IEICE TRANS. COMMUN., VOL.E93–B, NO.11 NOVEMBER 2010

INVITED PAPER *Special Section on Fundamental Issues on Deployment of Ubiquitous Sensor Networks*

Feedback Control-Based Energy Management for Ubiquitous Sensor Networks

Ting ZHU[†a)], **Ziguo ZHONG**[††], **Yu GU**[†††], **Tian HE**[††], ***and*** **Zhi-Li ZHANG**[††], ***Nonmembers***

SUMMARY Slow development in battery technology and rapid advances in ultra-capacitor design have motivated us to investigate the possibility of using capacitors as the sole energy storage for wireless sensor nodes to support ubiquitous computing. The starting point of this work is TwinStar, which uses ultra-capacitor as the only energy storage unit. To efficiently use the harvested energy, we design and implement feedback control techniques to match the activity of sensor nodes with the dynamic energy supply from environments. We conduct system evaluation by deploying sensor devices under three typical real-world settings — indoor, outdoor, and mobile backpack under a wide range of system settings. Results indicate our feedback control can effectively utilize energy and ensure system sustainability. Nodes running feedback control have longer operational time than the ones running non-feedback control.

key words: *wireless sensor networks, energy harvesting, ultra-capacitor, sustainable*

1. Introduction

With the increasing need for cyber-physical interaction, sensor networks have evolved into a key technology for long-term and ubiquitous applications such as target tracking [1], habitat monitoring [2], and scientific exploration [3]. Due to the low-cost and small size requirements, sensor nodes deployed in ubiquitous applications (e.g., MICA series) are normally equipped with limited power sources (e.g., 2200 mAh at 3 V for two AA batteries), and hazardous or inaccessible environments preclude manual battery replacement. Without renewable energy resources, a deployed sensor node can sustain only a few hours at 100% duty cycle.

Clearly, existing energy conservation techniques can only partially alleviate the conflict between system lifetime and performance. To ensure ubiquity and sustainability, environmental energy harvesting [4]–[6] has been pursued by the research community as the right solution for long-term and ubiquitous applications. To store harvested energy, researchers have proposed using rechargeable batteries [7], such as NiCAD, NiMH, or Li-ion, or a combination of capacitors and rechargeable batteries [8]. Although these solutions have been proven effective in prolonging the lifetime of the sensor nodes, there is still room for further improvement, because (i) rechargeable batteries have limited recharge cycles due to cyclic memory and crystalline formation (e.g., a Li-ion battery has 500 cycles, NiMH 300 cycles) and (ii) sophisticated recharging circuits and electrochemical conversion could reduce energy efficiency to as low as 6% [9].

Compared with rechargeable batteries, capacitors possess a set of advantages: they (i) have more than 1 million recharge cycles; (ii) have predicable remaining energy independent of discharge modes; (iii) are robust to temperature changes, shock, and vibration; and (iv) have high charging and discharging efficiency. Furthermore, recent advances in ultra-capacitor technology make it possible to use ultra-capacitors as the *only* energy storage device. For instances, a research group at MIT [10] has announced nanotube-based ultra-capacitors, which can provide energy storage densities comparable to those of batteries. In 2006, a U.S. patent [11] was issued for an electrical energy storage unit using an ultra-capacitor that has an energy/weight value of about 342 W·h/kg, twice that of Li-ion batteries. The largest capacitance currently available on the market is 3,000 F. Powered by this kind of capacitor, sensor nodes can work for more than 527 days under a 1% duty cycle with a single initial charge of the capacitor. It is highly possible that in the near future we will witness a paradigm shift from a battery-based to an ultra-capacitor-based design for all kinds of embedded devices. Therefore, it is essential to explore this frontier in advance.

Unlike previous research that focuses mainly on hardware design, the driving idea of this work is *sustainability* and *ubiquity* — a holistic approach towards energy equilibrium in WSN. A lightweight predictor is designed to estimate the life time of a sensor node based on the available environmental energy and the remaining energy inside the ultra-capacitor. A feedback control is then used to suggest an appropriate duty-cycle to the application layer, based on the gap between the predicted and targeted lifetimes. Instead of *conserving* as much energy as possible, the design objective of our work is to *consume* as much energy as possible while maintaining operational sustainability and ubiquity. Specifically, the major contributions of this work are as follows:

- To our knowledge, this is the first in-depth work to investigate the ultra-capacitor-based sustainable and ubiquitous sensor network design. The feedback control allows a node to precisely achieve an equilibrium

Manuscript received May 18, 2010.
Manuscript revised June 28, 2010.

[†]The author is with the Department of Computer Science at the University of Massachusetts Amherst, U.S.A.
[††]The authors are with the Department of Computer Science and Engineering at the University of Minnesota, U.S.A.
[†††]The author is at the Singapore University of Technology and Design, Singapore.
a) E-mail: tzhu@cs.umn.edu

DOI: 10.1587/transcom.E93.B.2846

between energy supply and demand.

- We design the battery-less energy efficient hardware platform and deploy our devices under diversified real-world environments including indoor, outdoor as well as mobile backpack.

The rest of the paper is organized as follows: Sect. 2 describes the motivation behind a battery-less design for sensor networks. Section 3 gives an overview of the design architecture. Hardware design, modeling, control, and adaption are presented in Sects. 4, 5, 6, and 7, respectively. System implementation and evaluation are detailed in Sect. 8. Related work is discussed in Sect. 9. Finally, Sect. 10 concludes the paper.

