Quality and Energy Aware Data Acquisition for Activity and Locomotion Recognition

Nguyen Hoang Viet, Ellen Munthe-Kaas, Thomas Plagemann
Department of Informatics
University of Oslo
Oslo, Norway
{*viet, ellenmk, plageman*}@*ifi.uio.no*

***Abstract*—With the advent of wireless sensors and pervasive environments, autonomic human activity recognition has received substantial attention in research. In such environments, many sensors are deployed on each object with the purpose to collect sufficient data to recognize the activities of the object. To perform activity recognition, low-level data streams from the sensors are combined at the sink. A key challenge is to recognize efficiently and with high accuracy the object's activities based on the low-level sensor data. However, there is a trade-off between high accuracy and efficiency, caused by the cost of delivering data samples from sensors to the sink. The challenge is to determine sampling rates that satisfy the required accuracy and minimizes the communication cost. We formalize this problem of choosing sampling rates that satisfy the required accuracy and minimize the communication cost. We formalize this problem as an integer programming problem and solve it by using Lagrangian relaxation with branch-and-bound method. Evaluation results with a publicly available dataset demonstrate the potential applicability of our approach.**

***Keywords*-Quality of Information, Activity Recognition, Energy Consumption.**

I. INTRODUCTION

There has been a rising attention towards human activity recognition using on body, object-placed or ambient sensors, for application areas like health-care, assistive technologies, manufacturing or gaming. These applications apply machine-learning techniques to classify signals gathered from multiple sensors with different modalities. This requires to process with high dimensional, noisy multimodal streams of data with large variability caused by changes in subject's behavior. Therefore, several challenges arise at the different processing stages ranging from feature selection, classification, to decision fusion at a sink.

Due to the limitations of sensors, such as limited battery, bandwidth, and measurement quality, the reported information is usually a distorted version of the ground truth. To estimate this distortion we use a metric, i.e. Quality of Information (QoI), to measure the goodness of the high-level activity recognition derived from low-level data streams from the sensors. The inherent noise in the environment can lead to insufficient QoI. Therefore, it is important to be able to control the QoI. There are two basic approaches to increase the QoI: (1) to increase the number of sensors: this is likely to increase the QoI, and will never decrease it; (2) to use different types of sensors which are sensitive to different types of noise. However, pervasive systems are battery-powered and the total energy consumption in these systems is equal to the sum of the energy used for computation, sensing, and communication. There is a trade-off between energy consumption and QoI from the sensors. Therefore, it is critically important to select for every sensing task a set of sensors that minimizes the cost, yet satisfies the user-specified QoI in a timely and efficient manner. The accuracy is a function of not just a chosen sensor but also the number of collected samples. Given the need to collect information from sensors embedded in pervasive computing environments, a promising approach to reduce the energy usage of sensors relies on reducing the numbers of samples collected at each sensor, or sending a reduced number of samples to the sink. However, it is important to assure that sending a reduced number of samples still leads to the minimum required QoI.

To get more insight into the core idea of adapting the number of samples to gain a user specified QoI, we consider a concrete example of a system made for activity and locomotion recognition in daily life: a person wakes up, prepares and eats breakfast, cleans up the table and dishes. Acceleration, gyroscope, ultrasonic and magnetic field sensors are attached to the person's body. These sensors have different modalities and are concurrently used to recognize an individual activity (e.g. getting up, breakfast preparation, opening the drawer, fetching a bread) and a mode of locomotion (e.g. walking, standing, running). If we want the system to concurrently recognize locomotion with 90% accuracy and an individual activity with 80% accuracy, we will find a solution with the specific number of samples of data from each sensor to contribute to the final fusion. An activity can be identified by the characteristics of the data collected during the activity. The more samples collected periodically, the higher the resolution of the data, which as a consequence leads to higher accuracy of recognition. By

978-1-4244-9529-0/13/$31.00 ©2013 IEEE

increasing or decreasing the number of samples from each sensor, the system can achieve the user specified QoI.

