

Optimizing Battery Lifetime-Fidelity Tradeoffs in BSNs using Personal Activity Profiles

Jeff S. Brantley¹, Adam T. Barth², John Lach¹

¹Charles L. Brown Department of Electrical and Computer Engineering
University of Virginia, Charlottesville, VA, USA 22903-3051

²Wireless Health Interactive, LLC, Vienna, VA, USA 22182
{jb7fx,atb4c,jlach}@virginia.edu

ABSTRACT

Body sensor networks (BSNs) often operate in dynamic environments, with the collected data profiles—and the resulting importance of data—varying throughout the system runtime. Therefore, the potential for power reduction fluctuates with changing user behavior, creating a dynamic battery lifetime-fidelity relationship that is subject to variations throughout the battery lifetime corresponding to an individual’s daily activities—past, present, and future.

This paper explores the potential for optimizing the tradeoff between meeting a desired battery lifetime and maximizing system fidelity through run-time adaptation of sensor acquisition (duty cycling) and profile-based predictions of an individual’s future activities. A “personal activity profile” describes the expected behavior of an individual and is used to inform the desired battery discharge characteristics over time. Using walking activity traces collected from three human subjects wearing Fitbit[®] trackers over several months in order to develop such activity profiles, the approach is demonstrated in simulation based on an analytical power model for an inertial BSN platform incorporating recent sensors. Results show improvements over statically setting a duty cycle for constant power consumption with respect to ideally setting the duty cycle based upon *a priori* knowledge of activities of interest throughout the system lifetime.

Categories and Subject Descriptors

J.3 [Life and Medical Sciences]: Health, Medical information systems

General Terms

Reliability

Keywords

body sensor networks, battery lifetime

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

BodyNets '12 Oslo, Norway

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

1. INTRODUCTION

Optimizing the tradeoff between battery lifetime and system fidelity is central to realizing the potential of body sensor networks (BSNs). One central challenge to this tradeoff is that, for many applications, energy consumption and data quality depend on the behavior and activities of the wearer. For example, given some available control setting (such as on-node data reduction through lossy compression, sampling rate adjustment, or wholesale duty cycling), the tradeoff between the fidelity level and remaining battery lifetime depends on how often the activities of interest will be performed. If the projection is too high, the fidelity level will be set unnecessarily low, leaving additional battery life on the table come re-charge. Conversely, if the projection is too low, the fidelity level will be set too high, expending the battery before the projected re-charge, leaving activities of interest entirely uncaptured. In order to properly optimize such a battery lifetime-fidelity tradeoff, it is necessary to predict and adapt to future dynamics over the course of operation, informed by past observations.

This work explores the potential of such an approach in the context of a gait monitoring application scenario (using 6 degrees-of-freedom motion capture), leveraging personal daily activity profiles and feedback to improve the battery lifetime-fidelity tradeoff. To illustrate, the focus is placed on the variability of the amount of data of interest—that is, the amount of time spent walking. During these periods of interest, the node selects a short-term power-fidelity tradeoff by setting a duty cycle of operation (sensor acquisition and radio transmission), but otherwise the node stays in a low-power mode during periods of non-interest. The goal is to capture the data at the highest allowable duty cycle while satisfying a battery life requirement, or, more generally, to give “equal opportunity” or equal fidelity to all data of interest within the specified monitoring period.

Daily walking traces were collected on three subjects for 133, 126, and 68 days, respectively, using Fitbit[®] trackers. Simulations were performed based on an analytical power model for an inertial motion capture system using recent components, comparing the proposed approach against statically setting a duty cycle for constant power consumption and ideally setting the duty cycle based upon *a priori* knowledge of the amount of walking remaining before re-charge.

The remainder of this paper is organized as follows. Section 2 provides background about the application scenario and related work. Section 3 describes the approach and methods for profiling a user’s daily activity to estimate expected activity over the battery lifetime and adjust acquisi-

tion duty cycling to meet the desired requirements. Section 4 describes the methodology used for evaluation of the proposed technique and the experimental setup. Section 5 describes the results of the proposed scheme compared to static settings and the ideal (“oracle”) case. Section 6 explains the conclusions and ideas for future research.

2. BACKGROUND

As described above, this work targets an example application involving continuous, longitudinal gait monitoring. Sensor acquisition effort is focused on periods of walking or, more generally, non-sedentary periods, which are of primary interest for gait analysis and activity monitoring. This example is motivated in part by ongoing research investigating the use of activity level and gait analysis from 6 degrees-of-freedom inertial BSNs to assess early warning signs for fall risk in the elderly in a retirement facility or homecare setting with the goal of intervening *before* fall events occur. It is important to capture high-precision inertial data continuously for an individual, including throughout the night (nighttime falls are common), to get accurate fall-risk estimates. The BSN nodes may be swapped periodically—daily for older platforms, but conceivably weekly for newer, more energy-efficient platforms—by nursing staff during one of the organized meals or a home visit, thus requiring a battery lifetime to cover the period between swapping or recharging.