2. Motivation

This section describes the motivation for our work by revealing the uncertainty in battery modeling and identifying challenging issues in ultra-capacitor-based design. In the empirical study, we used different types of batteries.

An essential requirement for energy management is estimating remaining energy accurately. In battery-based designs, remaining energy would be estimated by establishing battery models. Researchers [12] have demonstrated the possibility of estimating remaining energy with an ±3% error, assuming a fixed discharge rate, temperature, and battery type. However, it would fail to apply to sensor network settings in which the workload is driven by unpredictable events/workload, and changing environmental factors (e.g., temperature) affects electrochemical reaction significantly. In contrast, the remaining energy inside an ultra-capacitor can be precisely estimated according to its voltage.

Impact of Workload: To investigate the impact of changing workloads on the remaining energy estimation in battery-based designs, we conducted a series of experiments using four types of workload, as shown in Table 1. These synthetic workloads are generated by choosing different sets of active components (LED only vs. all components) and workload (1% vs. 99%). By periodically measuring the voltages, we can get the battery discharging characteristics under these four workloads, as shown in Fig. 1. Figure 1 confirms that even for batteries of the same type, remaining energy differs with different types of workloads. For example, Fig. 1(c) shows that at the 2.5 V, two NiMH batteries can supply 17.8 kJ at the 1%-LED workload, but only 9.8 kJ at the 99%-All-On workload. Clearly, this big uncertainty would lead to ineffective energy control.

Impact of Battery Type: We also conducted a series of experiments to investigate the consistency of battery models over different battery types. In this study, three different types of the batteries are investigated: Li-ion, NiCAD and NiMH. Comparing Figs. 1(a), 1(b), and 1(c) indicates that different types of the batteries have different discharge characteristics and different capacities. In other words, there is no single battery model that can be applied to different battery types. For example, when a sensor node runs under the 1%-LED workload, remaining energy under 2.7 V is 3.8 kJ, 10.55 kJ, and 22.16 kJ for Li-ion, NiCAD, and NiMH batteries, respectively.

Figure 1 also shows that the discharge curves for different batteries exhibit deep slopes within a certain voltage range, in which a small error in voltage reading would be amplified into a large error in estimating remaining energy. For example, as shown in Fig. 1(b), an 0.1 V error would be translated into a 200% energy difference between 2.4 V and 2.5 V readings.

3. Overview of System Design

The empirical study in Sect. 2 indicates that uncertainty in

Table 1 Workload patterns.

Workload Pattern	Avg. Current (mA)	Peak Current (mA)
1% LED On	3.08	10.77
1% All On	3.11	30.05
99% LED On	10.75	10.75
99% All On	30.03	30.03

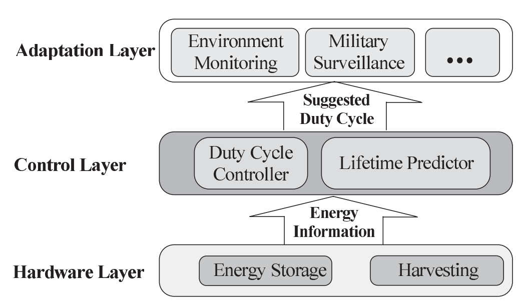

Fig. 2 Overview of system architecture.

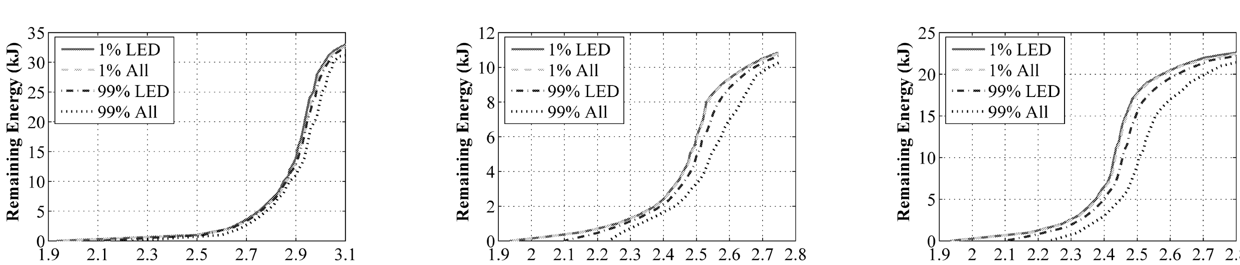

(a) Li-ion 3100 mAh (b) NiCAD 1100 mAh (c) NiMH 2500 mAh

Fig. 1 Battery discharge characteristics under four types of workload.

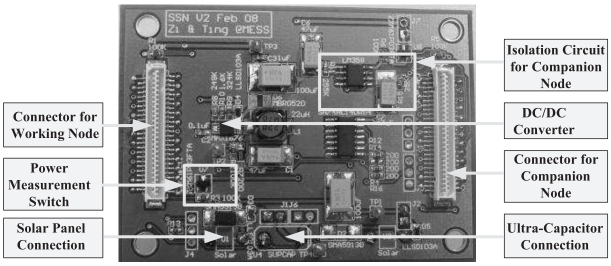

Fig. 3 TwinStar platform.

battery modeling is a critical hurdle that hinders designers exploiting the sustainability of battery-based design. To address the hurdle, we propose a three-layer ultra-capacitor-based design as shown in Fig. 2.

