

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

3-2011

An Energy Efficient Quality Adaptive Multi-Modal Sensor Framework for Context Recognition

Nirmalya ROY

Institute for Infocomm Research, Singapore

Archan MISRA

Singapore Management University, archanm@smu.edu.sg

Christine JULIEN

University of Texas at Austin

Sajal K. DAS

University of Texas at Arlington

Jit BISWAS

Institute for Infocomm Research, Singapore

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Software Engineering Commons](#)

Citation

ROY, Nirmalya; MISRA, Archan; JULIEN, Christine; DAS, Sajal K.; and BISWAS, Jit. An Energy Efficient Quality Adaptive Multi-Modal Sensor Framework for Context Recognition. (2011). *2011 IEEE International Conference on Pervasive Computing and Communications (PerCom): Seattle, 21-25 March: Proceedings.* 63-73.

Available at: https://ink.library.smu.edu.sg/sis_research/1660

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

An Energy Efficient Quality Adaptive Multi-Modal Sensor Framework for Context Recognition

¹Nirmalya Roy, ²Archan Misra, ³Christine Julien, ⁴Sajal K. Das and ¹Jit Biswas

¹Institute for Infocomm Research, Singapore, ²Applied Research, Telcordia Technologies, New Jersey

³The University of Texas at Austin, ⁴The University of Texas at Arlington

Email: {nroy, biswas}@i2r.a-star.edu.sg, archan@research.telcordia.com, c.julien@mail.utexas.edu, das@uta.edu

Abstract— Proliferation of mobile applications in unpredictable and changing environments requires applications to sense and act on changing operational contexts. In such environments, understanding the *context* of an entity is essential for adaptability of the application behavior to changing situations. In our view, context is a high-level representation of a user or entity's state and can capture activities, relationships, capabilities, etc. Inherently, however, these high-level context measures are difficult to sense directly and instead must be inferred through the combination of many data sources. In pervasive computing environments where this context is of significant importance, a multitude of sensors is already being embedded in the environment to provide streams of low-level sensor data about the environment and the entities present in that environment. A key challenge in supporting context-aware applications in these environments, therefore, is supporting energy-efficient determination of multiple (potentially competing) high-level context measures simultaneously using data from low-level sensor streams. In this paper, we first highlight the key challenges that distinguish the multi-context determination problem from single context determination and then develop our framework and practical implementation to account for them. Our model captures the tradeoff between the accuracy of estimating multiple context measures and the overhead incurred in acquiring the necessary sensor data. Given a set of required contexts to determine, we develop a multi-context search heuristic to compute both the best set of sensors to contribute to context determination and parameters of the sensing tasks. Our algorithm's goal is to satisfy the applications' specified needs for accuracy at a minimum cost. We compare the performance of our heuristic approach with a brute-force approach for multi-context determination. Experimental results with SunSPOT sensors demonstrate the potential impact of this approach.

I. INTRODUCTION

The integration of sensor networks into pervasive computing is enabling detailed monitoring of the network, the environment, and individuals (users) in the environment. This increase in sensing capability opens the question of how to efficiently and reliably convert streams of low-level sensed data into high-level abstractions called *contexts*. As just one example in this vast application space, remote medical monitoring, especially of elderly individuals and chronically ill patients, is widely perceived as a vital component of future health-care. Using automated and rich analytics to extract medically significant events from raw sensor data, remote monitoring can transform both the quality and cost of health-care delivery. In particular, we are already beginning to witness commercial activity centered on remote monitoring within "smart assisted-living homes," using a combination of body-worn medical and non-medical sensors (e.g., SpO₂ monitors and accelerometers) and *in-situ* sensors (e.g., thermal and motion detectors embedded in the home) [1].

Energy-efficient operation is critical in pervasive computing environments. Communication and energy are inherently intertwined, as communication is one of the most energy-costly tasks performed in pervasive computing networks. Given the

need to collect information from sensors embedded in pervasive computing environments, a promising approach for reducing the communication overhead relies on understanding the degrees of inaccuracy and imprecision that applications in these environments can tolerate. That is, as one example, we would like to enable sensors to change the ways in which they report information (specifically, reducing the frequency with which they communicate sensed data) if the application using those data can tolerate the change. Different applications require their context measures to be estimated to provide varying levels of accuracy (i.e., statistical confidence). Therefore, these context-aware pervasive computing environments require an approach that can dynamically relate applications' tolerances of imprecision to parameters for determining the behavior of the underlying sensor streams.

Pervasive computing environments typically contain many sensors of different types contributing to the determination of the same high-level context measure, albeit with differing degrees of accuracy. Multi-modal sensing, or the fusion of these disparate data streams to infer a single context metric, is an important method for context recognition in that it can help to account for the error of individual sensing stream [2]. Two additional observations further drive our work: 1) *the accuracy of inferred context increases with the use of a progressively larger sensor set*, and 2) *there is tradeoff between the energy overhead of sensing and the achievable quality of the sensed data*. The quality of the inferred context is a function not just of the chosen sensors but also of the permitted inaccuracy in the sensed values. Using these observations, we have defined *Quality-of-Inference (QoINF)* as the *error probability in estimating a context state, given the imprecision in the values of the contributing sensors*. In our previous work [12], we have formally related an application's specified QoINF value for a given context metric to both the set of sensors used to infer that metric and the imprecision allowed in each sensor in the set. We refer to the latter as the sensor's *tolerance range*. A tolerance range based approach allows us to transform data collection from the conventional sink-initiated polling data collection into a more energy efficient source-initiated event driven framework. Specifically, a sensor with a specified tolerance range reports new values to the sink only if they deviate from the previously reported value by more than the specified tolerance. It has previously been shown that a larger tolerance can result in a significant reduction in communication cost [14].




Fig. 1. Impact of Choice of Different Sensor Subsets on $QoINF$ (no consideration of tolerance range)

To demonstrate the interdependency among sensor selection streams in the multi-context recognition problem, we use an example from an Ambient Assisted Living (AAL) environment, as shown in Fig. 1. For example, in isolation, the single context *moving person* can be determined to 94% accuracy using the fusion of two sensor streams: accelerometer and ultrasonic. However, when the same system is asked to simultaneously determine the context *lying on bed*, the force sensitive resistors (FSR) is best used to aid in determining the new context, while *moving person* can be determined through combining the accelerometer and FSR, albeit at the reduced accuracy of 92%. As this example demonstrates, individual solutions to the single context optimization problem may not be the same solutions when we consider the joint determination of multiple contexts when the underlying sensor streams are shared. Fig. 1 does not explicitly relate the sensors' tolerance ranges to the achieved quality of context determination. These tolerance ranges add another dimension to this problem. In this example, however, it is easy to see that the estimation of the *moving person* context will be less accurate in both cases if the accelerometer sensor tolerance range is $\pm 40\%$ (indicating that the true reading may be up to 40% higher or lower than the reported value) as opposed to a tolerance range of $\pm 10\%$.

