the 120 mWh (= 79.2 mWh / 66% NiMH charge-discharge efficiency) per day requirement of our application.

4.4 Input Regulator

In selecting the input regulator, the important parameters are the operating range of the solar panel and batteries and the method and logic used to charge the battery. In our

Table 3. Different types of energy storage elements for micro-solar power systems.

Type	Lead Acid	NiCd	NiMH	Li-ion	Li-polymer	Supercap
Make	Panasonic	Sanyo	Energizer	Ultralife	Ultralife	Maxwell
Model No.	LC-R061R3P	KR-1100AAU	NH15-2500	UBP053048	UBC433475	BCAP0350
Characteristics of a single storage element						
Nominal voltage	6.0 V	1.2 V	1.2 V	3.7 V	3.7 V	2.5 V
Capacity	1300 mAh	1100 mAh	2500 mAh	740 mAh	930 mAh	350 F
Energy	7.8 Wh	1.32 Wh	3.0 Wh	2.8 Wh	3.4 Wh	0.0304 Wh
Weight energy density	26 Wh/Kg	42 Wh/Kg	100 Wh/Kg	165 Wh/Kg	156 Wh/Kg	5.06 Wh/Kg
Volume energy density	67 Wh/L	102 Wh/L	282 Wh/L	389 Wh/L	296 Wh/L	5.73 Wh/L
Weight	300 g	24 g	30 g	17 g	22 g	60 g
Volume	116.4 cm^3	8.1 cm^3	8.3 cm^3	9.3 cm^3	12.8 cm^3	53.0 cm^3
Self-discharge (per month)	3% - 20%	10%	30%	$< 10\%$	$< 10\%$	5.9% / day
Charge-discharge efficiency	70% - 92%	70% - 90%	66%	99.9%	99.8%	97% - 98%
Memory effect	No	Yes	No	No	No	No
Charging method	trickle	trickle/pulse	trickle/pulse	pulse	pulse	trickle

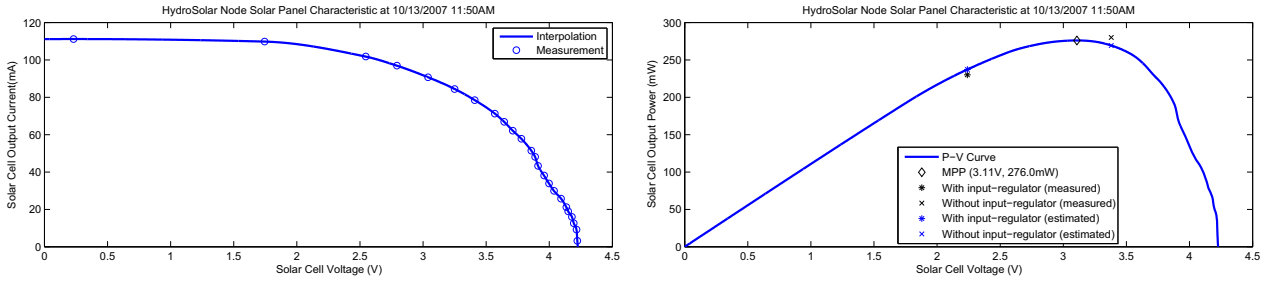

Figure 7. Current-Voltage and Power-Voltage performance of the Silicon Solar 4V-100mA solar panel.

design, we choose to trickle charge the batteries because it requires only a simple circuit and no software control.

In our initial design of the HydroSolar board, we used an input regulator to limit the voltage to the battery. However, we observed that the existence of the input regulator forced the solar panel to operate at a point far from its MPP. Not using the input regulator results in significantly more energy harvested from the solar panel because the input impedance of the regulator is less than that of the battery – see Figure 7(b). In addition to this increase, energy is no longer consumed by the input regulator, which empirically has about a 60% efficiency factor. This substantial gain in total system energy as well as efficiency led us to remove the input regulator from our design; removing the input regulator is only an option because the operating voltage of the solar panel matches the charging voltage of the batteries.