In this paper, we propose a model to control the QoI by adapting the number of collected samples. Continuing our previous work [1], we consider another aspect of system-guaranteed QoI and make two main contributions. First, we formalize the problem of collecting samples from selected sensors in a way that satisfies the minimum QoI of the concurrent activities. We formalize the problem in terms of an integer-programming problem and solve it using Lagrangian relaxation and branch–and–bound method. Second, we develop and evaluate a practical method for the calculation of the QoI function using the F-measure function that maps the number of samples to the QoI for each sensor and uses Bayesian formulation to fuse the total resulting QoI of the achieved recognition. By using a publicly available dataset [2], we show the applicability of our approach in a practical demonstration.

The rest of the paper is organized as follows: Section II defines the sampling and classification process, and proposes a joint optimization problem. Section III describes our Lagrangian relaxation and branch-and-bound algorithm to solve the resulting integer-programming problem. Evaluation is presented in Section IV. Section V reviews related work, followed by the conclusion in Section VI.

II. Problem Formulation

A. Sampling Rate and Sending Rate

Energy-efficient operation is critical in pervasive computing environments. For battery-powered sensors, the energy consumption is determined by the amount of computation, sensing and communication. Reducing energy consumption at the sensors can be achieved through reducing the number of samples collected at the sensors or sending a reduced number of samples to the sink. We aim to determine the minimum number of samples that must be sent to the sink in order to achieve the minimum QoI. The sampling rate defines the number of samples per unit of time taken from a continuous signal to make a discrete signal. The maximum sampling rate is determined by the A/D converter on the sensor. We define the sending rate as the number of samples sent by the sensor, out of 100 generated samples. We use a mask for each 100 samples of data to determine which samples will be sent to the sink. If the sink requires a higher accuracy, it can increase the sending rates by pulling more samples.

B. The Classification Process

We assume a model in which samples from different sensors are transmitted to a sink. At the sink, features from each data stream are extracted and fed into a classifier. We propose to perform fusion at the decision level. This has many advantages, such as having the same representation.

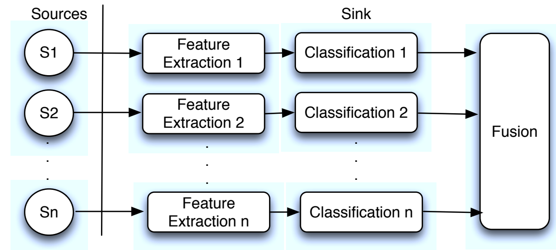

Figure 1. Schematic representation of our classification process

At this level, the values from all sensor streams are probability values, and are therefore independent of the chosen features and sensor data. At the other levels, the values are inhomogeneous, originating from different types of sensors. The schematic representation of our method is shown in Figure 1. For classifiers, we can use any simple and standard classification methods.

C. Sensor QoI Function and Sensor Fusion

The QoI provided by a sensor network is related to the accuracy of the entire classifier process, from collecting data to delivering the final result to the user. The higher the accuracy, the higher the QoI provided. In the literature, different metrics are used for QoI with respect to timeliness and confidence. Bisdikian [3] analyzes the impact of signal and system parameters on QoI, such as the sampling rate, while Hossain et al. [4] propose to use certainty, accuracy, timeliness and integrity as QoI attributes and model them statistically. In our paper, the QoI at each sensor reflects the recognition accuracy as a function of the sending rates. The overall QoI is then computed from the sensor QoIs using a Bayesian formulation.

1) Sensor Fusion: We fuse data streams from all selected sensors to make the final decision. The overall QoI of recognizing a specific acitivity is computed using a form of Bayesian formulation [5], which is given as:

$$QoI(\theta) = \frac{\prod_{s_i \in \theta} qoi(s_i)}{\prod_{s_i \in \theta} qoi(s_i) + \prod_{s_i \in \theta} (1 - qoi(s_i))} \quad (1)$$

$$\text{with } 0.5 \leq qoi(s_i) \leq 1$$

where θ is an arbitrary subset of the sensors. This formulation assumes that the qoi of each sensor is statistically independent of the others. It satisfies the key properties that a $QoI(.)$ function must have:

-The value of the function always falls between 0 and 1.

-By requiring that each $qoi(s_i) \geq 0.5$, the function increases monotonically, i.e., incorporating data from additional sensors does not decrease the QoI.