The problem of optimizing fidelity over a battery life has become of interest recently, with respect to mobile phones as well as BSNs, both of which have limited energy capacity and are subject to variable energy demand. In particular, related problems have been studied under assumption of Markovian user state transitions using a Markov Decision Process framework. On mobile phones, such approaches were used in conjunction with delaying update of state knowledge based on the probability of a state change [8] or choosing when to synchronize e-mail [5]. Others have applied the approach for health monitoring scenarios accounting for measures of health/criticality [7] or available energy for harvesting [6]. These heavyweight methods require extended offline computation and, moreover, may require more RAM (e.g., 128KB [7]) for storing decision tables than a typical low-power microcontroller, such as the TI MSP430F1611, is likely to have. In contrast, this paper presents a simpler, more straightforward approach for the purpose of exploring the potential gains of leveraging activity profiles for improved fidelity-battery lifetime performance.

3. APPROACH

As described above, this approach centers around a scenario in which the BSN node is only operated during times of detected activity (i.e., the wearer is walking), and otherwise the node is in a low-power mode monitoring for walking to begin again. This means that the achievable balance of lifetime and fidelity is dependent upon the amount of walking occurring in the deployment period, which is not known *a priori*. Thus, we present an analysis of the ideal operation (if one *did* have *a priori* knowledge), followed by an explanation of the proposed technique, whose rationale is informed by the ideal operation. An example application-agnostic fidelity metric is presented for use within this work, although more appropriate metrics better informed by a specific application should be substituted.

3.1 Ideal Operation

Given the above-stated goal of capturing data from all times of the day at an equal, maximal duty cycle, the ideal operation, given *a priori* knowledge, would be as follows. The starting battery energy, E_{batt} , should be spread evenly over the total amount of time, W spent walking during the day. That is, the preferred average power during periods of walking activity is

$$P_{avg} = \min\left(\frac{E_{batt}}{W}, P_{active}\right), \quad W \neq 0 \quad (1)$$

where P_{active} is the typical power when no duty-cycling is applied. The corresponding ideal setting of the duty cycle, d_{ideal} , is then simply

$$d_{ideal} = \frac{P_{avg}}{P_{active}} \quad (2)$$

That is, ideally, a fixed duty cycle would be chosen for the entire day, given knowledge of E_{batt} and W . However, since W is not known at the start of the day, any static setting will tend to be suboptimal, either exhausting the battery early or using a suboptimal duty cycle. Thus, we seek a more dynamic approach that makes adjustments in response to the actual amount of activity observed throughout the day.

3.2 Activity Profiling

In order to adapt to the wearer’s true amount of activity throughout the day, it is necessary to develop predictions, for different times throughout the day, of the amount of activity still to occur. Thus, we develop the notion of a *walking profile* as follows. Dividing the day into time steps $k = 1, 2, \dots, N$, let W_k denote the total amount of active time remaining at time step k , and let w_k be the amount of active time within a single time step k . W_k and w_k are related by the following:

$$w_k = W_k - W_{k+1} \quad (3)$$

and

$$W_k = \sum_{m=k}^N w_m \quad (4)$$

The activity profile, then, is a vector $\mathbf{W} = (W_1, W_2, \dots, W_N)^T$, which is simply a sequence of predictions of the total remaining active time at each time stage.

Figure 1 shows example profiles for 5 days. The profile begins at midnight at a high number, indicating the walking-to-go amount, in seconds, and remains constant until the wearer engages in the activities of the day. It then decreases steadily until the end of the day. On some days, for example, the subject may take a noticeably long walk around noon, causing that day’s initial value to start high, but nearer the end of the day, the profiles converge closer to one another, approaching 0 s of walking left at the end of the day.

The goal of activity profiling is to develop an *estimate profile*, $\hat{\mathbf{W}}$, based on actual profiles observed for previous days. The approach for training such an estimate profile is discussed later in Sec. 3.4. First, the usage of this profile during deployment is examined.

3.3 Profile-Informed Feedback

Once an estimate profile $\hat{\mathbf{W}}$ is developed for the day under test, it is used to employ a feedback algorithm. At each time