4.5 Output Regulator

The key criteria for choosing an output regulator are the operating ranges of the batteries and the load, as well as the efficiency of the regulator over the range of the load. With our choice of 2 NiMH AA batteries, the nominal voltage of the energy storage is 2.4V so a boost converter is required to match the 2.7-3.6V operating range of TelosB motes (Table 1(e)). The output regulator also has the important responsibility to provide a stable supply voltage to ensure the fidelity of sensor data. Though DC-DC converters introduce high-frequency noise from the switching process into the output signal, the amplitude of the noise does not negatively affect the sensor readings. If noise were a critical factor, either a low-pass filter or a higher voltage energy supply in combination with a linear drop out (LDO) regulator could be used instead.

We chose the LTC1751 regulator, which had an efficiency of around 50%. It requires very few discrete parts and has low, constant switching noise. However, as we learned how optimistic our capacity planning was in the forest watershed deployment (explained in Section 5.2), we would review our choice of output regulator. Table 4 shows the efficiency of a few suitable components at relevant output currents.

5 Evaluating the Design

To evaluate our model and design, we deployed two test networks of nodes with the HydroSolar subsystem. In both

Table 4. Power efficiency of a few 3.3V DC-DC boost converters.

V_{out}=3.3V	I_{out} 0.1mA	I_{out} 1mA	I_{out} 10mA
LTC1751 (V_{in}=2.75V)	55%	60%	60%
TPS61201 (V_{in}=2.4V)	45%	75%	80%
MAX1724 (V_{in}=2.5V)	78%	80%	82%

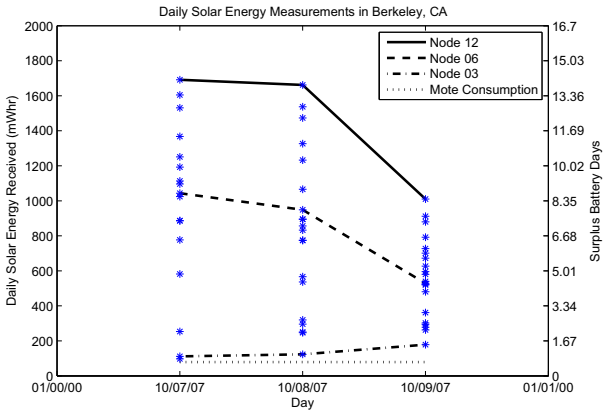

Figure 8. Scatter plot of solar energy received in the urban neighborhood deployment. Three representative nodes are highlighted.

cases, we used the same Primer Pack/IP gateway server and node application software as described in Section 3 with a combination of weather and routing nodes.

5.1 A Sensor Network in an Urban Neighborhood

The purpose of our first deployment was to confirm that nodes could sense, charge, and operate continuously for a period of days, as well as assess whether the model we developed accurately estimated the generation and consumption of energy in a variety of solar conditions. We deployed 22 nodes in an urban neighborhood in Berkeley; nodes were placed in varied locations, including on a house gutter, in and under trees, among shrubbery, and in a grassy yard. To emulate the situation in the forest watershed, we placed them in the vicinity of significant obstructions and varied the orientation of the solar panels: some were flat while others faced south, east, and west at a 45 degree inclination.

The range of daily solar energy via P_{sol} by each node over a period of three days can be seen in Figure 8. The lines on the graph show the behavior of the node that received the highest (Node 12), median (Node 06), and lowest (Node 03) amount of solar energy. The fourth line on the graph shows a constant 79.2 mWh break-even point. The first day (10/07/2007) was a fairly sunny day, resulting in the widest distribution of received solar energy (roughly 100-1700 mWh). However, as the days became cloudier, the variance of the distribution lessened; nodes at the high end of the distribution received slightly more than half the solar energy when cloudy compared to a sunny day. Interestingly, nodes on the lower end of the distribution received *more* solar energy on cloudier days; this is presumably because the diffusion of light caused by the layer of clouds scatters the light source and *enhances* the opportunity of the normally-occluded solar panel to harvest solar energy.[1] Nonetheless, every node harvests a surplus of energy on both sunny and cloudy days; the number of surplus battery days this energy creates is also in Figure 8. Surplus battery days are calculated by multiplying the surplus of energy flowing into the battery by the charge-discharge efficiency (66%) and dividing by the daily consumption (79.2 mWh).