2) Sensor QoI Function: Let $R_\theta = \{r_1, r_2, ..., r_k\}$ be the collection of sending rates for a selected sensor subset θ that is used for recognizing a single activity. The sending rates are required to be integers since they also represent the numbers of samples in the mask, which need to be sent to the sink. While finding a function that maps the sending samples to QoI requires exponential time, some function forms have proved to be more efficient in fitting inputs to outputs. In our practical case, a QoI function for each individual sensor s_i is represented by an inverse exponential distribution,

$$qoi(s_i) = \alpha_i exp(-\frac{\beta_i}{r_i}) \tag{2}$$

where α_i and β_i are sensitivity constants for sensor s_i. The constants are determined using curve fitting through the least square regression technique. A larger α_i indicates a higher contribution from sensor s_i. In order to meet the restrictions $0.5 \leq qoi(s_i) \leq 1.0$, each sending rate r_i must be chosen from an interval $r_{min} \leq r \leq 100$ where r_{min} can be found using Equation 2 with $qoi(s_i) = 0.5$.

D. The Cost Model

The cost of recognizing a single activity depends on both the sensor sending rate 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 in the uplink path from the sensor to the sink. The cumulative cost function of the set of selected sensors θ is given by:

$$Cost_\theta = \gamma \sum_{s_i \in \theta} h_i r_i \tag{3}$$

where γ is a scaling constant and h_i is the hop count.

E. Joint optimization for concurrent activity detection

In pervasive environments, activity recognition needs to support not only a single activity recognition but also a multiple activity recognition with complex patterns, such as consecutiveness and concurrency. Concurrent activities are activities that happen at the same time, but are reported potentially by different underlying sensors. We formulate the simultaneous recognition of the activities as a multi-objective optimization problem. Given a set of sensors θ, the optimization problem is to choose the sending rate values r_1, r_2,...,r_θ such that, when used together to detect the concurrent activities, we minimize the total cost while ensuring the user-specified accuracy of QoI for each activity:

$Cost(\theta, r_i)$, subject to: $QoI_j(\theta) \geq (QoI_j)_{min}$ with $r_i \in \{(r_i)_{min}, 100\}$ and r_i are integers.

The above multiple-objective optimization of (θ, r_i) is applied to J concurrent activities, where $j \in [1..J]$. There are different sensitivity constant values α_{ij} and β_{ij} for each sensor s_i and concurrent activity j. Given that QoI functions for all sensors and all activities have the same form but with different constant values, we get:

min $\gamma \sum_{s_i \in \theta} h_i r_i$, subject to for all $j \in [1..J]$, $\frac{\prod_{s_i \in \theta} \alpha_{ij} exp(-\frac{\beta_{ij}}{r_i})}{\prod_{s_i \in \theta} \alpha_{ij} exp(-\frac{\beta_{ij}}{r_i}) + \prod_{s_i \in \theta}(1-\alpha_{ij} exp(-\frac{\beta_{ij}}{r_i}))} \geq (QoI_j)_{min}$ with $r_i \in \{(r_i)_{min}, 100\}$ and r_i are integers.

III. Proposed Method

Our problem is an integer programming problem that is classified as NP-complete [6]. To solve our problem we use Lagrangian relaxations to obtain bounds in branch–and–bound algorithms for integer programming.

A. Branch and Bound

The branch–and–bound algorithm is a divide and conquer approach that dynamically constructs a search tree, each node of which represents a sub-problem. We summarize the steps of the branch–and–bound method for finding the optimal integer solution as follows.

1) Find the optimal solution with the integer restrictions relaxed with Lagrangian Relaxation.
2) Build a feasible solution tree. At the first node, the lower bound is the relaxed solution $Cost((r_{relaxed})_i)$ and the upper bound is the rounded-up integer solution $Cost(\lceil (r_{relaxed})_i \rceil)$. The optimal integer solution will be selected between these two bounds.
3) Create two new subsets to eliminate the fractional part of the solution value. Two branches are produced at the variable r_{lf} with the lowest fractional part: one for $r_{lf} \leq \lfloor (r_{relaxed})_{lf} \rfloor$ and the other for $r_{lf} \geq \lceil (r_{relaxed})_{lf} \rceil$.
4) Solve the relaxed integer programming model as Lagrangian relaxation with the new constraints added at each of these nodes.
5) The relaxed solution is the lower bound at each node, and the existing minimum integer solution at any node is the upper bound.
6) If the process produces a feasible integer solution with the lowest lower bound value of any ending node, the optimal integer solution has been reached. If a feasible integer solution is not found, we branch from the node with the lowest lower bound and repeat from step 3.