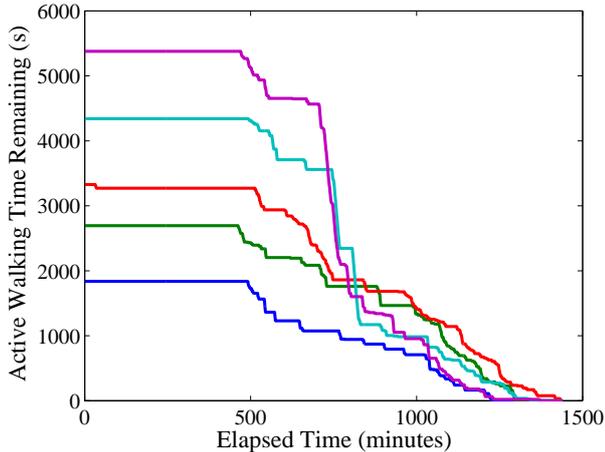


Figure 1: Example W profiles from five sample days

step k , the node calculates its desired average power, P_k in a similar fashion as Equation 1 but using the estimate \hat{W}_k instead of W_k , which is not known *a priori*:

$$P_k = \begin{cases} \min(E_k/\hat{W}_k, P_{active}), & \hat{W}_k \neq 0 \\ P_{max}, & \hat{W}_k = 0 \end{cases} \quad (5)$$

which gives a corresponding duty cycle

$$d_k = \frac{P_k}{P_{active}} \quad (6)$$

Applying the chosen P_k for the k th time step depletes the energy in proportion to the amount of walking that occurs, w_k , leading to the following recursion for battery energy remaining at time k :

$$E_{k+1} = E_k - P_k w_k \quad (7)$$

Since the actual amount of active time, w_k , is random, the remaining energy at time $k+1$ is itself a random quantity, being a function of the previous w_1, w_2, \dots, w_k . Thus, the recalculation of desired power P_{k+1} for time step $k+1$ —based on E_{k+1} and prediction \hat{W}_{k+1} —constitutes a feedback loop by which the node adjusts to the actual behavior (random disturbance) of the wearer over the course of a day.

In deploying this method, the BSN node must be aware of its current remaining energy, E_k , in order to calculate the proper duty cycle setting. Practically speaking, one can either assume the power P_k is deterministic or, for greater accuracy, track the energy consumption through the use of coulomb counting [4] or current sensing [2].

3.4 Profile Training

It is desirable to develop an estimate profile $\hat{\mathbf{W}}$ which optimizes some expected cost or utility achieved by the system. Let $g(\mathbf{d}, \mathbf{w})$ denote such a utility—or rather, fidelity—function, where $\mathbf{d} = (d_1, \dots, d_N)^T$ and $\mathbf{w} = (w_1, \dots, w_N)^T$. Given the initial battery charge (E_1) constraint, we have a constrained optimization problem:

$$\max g(\mathbf{d}, \mathbf{w}) \quad s.t. \quad \sum_{m=1}^N d_m w_m P_{max} \leq E_1 \quad (8)$$

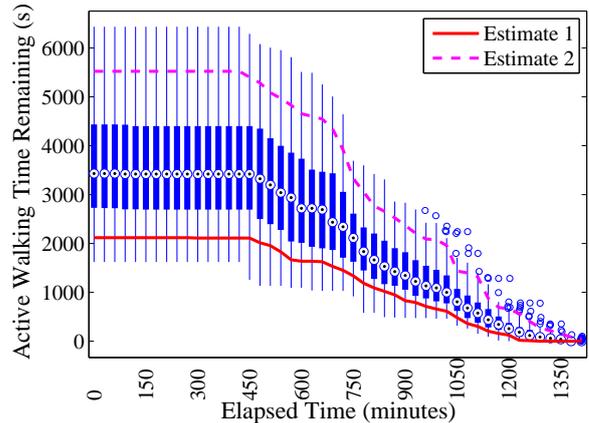


Figure 2: Example candidate profiles generated from the sample days.

where the d_m result from the feedback method, and the w_m are the (random) active times in each stage m . For the purposes of demonstrating of the usefulness of the profiling technique, we do not attempt to model the w_m probabilistically or delve deeply into optimization strategies within this work, but rather employ a simple heuristic approach, as follows.

Given D days of previous observations, we develop D candidate estimate profiles ($\hat{\mathbf{W}}^{(i)}, i = 1, \dots, D$) and choose the one which maximizes the average value of g over the D days. The candidate profiles, however, are not simply the actual profiles $\mathbf{W}^{(i)}$ previously observed on days $1, \dots, D$. Rather, for each time step k , $\hat{W}_k^{(i)} = W_k^{(r_i)}$, where r_i indicates rank in sorted order at time k . That is, the i th candidate profile's estimate at time k is the (i/D) th percentile among $W_k^{(1)}, W_k^{(2)}, \dots, W_k^{(D)}$.

This concept is illustrated in Figure 2, which depicts each time stage as a boxplot showing the distribution of W_k across previously observed days. Example candidate profiles are drawn as lines on the figure (one in the bottom quartile and one in the top quartile). The small circles indicate outliers, which may cause large jumps at those times among the highest-rank candidate profiles, but the training process should help to rule out those candidates if they are too extreme.