Our contributions in this paper build on the foundation of our previous formalization [12] of minimum-cost context information using a generic function that captures the relationship between $QoINF$ and the set of sensors (and their assigned tolerance ranges) as reviewed in Section II. Specifically, we make the following novel contributions (described in Section III) and its practical evaluation (presented in Section IV):

- We formalize the problem of multi-context estimation to address the inevitable variety of context streams available in emerging applications.
- We design a low-complexity Lagrangian-based heuristic algorithm to approximate the selection of the best set of sensors and the appropriate tolerance ranges to maintain the specified $QoINF$ s for the multiple contexts simultaneously at a minimum cost.
- We develop and evaluate a practical method for the $QoINF$ functions for multi-context estimation and empirically demonstrate the applicability of the approach.

To evaluate our approach, we articulate a brute-force algorithm for multi-context optimization and compare its time complexity with our heuristic. We further compare the two approaches using statistical regression. We demonstrate how a $QoINF$ function can be computed in practice, lending support for incorporating intelligent adaptation of context quality for conserving energy expense in real world pervasive computing environments

(Section IV). We provide a discussion of issues related to our model (Section V) and other related work (Section VI), followed by the conclusion (Section VII).

II. CONTEXT INFERENCE MODEL

We begin with an overview of model of $QoINF$, that we developed in our prior work to determine a single context variable from multiple underlying sensor streams [12]. The work in this paper extends this to support more challenging multi-context recognition based on multiple underlying sensing streams. In either of the single or multi context estimation problems, determination of a specific context attribute may be viewed as an inference obtained by fusing the values from multiple sensor streams.

A. Uncertainty in the Tolerance Ranges

The model of context inference and associated $QoINF$ functions can capture complex relationships between context inferring quality and the choice of sensors used to infer that context. Such a model does not, however, completely capture the benefits that can be garnered from using continuous event-driven monitoring as described in our problem statement. In event-driven monitoring, sensors can be assigned tolerance ranges, allowing them to only report changes that fall outside of the tolerance range. Given different tolerance ranges for a variety of sensors, different sensors' values may be known to different degrees of precision. Previous work has demonstrated that one may achieve a significant reduction in the communication and energy overheads of sensing by providing small non-zero tolerance ranges to individual sensors [3], [6].

Let $Q_\theta = \{q_1, q_2, \dots, q_\theta\}$ be the collection of tolerance ranges for the selected sensor subset θ , used to infer a single context metric independently. We represent the error associated with this choice of sensors through the following modified dependency relationship, which explicitly relates our $QoINF$ function to the choice of not just the sensors but also Q_θ , the permitted inaccuracy in the value of each sensor:

$$QoINF_C(\theta, Q_\theta) = 1 - \sum_{x \in \Lambda_C} err_C(x, \{(s_i, q_i) : s_i \in \theta, q_i \in Q_\theta\}) \quad (1)$$

where $err_C(\cdot)$ is the probability of error, given accurate readings from the sensor subset θ and tolerance range q_i , when the value for the context variable C is actually x . Note that this formulation computes the average error, summing across all possible values of the context variable C . This implicitly assumes that all context states are equally likely.

B. Cost Model

One of our primary objectives is to reduce the total energy consumption due to communication in the pervasive computing network, while ensuring the application's minimum $QoINF(\cdot)$ bound for a given context metric. The cost incurred in acquiring one of the sensor streams needed to update the context variable is basically a function of both the sensor s_i 's tolerance range (q_i) and the multi-hop transmission cost from the sensor to the sink. We assume that this cost is a linear function of the number of hops (h_i) in the uplink path from sensor s_i to the sink, given a

network in which links have identical transmission power. Intuitively, the cost and tolerance range should be inversely related: if the tolerance range is large, it is less likely for a value to fall outside the permitted range, making communication less frequently necessary. In the absence of any temporal correlation among sensor's samples, we assume the underlying data samples evolve as a random walk, and the $cost \propto \frac{h_i}{q_i^2}$ [6]. In this case, the resulting cumulative cost function for the set of selected sensors, θ , is given by: $COST(\theta, q_\theta) = \kappa \sum_{s_i \in \theta} \frac{h_i}{q_i^2}$, where κ is a scaling constant and h_i is the hop count.

C. Choice and Characteristics of a $QoINF$ Function

While a completely arbitrary $QoINF(.)$ function requires a brute-force search for a solution, there are certain forms of $QoINF(.)$ that prove to be more tractable and lend themselves to more efficient optimization heuristics. A particularly attractive case occurs when the i^{th} sensor's individual $qoinf(.)$ is represented by an Inverse-Exponential distribution of the form

$$qoinf(i) = 1 - \frac{1}{\nu_i} \exp\left(\frac{-1}{\eta_i q_i}\right) \quad (2)$$

where η_i and ν_i are sensitivity constants for sensor s_i . A larger value of ν_i indicates a lower contribution from sensor s_i to the inference of context C . Moreover, for a selection of sensors θ , the resulting $QoINF(.)$ function is modeled as:

$$QoINF_C(\theta) = 1 - \prod_{s_i \in \theta} (1 - qoinf_C(i)) \quad (3)$$

This formulation assumes that the estimation error of each sensor is statistically independent of the others. The above equation satisfies three of the key properties a $QoINF(.)$ function must have: (i) the value of the function always falls between 0 and 1; (ii) $QoINF_C(.)$ is non-decreasing in θ in the sense that, incorporating data from an additional sensor does not lead to a reduced $QoINF$ value; and (iii) as a sensor's tolerance range increases towards infinity, the quality of the contribution of that sensor decreases towards 0. As shown in Fig. 1, the context of *lying on bed* may be computed with an inferencing accuracy of 0.9 (i.e., with a 10% error rate) using data from FSR sensor, but only with 0.7 accuracy using data from a high-quality accelerometer sensor. However, by fusing the data available from FSR and accelerometer, we can achieve an inferencing accuracy of 0.97 ($1 - [(1 - 0.9) \times (1 - 0.7)]$) (i.e., only a 3% error rate, calculated based on Eqn. 3).

III. JOINT OPTIMIZATION OF MULTIPLE CONTEXTS

Pervasive computing environments that support multiple context-aware applications entail multiple needs for context information, necessitating simultaneous determination of multiple context metrics from shared underlying sensor streams. Recognizing multiple contexts simultaneously from underlying sensor data streams requires selection of the best subset of the sensors along with their optimal tolerance ranges in a way that satisfies the applications' $QoINF$ functions for each of the needed context metrics with a minimum incurred overhead. As sensor networks become ubiquitous, they will increasingly be treated as a platform for multi-modal sensing for multi-context recognition problems. This diverges from the current paradigm of

designing an individual context recognition framework from a specified set of sensors. However, sharing the same set of sensors across multiple contexts may result the reduction of an individual context metric's $QoINF$ accuracy (in comparison to the $QoINF$ that is achievable for that context metric in isolation). To gain a deeper understanding, we formalize the simultaneous determination of multiple contexts as a multi-objective optimization problem. In particular, we propose a linear time sensor selection multi-context search heuristic algorithm to compute the best set of sensors to be chosen and their associated tolerance ranges. We also outline a brute-force algorithm for multiple context recognition and compare the time complexity to our heuristic.