The hardware layer uses an energy harvesting circuit (e.g., solar panel or wind generator) to harvest the environmental energy and uses an ultra-capacitor as the *only* energy storage unit to power sensor nodes. The hardware platform is shown in Fig. 3. This layer also (i) monitors the remaining energy inside the ultra-capacitor and (ii) samples the energy harvesting rate from the energy harvesting circuit. These two pieces of information are then fed into the control layer.

The control layer predicts the lifetime of a node based on the energy harvesting and consumption rates. The difference between the predicted lifetime and user-specified lifetime is used as the control signal to the duty cycle controller, which *suggests* a certain percentage of duty cycle change to the application. The control layer deals with two conditions: energy oversupply and undersupply. In the case of oversupply, the controller suggests the running program to maximize the energy utilization. In the case of undersupply, the controller signals the running program to reduce the rate of energy consumption so as to ensure the minimal but critical activities of the sensor node.

In the adaptation layer, a running application changes its working schedules according to the suggested energy budget. This type of adjustment can improve the quality of applications during the energy-rich stage and maintain the minimum application quality during the energy deficit stage. For example, nodes can increase their sampling rate to provide higher accuracy when they have surplus energy. These nodes can also keep the minimum sampling rate when they are in energy deficit stage. We note that designs for the adaption layer are highly diverse. Instead of using a case-by-case approach, we propose a generic adaptation technique by treating duty-cycle as first-class resource that can be allocated, scheduled and adjusted in Sect. 7.

4. Hardware Design of TwinStar

TwinStar is an add-on power board that harvests energy from environments and uses an ultra-capacitor as the only energy storage to overcome the intrinsic limitations of batteries (e.g., energy uncertainty, limited recharge cycles, low conversion efficiency, and environment unfriendliness). It consists of (i) the solar panels and peripheral circuit for energy harvesting; (ii) the power measurement switch; (iii) the ultra-capacitor-based energy storage; and (iv) the smart power supply circuit with a DC/DC converter for powering the working node attached to the TwinStar board. The corresponding printed circuit board is shown in Fig. 3. Due to space constraints, we only explain a few unique features of the TwinStar design in the rest of the section.

4.1 Smart Power Supply Circuit

To accommodate fluctuating ambient energy, the voltage of a power supply shall be stabilized. We apply a high-efficiency switch-type DC/DC converter for providing a stable power supply for the working node. In the battery-based approaches, normally a DC/DC converter is powered directly by a battery, assuming enough energy is left in the battery [13]. However, it would be problematic for the ultra-capacitor-based design, especially under extremely low ambient energy situations due to the *zero-energy bootstrap* problem: Initially, the voltage of the ultra-capacitor is around 0 and all the system components do not work including the DC/DC converter. When the voltage of the ultra-capacitor reaches a threshold level, the converter enters a warming up stage, during which the high-efficiency switch-type converter needs to excite the external inductor into oscillations with high transient current (at the level of 10 mA). If the scavenged energy is not sufficient to support the converter to finish this stage, the voltage of the ultra-capacitor drops below the threshold level and the converter fails to boot up. Such failure repeats itself, wasting energy harvested from the environment.

To address this zero-energy boot-up challenge, a smart control circuit is designed to automatically boot-up and shut down the DC/DC converter. The basic idea is that during the charging state, the DC/DC converter is kept shutdown by the smart control circuit until the ultra-capacitor accumulates sufficient energy and reaches a voltage much higher than the DC/DC converter's input voltage threshold (1.1 V); during the discharging period, the DC/DC converter is turned off when the voltage of the ultra-capacitor approaches the minimum input voltage (0.7 V). In this way, high boot-up current stage and energy waste can be avoided.

However, achieving this kind of smart control needs to overcome a hidden challenge: there is no appropriate power supply for the smart control circuit itself before enabling the DC/DC converter. In our design, we address all the above challenges by proposing a novel dual solar panel solution (a small boot solar panel and a big main panel) together with a Schmitt trigger-based control circuit. The design is illustrated in Fig. 4.

As shown in Fig. 4, a small boot solar panel is applied to power the Schmitt trigger module that is used to control the on/off status of the DC/DC converter. The small panel charges a small capacitor (C1, which is 47 μF) and keeps the voltage of C1 higher than 2.5 V so as to power the Schmitt trigger, which is an extremely low power device with power consumption less than 10 μW. Therefore, even with a low

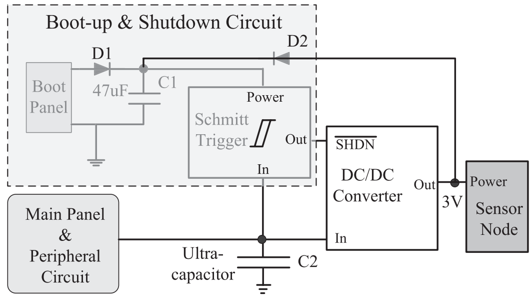

Fig. 4 Boot-up and shutdown control circuit.

environmental energy supply, the smart control circuit can still work.

The main panel is the one harvesting energy for the ultra-capacitor. Note that when the DC/DC converter starts to work, its output also feeds back to power the Schmitt trigger, avoiding shutting down the DC/DC converter when there is no environmental energy but the ultra-capacitor is still energy-rich.