The candidate profile which, when simulated with all D days previously observed, maximizes the average utility g over all days is chosen. That is,

$$\max_i \frac{1}{D} \sum_{m=1}^D g(\mathbf{d}^{(m)}, \mathbf{w}^{(m)}) \quad (9)$$

is used to choose the “best” profile, $\hat{\mathbf{W}}^{(i)}$

3.5 Fidelity Metric

While the framework described previously can be used to map to a particular set of utility/fidelity measures and/or constraints, for the remainder of this work, we will consider a specific fidelity metric designed to capture the objectives that were expressed qualitatively until now. (Ideally such a metric would be informed directly by the application, but for now we choose a more agnostic metric.) Recall, we wish to

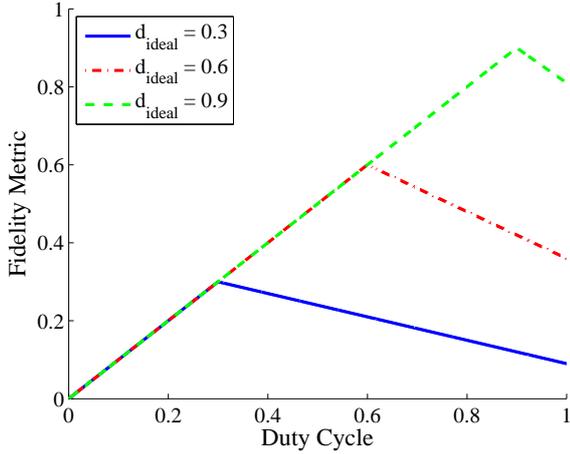


Figure 3: Response of the utility measure as a function of duty cycle, given a fixed choice of duty cycle for the day

capture the active periods at all times of the day at an equal, maximal duty cycle. Therefore, a utility metric should reward a high duty cycle while penalizing variations in the choice of duty cycle over the day. If the battery is exhausted prematurely, subsequent active times are treated as operating at a duty cycle of 0, manifesting as extra variation in contrast to earlier periods of nonzero duty cycle.

Again, we emphasize that choice of fidelity metric may be changed to suit one’s needs. For this work, we select the following metric for a single day’s operation:

$$v = g(\mathbf{d}, \mathbf{w}) = \text{mean}(\mathbf{d}) - \text{Var}(\mathbf{d}) = \frac{1}{W_1} \left(\sum_{k=1}^N d_m w_m - \sum_{k=1}^N w_m (d_m - \mu)^2 \right) \quad (10)$$

This function was chosen, in part, to give a reasonable response for the fixed-duty-cycle baseline approach. The fidelity v is maximized when $d = d_{ideal}$ (Equation 2), but decreases linearly as d moves away from the d_{ideal} . This is illustrated in Figure 3 for three possible values of d_{ideal} . When the duty cycle is too conservative ($0 \leq d \leq d_{ideal}$), the $\text{Var}(\mathbf{d})$ term is 0, leaving the linear function $v = \text{mean}(\mathbf{d}) = d$. When the duty cycle is too aggressive ($d_{ideal} \leq d \leq 1$), the mean saturates at d_{ideal} while the $\text{Var}(\mathbf{d})$ term grows due to periods of missed data (effective duty cycle of zero), causing a linear decrease in v (specifically, $v = d_{ideal} - d_{ideal}(d - d_{ideal})$).

4. EVALUATION METHODOLOGY

The approach described above is evaluated in simulation based on sample profiles collected from three human subjects over several months. An energy model was developed for calculating idle and active power figures for the system, and step count data from a Fitbit® tracker were used as a proxy for the profiles that would normally be collected directly with the inertial BSN node running over multiple days.

4.1 Power Modeling

The various energy-consuming components of a BSN node can be included in a power model to predict the average op-

erating power for collecting and sending sensor data. An analytical power model allows for power-fidelity analysis to be done off-line, and it can be easily modified for other hardware platforms or components. The total average power needed to collect and transmit data on a BSN node can be broken into contributions from the various board-level hardware components: the microcontroller (P_{MCU}), the radio (P_{radio}), and the sensors (P_{sensor}) as shown in (11).

$$P_{SYS} = P_{sensor} + P_{radio} + P_{MCU} \quad (11)$$

This work attempts to model a custom 6 degrees-of-freedom inertial sensor node based upon the TEMPO inertial measurement BSN node [3] with newer sensors and radio components substituted. Specifically, the Analog Devices ADXL345 tri-axial digital accelerometer was chosen for its low-power sensing mode, and the Invensense MPU-6000 was chosen as the lowest-power available sensor providing a tri-axis gyroscope. In its low-power sensing mode, the ADXL345 consumes only $50\mu\text{A}$ at a 100 Hz sampling rate; in its higher power (lower-noise) mode, it consumes $140\mu\text{A}$ at the same sampling rate. The MPU-6000 consumes 3.6 mA in active mode, thus dominating the sensing power when turned on.

