

Energy-Efficient Sampling Schedules for Body Area Networks

Vishwa Goudar and Miodrag Potkonjak
Computer Science Department
University of California, Los Angeles
Los Angeles, USA

Abstract— We present a novel sampling method to overcome the tradeoff between sensing fidelity and energy-efficiency in the context of localized sensor arrays used by Body Area Networks (BANs). Prior research has tackled this tradeoff as a coverage problem, wherein a subset of sensors must cover the sensor field. Instead, we formulate it as a power-constrained sampling problem, limiting the number of samples taken per epoch to produce schedules with enhanced coverage and energy savings. This formulation capitalizes on the periodic nature and the strong spatio-temporal interactions that are innate to BAN sensor samples. Our algorithm produces schedules with over 170% in energy savings with increased sensor coverage that yields up to a 41% improvement in diagnostic estimates.

I. INTRODUCTION

Major advances in Wireless Sensor Network (WSN) performance and application support services have wielded a host of scientific and engineering applications with robust, flexible and scalable remote monitoring and control capabilities. While early applications enlisted them to static deployment over large areas, recent ones are a product of novel deployment configurations that open WSNs up to new application domains.

For example, they are being retooled as smaller networks of wearable sensors and devices known as Body Area Networks (BANs), with applications in the medical, sport, military and security domains. For example, the field of wireless health aims to improve healthcare and medical services via BANs that continuously and unobtrusively monitor patient activity/behavior. These systems endow healthcare providers and medical researchers access to previously unavailable, temporally and diagnostically rich measurements of validated medical metrics.

However, widespread adoption of BANs has been stymied by usability issues. The validity and reliability of measured metrics under noisy conditions and in varied environments is crucial to adoption. Furthermore, BANs are often mobile and comprised of localized sensor arrays powered by a single battery. Such systems must be able to function unattended for extended periods, while staying light-weight and maintaining a small form factor so as not to hinder/disrupt normal subject activity. Since energy usage affects the battery size and therefore the form factor, it is crucial to adoptability.

This gives rise to a tradeoff between lowering energy consumption that requires sampling fewer sensors at a time, and sustaining a high quality of medical measurements. In this paper, we present a sampling algorithm that raises sampling efficiency to tackle this tradeoff. It abides by pre-specified power thresholds to procure at most k samples at a time while maximizing their diagnostic relevance. This formulation is based on, (i) the localized scope of BAN sensor arrays that engenders strong spatio-temporal interactions, thereby allowing high-fidelity sample inference, and, (ii) the periodic nature of the measurements, such as plantar pressure, ECG, and respiration.

II. RELATED WORK

Among the host of engineering considerations in the design and deployment of WSNs, energy-efficiency has been identified as the most important performance criterion [8]. Although a subclass of WSNs, BANs offer a distinct set of challenges and opportunities [1], with medical applications ranging from gait analysis [11] to emotional health monitoring [13]. However, energy remains a crucial resource for BANs [2]. Related optimizations span communication [14], sensing coverage [12], etc.

Prior research in the BAN community to have studied the tradeoff between sensing fidelity and energy-efficiency, propose to mitigate it either through energy scavenging [4][18][17], or by posing it as a coverage problem wherein a subset of sensors is selected for its ability to jointly predict the measurements at all sensors [3]. In [9][6] [15] and [16], the authors propose algorithms to further reduce the “covering” subset size while improving sensing fidelity by defining the fidelity only in terms of diagnostic accuracy. However, by taking advantage of the strong spatio-temporal relationships and periodicity in the data, we are able to pose the problem as one of sample selection thereby generating a richer set of measurements under similar power constraints.

III. PRELIMINARIES

A. Human Balance Monitoring via Hermes

We present our algorithm in the context of a human balance monitoring system known as Hermes [10] (Fig. 1a)

This work was supported in part by the *NSF* under awards CNS-0958369, CNS-1059435, and CCF-0926127, and is based upon research performed in collaborative facilities renovated with funds from the *National Science Foundation* under Grant No. 0963183, an award funded under the American Recovery and Reinvestment Act of 2009 (ARRA).