5. System Energy Modeling

System modeling is an essential step for effective control. In this section, we introduce three models: the energy harvesting model (Sect. 5.1) and the energy consumption model (Sect. 5.2). To achieve effective energy management, we need to balance the interaction among these models into an equilibrium state.

5.1 Energy Harvesting Model

The energy harvesting model is built online using the measurement switch. The total energy E_E harvested during the time interval $[(k-1)T, kT]$ is:

$$\Delta E_E[kT] = \int_{(k-1)T}^{kT} I_{cap}(t) \cdot V_{cap}(t)dt \quad (1)$$

Here $V_{cap}(t)$ is the ultra-capacitor's voltage at time t. Both $I_{cap}(t)$ and $V_{cap}(t)$ can be measured by our ultra-low-power measurement switch. T is called the sampling *period* and k is called the *sampling instant*.

5.2 Energy Consumption Model

The energy consumption of the sensor node is determined by three parameters: average active mode current I_{active}, average sleep mode current I_{sleep} and duty cycle D (the percentage of active time). The energy consumption during the time interval $[(k-1)T, kT]$ can be represented as:

$$\Delta E_C[kT] = V_{node} \cdot (I_{active} \cdot D + I_{sleep} \cdot (1-D)) \cdot T \quad (2)$$

The average power consumption at the end of time kT can be represented as follows:

$$P_{C_{AVG}}(kT) = \frac{\Delta E_C[kT]}{T} \quad (3)$$

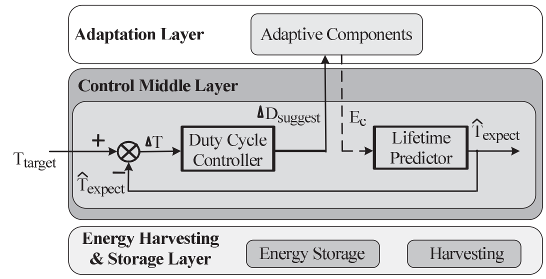

Fig. 5 Control layer design overview.

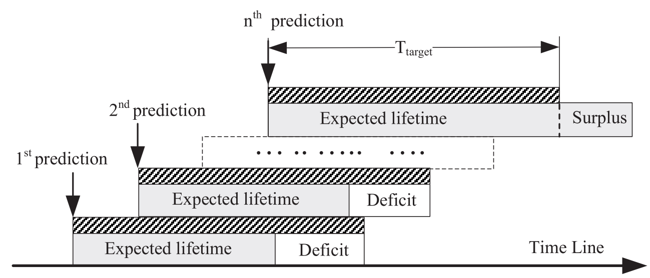

Fig. 6 Sustainability.

6. Control Layer Design

The design objective of the control layer is to *consume as much energy as possible while maintaining sustainability*. Sustainability T_{target} is defined as the duration a node must survive without ambient energy. T_{target} is a configurable parameter, set by users according to the environment in which sensors are deployed. For example, in energy-rich environments such as deserts, users can set a smaller T_{target} to consume energy aggressively. On the other hand, in energy-poor and unpredictable environments, a larger T_{target} is desired to ensure the aliveness of sensor nodes. Once T_{target} is chosen, it is used as the set point in the feedback control based design as shown in Fig. 5.

As shown in Fig. 6, to maintain sustainability, the lifetime predictor periodically predicts the expected lifetime $\widehat{T}_{expect}$ and compares $\widehat{T}_{expect}$ with the target lifetime T_{target}. In the case of energy deficit (i.e., $\widehat{T}_{expect} < T_{target}$), a node is expected to run out of energy after $\widehat{T}_{expect}$ seconds, if it maintains current level of activity. In the case of energy surplus (i.e., $\widehat{T}_{expect} > T_{target}$), a node can increase the activity accordingly for application performance improvement. As shown in Fig. 5, the change of activity is achieved by using the difference between $\widehat{T}_{expect}$ and T_{target} as the input to the duty cycle controller, which suggests a certain percentage of duty cycle change to the adaptation layer. Unlike conventional control designs, the duty cycle is *suggested* by the control layer to the adaptation layer. The adaptation layer adjusts the duty cycle based on the application's requirement. This allows a more flexible design of the energy adaptation layer. For example, a node should not stop sending critical control messages disruptively simply because the control layer suggests to reduce duty cycle due to the brief drop in energy supply. Short-term mismatching is tolerable because the ultra-capacitor can serve as a buffer.

The control layer contains two major components: the lifetime predictor and the duty cycle controller, as shown in Fig. 5. These are described in the following subsections.

6.1 Design of the Lifetime Predictor

To ensure that sensor nodes can survive dark periods during which no ambient energy is available (e.g., night), it is essential to predict lifetime conservatively based on (i) the remaining energy, and (ii) the energy consumption rate. More specifically, the prediction goes through the following steps:

Step 1: Measure Remaining Energy

At the end of every T seconds, the control layer periodically monitors the remaining voltage $V_{cap}(t)$ of the ultra-capacitor. Therefore, at the end of kT, the remaining energy $E_R(kT)$ can be obtained by using the following equation:

$$E_R(kT) = \frac{1}{2}CV_{cap}^2 \tag{4}$$

Step 2: Calculate the Power Consumption during Time Interval $[(k-1)T, kT]$

Since the sensor node is powered by the ultra-capacitor, the energy consumed by the sensor node during $[(k-1)T, kT]$ equals to the environmental energy gained during $[(k-1)T, kT]$ plus the difference of the remaining energy at time $(k-1)T$ and kT. It can be calculated as follows:

$$\Delta E_C[kT] = \Delta E_E[kT] + E_R((k-1)T) - E_R(kT) \tag{5}$$

The average power consumption during the time interval $[(k-1)T, kT]$ can be calculated as follows:

$$P_{C_{AVG}}[kT] = \frac{\Delta E_C[kT]}{T} \tag{6}$$

Step 3: Predict the Expected Lifetime $\widehat{T}_{expect}$

We assume that energy consumption in the past is an approximation of the future, therefore the expected lifetime can be estimated as follows:

$$\widehat{T}_{expect} = \frac{E_R(kT) - E_{min}}{P_{C_{AVG}}[kT]} \tag{7}$$

Where E_{min} is the minimum energy needed to sustain a node's basic operations.

6.2 Design of the Duty-Cycle Controller

The energy model in Sect. 5.1 is an approximation of reality with random noises in measurements and uncertainties in future energy consumption rates, therefore we should treat energy management in sustainable sensor networks as a dynamic process and use control theory to converge and stabilize the performance of such a system to a desired set point. In this section, we describe the design of the controller in the system. The controller indicates the suggested change of duty cycle $\Delta D_{suggest}$ to the adaptation layer based on the remaining energy inside the ultra-capacitor and the time difference $T_{diff} = \widehat{T}_{expect} - T_{target}$. Here we use a P-controller, considering the fact that the bigger the difference T_{diff}, the larger ΔD is needed to adjust.

The major issue for the P-controller design is to choose an appropriate control gain P, so that the sensor node can increase the amount of energy consumed and still meet the lifetime goal of T_{target}. Intuitively, if the sensor node increases its duty cycle at the ultra-capacitor's high voltage stage, it can utilize more energy before the next available environmental energy comes. To illustrate the design, consider two scenarios:

Energy Surplus: In the case of an energy surplus, $\widehat{T}_{expect}$ is larger than T_{target}, we get a positive T_{diff} value. To converge $\widehat{T}_{expect}$ to the set point T_{target}, a positive duty-cycle change that can increase the power consumption $P_C(t)$ needs to be suggested to the adaption layer. Clearly, it is desirable to increase more duty-cycle when the capacitor's voltage value is high. Therefore, in the case of an energy surplus, we used the following adaptive P-controller:

$$\Delta D_{suggest} = G_+ \cdot P(V_{cap}) \cdot T_{diff} \tag{8}$$

where G_+ is a coefficient for control gain adjustment in the case of an energy surplus. $P(V_{cap})$ is a function of V_{cap}.

Energy Deficit: In the case of an energy deficit, $\widehat{T}_{expect}$ is smaller than T_{target}, and we get a negative T_{diff} value. To converge $\widehat{T}_{expect}$ to the set point T_{target}, we shall reduce the power consumption $P_C(t)$ by suggesting a negative duty cycle change to the adaption layer. It is desirable to keep a high duty cycle when the capacitor's voltage is high. Therefore, in the case of an energy deficit, we would like to reduce the duty cycle slowly in the ultra-capacitor's high voltage stage. The following adaptive P-controller is used:

$$\Delta D_{suggest} = \frac{G_-}{P(V_{cap})} \cdot T_{diff} \tag{9}$$

where G_- is a coefficient for control gain adjustment in the case of an energy deficit.

7. Adaptation Layer Design

The controller layer provides suggestive duty cycle as a feedback to the adaptation layer where the activity of application-level components can be adjusted. In our design, we opt to use duty-cycle-based adaption in communication, sensing, and actuation modules. The high-level idea is to design system components with a tunable duty cycle parameter for each type of activity. This parameter is enforced through a mechanism similar to a token-bucket in network traffic control — i.e., each component is assigned a certain number of *active instances* per period of time. Specifically, let the active instance of a sensor node be a tuple (t, d), where t is the start time of the active instance and d is the corresponding duration of this active instance, the working schedule of a node can be represented as $\Gamma = \{(t_1, d_1), (t_2, d_2), ...\}$. Essentially, Γ stores all active instances for a node. For most of the applications, the working schedules of sensor nodes are periodic. For example, a periodic working schedule $\{(1, 3), (101, 3), (201, 3), ...\}$ represents the sensor node is active for 3 active durations every 100 units of time, The *duty cycle* of a sensor node, therefore, is the ratio between the total number of active durations within a period and the time duration of the period. In the above example, the duty cycle is 3%. Duty-cycle-based adaption is generic enough to be applied in many protocols.

With the suggested duty-cycle from the control layer, a node can do either component-level or node-level adaptation [14]. *Component-level adaptation* is achieved by (i) increasing and reducing the number of active instances and/or (ii) extending or shrinking the duration. *Node-level adaptation* is done by adjusting duty-cycle parameters either proportionally or based on priority.

8. Evaluation

In this section, we evaluate system performance of a single TwinStar node under different types of environmental energy patterns. We compare our system working in Feedback Control (FC) mode with the same system but in a Non-Feedback Control (NFC) mode. The system in the non-feedback control mode works under constant duty cycle for both the energy surplus and energy deficit scenarios.