The radio consumption was modeled as a constant energy per bit, E_{bit} , assessed to be approximately equal to $2.83\mu\text{J}$ using values from a common 2.4 GHz transceiver capable of transmitting at 250 kbps [1]. The radio values assume that the transceiver buffers the entirety of a message before sending it across the body area channel. So higher compression ratios result in fewer messages being sent (as opposed to shorter messages) which is desirable given the significant overhead of sending a message. The average power of the radio is expressed in terms of the average bitrate, f_{bit} :

$$P_{radio} = f_{bit} E_{bit} \quad (12)$$

The average power consumption of the microcontroller is broken in contributions from active mode (P_{AM}) and low-power mode (P_{LPM}),

$$P_{MCU} = \alpha_{AM} P_{AM} + \alpha_{LPM} P_{LPM} + E_{LPM,trans} \quad (13)$$

with α_{AM} and α_{LPM} defined as follows:

$$\alpha_{AM} = (t_{proc} + t_{send})/t_{total} \quad (14)$$

$$\alpha_{LPM} = (t_{total} - t_{proc} - t_{send} - t_{LPM,trans})/t_{total} \quad (15)$$

are relative on-time factors which scale the raw power figures to average power. The time amounts are relative to some epoch of time t_{total} in which the MCU burst-reads sensor data and updates its walking detection algorithm (t_{proc}), and—if the wearer is walking—sends the data to the radio (t_{send}). The time to switch into low-power mode and back is denoted $t_{LPM,trans}$, while the finite energy required for switching to LPM and back is $E_{LPM,trans}$.

All of these values are either known or found in the microcontroller documentation and datasheets except for t_{proc} and t_{send} . t_{proc} is directly related to the number of clock cycles needed to read in the sensor values and perform walking detection, and t_{send} is directly related to the sampling rate. An example of a simple walking detection algorithm would consist of computing the vector magnitude of each 3-axis accelerometer sample, periodically calculating its variance over a recent window, and comparing to a pre-determined threshold. Relative to a one-second epoch, at a 100-Hz sampling rate, this can be done in roughly 11.9 ms

Model Parameter	Value
$P_{accel,low}$	165 μ W
$P_{accel,high}$	462 μ W
P_{gyro}	11.9 mW
f_{bit}	9600 bps
E_{bit}	2.83 μ J
P_{AM}	8.83 mW
P_{LPM}	8.56 μ W
E_{LPM_trans}	300 nJ
t_{total}	1 s
$t_{proc,idle}$	13.8 ms
$t_{proc,active}$	16.2 ms
t_{send}	4.8 ms
t_{LPM_trans}	3 ns
P_{active}	39.7 mW
P_{idle}	296 μ W

Table 1: Energy model parameters

on a TI MSP430F1611 at 3.69 MHz. This microcontroller power model was validated by measuring current consumption through a sense resistor for a TIMSP430F1611 and was shown to give values within 3% of those measured over a sweep of processing cycles/epoch.

A summary of relevant parameters in the power model for the BSN node model used in this work is given in Table 1.

4.2 Fitbit® Profiles

In a realistic deployment using this technique, one would derive the observed walking profile $\mathbf{W}^{(i)}$ for a given day, i , directly using the capabilities of the sensor node, which would use its accelerometer and a walk detection algorithm to detect periods of activity and log them in local RAM or flash memory. For example, a relatively high-granularity walking profile consisting of one value—the number of seconds active—for each minute of the day, would require 1440 bytes, which is not unreasonable for an embedded MCU such as the TIMSP430F1611 used in our node model.

However, for the convenience of experimentation, we have used the Fitbit® tracker as a proxy for the BSN node to collect sample profiles from three subjects over extended periods of time. The Fitbit® is a commercial system consisting of a small clip-on device that tracks daily activity levels. It reports the total number of steps taken for each minute of the day. This was multiplied by a sample cadence value for the subject to get an estimate of the time spent walking for each minute. Admittedly, this is only an estimate, as the cadence cannot be assumed constant. When the cadence is clearly underestimated, resulting in an apparent active time of 60s for a given minute, this value is reduced to 60s. In general, we expect that these profiles are a reasonable representation of the relative intra-day and inter-day patterns for the wearers and illustrate the value of the proposed approach, and are thus satisfactory for our initial analysis.

Three volunteer subjects (adult male) wore Fitbit® trackers, each beginning on a different date, resulting in 133, 126, and 68 days of tracking, respectively. The subjects worked 40-50 hours a week, sitting down for the majority of the time. Outside of work, the subjects went about their daily lives, which consisted of typical daily activities (housework, exercise, television watching, grocery shopping, computer work, etc.). The distributions of total daily walking time for each

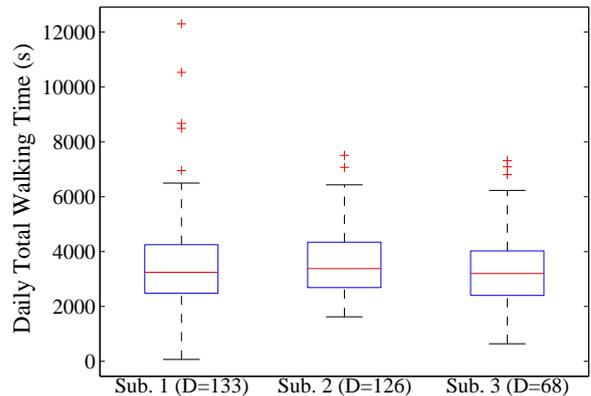


Figure 4: Distribution of total daily walking time, for each subject. D indicates the total number of days in the data set for the particular subject.

subject are indicated by boxplots in Figure 4. Subject 1 in particular exhibited the largest variation, including multiple significant outliers (indicated by ‘+’ symbols).