A. Multi-Context Optimization Problem

Given a set θ of sensors, the optimization problem for a single context variable is to choose the sensor data values of $q_1, q_2, \dots, q_\theta$, that, when used together to infer the value of a context variable C , minimize the total cost while ensuring the application-specified minimum accuracy level of $QoINF$:

$$\min COST(\theta, q_\theta) \text{ subject to: } QoINF_C(\theta) \geq QoINF_{min}$$

The above joint optimization of (θ, q_i) applies to the efficient determination of a single context variable C , which we have investigated in [12]. Considering the problem of simultaneously determining l separate context variables, $\{C(1), C(2), \dots, C(l)\}$, we propose the following Lagrangian Optimization problem:

$$\text{minimize } \sum_{s_i \in \theta} \frac{h_i}{q_i^2} + \sum_{l=1}^L \lambda_l [QoINF_{C(l)}(q_1, q_2, \dots, q_\theta) - QoINF_{min}^l] \quad (4)$$

There are different sensitivity constant values ν_{il} and η_{il} for each of the context l being sensed as each context has a different sensitivity to the tolerance range of a sensor s_i . As the tolerance range q_i is identical across all contexts, a sensor will report only when its new value diverges from the previous value by q_i . If the individual $QoINF_{C(i)}(.)$ functions have the inverse exponential form, then the collective optimum values of $\{q_i\}$ may be explicitly computed, as shown below. Considering the $QoINF$ functions for all of the context metrics are the same but with different sensitivity constant values, we get:

$$\text{minimize } \sum_{s_i \in \theta} \frac{h_i}{q_i^2} + \sum_{l=1}^L \lambda_l \left[1 - \prod_{s_i \in \theta} \left[\frac{1}{\nu_{il}} \exp\left(\frac{-1}{\eta_{il} q_i}\right) \right] - QoINF_{min}^l \right] \quad (5)$$

Lemma 1: If the $QoINF(.)$ functions for the set of l contexts and the set of sensors θ satisfies Equations (2) and (3), then the optimal choices of q_i that minimize the cost function for context l follow the relationship:

$$\frac{\frac{2h_1}{q_1^2}}{\sum_{l=1}^L [\log(1 - QoINF_{min}^l) + \log(\nu_{1l})]} = \frac{\frac{2h_1}{q_1^2}}{\sum_{l=1}^{L-1} [\log(1 - QoINF_{min}^l) + \log(\nu_{1l})]}$$

$$\begin{aligned} & \frac{\frac{2h_1}{q_1^2} + \frac{2h_2}{q_2^2}}{\sum_{l=1}^L [\log(1 - QoINF_{min}^l) + \log(\nu_{1l}) + \log(\nu_{2l})]} - \\ & \frac{\frac{2h_1}{q_1^2} + \frac{2h_2}{q_2^2}}{\sum_{l=1}^{L-1} [\log(1 - QoINF_{min}^l) + \log(\nu_{1l}) + \log(\nu_{2l})]} = \dots \\ & = \frac{\frac{2h_1}{q_1^2} + \frac{2h_2}{q_2^2} + \dots + \frac{2h_\theta}{q_\theta^2}}{\sum_{l=1}^L [\log(1 - QoINF_{min}^l) + \sum \log(\nu_{\theta l})]} - \\ & \frac{\frac{2h_1}{q_1^2} + \frac{2h_2}{q_2^2} + \dots + \frac{2h_\theta}{q_\theta^2}}{\sum_{l=1}^{L-1} [\log(1 - QoINF_{min}^l) + \sum \log(\nu_{\theta l})]} \end{aligned}$$

and the optimal value of q_i (the tolerance range to assign to sensor $s_i \in \theta$) across all contexts is given by:

$$\hat{q}_i = \frac{h_i * \left[\sum_{l \in L} \sum_{s_i \in \theta_l} \frac{1}{h_i * \eta_{il}} \right]}{\sum_{l \in L} [\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta_l} \log(\nu_{il})]} \quad (7)$$

and the minimal cost to achieve the specified inference accuracy using these q_i is:

$$COST(\theta) = \frac{\left(\sum_{l \in L} [\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta_l} \log(\nu_{il})] \right)^2}{\sum_{l \in L} \sum_{s_i \in \theta_l} \frac{1}{h_i * \eta_{il}^2}} \quad (8)$$

Proof: Proof of Lemma 1: The multi-context optimization problem can be stated as:

$$\min \sum_{s_i \in \theta} \frac{h_i}{q_i^2}$$

subject to L different constraints, where the l^{th} constraint is:

$$1 - \prod_{s_i \in \theta} \frac{1}{\nu_{il}} \exp\left(-\frac{1}{\eta_{il} q_i}\right) \geq QoINF_{min}^l \quad (9)$$

Taking the logarithm of each constraint and setting up the Lagrangian, we obtain:

$$\sum_{s_i \in \theta} \frac{h_i}{q_i^2} + \sum_{l \in L} \lambda_l * \left(\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}) + \sum_{s_i \in \theta} \frac{1}{\eta_{il} q_i} \right) \quad (10)$$

Taking derivative of these with respect to each q_i , we get:

$$\sum_{s_i \in \theta} \frac{2h_i}{q_i^3} + \sum_{l \in L} \sum_{s_i \in \theta} \frac{\lambda_l}{\eta_{il} q_i^2} = 0 \quad (11)$$

which yields:

$$q_i = \frac{2h_i}{\sum_{l \in L} \frac{\lambda_l}{\eta_{il}}} \quad (12)$$

Further taking derivatives of the Lagrangian with respect to λ_l , we have:

$$\sum_{l \in L} \left(\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}) + \sum_{s_i \in \theta} \frac{1}{\eta_{il} q_i} \right) = 0. \quad (13)$$

In order to derive the values of q_i , we solve together the set of L equations. (Eqn. 13) and set of $|\theta|$ equations. (Eqn. 11). Rewriting Eqn. (11) as follows:

$$\sum_{s_i \in \theta} \frac{2h_i}{q_i^2} + \sum_{l \in L} \lambda_l \sum_{s_i \in \theta} \frac{1}{\eta_{il} q_i} = 0 \quad (14)$$

and replacing $\sum_{s_i \in \theta} \frac{1}{\eta_{il} q_i}$ using Equation (13), we get:

$$\begin{aligned} (6) \quad & \sum_{s_i \in \theta} \frac{2h_i}{q_i^2} - \sum_{l \in L} \lambda_l \left(\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}) \right) = 0 \\ & \Rightarrow \sum_{l \in L} \lambda_l = \frac{\sum_{s_i \in \theta} \frac{2h_i}{q_i^2}}{\sum_{l \in L} (\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}))} \end{aligned}$$