Figure 1. (a) The Hermes smart shoe platform [10] (left); (b) Entropy in maximum amplitude at each sensor for a subject (right).

that measures foot plantar pressure via an array of 99 passive resistive pressure sensors purposefully placed at locations specified by the Pedar plantar mapping [5]. The sampling unit for the platform samples each sensor at 60Hz with data collected through a 16-bit analog-to-digital converter. While Hermes is certainly capable of continuous sampling at all sensors, medical practitioners only require a set of diagnostic metrics for each stride taken by the subject, which may not require all samples for the stride. In this paper, we focus on one such metric. The maximum amplitude metric has been identified as a causative factor in the development of diabetic plantar ulcers [7] and may be measured as the average maximum pressure over all sensors for a stride. Based on this characterization, it is clear that the only relevant sample of each sensor for a stride is the one that provides its maximum pressure over that stride.

B. Algorithmic Motivation

A stride is defined as a single sequence of functions of one limb and is divided into the stance phase, when the limb is in contact with the ground and maximum amplitude is achieved, and the swing phase, when the limb is in the air for advancement. We exploit the periodic nature of a stride to construct a repeatable sampling scheme by dividing the stance phase into a number of *epochs* at which we would like plantar pressure to be sampled at various sensors.

To maximize the information gained with a few samples, we take advantage of the spatio-temporal redundancy inherent to localized sensor arrays. This is achieved with simple linear regression models, similar to [9], that are used to infer multiple readings from a single sample when their mutual information is high.

Fig. 2 shows the Spearman rank-order correlation coefficients between one sensor, s , and all 99 others, ordered by their physical distance from s , for each epoch of the stance phase from a dataset collected by the Hermes platform. While it is clear that, in general, nearby sensors are well correlated with the sensor s compared to distant ones, we observe that the level of correlation varies significantly across epochs. We assert that these observations generalize across most sensors in all our datasets owing to a variation in the spatial relationships across sub-phases of the stance phase.

Our approach to maximizing the quality of diagnostic information extracted, while limiting the maximum number

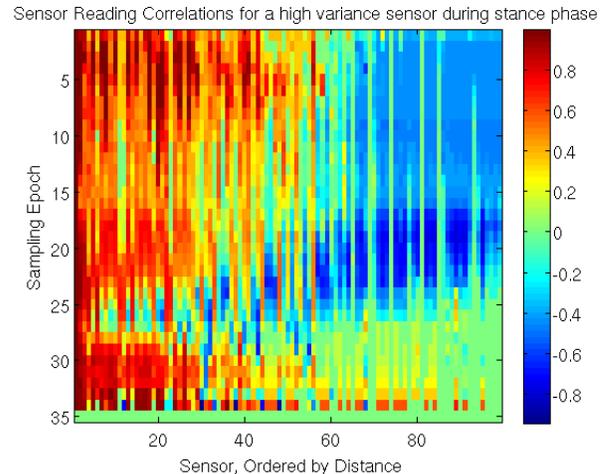


Figure 2. Spatio-Temporal correlations of plantar pressure measurements between a single sensor and all others on a Hermes shoe, for each epoch of the stance phase for a subject.

of sensors, k , sampled at an epoch is guided by the following directives.

- 1) Inferring readings with the most contextually relevant information, which corresponds to measurements at sensors with the most *entropy*, at epochs when they are likely to achieve maximum amplitude. This latter property is interpreted from a distribution (*dist*), over all epochs, of the likelihood that a sensor will observe its maximum amplitude, and is constructed from a training dataset.
- 2) Utilizing spatio-temporal relationships for a given epoch, to infer readings at un-sampled sensors with high fidelity. Here we also exploit the observation in [3], that time shifting the signals from 2 sensors greatly improves their correlation. *senPred* captures those sample pairs that exhibit low relative linear regression rms error.
- 3) Covering sensors, via measurement or inference, over a majority of the epochs when they are most likely to achieve maximum amplitude.
- 4) Forfeiting concurrent coverage of sensors whose maximum amplitude is strongly correlated.

The benefits of these directives are highlighted by Fig. 3. From the likelihood distribution of 4 representative color-coded sensors over 8 epochs (Fig. 3a), we see that the narrow width of the distributions allows us to gain information by forgoing inference of a sensor's reading at most epochs which creates space for others. For example, s_4 need be sampled only at epochs e_5 thru e_7 leaving space for others to be sampled at e_1 thru e_4 . It also obviates sampling at irrelevant epochs, such as e_8 , thereby improving energy-efficiency.