4.3 Simulation

The profiling approach is explored via simulation in MATLAB. The profiles collected from Fitbit® are partitioned into a training set (60%) and a testing set (40%) for the subject, and the training method described in Section 3.4 is applied to the training set using the fidelity metric developed in Section 3.5. To ensure fair comparison with the static (fixed-duty-cycle) baseline approach, the static estimate \hat{W} is analogously developed by training. Note that this is a scalar estimate used once at the beginning of the day to select a static operating point.

One parameter of particular interest in simulation is the starting battery capacity, E_{batt} . The approach in this work assumes that the battery is not large enough for the node to run at full duty cycle for all days’ walking amounts, but not so small that even an ideal control scheme would show minimal improvement. Thus, the battery is studied over a range of interest and its effect is explored in the final results. Note: given the relatively low power model parameters assumed here (Section 4.1), we assume the deployment would target a weekly, rather than daily, recharge period, so E_{batt} reflects $(1/7)^{th}$ of the battery pre-allocated to a given day, although one could imagine adapting the approach to use a 7-day, rather than 1-day, horizon.

During deployment, the node itself would be responsible for collecting the daily walking profiles in addition to the primary task of capturing periods of walking activity. For this analysis, we reserve a portion of the battery capacity, $E_{reserve}$, sufficient to operate the node in idle mode for the duration of the observation period. That is, $E_{reserve} = P_{idle}T_{target}$ (where we assume $T_{target} = 24$ hr). Therefore, the value used for the starting energy (as in Section 3.3) is $E_1 = E_{1,actual} - E_{reserve}$.

5. EXPERIMENTAL RESULTS

A simulation was performed as described in Section 4 using sample Fitbit profile data from a subject who wore the device for a total of 126 days. Figure 5 illustrates the per-

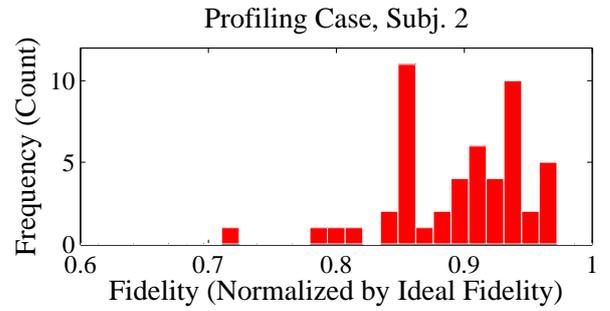
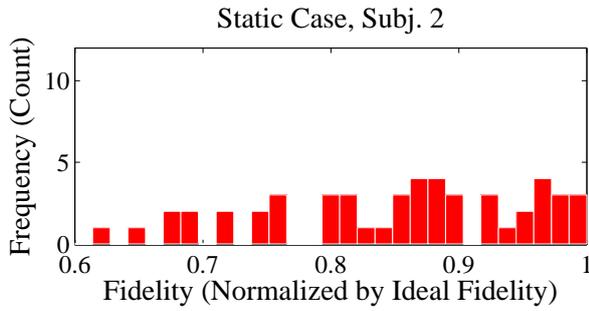


Figure 5: Distribution of fidelity, normalized to the ideal fidelity ($v = d_{ideal}$), provided by an oracle, for subject 2, for the *static-choice baseline* method (for $E_{batt} = 82.7$ J).