8.1 Evaluation Metrics and Baseline

The key advantage of the feed-back control design is efficiently using energy before the available environmental energy comes. We use the following metric to evaluate the performance of our system:

Availability (A): Availability is the ratio of the systems operational time to the entire experiment time. A system is operational when the voltage of the capacitor is above the DC/DC converter's minimum voltage (i.e., 0.7 V for our system). Under a given energy harvesting pattern, the larger the *A* value, the longer the system is operational. In other words, a larger *A* value indicates the system is more reliable.

8.2 Different Deployment Scenarios

We collected the solar charging data and simulated our system behavior under three different scenarios: outdoor, indoor, and mobile backpack, as shown in Figs. 7, 8 and 9, respectively. These scenarios are carefully selected to represent a wide range of energy harvesting patterns: (i) periodical and vibrated energy for the outdoor environment, (ii) periodical and constant energy for the indoor environment, and (iii) dynamic and highly unpredictable energy in the mobile environment.

8.3 Outdoor Experiment

The outdoor scenario represents such potential applications as environment monitoring [2] and scientific exploration [3]. In the outdoor experiment, we deployed solar panels outside of a fifth-floor apartment, as shown in Fig. 7(a). The total time duration of this experiment was 76 hours. Figure 7(b) shows the energy harvested.

Outdoor

(a) Experimental Site

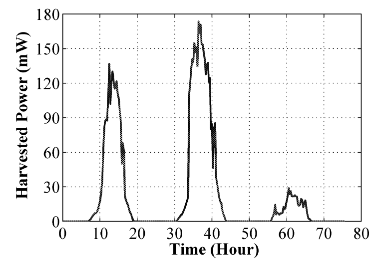

(b) Harvested Power

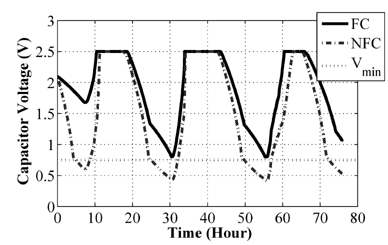

(c) Capacitor Voltage

Fig. 7 Outdoor.

Figure 7(c) compares the capacitor's voltage among the systems with feedback control and non-feedback control. Clearly, the feedback control system allows nodes to consume as much energy as possible while still maintains the sustainability of the system. As a result, during the whole experiment period, the capacitor's voltage for the feedback control system was always above the minimum voltage to maintain liveness, which is 0.7 V. However, the non-feedback control system stopped working for 4 times. It needs to wait for the available environmental energy to charge the capacitor to boot it up again. As shown in Fig. 10, the feedback control system has the highest availability value 100%, while the non-feedback control system only has the availability value 61.2%.

8.4 Indoor Experiment

The indoor scenario represents such potential applications as facility management, structure monitoring. As shown in

Indoor

(a) Experimental Site

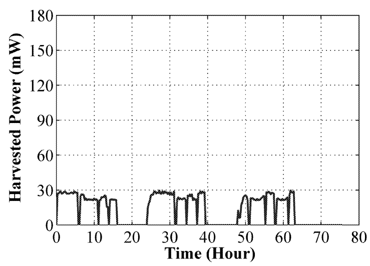

(b) Harvested Power

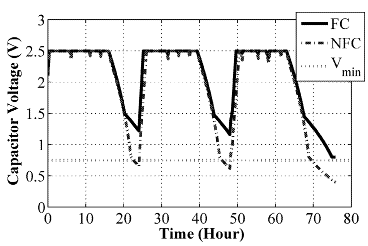

(c) Capacitor Voltage

Fig. 8 Indoor.

Mobile Backpack

(a) Experimental Site

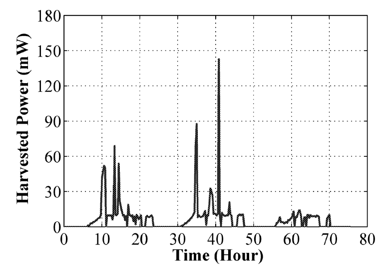

(b) Harvested Power

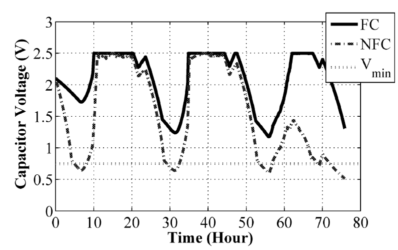

(c) Capacitor Voltage

Fig. 9 Mobile backpack.

Fig. 8(a), our system was deployed under the overhead light in our lab. The total time duration for this experiment was 72 hours. The light was turned on in the morning, when people came to the lab, and turned off in the middle of the day or during the night when no one was inside the lab.

Figure 8(b) shows the energy harvested by the system over the 72-hour period. The fluctuation of the energy level was due to the turned on or off of the neighboring overhead lights. Compare with the outdoor experiment, the voltage curves of feedback control system and non-feedback control system are very close to each other. This indicates that the non-feedback control system works better in the environment with relatively constant available energy. Although the non-feedback control system still run out of energy, the total time that the system was not operational is smaller than in the outdoor experiment. As a result, the availability for the non-feedback control system is 82.3%, while the feedback control system's availability is still 100% (shown in Fig. 10).

8.5 Mobile Backpack Experiment

We also deployed our system on a backpack, as shown in Fig. 9(a). It was carried by a graduate student every day for 3 days. During the night and in the early morning, the backpack was placed near a living room window to harvest environmental energy through the window. During the day, the graduate student carried it to attend outdoor and indoor activities. This scenario represents such potential applications as assisted living and human-centric sensing [15]. The total time duration of this experiment was 76 hours.