For the base case when $\lambda = 1$,

$$\lambda_1 = \frac{\sum_{s_i \in \theta} \frac{2h_i}{q_i^2}}{\log(1 - QoINF_{min}^1) + \sum_{s_i \in \theta} \log(\nu_{i1})} \quad (15)$$

Similarly for λ_2 ,

$$\begin{aligned} \lambda_2 &= \sum_{l=1}^2 \lambda_l - \sum_{l=1}^1 \lambda_l \\ \lambda_2 &= \frac{\sum_{s_i \in \theta} \frac{2h_i}{q_i^2}}{\sum_{l=1}^2 (\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}))} \\ &\quad - \frac{\sum_{s_i \in \theta} \frac{2h_i}{q_i^2}}{\log(1 - QoINF_{min}^1) + \sum_{s_i \in \theta} \log(\nu_{i1})} \\ &\quad \vdots \end{aligned}$$

and for λ_L ,

$$\begin{aligned} \lambda_L &= \sum_{l=1}^L \lambda_l - \sum_{l=1}^{L-1} \lambda_l \\ \lambda_L &= \frac{\sum_{s_i \in \theta} \frac{2h_i}{q_i^2}}{\sum_{l=1}^L (\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}))} \\ &\quad - \frac{\sum_{s_i \in \theta} \frac{2h_i}{q_i^2}}{\sum_{l=1}^{L-1} (\log(1 - QoINF_{min}^l) + \sum_{s_i \in \theta} \log(\nu_{il}))} \end{aligned}$$

In addition to finding the minimum cost for a given θ , we also need to determine the best subset of sensors, θ , that minimizes the overall update cost across all the contexts. Clearly, one brute-force approach is to iterate through all possible combinations, computing $COST(\cdot)$ for each combination of sensors for all contexts. ■

B. Brute-Force Sensor Selection for Multiple Contexts

Fig. 2 presents the brute-force algorithm for multiple contexts. The brute-force algorithm forms all possible combinations of sensors from the set S (i.e., it generates the power set of S). The algorithm iterates over all the contexts starting

```

Procedure Brute_Force_Multi_Fusion (input set  $S$ ,  $QoINF_{min}^l$ 
 $\forall l = (1, \dots, L)$ )
1. Initialize empty set of sensors;  $\theta = \phi$ ;  $\widehat{COST}(\theta) = 0$ ;
 $MinCost(l) = \infty$ ;
2. Find the power set of  $S$  //Assume  $\theta \subseteq \mathcal{P}(S)$ 
3. For ( $i = 1$ ;  $i < |\mathcal{P}(S)|$ ;  $i++$ )
4.   For ( $l = 1$ ;  $l \leq L$ ;  $l++$ )
5.     Compute the update cost  $\widehat{COST}(\theta)$  for  $QoINF_{min}^l$ 
6.     if  $(\widehat{COST}(\theta) - MinCost(l)) < 0$ 
7.        $MinCost(l) = \widehat{COST}(\theta)$ .
8.   End-For
9. End-For
10. return  $\{\theta, \widehat{COST}(\theta)\}$ .

```

Fig. 2. Sensor Set Selection Brute Force Algorithm for Multiple Contexts

from a specific combination of sensors to determine its cost and application-specified $QoINF$ satisfiability. In this way, it determines the the minimal possible cost ($\widehat{COST}(\cdot)$) with a particular subset of sensor which might be good for a subset of contexts. But our goal is to determine an unified set of sensors across all the contexts simultaneously with the minimal cost and achievable $QoINF_{min}^l$. We repeat this for all combinations of sensors and find out the optimal update cost $\widehat{COST}(\theta)$ and the optimal subset of the sensor θ which hold across all the contexts. However, because this brute-force search is excessively inefficient for determining the best subset of sensor in practical settings, we have developed a selection heuristic for multiple contexts.

C. Sensor Selection Heuristic for Multiple Contexts

Here we propose a selection heuristic for multiple contexts $\{C(1), C(2), \dots, C(L)\}$. The heuristic is based on the observation that the additional cost (based on Eqn. 8) of adding a sensor s_i to an existing set θ is dependent on the term $\log^2(\nu_{il}) * h_i * \eta_{il}^2$. This can be derived from Eqn. (8) by considering the limiting case of the $QoINF_{min}^l$ value. This term can be generalized so that each sensor can have a different sensitivity and hop count factor for each context $C(l)$; $\forall l = (1, \dots, L)$ as $[\log^2(\nu_{il}) * h_i * \eta_{il}^2]$; that is we can allow each sensor to be in a different distance from each of the l context fusion centers, and we can account for the fact that decreased quality in an individual sensor stream can have different degrees of impact on the determination of different contexts. Since a lower value of this term indicates a greater preference for selecting a sensor, the selection heuristic sorts the available sensor set S in ascending order of this term for each context generating L sorted lists, \mathfrak{R} . Our goal is to incrementally select a subset of sensors, iteratively considering additional context metrics and adjusting the selected set to continue to satisfy the growing set of considered context metrics. We first select a single context metric ($C(1)$) and find the subset of sensors from S such that $C(1)$'s $QoINF$ function can be satisfied with least cost. We then incrementally consider additional context metrics (i.e., $C(2)$ is considered next) and determine whether the set of sensors selected for determining $C(1)$ can also satisfactorially determine $C(2)$. If not; sensors are added to the selected set, enabling determination of the newly added context metric given its $QoINF$ function at the least cost. This process continues until all of the contexts have been considered. Fig. 3 shows the pseudocode for this search heuristic for multiple contexts.

```

Procedure Heuristic_Multi_Fusion (input set  $S$ ,  $QoINF_{min}^l$ 
 $\forall l = (1, \dots, L)$ )
1. Initialize empty set of sensors;  $\theta = \phi$ ;  $MinCost(l) = \infty$ ;
2. Sort the sensor set  $S$  into a list  $\mathfrak{R}_l$ ,  $\forall l = (1, \dots, L)$  in
increasing order of sensitivity  $\log^2(\nu_{il}) * h_i * \eta_{il}^2$ 
3. For ( $l = 1$ ;  $l \leq L$ ;  $l++$ )
4.   For ( $i = 1$ ;  $i \leq |S|$ ;  $i++$ ) // cardinality of set  $S$ 
5.      $\theta = \theta + \mathfrak{R}_l(i)$ ; //set-theoretic addition
6.   For ( $l = l + 1$ ;  $l \leq L$ ;  $l++$ )
7.      $\mathfrak{R}_l = \mathfrak{R}_l - \theta$ ; //set-theoretic subtraction
8.   Compute the optimal update cost  $\widehat{COST}(\theta)$ 
for  $QoINF_{min}^l$ 
9.   if  $(\widehat{COST}(\theta) - MinCost(l)) > 0$  OR below threshold
break;
10.  else  $MinCost(l) = \widehat{COST}(\theta)$ .
11.  End-For
12.  End-For
13.  End-For
14. End-For
15. return  $\{\theta, \widehat{COST}(\theta)\}$ .