Looking at the daily graph of solar current and voltage experienced at each of the three representative nodes on a sunny day – shown in Figure 9 – we can see the variations in available solar energy inputs among nodes throughout a day. Nodes that generated very little solar energy still had a solar panel voltage above 3 volts for the light portion of the day. This voltage is limited by the load – in this case, the batteries. Thus, the solar voltage exhibits near binary behavior between 0 volts when there is no incident light and its maximum voltage (as dictated by its load) any time between dawn and dusk. Additionally, these current graphs are plotted alongside the astronomical model described in Section 2 as a basis for comparison. The solar profile in each case fits the astronomical model except for discrepancies caused by shadows from buildings and trees, non-optimally directed panels, or cloudy days. For example, in the current graph for each of the nodes, for various periods the panels are obstructed and the current falls significantly. Also, the panel on Node 06 only receives high current in the afternoon sun in accordance with the panel facing west. The sporadic pattern of the solar energy received throughout the day has implications for the daily power cycle introduced in Figure 2 as well; the progression through the daily model may instead oscillate among the recharge, saturation, and discharge phases during the daylight hours.

The urban neighborhood deployment demonstrated that even nodes with severe arboreal and other occlusions received enough sunlight to sustain operation; that is, the nodes in the most shade still received at least 30 minutes of sunlight on both sunny and cloudy days validating the prediction of our model and making us (falsely) confident that our design would succeed in the forest watershed.

[1] This effect is most pronounced in this figure (the solar energy *doubles* on a cloudy day for Node 03), but appeared in other observations as well.

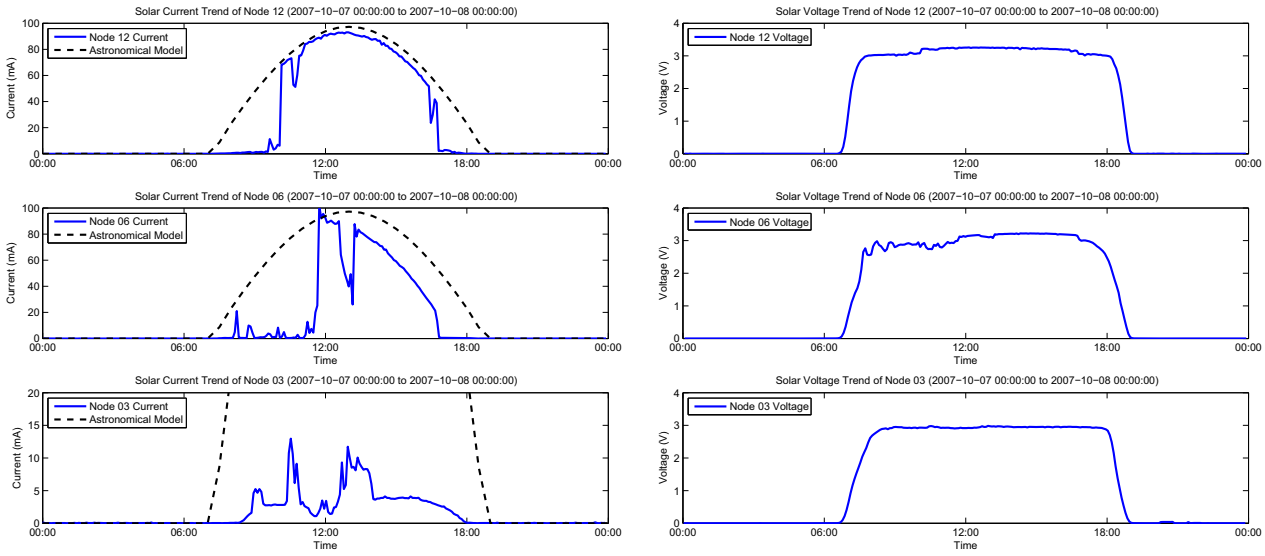

Figure 9. Comparison of solar panel output current and voltage on a sunny day for the urban neighborhood deployment. Notice the differences in scale of the graphs.