Figure 9(b) shows the energy harvested by the system over 76 hours. Compared with the outdoor and indoor cases, the energy harvested in the backpack experiments was very bursty, with the high peak corresponding to outdoor activity and the flat part corresponding to indoor activity. Figure 9(c) shows again that the feedback control system maintain the aliveness of the system all the time, while the non-feedback control system run out of energy for 4 times. The availability of the feedback control system is 100% in this scenario. It is 37% more than the availability of the non-feedback control system (shown in Fig. 10).

8.6 Summary

We have evaluated our system under different environmental energy patterns. In all these scenarios, our feedback control

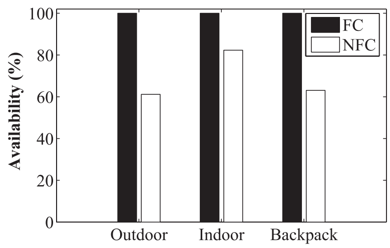

Fig. 10 Availability for 3 scenarios.

system outperforms the non-feedback control system. The key observations are: (i) the system working under feedback control never runs out of power (note: minimum voltage to maintain liveness is 0.7 V); and (ii) the feedback control design is more responsive to the influx of energy.

9. State of the Art

Energy harvesting is the conversion of ambient energy into usable electrical energy. Several technologies have been developed for extracting energy from the environment, including solar [8], wind [16], and kinetic energy. With these energy-harvesting technologies, researchers have designed various types of platforms to collect ambient energy from human activity or environments [8], [16]. Notable ones designed specially for sensor networks are Heliomote [7], Prometheus [8], and AmbiMax [16]. According to the type of energy storage used, these platforms can be separated into two categories: (i) rechargeable battery-based platforms [7], and (ii) designs combining ultra-capacitors and rechargeable batteries [8].

• In rechargeable battery designs, such as Heliomote [7], the energy harvesting panel is directly connected to its battery. As the primary energy storage device, the rechargeable battery is charged and discharged frequently, leading to low system lifetimes due to the physical limitation on the number of recharge cycles.

• In designs that combine ultra-capacitors and rechargeable batteries such as Prometheus [8], the solar energy is first stored in the primary energy buffer, which is one or multiple ultra-capacitors. The rechargeable batteries are then used as the secondary energy buffer. This design inherits both the advantages and limitations of batteries and capacitors. For examples, it is difficult to predict remaining energy because of the inclusion of batteries, and the lifetime of energy storage subsystem is decided by the shelf time of batteries (in the order of a few years).

Energy transmission technology such as WiTricity provides potential solution for resolving energy problem in sensor network. Different from previous works, this work investigates the frontiers of capacitor-only energy storage design and studies not only hardware designs but also related software control techniques.

Energy conservation is an intensively studied area. Many solutions have been proposed at different layers, including high-efficiency hardware design [17], [18], network layer design [19]–[21], and operating system [22]. However, only a few works have focused on energy adaptation [13], [23]. The most closely related work for energy adaption is Eon [13], which is the first programming language for energy-adaptive applications in wireless sensor networks. Using Eon, programmers can easily annotate flows in energy states. Eon's automatic energy management uses these annotations to change application behavior adaptively. Since Eon focuses on language and the adaption layer, it is highly complementary to the hardware and controller design in this work.

10. Conclusions

Slow development in battery technology and rapid advances in ultra-capacitor design have motivated us to investigate the possibility of using capacitors as the sole energy storage for ubiquitous wireless sensor networks. In this work, we build TwinStar, an add-on energy efficient hardware platform that uses energy harvesting circuits (e.g., solar panels or wind generators) to harvest the energy from the environment and employs an ultra-capacitor as the only energy storage unit to power the sensor node.

To promote the capacitor-only-design and ensure long-term sustainability and ubiquity, we propose a feedback-based approach that suggests an appropriate duty-cycle change to the adaption layer, based on the gap between the predicted and targeted lifetimes. The controller increases activity when the capacitor's voltage is high. This type of control allows us to utilize the energy before the next available environmental energy comes. We invested significant amount effort to evaluate our design by deploying the devices in multiple real-world settings: (i) indoor, (ii) outdoor, and (iii) mobile backpack. The results indicate the effectiveness of the feedback control design compared with a design that with non-feedback control.

Acknowledgements

This work was supported in part by NSF grants CNS-0917097 and CNS-0845994. It was also supported by NSF under grant #1019343 to CRA for the CI Fellows Project. We thank Prof. Yoshito Tobe and the reviewers for their valuable comments and support.