B. Lagrangian Relaxation

A standard approach for relaxing integer programming problems is to omit the integrality restriction. In our case, this means replacing the integer sending rates with real values. Based on this we propose the following Lagrangian Optimization problem for the steps 1 and 4:

$$minimize\ \gamma \sum_{s_i \in \theta} h_i r_i + \sum_{j=1}^{J} \lambda_j \left[\frac{\prod_{s_i \in \theta} \alpha_{ij} exp(-\frac{\beta_{ij}}{r_i})}{\prod_{s_i \in \theta} \alpha_{ij} exp(-\frac{\beta_{ij}}{r_i}) + \prod_{s_i \in \theta}(1-\alpha_{ij} exp(-\frac{\beta_{ij}}{r_i}))} - (QoI_j)_{min} \right]. \tag{4}$$

To solve this we take the logarithm of each constraint and take derivatives of the Lagrangian with respect to each r_i and λ_j.

IV. Evaluation

A. Experiments

To examine the validity of our approach and understand the relationship between sending rates and QoI, we experiment with an emulator that mimics the samples that a sensor would have reported, given the trace and the sending rates. The classification result is compared to the ground truth. This trace-driven approach allows us to resemble real-world evaluations in a controllable and reproducible manner.

For demonstrating our approach, we use the Opportunity dataset [2] as the input for the emulator. The data was recorded in a highly instrumented environment, including 72 sensors of 10 modalities in 15 wireless and wired networked sensor systems while the subjects are performing morning activities. We use a subset of 6 body-worn accelerometer sensors in the lower part of the right upper arm (RUA_), the left hand (LH), the right lower arm (RLA), the right upper knee (RKN^), the left upper arm (LUA) and the right upper arm (RUA) from the dataset, corresponding to 4 subjects. The aim is to recognize the subjects' gestures and modes of locomotion.

We use Quadratic Discriminant Analysis (QDA) as a supervised learning technique to train the classifier with data that was manually labeled for modes of locomotion, gestures and high-level activities. The data from three axes of each sensor is fed to one classifier. Experiments are performed using both the mean and the variance of the activity time window. We tested our method on one of the four subjects using 10-fold cross-validations.

B. Performance and Energy Consumption Measurement

The QoI can be evaluated by measuring the accuracy of the classifier. There exist different ways to measure the performance of activity recognition. The simplest is the accuracy (i.e. correctly predicted samples/the total number of samples), which is highly affected by the sample distribution across activity classes. Alternatively, the F-measure, taking into account the precision and recall for each class, can give a better assessment of performance. Furthermore, to counter the class imbalance, classes can be weighted according to their sample proportion:

$$F_1 = \sum_i 2w_i \frac{precision_i \cdot recall_i}{precision_i + recall_i} \tag{5}$$

where i is the class index and w_i is the proportion of samples of class i: $w_i = n_i/N$. The ground truth for the measurement comes from annotated activity labels provided by the dataset.

We evaluate the relationship between sending rates, QoI and energy consumption. The power consumers of a sensor node are mainly computing, sensor, and wireless modules. Sample collection mainly affects power consumption of the sensing and wireless modules, but in this paper we focus on the power consumption of the wireless module, i.e., the sensor saves energy by sending a reduced number of samples by request from the sink. If the sink needs more or less samples, it will issue another request with an updated sending rate. We consider for our experiment a general one-hop Body Sensor Network with the star topology. We calculate the total energy consumption E based on the number of transmitted and received packets.