5.2 A Sensor Network in a Forest Watershed

The blend of solar profiles seen by the nodes in the forest watershed was far less diverse than the urban neighborhood as shown in Figure 12. Most of the nodes received no more than 50 mWh of energy on any of the days of the deployment. Just as in Figure 8, the lines represent nodes chosen to show the range of the solar distribution. However, in Figure 12, the middle line represents the second-best performing node (not the median) and the lowest line is for a node representative of those that are receiving very limited energy. The stunning difference between the two deployments is how much less solar energy was harvested in the forest watershed – the best-performing node on a sunny day in the forest did not receive as much solar energy as the median node on a cloudy day in the urban neighborhood. Additionally, Angelo 02 (and other sun-starved nodes like it) harvested less than the node consumption *each day*. This daily energy deficit results in a negative number of surplus battery days. It is important to note that these nodes are experiencing different degrees of sun starvation – some are only consuming about half a day's worth of battery energy daily, while others are consuming a full day's worth of energy daily. Still, a majority of the nodes were not receiving sufficient solar energy to operate sustainably, causing a finite lifetime for the network.

What was the cause of such critical energy shortages? Figures 10 and 11 show the solar current and voltage of the three representative nodes on a sunny and overcast day, respectively. The solar voltages exhibit the familiar binary behavior in both cases. The solar currents noticeably suffer on the overcast day, but the heavily shaded node slightly improves its energy harvesting. Perhaps the most important observation is how spiky the solar profile is for the nodes that receive reasonable amounts of solar energy.

It appears that the primary limitation of available solar energy in the forest context is not the amount of light, but the speckled nature of the light that is present. Rarely is the spot of light that falls on even our small panels large enough to illuminate the entire panel. Overcast days diffuse the shadows, reducing the spotting. An individual solar cell produces about 0.5 volts, so several are placed in series within the panel to provide a useful output voltage. For example, our panels have a chain of eight cells in series. The current of the cell is determined by its area, and cells can be interconnected in various serial-parallel networks. The problem is that when a single cell in a serial chain is not well-illuminated, it limits the current flow through the entire chain. A simple experiment connecting panels in serial or parallel confirms this behavior. Thus, enlarging the panel does not necessarily increase the power output in speckled light. Instead, many small panels should be connected in a highly parallel configuration. Large residential and commercial arrays have this character because of the sheer number of panels involved. We are not aware of any such array structures for micro-solar panels.

Increasing the battery size also has surprising implications. With the low daily consumption of a well-engineered environmental monitoring application, it is reasonable to

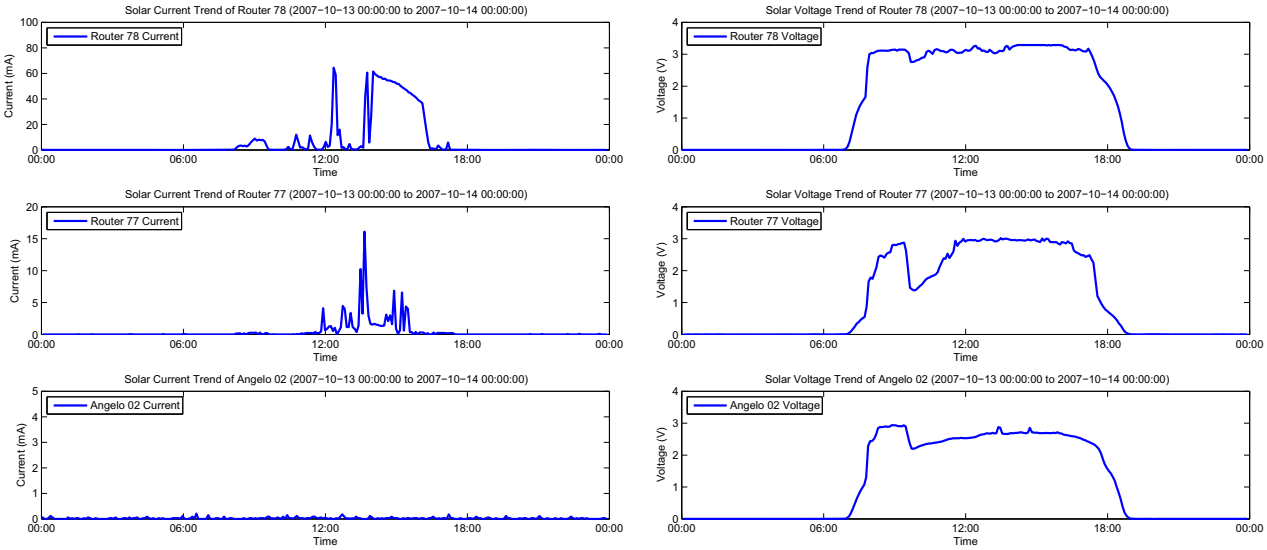

Figure 10. Comparison of solar panel current and voltage on a sunny day (10/13/2007) in the forest watershed deployment. Notice the differences in scale of the graphs.