```

Fig. 3. Sensor Set Selection Heuristic Algorithm for Multiple Contexts

D. An Example and Computational Complexity

To understand our proposed algorithm and its time complexity we consider a simple example. Consider a set of five sensors $S = \{1, 2, 3, 4, 5\}$ used to determine three different contexts $C(1), C(2)$ and $C(3)$; $\forall l = (1, \dots, 3)$. Our goal is to determine the subset of sensors that minimizes the overall update cost. We sort the set of available sensors based on $[\log^2(\nu_{il}) * h_i * \eta_{il}^2]$ and generate three different lists (one for each context) as follows: $R_{C(1)} = \{1, 2, 4, 5, 3\}$; $R_{C(2)} = \{5, 4, 3, 2, 1\}$ and $R_{C(3)} = \{3, 2, 4, 1, 5\}$. For list $R_{C(1)}$ assume the optimal choice of sensor set θ is $\{1, 2\}$ with minimum cost, these two sensors together can achieve the specified $QoINF$, and including any more sensors increases the update cost. After selecting sensor 1 and 2, we want to see what other sensor we need to add to the sensor set (if anything) to satisfy context $C(2)$ with the required $QoINF$ accuracy. In each step we exclude the already chosen sensor from rest of the lists and thus in this case generate new list $R'_{C(2)} = \{5, 4, 3\}$ and $R'_{C(3)} = \{3, 4, 5\}$ with fewer sensors, remaining that are available to support these additional context metrics. Considering the modified next list $R'_{C(2)}$, assume we have optimal set θ as $\{1, 2, 5\}$ with minimal cost. So the new list $R''_{C(3)}$ becomes $\{3, 4\}$. Considering $R''_{C(3)}$, assume we have the optimal sensor set $\theta = \{1, 2, 5, 3\}$ with minimal cost. This example clearly states how our heuristic works with a linear running time. Thus the time complexity of the heuristic is $O(n)$.

In contrast, in the case of the brute-force search algorithm, we have to iterate over the power set of S , which, in this example, contains $2^5 - 1$ elements, across three different contexts $C(1), C(2)$ and $C(3)$. Starting from a singleton sensor set we assume that sensor set $\{1, 2\}$ is good for context $C(1)$, but not satisfactory for context $C(2)$ and $C(3)$ simultaneously. After iterating all set of three sensor across all contexts we find sensor set $\{1, 2, 5\}$ simultaneously satisfies context $C(1)$ and $C(2)$, but not context $C(3)$ with the required conditions. Next iterating all combinations of four and five sensor across all contexts, we find out optimal sensor set $\theta = \{1, 2, 5, 3\}$ which holds across the three contexts with minimal cost and satisfies $QoINF$ accuracy. Obviously, the time complexity of the brute force algorithm grows exponentially as the number of sensors increases.

IV. EXPERIMENTAL SETUP AND RESULTS

To empirically examine the validity of our approach and understand the interplay between tolerance ranges ($\{q_i\}$) and interfering error, we have experimented with SunSPOT sensors (Sun Small Programmable Object Technology devices [13]). The SunSPOT sensor board contains a 3-axis accelerometer, a light sensor, and a temperature sensor. For our studies, we have utilized readings from the accelerometer (to estimate motion and orientation), light sensor, and a Dual-axis Mems Gyro sensor to characterize multiple contexts (in this case *sitting*, *walking* and *running*). Fig. 4 shows the gyro sensor on the left and the wiring of the gyro with the SunSPOT sensor board on the right. We used this setup to evaluate a medical monitoring application that infers a patient's *activity* using a variety of sensor types.

A. Empirical Determination of Multiple Contexts

We used the SunSPOT's accelerometer to measure the tilt of the SunSPOT (in degrees) when the individual user was in three different context states: *sitting*, *walking* and *running*. From the collected samples, we computed the 5th and 95th percentile of the tilt readings corresponding to each state. Table I shows the resulting ranges in the accelerometer tilt readings observed for each of the three states. Such results indicate that there is indeed an observable separation in the ranges of the tilt values for the three different states. This suggests that the states can be distinguished with reasonable accuracy even under moderate uncertainty in the sensor's readings.

TABLE I
CALIBRATED ACCELEROMETER SAMPLE VALUES

Range (5 th – 95 th percentile) of Tilt Values (in degree)	Context State
85.21 to 83.33	<i>Sitting</i>
68.40 to 33.09	<i>Walking</i>
28.00 to –15.60	<i>Running</i>

Similarly, we used the SunSPOT light sensor to measure the light levels for different contexts. Intuitively, low values of light intensity may indicate a *sleeping* state, while higher values are likely to result in the *active* state. Table II shows the observed ranges for light values for these two states. The accuracy of context from the light sensor is, however, much lower, as users may often be inactive (e.g., *sitting*) in the light. We have also

TABLE II
LIGHT SENSOR SAMPLE VALUES (LUMEN)

Average Range of Light level (lumen)	Context State
LightSensor.getValue() = 10 to 50	Turned on → active
LightSensor.getValue() = 0 to 1	Turned off → sleeping

used the external gyro sensor to measure the variation of the rate of angular rotation when the user is walking (active), or sitting or sleeping (inactive). As the gyro is quite sensitive, we use a calibration function to measure the average angular rate variation range for both the active and inactive states, as shown in Table III.

B. Methodology

We emulated an activity monitoring scenario deployed by wellness management professionals to monitor the daily activities of a user on real collected traces from SunSPOT wireless sensors. The application specifies a quality-of-inference

TABLE III
GYRO SENSOR SAMPLE VALUES

Average Angular Rate of Rotation (degs/sec)	Context State
X-Rotational Variation (Xout) = 150 to 350 Y-Rotational Variation (Yout) = 150 to 350	Walking → active
X-Rotational Variation (Xout) = 2 to 8 Y-Rotational Variation (Yout) = 2 to 8	Sitting → inactive

accuracy, exercising our context inference and *QoINF* models. We recruited participants and instrumented them with SunSPOT sensors whose data was logged onto a laptop through a SunSPOT base station. We mounted the sensor on the wrist of each individual participant. To study the potential impact of our approach, we collected initial traces for the SunSPOT motion, light, and gyro sensors for five users, who engaged in a mix of three activities (*sitting*, *walking* and *running*) for a period of three days. Our users were all adults with no known serious medical conditions but with differing levels of physical fitness. We then used an emulator to mimic the samples that a sensor would have reported given the trace and a given q and compared the context inferred from the values reported by the emulation against the ground truth. This trace-driven event-based approach allows us to make meaningful comparisons, as the uncontrollable physiological and environmental variations would otherwise make it impossible to obtain the exact same data stream from two different sessions from the same user.