Furthermore, inference based on spatio-temporal correlations (Fig. 3c) enables the recreation of multiple readings from a single sample, including those at epochs other than the sampled one, vacating the necessity to sample at these other epochs. In our example, sample (s_1, e_4) predicts readings (s_3, e_3) and (s_3, e_4) . However, it should be clear that our approach is not suitable for measurements of behavior that is not periodic or does not have spatio-temporal redundancy.

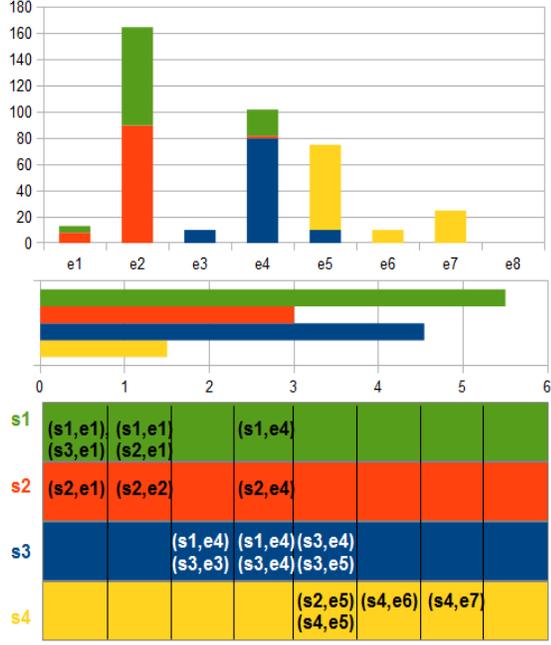


Figure 3. (a) Max. Amplitude Likelihood Distribution (top); (b) Max. Amplitude entropy (middle); (c) Spatio-temporal inference map that lists the set of samples that may be used to infer each reading (bottom).

IV. DIAGNOSTIC DRIVEN SAMPLE SELECTION

We now describe an offline learning approach that combines these considerations to produce a customized sampling schedule based on a training dataset comprised of a few strides of the subject. The resulting schedule may then be applied towards energy-efficient diagnostic estimation for the subject. Table I summarizes the power constrained sample selection algorithm (PCSS) that determines the sensors to be sampled at each epoch, as well as the readings that are inferred from them.

Based on our guiding directives described earlier, line 4 of the algorithm computes the value of a sample. It is measured by a *heuristic* that aggregates the merit of each measurement that may be accurately inferred or *covered* by the sample (directive 2). The merit of such an inferred reading is composed of the likelihood that the corresponding sensor will observe maximum amplitude at the inferred epoch (directive 1), along with the entropy (directive 1) and the coverage received thus far (directive 3) by the corresponding sensor. Line 4 also ensures that we don't attribute an inference to a sample if the inferred reading has already been covered. At each iteration, the algorithm schedules the sample with the best heuristic value. It terminates either when there are no remaining *free epochs* at which assignments may be made, or no merited assignments to make.

While the PCSS algorithm is driven by most of our guiding directives, it does not ensure that the selected sensors are well covered (directive 3). Put differently, each sensor that we decide to cover must be covered in most, if not all, epochs when they are likely to observe the metric, barring which we increase the likelihood of using erroneous values in the

TABLE I. POWER CONSTRAINED SAMPLE SELECTION (PCSS)

1: Input: <i>senPred</i> ; <i>distr</i> ; <i>entropy</i> ; <i>k</i>
2: Output: Sampling and Prediction Schedule
3: Do until no more coverage is possible
4: Compute heuristic for sampling each sensor at each free epoch upto <i>k</i> , based on non-covered samples
5: Schedule sensor sample with best heuristic value and track the readings it will cover (and infer)
6: End do
7: Return sampling and inference schedules

composition of the diagnostic and increase its error. Also, PCSS does not refrain from simultaneous covering sensors that are well correlated or *equivalent* (directive 4). These issues are mitigated by an iterative refinement strategy outlined in table II.

CIIR iteratively eliminates those sensors from consideration for coverage that have failed to meet a coverage threshold τ , via *senPred*, thereby reducing competition and improving coverage (lines 7 and 8). By calling the PCSS algorithm with stepwise increments of the maximum number of samples per epoch, *k*, CIIR is able to avoid concurrent coverage of equivalent sensors, by filtering out from *senPred* all sensors that are equivalent to each covered sensor (line 11). The CIIR algorithm terminates to produce the sampling and inference schedules. By sampling measurements at each stride based on this sampling schedule, the maximum amplitude metric is composed, based on a multiple linear regression model, from the inferred maximum amplitude values at each covered sensor.