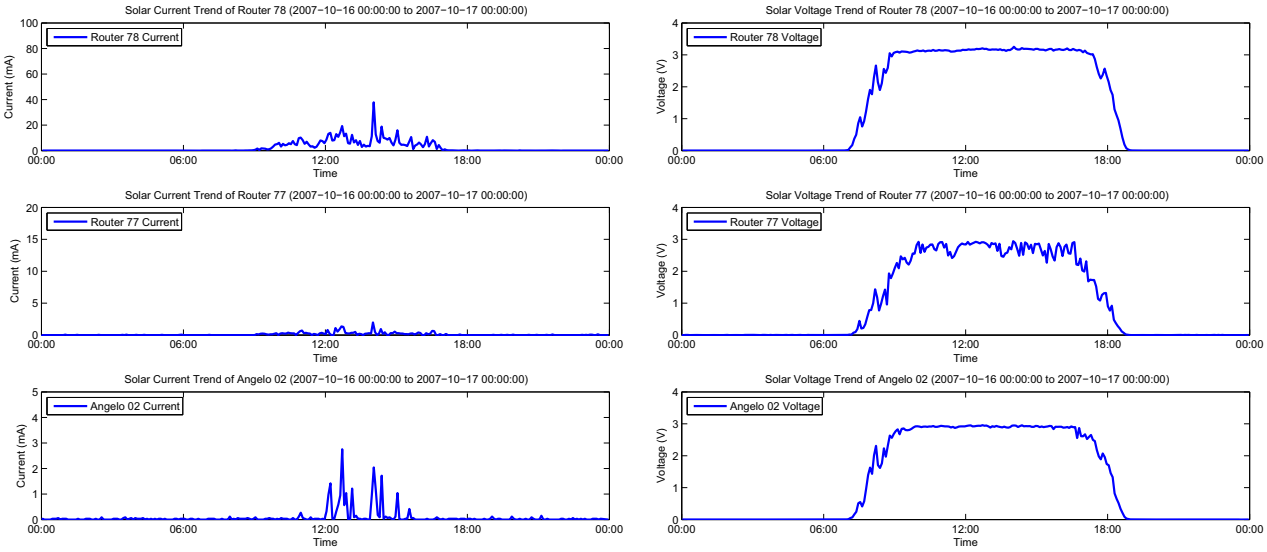

Figure 11. Comparison of solar panel current and voltage on an overcast day (10/16/2007) in the forest watershed deployment. Notice the differences in scale of the graphs.

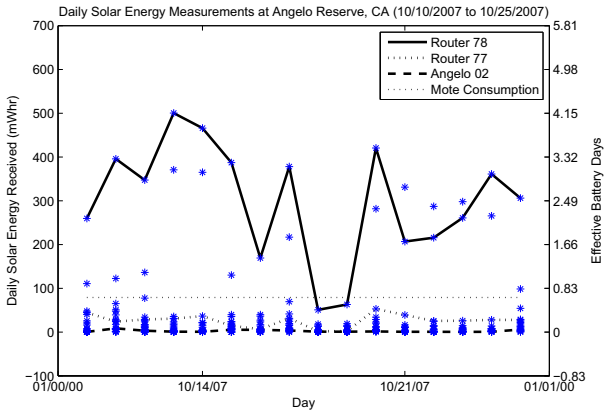

Figure 12. Scatter plot of solar energy received in the forest watershed deployment. Three representative nodes are highlighted.

size batteries to last for several seasons. In deciduous forests, this would allow nodes to store up all their energy after the leaves fall. Even in coniferous forests, it means that energy can be collected when the interaction of the canopy and the sun angle are most favorable.