C. Role of Tolerance Range on Energy and *QoINF* Accuracy

To define the power consumption of SunSPOT sensors for our application, we note the average power of the eSPOT board (95 mA) and application daughter board (400 mA) in run mode and wireless CC2420 radio transmission power (18 mA) [13]. Based on Eqn. (16) we find out the average power consumption of each of the sensor for a specific communication frequency.

$$\text{Power} = \text{Power of eSPOT board} + \text{Power of daughter board} +$$

$$(\text{Radio Transmission Power} \times \text{Communication Frequency}) \quad (16)$$

Figure 5 shows the sensor power consumption and the corresponding *QoINF* achieved (defined as $1 - \text{error rate}$), for different values of the tolerance range (q_m) for the motion sensor. Fig. 6 and Fig. 7 plot the corresponding values versus the tolerance ranges for the light and gyro sensor respectively. As these figures demonstrate, in general there is a continuous drop in the power consumption and the *QoINF* as q increases, for all three types of sensors. However, as seen in Fig. 5, for the motion sensor, the *QoINF* accuracy of $\approx 81\%$ is achieved for a q value of 20; moreover, using this tolerance range reduces the sensor power consumption by $\approx 63\%$ ($3.65 \times 10^4 \text{mA} \rightarrow 1.35 \times 10^4 \text{mA}$). This suggests that it is indeed possible to achieve significant savings in network energy expense, if one is willing to tolerate marginal degradation in accuracy. A similar behavior is observed in Fig. 6 for the light sensor where $q = 4$ incurs a 5% loss in accuracy vs. $\approx 62\%$ ($2.1 \times 10^4 \text{mA} \rightarrow 0.81 \times 10^4 \text{mA}$) reduction in power consumption. However, as the difference between the lumen ranges for *Active* versus *Sleeping* is only ≈ 10 (Table II), increasing q beyond ≈ 10 leads to a sharp fall in *QoINF*. Similarly, for the gyro sensor results shown in Fig. 7, increasing q leads to a decrease in power consumption though the *QoINF* accuracy remains constant $\approx 49\%$ after it crosses the




Fig. 4. Single Chip Dual-Axis Gyro Sensor with SunSPOT

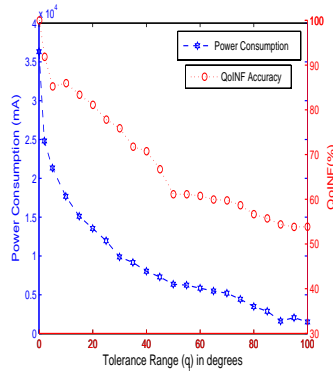


Fig. 5. Power Consumption & Accuracy vs. Tolerance for the Motion

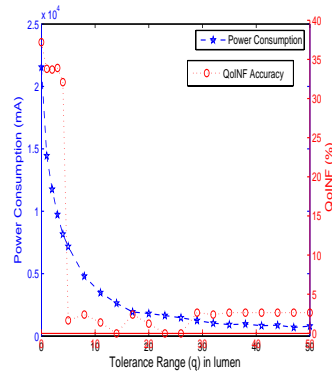


Fig. 6. Power Consumption & Accuracy vs. Tolerance for the Light




Fig. 7. Power Consumption & Accuracy vs. Tolerance for the Gyro

ranges for *Inactive* ≈ 8 . For $q = 8$, *QoINF* accuracy remains $\approx 49\%$, whereas power consumption got reduced to almost half ($3.0 \times 10^4 mA \rightarrow 1.6 \times 10^4 mA$).

From these results we can visualize the application of our multi-context model properly. Consider an application specifies $QoINF_{min}$ as $\approx 80\%$ and interested in recognizing all the contexts like (*sitting, walking and running*) simultaneously with an objective of reducing the energy cost as well. We can conclude that use of motion sensor would be a good choice in this case with a tolerance range of $q = 20$ with an achievable *QoINF* accuracy $> 80\%$ and with a cost (energy) reduction of $\approx 63\%$.

D. *QoINF* Accuracy & Sensor Overheads for Multiple Users

We also investigated the tradeoff between tolerance range and *QoINF* for five users in our experiments. The goal was to study the sensitivity of the tradeoff to individualized activity patterns; we focused on results from a single sensor (the SunSPOT motion sensor). Once we collected the traces, we replayed them through the emulator as before. In each session the five participants were engaged in a mix of three different activities (*sitting, walking and running*), again over the course of three days. Figs. 8 and 9 depict the variation, across users, in the communication overhead and inferencing accuracy, respectively, as a function of the tolerance range. There are clearly significant differences (especially in the *QoINF* accuracy) across users; in particular, user 2 has a much sharper drop in *QoINF* once the tolerance range exceeds 40 degrees.

The figures suggest that *personalization* of the *QoINF* function might be an important step towards maximizing the benefits of inference-quality aware sensing. However, even in the absence of personalization, the benefits from quality-aware context inference are significant. For example, if a tolerance range of 20 is applied to all users, the lower bound of the accuracy achieved is $\approx 71\%$ (for User 3); at the same tolerance range, the worst case (smallest) reduction in the reporting overhead is observed to be $\approx 60\%$ (for User 4).

E. Applicability of the Proposed Heuristic Algorithm

These initial results demonstrate that significant savings can be achieved by relaxing the tolerance of each sensor without compromising the accuracy of context estimation. To validate and quantify these benefits, we apply our results to our formal *QoINF* model. First we fit the *QoINF*(.) accuracy versus q




Fig. 8. Communication Cost vs. Tolerance for Motion Sensor and Multiple Users




Fig. 9. Accuracy vs. Tolerance for Motion Sensor and Multiple Users

curves from Figs. 5, 6 and 7 to the inverse exponential model of Eqn. (2). After obtaining the best parametric fit (using a linear least squares regressor), we use the heuristic to compute the q values for a target $QoINF_{min}$ value and then use additional traces to verify if this approach can provide the required accuracy.

We further compare the performance of our heuristic with the brute-force search technique. Recall that, while this technique still uses the least-square estimator to first compute the “best” inverse-exponential *QoINF* function, it uses an exhaustive search over the $2^n - 1$ possible combinations (in this case $2^3 - 1 = 7$) where n is the number of sensors to determine the best *QoINF* vs. q tradeoff. Table IV shows the estimated coefficients (η_i and ν_i) for each of the sensors for User 1, based on the empirical data. We then use the heuristic and brute-force algorithms to compute the optimal sensor set ($\hat{\theta}$) and associated tolerance ranges $Q(\hat{\theta})$ that minimize the communication overhead for a target *QoINF*.