References

[1] Z. Zhong, T. Zhu, D. Wang, and T. He, "Tracking with unreliable node sequence," INFOCOM'09, 2009.
[2] R. Szewczyk, A. Mainwaring, J. Anderson, and D. Culler, "An analysis of a large scale habit monitoring application," SenSys'04, 2004.
[3] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh, "Fidelity and yield in a volcano monitoring sensor network," OSDI'06, 2006.
[4] S. Meninger, J.O. Mur-Miranda, R. Amirtharajah, A. Chandrakasan, and J. Lang, "Vibration-to-electric energy conversion," ISLPED'99, 1999.
[5] T. Zhu, Z. Zhong, Y. Gu, T. He, and Z.L. Zhang, "Leakage-aware energy synchronization for wireless sensor networks," MobiSys'09,

2009.
[6] Z. Zhong, T. Zhu, T. He, and Z.-L. Zhang, "Leakage-aware energy synchronization on twin-star nodes," SenSys'08, 2008.
[7] A. Kansal, J. Hsu, S. Zahedi, and M.B. Srivastava, "Power management in energy harvesting sensor networks," TECS, vol.6, no.4, 2007.
[8] X. Jiang, J. Polastre, and D. Culler, "Perpetual environmentally powered sensor networks," IPSN'05, 2005.
[9] Suzanne Foster, "The energy efficiency of common household battery charging systems: Results and implications," Natural Resources Defense Council. http://www.efficientproducts.org/reports/bchargers/NRDC-Ecos_Battery_Charger_Efficiency.pdf
[10] R. Signorelli, J. Schindall, and J. Kassakian, "Carbon nanotube enhanced double layer capacitor," 15th International Seminar on Double Layer Capacitors and Hybrid Energy Storage Devices, 2005.
[11] Weir, et al., "Electrical-energy-storage unit (EESU) utilizing ceramic and integrated-circuit technologies for replacement of electrochemical batteries," United States Patent 7033406, 2006.
[12] R. Rao, S. Vrudhula, and D.N. Rakhmatov, "Battery modeling for energy-aware system design," Computer, vol.36, no.12, pp.77–87, 2003.
[13] J. Sorber, A. Kostadinov, M. Garber, M. Brennan, et al., "Eon: A language and runtime system for perpetual systems," SenSys'07, 2007.
[14] Y. Gu, T. Zhu, and T. He, "ESC: Energy synchronized communication in sustainable sensor networks," ICNP'09, 2009.
[15] S.B. Eisenman, E. Miluzzo, N.D. Lane, R.A. Peterson, G.S. Ahn, and A.T. Campbell, "The bikenet mobile sensing system for cyclist experience mapping," SenSys'07, 2007.
[16] C. Park and P.H. Chou, "AmbiMax: Autonomous energy harvesting platform for multi-supply wireless sensor nodes," SECON'06, 2006.
[17] D. Lymberopoulos and A. Savvides, "XYZ: A motion-enabled, power aware sensor node platform for distributed sensor network applications," IPSN'05, 2005.
[18] Y. Sun, T. Zhu, Z. Zhong, and T. He, "Energy profiling for mPlatform," SenSys'09, 2009.
[19] T. Zhu, Z. Zhong, T. He, and Z.L. Zhang, "Exploring link correlation for efficient flooding in wireless sensor networks," NSDI'10, 2010.
[20] T. Zhu and M. Yu, "A secure quality of service routing protocol for wireless ad hoc networks," GLOBECOM'06, 2006.
[21] T. Zhu and M. Yu, "A dynamic secure QoS routing protocol for wireless ad hoc networks," Sarnoff'06, 2006.
[22] C.C. Han, R. Kumar, R. Shea, E. Kohler, and M. Srivastava, "A dynamic operating system for sensor nodes," MobiSys'05, 2005.
[23] J. Flinn and M. Satyanarayanan, "Managing battery lifetime with energy-aware adaptation," ACM Trans. Comput. Syst., vol.22, no.2, pp.137–179, 2004.

Ting Zhu is currently a CI Fellow in the Department of Computer Science at the University of Massachusetts Amherst. He received his Ph.D. degree in computer science from the University of Minnesota. He received his B.S. and M.E. degrees from the Department of Information Science and Electronic Engineering at the Zhejiang University in 2001 and 2004, respectively. His research interests include wireless networks, renewable energy, embedded systems, distributed systems, and security. He is a member of the IEEE, ACM, and USENIX.

Ziguo Zhong received the B.E. in radio engineering, the M.E. in signal processing, both from Southeast University, Nanjing, China, in 2003 and 2005, respectively, the M.S. degree in computer science from University of Minnesota, Minneapolis, MN, in 2008. He is currently pursuing the Ph.D. degree at the University of Minnesota. His research interests include wireless sensor networks, computer networking and signal processing.

Yu Gu is currently an assistant professor at the Singapore University of Technology and Design. His research includes Networked Embedded Systems, Wireless Sensor Networks, Cyber-Physical Systems, Wireless Networking, Real-time and Embedded Systems, Distributed Systems, Vehicular Ad-Hoc Network and Stream Computing Systems. Yu Gu is a member of ACM, IEEE and SIAM.

Tian He is currently an assistant professor in the Department of Computer Science and Engineering at the University of Minnesota-Twin City. He received the Ph.D. degree under Professor John A. Stankovic from the University of Virginia, Virginia in 2004. His research includes wireless sensor networks, intelligent transportation systems, real-time embedded systems and distributed systems, supported by National Science Foundation and other agencies. Dr. He is a member of ACM and IEEE.

Zhi-Li Zhang received the B.S. degree in computer science from Nanjing University, China, and his M.S. and Ph.D. degrees in computer science from the University of Massachusetts. In 1997 he joined the Computer Science and Engineering faculty at the University of Minnesota, where he is currently a Full Professor. He has held visiting positions at Sprint Advanced Technology Labs; IBM T.J. Watson Research Center; Fujitsu Labs of America, Microsoft Research China, and INRIA, Sophia-Antipolis, France.