TABLE II. COVERAGE IMPROVING ITERATIVE REFINEMENT (CIIR)

1: Input: <i>senPred</i> ; <i>distr</i> ; <i>entropy</i> ; <i>k</i>
2: Output: Sampling and Prediction Schedule
3: for <i>k'</i> in 1: <i>k</i> do
4: <i>senPred'</i> = <i>senPred</i>
5: while τ_{cov} spans the range $[\tau_{mincov}, \tau_{maxcov}]$ in small increments do
6: Run PCSS with input <i>senPred'</i> , <i>distr</i> , <i>entropy</i> and <i>k'</i> and store the returned schedules temporarily
7: Remove samples from schedules corresponding to predicted sensors whose coverage is less than τ_{cov}
8: Restrict <i>senPred'</i> to only cover remaining sensors
9: end while
10: Copy temporary schedules to permanent ones
11: Restrict <i>senPred</i> to only cover sensors that are not covered or equivalent to those that are covered
12: end for
13: Return sampling and inference schedules

I. RESULTS

We discuss the performance of our algorithm in the context of 5 plantar pressure datasets obtained from distinct subjects with distinct gait profiles, wearing the Hermes shoe while walking. We divide each dataset into 2 subsets, ascribing 80% for training and 20% for testing. All the results discussed here are for the testing subsets wherein a distinct sampling schedule was constructed from each training

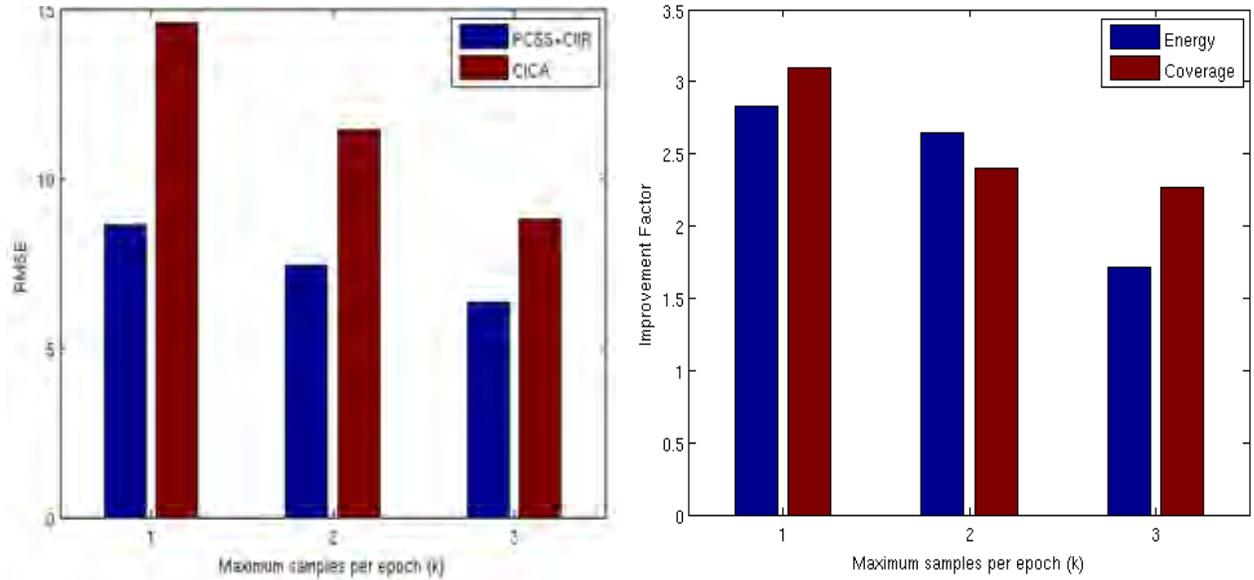


Figure 4. (a) RMSE of max amplitude diagnostics for our algorithm vs. CICA (left); (b) Sensor Coverage and Energy Savings Improvement Factors for different k (right).

dataset. We compare our algorithm's energy savings and diagnostic accuracy, to the CICA-based algorithm in [9].