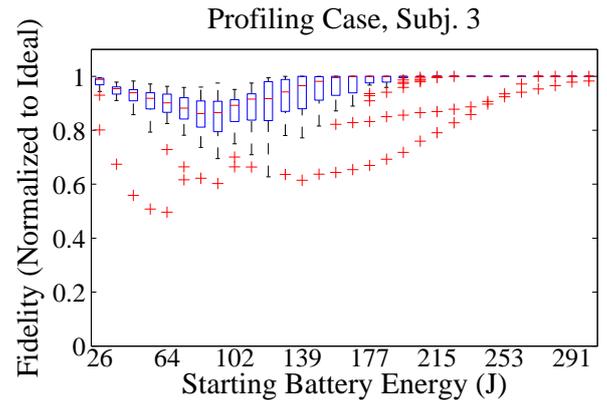
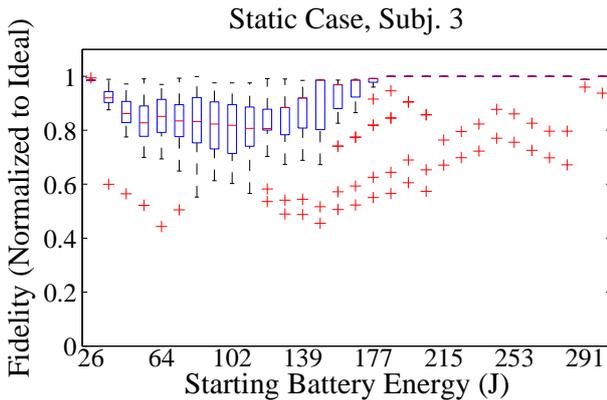
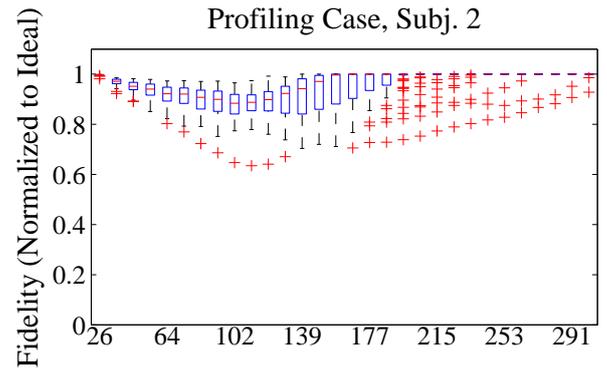
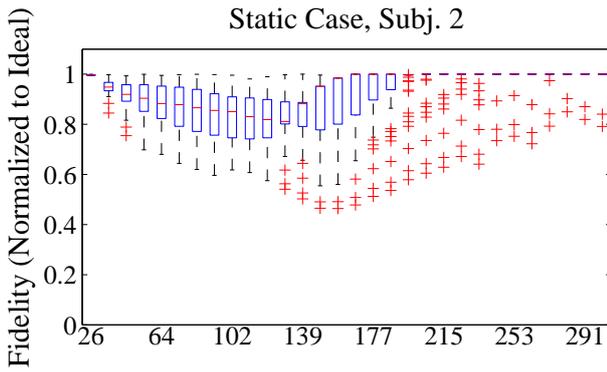
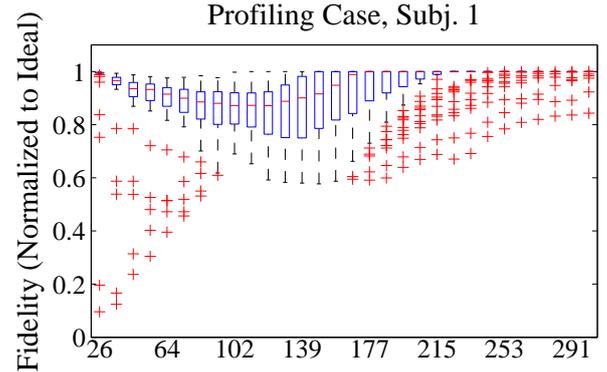
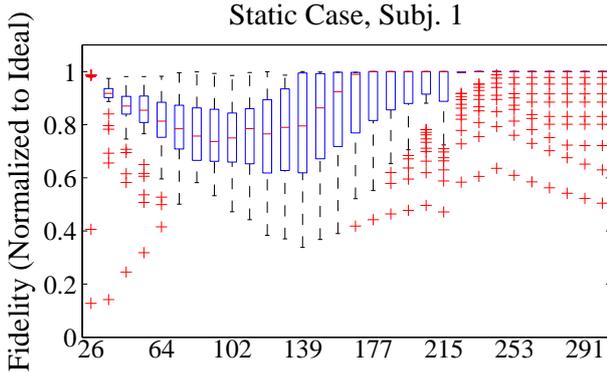


Figure 6: Distribution of fidelity, normalized to the ideal fidelity ($v = d_{ideal}$), provided by an oracle, as a function of initial battery capacity for the day. Boxplots indicate the median, quartiles, and outliers (denoted by '+' symbols) for each battery value. Results are shown for both the *static-choice baseline* case as well as the proposed *profiling/feedback* approach.

formance of the static baseline approach and the proposed profiling technique. In each case, the resulting fidelity score for each day has been normalized by the ideal (maximum possible, as selected by an “oracle” with *a priori* knowledge) fidelity for that day. Thus Figure 5 shows a histogram of such normalized fidelity—for both the baseline static case and the profiling technique—over all days in the test set, for a particular battery size at the start of all days.

It can be seen that the distribution of normalized fidelity tends to be both more concentrated and tends toward higher values. The plots in Figure 6 show the same information for the baseline and profiling techniques, respectively, as a function of starting battery capacity. That is, the histograms from the previous figures become boxplots in the latter figures (one boxplot for each battery size). Again, the distribution of normalized fidelity tends to exhibit less variation and to tend more toward the ideal than in the static case.