Additional improvements are possible through utilizing more efficient regulators with somewhat more complex circuit requirements. Exploration of novel collectors and storage profiles for important solar-challenged environments will drive further improvements in the models as well as the physical design.

6 Related Work

In an effort to support sustainable sensor networks, several research groups have developed micro-solar power subsystems. Heliomote [17], which consists of a solar panel, NiMH battery, and a boost converter for controlling load supply voltage, demonstrated sustainable operation of a single mote-based node with a 20% fixed duty cycle in a week-long experiment. Though the components used by Heliomote are similar to the HydroSolar board and many design decisions were reached similarly, the design in the paper is not driven by a realistic application and the evaluation of the paper is limited to a single node with copious available solar energy. Kansal *et al.* [10] showed an analytical model of micro-solar power systems. Using mathematical analysis, they showed how each component of a micro-solar power system should be related for sustainable operation. Additionally, the authors introduce an algorithm for varying the duty cycle based on the available solar energy and evaluate it mathematically. The paper also includes empirical results of a single Heliomote with a 40% fixed duty cycle sustained for over two months during the summer in Los Angeles. In comparison, our model augments this work by considering solar energy input variations by using an astronomical model with occlusion effects and the efficiency implications of using non-ideal regulators. Furthermore, our system is evaluated with a real application in a variety of challenging solar environments.

Prometheus [9] consists of a solar panel, a two-tier storage hierarchy of supercapacitors and a Li-ion battery, and software-controlled battery charging. While Li-ion batteries have higher discharge efficiency than NiMH batteries and the use of tiered storage improves the battery lifetime, its use of software-controlled charging can be problematic. This was evident in Trio [8], which used Prometheus for a long-term outdoor deployment. When charging logic on the mote did not work properly, the battery was not charged even with sufficient solar radiation.

ZebraNet [24], whose energy harvesting nodes are composed of solar panels, a Li-ion battery, and a boost converter for battery charging, was deployed for outdoor habitat monitoring. Application requirements (GPS sensors and long-range radios) dictated power consumption 15 - 30 times higher than a mote device, leading to a design focus of minimizing the duty cycle of high energy components. ZebraNet developed application-driven hardware for solar energy harvesting and considered capacity needs and the effects of solar cell shading; it represents a single point in the design space that could have been formulated using our model.

Everlast [19], which consists of a solar panel, buck converter, supercapacitor, and step-up regulator, is designed with two key points: first, a larger number charge-discharge cycles is possible by using a supercapacitor instead of batteries as the energy storage; second, the operating point of the solar panel is continually optimized by using maximum power point tracking (MPPT). While MPPT does help increase the solar energy input into the system, the MPPT method used requires control by the MCU and little is discussed on the efficiency and energy consumption of the entire system including the two regulators.

Fleck [6] nodes have an energy subsystem consisting of a solar panel, NiMH battery, and a boost converter for controlling the load supply voltage. Like Trio, Fleck improved micro-solar power sensor nodes by demonstrating long-term and large-scale outdoor deployments. However, the system was designed to work only in ample sunlight and had limited consideration for other solar inputs. Fleck presents another specific design that could be represented using our model.

7 Conclusion

We began this work with the goal of creating the power subsystem for a microclimate sensor network for studies of

hydrological cycles in forest watersheds. To explore the design space of micro-solar power systems, we created a model for each of the constituent components and calculated that half an hour of sunlight per day is an appropriate requirement for these nodes to operate perpetually. This approach enabled us to provision our system specifically for the application load we expected, including a low-power multi-hop networking stack, a critical component for building large-extent, low-duty-cycle, and highly-scalable sensor networks. Then, we designed our solar-energy harvesting module based on the energy budget predicted by an astronomical model of the sunlight we could expect to see at our deployment location. In addition, we augmented our system with circuit monitoring capabilities to enable further analysis of performance and iterative improvements to guide future design of micro-solar power subsystems. In a series of deployments of the HydroSolar board we created in accordance with our model, we discovered that our prediction of available sunlight was accurate for an urban neighborhood setting, yet highly optimistic for a forest watershed. With this empirical observation, we refined our model and identified potential solutions to the challenge of designing a node that could operate indefinitely in forested or otherwise solar-challenged environments. Our experience provides insight into the unique issues that arise from designing micro-solar power systems as opposed to the more familiar realm of macro-solar power systems.