$$E = (B_D + B_H) \cdot (N_{TX} \cdot P_{TX} + N_{RX} \cdot P_{RX}) \cdot T \tag{6}$$

where B_D and B_H are respectively the number of data and header bits in a packet, N_{TX}, N_{RX} are the transmitted and received numbers of packets, T is the bit time in seconds (or the reversed data rate), P_{TX} is the power of the transmitter in mW, and P_{RX} is the power of the receiver in mW. In the dataset, the data was collected using two wireless protocols: Bluetooth and Zigbee. Without loss of generality, we calculate the energy consumption using the Zigbee standard [7], which considers IEEE 802.15.4-compatible transceivers operating at 2.4 GHz with a data rate of 250 kbps, $P_{TX} = 35mW$, $P_{RX} = 35mW$, $B_D = 27$ bytes and $B_H = 16$ bytes.

C. QoI and Sensor Communication Overheads

First, we evaluate the trade-off between QoI, energy consumption and sending rate. We use the emulator to analyze two main tasks: an activity recognition (24 classes of activities, e.g. Open Fridge, Close Door, Reach Cup...) and a locomotion detection (Stand, Walk, Sit, Lie). Figures 2–4 show the power consumption and the corresponding QoI for recognition of activities for different sending rates for the accelerometer sensors in RKN_, LH, LUA. The results for sensors RLA, RKN^and RUA are similar. As figures demonstrate, in general there is a continuous increase in accuracy and energy consumption as the sending rate increases for all of the sensors. Table I shows the estimated coefficients (α_i and β_i) for each sensor of Subject 1 in the activity recognition task. The QoI functions fit quite well to the data from the dataset. The total energy consumption increases linearly as opposed to the nonlinear increase in the QoI function. This suggests that as energy consumption increases, the system will get a decreasing gain in accuracy. For instance, for activity recognition at the sending rate of 20%, we get QoI of 60% instead of 70% for the sending rate of 100% (Figure 3). We can conclude that if we want to sacrifice 10% of accuracy, we can reduce the cost by 80%.

Figure 5 shows the corresponding values for locomotion activity recognition for the sensor in LUA. The figure shows that the data fluctuates slightly, but in general the data follows the distribution function. Thus, we can use the QoI function for the optimization in the next steps. The coefficients (α_i and β_i) for locomotion recognition for each sensor of Subject 1 are listed in Table II.

D. QoI and Sensor Communication Overheads for Multiple Subjects

We continue to investigate the trade-off between the sending rate and QoI for the four subjects. Using a single sensor, we focus on the sensitivity of the trade-off to individual activity patterns. Using our emulator, we replayed the four subjects with different activity recognition and different sending rates. Figure 6 depicts the variation in the communication overhead and the accuracy of the accelerometer sensor in the RUA of the different subjects. There are significant differences across the subjects. In particular, Subject 2 has the lowest accuracy for his activity when using an accelerometer on the RUA. The difference from the three other subjects is about 20%. In conclusion, we can say that the QoI highly depends on the personalization of the subject.

E. QoI and Sensor Overheads for Concurrent Activities

As mentioned, we have investigated QoI and Energy Consumption versus sending rates for two main tasks: activity recognition and modes of locomotion. We used our optimization approach to find a set of sensors and their sending rates to minimize the cost while maintaining a QoI of at least the specified objective QoI_{min}.

First, we use the optimization to find how many samples each sensor has to contribute to the final fusion task of the activity recognition. We notice that in order to guarantee that the QoI for each sensor is larger than 0.5, the lower bounds for sampling rates are $\{7, 9, 7, 41, 17\}$ for the sensors {RUA_, LH, RLA, RKNˆ, LUA} . In Figure 7, we plot the minimum Energy Consumption for the activity recognition for different values of QoI_{min}. The total energy consumption increases only marginally as QoI_{min} increases from 50% to 80%, and sharply increases as QoI increases from 80% to 94%. This result also supports our previous observation that if we are willing to sacrifice a small amount of accuracy, we can get a large reduction in energy consumption in return.