TABLE IV

η_i AND ν_i VALUES OBTAINED BY CURVE FITTING

Sensor	η_i	ν_i
Motion	0.0512	1.8474
Light	3.6553	1.2288
Gyro	17.2512	1.8488

Fig. 10 compares the cumulative theoretical computational cost of the heuristic and brute-force search methods to find the minimum cost for different values of $QoINF_{min}$. The figure illustrates that the heuristic solution always incurs lower cumula-

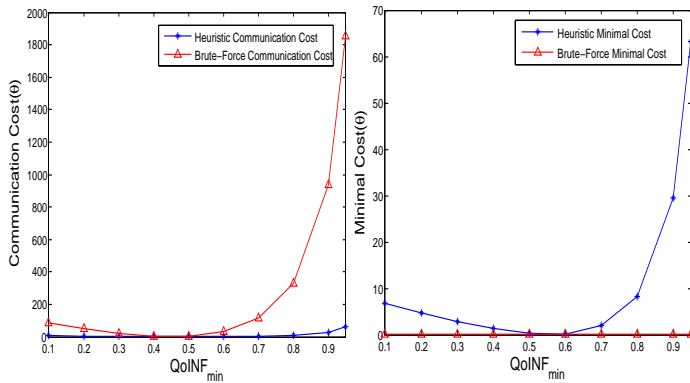


Fig. 10. Cumulative Computational Costs of Heuristic and Brute-Force




Fig. 11. Heuristic & Brute-Force Minimal Cost Comparison

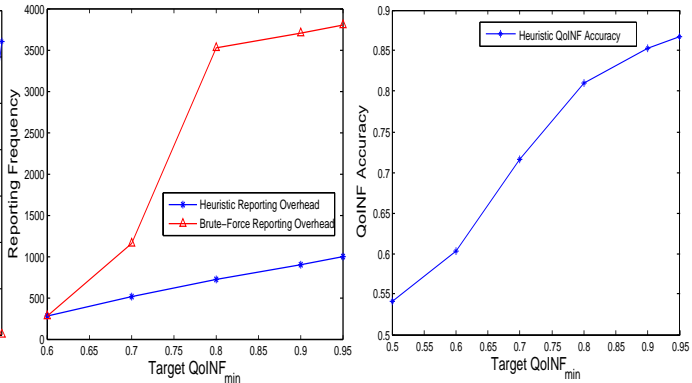


Fig. 12. Heuristic & Brute-Force Communication Overhead Comparison

tive computational cost than the brute-force search. It converges much more quickly than the brute-force when the target $QoINF$ value is high or moderately high (≥ 0.6). At higher values of $QoINF$, our heuristic search can result in a substantial decrease in computational cost (about 97% for $QoINF=0.9$), compared to the brute-force approach.

We also plot the minimum cost for heuristic and brute-force search for different values of $QoINF_{min}$ in Fig. 11 where the brute-force minimum cost outperforms the heuristic in all cases as it searches the entire search space to find the minimal cost. We find this cost from Fig. 10, but here instead of cumulative search cost we just plot the minimum cost point. Particularly, we also note in our experiments, in case of heuristic the optimal set θ remain as $\{motion\}$ sensor until $QoINF_{min} < 0.70$ as the inclusion of other sensor incur an increase in cost. Beyond that it selects both motion and light sensors. Whereas in case of brute-force at the minimum cost point, the optimal subset of sensor θ remains $\{motion\}$ sensor until $QoINF_{min} \leq 0.50$. The optimal set θ becomes $\{motion, light\}$ when $0.50 < QoINF_{min} < 0.70$ and beyond that $QoINF_{min}$, optimal set θ becomes $\{motion, light, gyro\}$.

More importantly, we compare the behavior of the two search algorithms, in terms of their update overhead (cost) and their ability to accurately attain the target $QoINF$ objective at the minimum cost. For this purpose, for a given $QoINF_{min}$, first we calculate the tolerance range q for all sensors and the minimum cost point, followed by the determination of the optimal subset θ associated with that minimum cost. Using the determined θ and individual q values, we find out the update overhead. Fig. 12 plots the update overhead against that observed empirically when the sensor data is subject to the θ , $Q(\theta)$ values suggested by the Heuristic or Brute-Force search algorithms. Heuristic always stops if there is an increase in cost on adding a sensor to θ , but brute-force keeps on adding sensor and compares all the cost value to find the minimum one. Leveraging all the options empower the brute-force to form a larger optimal subset θ in comparison with the heuristic as we previously observed. Thus brute-force update overhead becomes more than the heuristic as shown in Fig. 12. Fig. 13 plots the empirically observed $QoINF$ values for the Heuristic algorithm verses the target $QoINF$ values. We can see that the heuristic approach is able to approximate the target objectives fairly well; in particular, the inference accuracy observed by the heuristic is no more

than 5% lower than the target $QoINF$ value. These plots provide compelling evidence that a model-based, heuristic approach towards $QoINF$ -aware sensing can meet an application's $QoINF$ objectives, while significantly reducing the sensor communication overhead and energy expense.

V. ONGOING WORK

In this section, we highlight the open issues with our proposed model-based $QoINF$ -aware sensing mechanism. We are addressing these issues in our ongoing work.

A. Parametric Estimation of $QoINF(.)$ Coefficients

One of the main challenges in the application of our suggested formalism is the establishment of appropriate $QoINF_C(.)$ functions for specific context variables. Indeed, much of the work on utility-based context models has failed to achieve the desired impact due to the difficulty of computing useful utility functions. To overcome this challenge, we have used *statistical regression* to obtain the best coefficients for a pre-selected $QoINF(.)$ function from the empirically observed data. However, the appropriate initial choice of a suitable $QoINF$ function is itself an open question. While ideally, we would like a common $QoINF$ function to model multiple sensor-to-context mappings, it is possible that the appropriate $QoINF$ function be different for different contexts. Another challenge relates to how to vary the tolerance ranges $q_i(.)$ to obtain the full specification of the $QoINF(.)$ function. In many pervasive computing environments (e.g., medical contexts), it is necessary to ensure that context estimation errors do not become too large. While occasional errors in deducing certain contexts may be tolerable, it is necessary to ensure that the errors do not occur at critical times. Accordingly, we believe that *incremental learning techniques* will be useful. Such adaptive techniques progressively probe the state space of $\{q_i\}$ values, while ensuring that high tolerance ranges for certain sensors are only specified when alternative sensors are unlikely to provide the required accuracy.

B. Discussions

In the following we discuss some of the interrelated issues and underpinnings of our $QoINF$ -aware model.

Usability: The minimum $QoINF$ level can be specified for a particular application by the external stake holders such as wellness professionals. But based upon the situational requirements, multiple optimization cycles can be triggered with a higher or

lower $QoINF$ level. This can be represented using traditional database queries [3]. In the future we would like to investigate active learning techniques and feedback control systems to make this $QoINF$ specification more dynamic and adaptive in nature.