Here, we measure accuracy in terms of the Root Mean Square Error (RMSE) in the maximum amplitude metric, where the RMSE is taken jointly over all datasets. Fig. 4a plots the observed RMSE of the 2 algorithms, for values of k between 1 and 3. Our diagnostic improvement over CICA ranges between 28% and 41%. Note here, that the CICA-based algorithm will sample the same k sensor(s) at every epoch. While the quality of the metric for these sensors would be perfect, we maintain that in 4 out of the 5 datasets, our algorithm's metric prediction accuracy, as described by its coefficient of determination, is above 0.9 for all covered sensors. Energy is saved by our algorithm when fewer than k samples are taken at an epoch. Calculated as a weighted average over all datasets, with the weights based on the number of strides per dataset, our energy savings over CICA vary between 70% and 178% (Fig. 4b), rising to 263% for one of the datasets when k is 2.

Our ability to produce more accurate diagnostics with less energy can be chalked up to improved sensor coverage. The sensor coverage factor is the ratio of the total number of covered sensors, to k and is also calculated as a weighted average over all datasets. We are able infer metric values at over twice the number of sensors sampled by CICA (Fig. 4b). Finally, we observe that the covered sensors are very informative – the average entropy of the covered sensors as a percentage of the maximum ranges between 77% and 96%.

REFERENCES

- [1] Hanson M.A et al. "Body area sensor networks: challenges and opportunities," *Computer*, vol.42, no.1, pp. 58-65, Jan. 2009.
- [2] Yan L. et al. "Energy comparison and optimization of wireless body-area network technologies," In Proc. ICST BodyNets 2007, Florence, Italy, pp.1-8, June 2007.
- [3] Noshadi, H. et al. "Energy optimization in wireless medical systems using physiological behavior," In Proc. ACM Wireless Health 2010, San Diego, CA, pp.128-136, Oct. 2010.

- [4] Wendt J. B. et al. "Spatiotemporal assignment of energy harvesters on a self-sustaining medical shoe," In Proc. IEEE Sensors 2012, in press.
- [5] Novel.de, Pedar, 2007, <http://www.novel.de/>
- [6] Wendt J. B. et al. "Semantics-driven sensor configuration for energy reduction in medical sensor networks," In Proc. IEEE ISLPED 2012, Redondo Beach, CA, pp. 303-308, July 2012.
- [7] Cavanagh P.R. and Bus S.A. "Off-loading the diabetic foot for ulcer prevention and healing," *JAPMA*, vol.100, no.5, pp.360-368, Sep 2010.
- [8] Culler, D. et al. "Guest editors' introduction: overview of sensor networks," *Computer*, vol.37, no.8, pp. 41- 49, Aug. 2004.
- [9] Wendt, J.B. and Potkonjak, M. "Medical diagnostic-based sensor selection," *IEEE Sensors*, pp.1507-1510, Oct. 2011.
- [10] Noshadi H. et al. "Hermes - mobile balance and instability assessment system," In Proc. BIOSIGNALS 2010, Valencia, Spain, pp.264-270, Jan. 2010.
- [11] Barth A.T. et al. "Longitudinal high-fidelity gait analysis with wireless inertial body sensors," In Proc Wireless Health 2010, ACM, San Diego, CA, pp. 192-193, Oct. 2010.
- [12] Ghasemzadeh H. et al. "Action coverage formulation for power optimization in body sensor networks," In Proc. IEEE ASPDAC 2008, Los Alamitos, CA, pp. 446-451, March 2008.
- [13] Dickerson R. et al. "Empath: a continuous remote emotional health monitoring system for depressive illness," In Proc. Wireless Health 2011, ACM, San Diego, CA, Article 8, Oct. 2011.
- [14] Ghasemzadeh H. and Jafari R. "Data aggregation in body sensor networks: a power optimization technique for collaborative signal processing," In Proc. IEEE SECON 2010, Boston, MA, pp. 439-447, June 2010.
- [15] Wendt J. B. et al. "Energy and cost reduction in localized multisensory systems through application-driven compression," In Proc. IEEE DCC 2012, Snowbird, UT, pp. 411, April 2012.
- [16] Noshadi H. et al. "Energy Optimization in Wireless Medical Systems Using Physiological Behavior", In Proc ACM/BMES Wireless Health, San Diego, CA, pp. 128-136, Oct 2010.
- [17] Rofouei M. et al. "Optimization Intensive Energy Harvesting", In Proc DATE, pp. 272-275, March 2012.
- [18] Goudar, V. and Potkonjak, M. "Dielectric Elastomer Generators for Foot Plantar Pressure Based Energy Scavenging," *IEEE Sensors*, Oct. 2012.