Second, we find the sending rates needed to recognize the two tasks at the same time: activity recognition with the sensor subset {RUA_, LH, RLA, RKNˆ, LUA} and the modes of locomotion with {RKNˆ, LUA}. As we see, the tasks share two sensors. We use two different constraints QoI_{min} for locomotion and activity recognitions, with the values 90% and 80% respectively. The optimization yielded the sensor sending rates $\{18, 15, 18, 41, 19\}$ which satisfy $QoI_{locomotion} = 97.82\%$, $QoI_{activities} = 80.31\%$ and $E = 5.6378\ mJ$. We conclude that our optimization approach successfully calculates the optimal sending rates while guaranteeing that the system can classify the different concurrent activities with the required specific accuracies.

V. Related Work

The trade-off between sensing and communicating overhead and the quality of reconstructed data in wireless sensor networks has been widely studied in the literature, e.g., [8], [9]. We focus on the two most relevant works, i.e. the relationship between the sampling rate and the accuracy of detecting transient events [10], and the connection between the quality of inference and the sensor's tolerance range [11]. These works and ours consider the trade-off between sampling rates and QoI and try to leverage the QoI function with sampling rates as input. However, [10] investigates event detection and [11] context recognition, while we investigate activity recognition. Zadehi et al. [10] suggested to build a model to capture the accuracy of event detection given a set of sensors. In contrast to this, we focus on joint optimal determination of sampling rates from a sensor subset. Roy et al. [11] deal with the quality of inference with the tolerance range or precision of data value as inputs. From our implementation experience, their work encounters the problem of mapping tolerance ranges to QoI if they use standard classification techniques. If it is necessary to pull more data to provide a higher accuracy, the approach cannot easily trace which samples were skipped. Our work does not have this limitation.

Sensor	RUA_	LH	RLA	RKNˆ	LUA
α_i	68.9089	70.2952	72.6076	54.8727	62.5572
β_i	1.9925	2.9803	2.5514	3.7677	3.7406

Table I
SENSITIVITY CONSTANTS OBTAINED BY CURVE FITTING FOR ACTIVITY RECOGNITION

Sensor	RKNˆ	LUA
α_i	92.2889	87.9588
β_i	0.6511	0.7219

Table II
SENSITIVITY CONSTANTS OBTAINED BY CURVE FITTING FOR LOCOMOTION RECOGNITION

VI. Conclusion

In this paper, we propose a model to control the QoI by adapting the number of samples collected from sensors. Based on this model, we minimize the communication cost of collecting the samples from each sensor, while guaranteeing that the user-specified QoI is achieved. We formalize this as an integer programming problem and solve it with branch–and–bound method with Lagrangian relaxation. Using an exponential distribution function, we formulate the QoI of each sensor as a function of its sending rate. Evaluation with traces from a public dataset demonstrates the significance of the trade-off between QoI and communication cost, e.g., 10% reduced accuracy gives 80% reduced communication cost. Thus, a considerable amount of energy can be saved by accepting marginally lower accuracy.

In this work, we only consider the exact solution using branch and bound. In future work, we will extend our

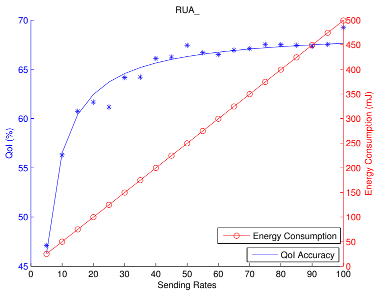

Figure 2. QoI and Energy Consumption vs. Sending Rates for Right Upper Arm Activity Recognition (RUA_)

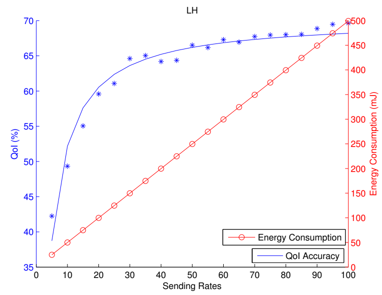

Figure 3. QoI and Energy Consumption vs. Sending Rates for Left Hand Activity Recognition (LH)

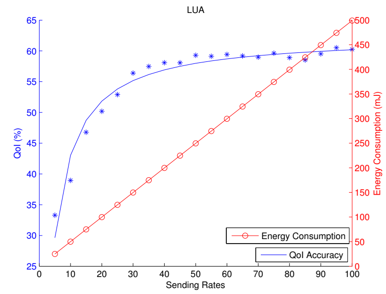

Figure 4. QoI and Energy Consumption vs. Sending Rates for Left Upper Arm Activity Recognition (LUA)