Rationale of the Distribution: The basic premise behind our $QoINF$ distribution (i.e., the negative exponential) is to determine the optimal subset of sensors with their associated tolerance ranges. The choice of this distribution has been empirically validated with real traces using the curve-fitting approach ensuring no less than 5% of targeted $QoINF$ accuracy. Similar work can be found in [10], where a decision fusion rule based on counting policy has been proposed based on a Poisson sensor distribution model. Another decision fusion algorithm for a large scale sensor networks is formulated in [8], where statistical dependence and independence of the sensors has been exploited.

Limitations: In our experiments, we collected data traces from the users simultaneously being in multiple context states and thus ran our emulator to determine the sensitivity factors using $QoINF$ and q curves through the least square regression technique. Thus we had only one set of sensitivity factors across all the contexts for each sensor. We are currently collecting data for each context state distinctly with a goal to determine the sensor sensitivity factor associated with each context state separately. We also intend to use a larger set of sensors to validate the proposed algorithm.

VI. RELATED WORK

The tradeoff between communication overhead and the quality of reconstructed data was first studied in [11], which envisioned the effect of tolerance ranges on the frequency of sink-initiated fetching vs. source-initiated proactive refreshing. The focus, however, was on snapshot queries and not on *continuously satisfying* the $QoINF$ bound of a long-standing subscription. The idea of exploiting temporal correlation across successive samples of individual sensors for reducing communication overhead for snapshot queries is addressed in [3], which used training data to parameterize a jointly-normal density function. While a precursor to our work, the focus there was on meeting the $QoINF$ requirements for a class of aggregation queries, whereas we focus on arbitrary relationships between a context variable and the underlying data. Entropy-based sensor selection heuristic algorithms are also proposed in [4], [9], [15]. In these literature, an information theoretic approach is taken for specific application scenarios, where the belief state of the target value is gradually improved by repeatedly selecting the most informative unused sensor until the required accuracy is achieved. The CAPS algorithm [6] is designed for long-running aggregation queries (such as $\{min, max\}$) and computes the optimal set of tolerance ranges for a given set of sensors that minimizes communication overhead while guaranteeing the accuracy of the computed response. In contrast, our objective is to compute *both* the best subset of available sensors and their tolerance ranges to achieve the desired accuracy for arbitrary context variables.

QoI (quality of information) has been studied before within the context of data collection and storage for database systems, with a focus on data consistency, completeness, currency, etc [11]. However, the development of formal models of quality of information (or inference) has not been extensively addressed in the context of sensor-generated data streams. Recent

work [16] has suggested the use of specialized QoI models to capture the accuracy of detecting transient events, given a set of sensors. However, this approach does not focus on the optimal joint selection of *sensors and their tolerance ranges*, which is the distinguishing characteristic of our work.

VII. CONCLUSIONS

We have presented a formal framework for energy-efficient determination of contexts in pervasive computing environments. The key idea is to express the accuracy of context estimation, for arbitrary contextual attributes, through quality of inference ($QoINF$) function that captures the dependence of estimation accuracy on the selected sensors and their specified tolerance ranges. Such a multi-context model in terms of the uncertainty range of the underlying sensor data has not been rigorously investigated in the past and holds promise for reducing the communication overhead of sensor data transmissions. We have outlined a multi-context search heuristic algorithm to solve the proposed optimization problem. Experimental traces collected in the laboratory setting demonstrate the significance of the quality-verses-communication cost tradeoff and establish that our proposed heuristic is able to provide close-to-optimal tradeoff between the $QoINF$ value and the communication overhead, at the cost of only modest computational requirements. Experiments using a combination of motion, light and gyroscopic sensors show that our model-based approach is able to provide a $QoINF$ value that is no more than $\approx 5\%$ lower than the desired target.

REFERENCES

- [1] Aperia Inspiration Communities, http://aperion.typepad.com/aperion-companies_weblog/aperion_health_news/index.html
- [2] Oliver Amft "On the need for quality standards in activity recognition using ubiquitous sensors", *In How To Do Good Research In Activity Recognition. Workshop in conjunction with Pervasive 2010.*
- [3] A. Deshpande, C. Guestrin, S. Madden, J. M. Hellerstein, W. Hong, "Model-based Approximate Querying in Sensor Networks", *Int'l Journal on Very Large Data Bases (VLDB Journal)*, 2005
- [4] E. Ertin, J. Fisher, and L. Potter, "Maximum mutual information principle for dynamic sensor query problems", *In Proc. IPSN03*, Apr 2003.
- [5] W. Heinzelman, A.L. Murphy, H.S. Carvalho and M.A. Perillo, "Middleware to Support Sensor Network Applications", *IEEE Network* 18, pp. 6-14, Feb. 2004.
- [6] W. Hu, A. Misra and R. Shorey, "CAPS: Energy Efficient Processing of Continuous Aggregate Queries in Sensor Networks", *Fourth IEEE Int'l Conference on Pervasive Computing and Communications (PerCom)*, 2006, pp. 190-199
- [7] A. Krause, A. Smailagic, D. Siewiorek, "Context-Aware Mobile Computing: Learning Context-Dependent Personal Preferences from a Wearable Sensor Array," *IEEE Transactions on Mobile Computing*, pp. 113-127, Feb 2006.
- [8] N. Katenka, E. Levina, and G. Michailidis, "Local vote decision fusion for target detection in wireless sensor networks," *Joint Research Conference on Statistics in Quality Industry and Technology*, 2006.
- [9] J. Liu, J. Reich, and F. Zhao, "Collaborative in-network processing for target tracking," *EURASIP JASP: Special Issues on Sensor Networks*, 2003(4):378391, Mar 2003.
- [10] R. Niu and P. K. Varshney, "Distributed detection and fusion in a large wireless sensor network of random size," *EURASIP Journal on Wireless Communications and Networking*, vol. 5, no. 4, 2005, pp. 462-472.
- [11] C. Olston and J. Widom, "Offering a precision performance tradeoff for aggregation queries over replicated data", *In Proceedings of VLDB*. Morgan Kaufmann Publishers Inc., 2000, pp.144-155.
- [12] N. Roy, A. Misra, S. K. Das and C. Julien, "Quality-of-Inference (QoINF)-Aware Context Determination in Assisted Living Environments", *Proc. of ACM SIGMOBILE Workshop on Medical-Grade Wireless Networks (WiMD 2009)* in conjunction with MobiHoc, May 2009
- [13] SunSpotWorld-Home of Project SunSPOT: www.sunspotworld.com
- [14] X. Tang and J. Xu, "Extending Network Lifetime for Precision-Constrained Data Aggregation in Wireless Sensor Networks", *Proc. of IEEE INFOCOM 2006*, Apr 2006.
- [15] H. Wang, K. Yao, G. Pottie and D. Estrin, "Entropy-based sensor selection heuristic for target localization", *Proc. of IPSN 2004*, Apr 2004.
- [16] S. Zahedi, M. B. Srivastava and C. Bisdikian, "A Framework for Quality of Information Analysis for Detection-oriented Sensor Network Deployments", *Proceedings of MILCOM 2008*, Nov. 2008 and Annual Conference of ITA 2008