However, there is significant variation among the three subjects. Subject 2, in particular, exhibits the most pronounced improvement, followed arguably by Subject 3, then Subject 1. Consulting Figure 4 from Section 4.2, we can see that, correspondingly, Subject 2 exhibits the least amount of variation in daily walking-time totals, followed in order by Subjects 3 and 1. This suggests that the proposed approach is more effective for a wearer whose routines are more regular, as one would expect.

It is worth noting that the fidelity metric used here, as explained in Section 3.5, exhibits a linear degradation in fidelity—for the static baseline case—as the fixed choice of duty cycle moves away from the ideal, which one might characterize as forgiving. An alternate formulation with a different fidelity function (e.g., quadratic degradation) could alter the results significantly.

6. CONCLUSION

This work presented a method for balancing competing needs: ensuring battery life—or rather, capturing all data of interest during a deployment period—and maximizing fidelity of captured data under uncertain amounts of load (periods of interest). In order to accomplish this, personal activity profiles were utilized to predict future user behavior, allowing online adjustment (through feedback) to actual behavior (and corresponding energy expenditure) to better balance both battery lifetime and fidelity. It was shown that, in a subject with reasonably regular behavior trends, the level of fidelity (determined here by the chosen duty cycle) can be increased and made more consistent across various days. Such approaches as this, which combine feedback with profile-based predictions, could help to better enable BSNs for use in longitudinal studies of continuous monitoring by reducing the necessary battery size and/or frequency of recharges.

Future work will seek to further develop a general profiling- and feedback-based battery lifetime-fidelity optimization methodology, while addressing the limitations of this initial approach. The methodology should support a variety of BSN platforms and applications, each of which implies a different power model, fidelity definition, and/or power-fidelity settings (other than duty cycling). In addition, for multi-day recharge periods, prediction and feedback accounting for the entire recharge period as a whole (rather than individual days) will be explored. Finally, additional exploration is needed in the probabilistic/statistical model-

ing of user behavior and optimization methods to produce more accurate profiles that result in higher fidelity, including incorporating or comparing with aspects of related methods described in Section 2. For instance, the observations of activity-so-far in a day (w_1, \dots, w_N) may well be predictive of future activity and thus could be used to update predictions about the remainder of the day on-the-fly, rather than only offline at the beginning of the day; this would be especially useful if the wearer’s profiles cluster into similar, yet separate, groupings. This research direction will facilitate pervasive adoption of BSNs by enabling them to intelligently adapt to system dynamics and resource availability.

7. ACKNOWLEDGMENTS

The authors would like to thank Dr. Stephen D. Patek (Systems & Information Engineering, UVA) for his helpful input. This work is supported by the National Science Foundation under grants ECCS-0901686 and CNS-1035771 and the Graduate Research Fellowship (DGE-00809128).

8. REFERENCES

- [1] Texas Instruments, CC2500 Low-cost low-power 2.4 GHzRF transceiver Datasheet, 2007.
- [2] L. Au, W. Wu, M. Batalin, D. McIntire, and W. Kaiser. Microleap: Energy-aware wireless sensor platform for biomedical sensing applications. In *IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 158–162, 2007.
- [3] A. T. Barth, M. A. Hanson, H. C. Powell Jr., and J. Lach. TEMPO 3.1: A body area sensor network platform for continuous movement assessment. In *International Conference on Body Sensor Networks (BSN)*, pages 71–76. IEEE, 2009.
- [4] A. T. Barth, M. A. Hanson, H. C. Powell Jr., and J. Lach. Online data and execution profiling for dynamic energy-fidelity optimization in body sensor networks. In *Proceedings of the 2010 International Conference on Body Sensor Networks, BSN '10*, pages 213–218, Washington, DC, USA, 2010. IEEE Computer Society.
- [5] T. L. Cheung, K. Okamoto, F. Maker, X. Liue, and V. Akella. Markov decision process (mdp) framework for optimizing software on mobile phones. In *Proceedings of the Seventh ACM International Conference on Embedded Software, EMSOFT '09*, pages 11–20, New York, NY, USA, 2009. ACM.
- [6] Y. He, W. Zhu, and L. Guan. Optimal resource allocation for pervasive health monitoring systems with body sensor networks. *IEEE Transactions on Mobile Computing*, 10(11):1558–1575, Nov. 2011.
- [7] A. Panangadan, S. M. Ali, and A. Talukder. Markov decision processes for control of a sensor network-based health monitoring system. In *Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence - Volume 3, IAAI'05*, pages 1529–1534. AAAI Press, 2005.
- [8] Y. Wang, B. Krishnamachari, Q. Zhao, and M. Annamaram. Markov-optimal sensing policy for user state estimation in mobile devices. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN '10*, pages 268–278, New York, NY, USA, 2010. ACM.