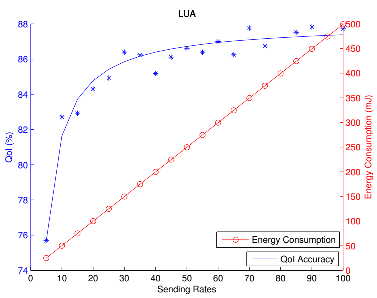

Figure 5. QoI and Energy Consumption vs. Sending Rates for Left Upper Arm Locomotion Recognition (LUA)

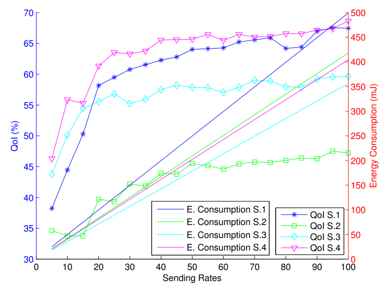

Figure 6. QoI and Energy Consumption vs. Sending Rates for Activity Recognition for Multiple Subjects (RUA)

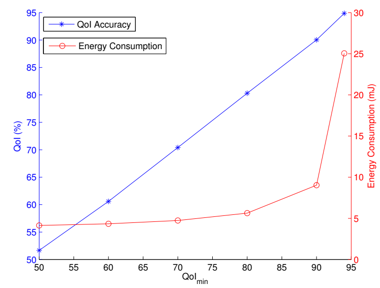

Figure 7. Optimization of the Energy Consumption for Different Objectives Qoi_{min}

approach with an approximation solution which relaxes any discreteness constraints on the sending rates and then rounds the solution to get more efficient in computation. There are also other open issues to solve, such as investigating how sending rates influence a complex query that consists of multiple predicates. These predicates include not only concurrent activity predicates but also And, Or, consecutive activity predicates. We will also investigate non-uniform sampling for a single query with different accuracies for different predicates.

ACKNOWLEDGMENT

This work is funded by the Norwegian Research Councils Program, through the Sirirus Project (No.2009/784).

REFERENCES

[1] N. H. Viet, E. Munthe-Kaas, and T. Plagemann, "Energy-balanced sensor selection for social context detection," in *PERCOM Workshops, 2012 IEEE*, March 2012.

[2] D. Roggen *et al.*, "Collecting complex activity data sets in highly rich networked sensor environments," in *7th INSS*, 2010.

[3] C. Bisdikian, "On sensor sampling and quality of information: A starting point," in *PERCOM Workshops, 2007.*

[4] M. A. Hossain, P. K. Atrey, and A. E. Saddik, "Modeling quality of information in multi-sensor surveillance systems," in *Proc. of the 23rd Inter. Conf. on Data Eng. Workshop*. IEEE, 2007.

[5] P. K. Atrey, M. S. Kankanhalli, and J. B. Oommen, "Goal-oriented optimal subset selection of correlated multimedia streams," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 3, no. 1, Feb. 2007.

[6] C. H. Papadimitriou, "On the complexity of integer programming," *Journal ACM*, vol. 28, no. 4, Oct. 1981.

[7] "Chipcon CC2420 Datasheet, Texas Instruments," 2007.

[8] C. Olston and J. Widom, "Offering a precision-performance tradeoff for aggregation queries over replicated data," in *26th VLDB*. Morgan Kaufmann Publishers Inc., 2000.

[9] W. Hu, A. Misra, and R. Shorey, "Caps: Energy-efficient processing of continuous aggregate queries in sensor networks," ser. PERCOM '06. IEEE, 2006.

[10] S. Zahedi and C. Bisdikian, "A framework for QoI-inspired analysis for sensor network deployment planning," ser. WICON '07. ICST.

[11] N. Roy, A. Misra, C. Julien, S. Das, and J. Biswas, "An energy-efficient quality adaptive framework for multi-modal sensor context recognition," in *PERCOM, IEEE 2011*, March.