Acknowledgments

This material is based upon work supported by the Keck Foundation under a grant for the HydroWatch Center as well as by the National Science Foundation under grants #0435454 ("NeTS-NR") and #0454432 ("CNS-CRI"). This work was also supported by a National Defense Science and Engineering Graduate Research Fellowship as well as generous gifts from the Hewlett-Packard Company, Intel Research, and California MICRO.

References

[1] Arch Rock Corporation, Primer Pack / IP. `http://www.archrock.com/downloads/datasheet/primerpack_datasheet.pdf`.

[2] HydroWatch Project. `http://hydrowatch.cs.berkeley.edu`.

[3] Sentilla Tmote Sky. `http://www.sentilla.com/pdf/eol/tmote-sky-datasheet.pdf`.

[4] P. Buonadonna, D. Gay, J. M. Hellerstein, W. Hong, and S. Madden. TASK: Sensor Network in a Box. *EWSN*, 2005.

[5] R. Cardell-Oliver, K. Smettem, M. Kranz, and K. Mayer. A reactive soil moisture sensor network: Design and field evaluation. *IJDSN*, Mar. 2005.

[6] P. Corke, P. Valencia, P. Sikka, T. Wark, and L. Overs. Long-Duration Solar-Powered Wireless Sensor Networks. *EmNets*, Jun. 2007.

[7] J. V. Dave, P. Halpern, and H. J. Myers. Computation of Incident Solar Energy. *IBM Journal of Research and Development*, 19(6):539–549, 1975.

[8] P. Dutta, J. Hui, J. Jeong, S. Kim, C. Sharp, J. Taneja, G. Tolle, K. Whitehouse, and D. Culler. Trio: Enabling sustainable and scalable outdoor wireless sensor network deployments. *SPOTS*, Apr. 2006.

[9] X. Jiang, J. Polastre, and D. Culler. Perpetual Environmentally Powered Sensor Networks. *SPOTS*, Apr. 2005.

[10] A. Kansal, J. Hsu, S. Zahedi, and M. B. Srivastava. Power Management in Energy Harvesting Sensor Networks. *ACM Trans. in Embedded Computing Systems*, 6(4), 2007.

[11] K. Langendoen, A. Baggio, and O. Visser. Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture. *WPDRTS*, Apr. 2006.

[12] P. Levis, N. Patel, D. Culler, and S. Shenker. Trickle: A Self-Regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks. *NSDI*, Mar. 2004.

[13] R. A. Messenger and J. Ventre. *Photovoltaic Systems Engineering, Second Edition*. CRC, 2003.

[14] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler. Transmission of IPv6 Packets over IEEE 802.15.4 Networks, September 2007. `http://tools.ietf.org/html/rfc4944`.

[15] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. *Sensys*, Nov. 2004.

[16] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power wireless research. *SPOTS*, Apr. 2005.

[17] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. Srivastava. Design Considerations for Solar Energy Harvesting Wireless Embedded Systems. *SPOTS*, Apr. 2005.

[18] V. Shnayder, M. Hempstead, B. Chen, G. Werner-Allen, and M. Welsh. Simulating the power consumption of large-scale sensor network applications. *ACM SenSys*, Nov. 2004.

[19] F. Simjee and P. H. Chou. Everlast: Long-Life, Supercapacitor-Operated Wireless Sensor Node. *ISLPED*, Oct. 2006.

[20] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, and D. Culler. An analysis of a large scale habitat monitoring application. *Sensys*, Nov. 2004.

[21] R. Szewczyk, E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin. Habitat monitoring with sensor networks. *Communications of the ACM*, 47(6):34–40, 2004.

[22] G. Tolle, J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong. A macroscope in the redwoods. *Sensys*, 2005.

[23] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh. Monitoring Volcanic Eruptions with a Wireless Sensor Network. *EWSN*, Jan. 2005.

[24] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi. Hardware Design Experiences in ZebraNet. *Sensys*, Nov. 2004.