

Energy Efficient Data Collection in Distributed Sensor Environments

Qi Han Sharad Mehrotra Nalini Venkatasubramanian
School of Information and Computer Science
University of California, Irvine, CA 92697
{qhan, sharad, nalini}@ics.uci.edu

Abstract

Sensors are typically deployed to gather data about the physical world and its artifacts for a variety of purposes that range from environment monitoring, control, to data analysis. Since sensors are resource constrained, often sensor data is collected into a sensor database that resides at (more powerful) servers. A natural tradeoff exists between the sensor resources (bandwidth, energy) consumed and the quality of data collected at the server. Blindly transmitting sensor updates at a fixed periodicity to the server results in a suboptimal solution due to the differences in stability of sensor values and due to the varying application needs that impose different quality requirements across sensors. This paper proposes adaptive data collection mechanisms for sensor environments that adjusts to these variations while at the same time optimizing the energy consumption of sensors. Our experimental results show significant energy savings compared to the naive approach to data collection.

1. Introduction

With the advances in computational, communication, and sensing capabilities, sensor enriched communication and information infrastructures have the potential to revolutionize almost every aspect of human life benefiting application domains. An integral component of such an infrastructure is a data management system that allows seamless access to data dispersed across a hierarchy of storage, communication, and processing units – from sensor devices where data originates to large databases where the data generated is stored and/or analyzed.

Designing a scalable data management solution to drive distributed sensor applications poses many significant challenges. Given the limited computational, communication, and storage resources at the sensors, a traditional distributed database approach in which sensors function as nodes in a distributed system might not be a feasible option. In order to facilitate complex query processing and analysis, data

might need to be migrated to repositories that resides at (more powerful) server(s). An alternative solution, where sensor data is continuously collected at a (logically) centralized database might also be infeasible. Since sensor readings may change very frequently/continuously, such environments are highly dynamic. Blindly transmitting the sensor updates to the server will impose severe network and storage overheads. Furthermore, since communication constitutes a major source of power drain [1] in battery-operated sensors, it would incur a very high energy cost.

The problem of effective data collection in highly dynamic environments has recently been studied in [4, 3]. The key observation is that a large number of sensor applications can tolerate a certain degree of error in data. Data imprecision, of course, impacts application quality. For example, in an application such as target tracking in a sensor network, error in sensor intensity readings may result in error in localizing the object. Similarly, the result of a query for average temperature in a given region may be imprecise due to data error. The communication overhead between the data producers and the server can be alleviated by exploiting the applications' error tolerance.

Motivated by [4], in this paper, we explore data collection protocols for sensor environments that exploits the natural tradeoff between application quality and energy consumption at the sensors. Modern sensors try to be power aware, shutting down components (e.g., radio) when they are not needed in order to conserve energy. We consider a series of sensor models that progressively expose increasing number of power saving states. For each of the sensor models considered, we develop quality-aware data collection mechanisms that enable quality requirements of the queries to be satisfied while minimizing the resource (energy) consumption.

2. Problem Formulation

In this section, we describe system and query models used in this paper and develop a formal characterization of the sensor data collection problem.

System and Query Model: Our system consists of a set of n sensors and one server(residing at a resource sufficient node) that maintains a database. For simplicity, we will assume that each sensor can communicate directly to the server. Our solution can serve as a building block for large scale distributed sensor system.

Each incoming query Q_i is associated with an accuracy constraint A_i indicating its tolerance to error in answer precision (We will explain the accuracy constraint and query answer accuracy later). Furthermore, a query has a latency bound D which requires that each query be answered within D time units.

Each sensor node has a processor with limited memory, an embedded sensor, an analog-to-digital converter, and radio circuitry. A micro-operating system controls each component. We only consider different radio modes and assume all the other components are always turned on. we consider three sensor states: active (a), listening (l) and sleeping (s). While the sensor is in the *active/listening* mode, the transmitter/receiver is on; when the sensor goes to the *sleeping* state, its radio is turned off completely. Currently, there exist two types of microsensor devices. One type (such as Berkeley Mica Mote [6]) has only has one radio (i.e., it either transmits or receives data); the other type (such as MIT μ AMPS node [7]) has two radios, and can transmit and receive data simultaneously. These two types of sensor models can both be represented by our abstract sensor models as in Table 1. Note that when a radio is in the idle mode, it is

Table 1: **Sensor states**

radio mode		sensor state
1-radio node	2-radio node	
Tx on, Rx off	Tx on, Rx on	active (a)
Tx off, Rx on		listening (l)
Tx off, Rx off		sleeping (s)

capable of detecting an incoming packet, but not in the process of receiving a packet. We classify this mode as the listening state since Tx is off and Rx is on. Also note that even in the sleeping state, the changes in sensor values can still be detected, since we assume the sensor and processor are always on. In this paper, we assume that a sensor node has two radios, and similar analysis can be applied for sensor nodes with only one radio.

Data Collection Framework: Previous work has shown that an effective approach to exploit the tradeoff between application quality and data imprecision is for the server to maintain an approximate value of the data whose divergence from the true value is guaranteed to be bounded by an error at any time. Specifically, let $S = \{s_1, \dots, s_n\}$ be the set of sensors. Each sensor hosts its exact value that may change frequently. For each $s_i \in S$, let v_i denote the value stored at sensor s_i . The approximation of v_i is represented by a range r_i with lower bound l_i and upper bound

$u_i: r_i = [l_i, u_i]$, which is stored in the database at the server. A query for the value of sensor s_i is answered in the format of a range with a lower and an upper bound, so the answer accuracy is defined by the range size $u_i - l_i$. The accuracy constraint A_i of query Q_i specifies the maximum acceptable width of the result.




Figure 1: **The Data collection process**

Whenever the sensor value v_i changes to v'_i , sensor s_i checks whether r_i is still a valid approximation for the new value. If v_i falls outside r_i , a new approximation of v'_i is sent to the server to update the database (This process is called *source-initiated update*). Otherwise, there is no need to transmit the update to the server, hence reducing communication overhead. Queries are executed over the cached ranges at the server. If the error tolerance of the query is larger than the data error, i.e., $A_i \geq u_i - l_i$, it is processed without any communication with the sensor. Otherwise, the approximation offered by the database is insufficient, the server may request the exact value from remote sensor. The sensor responds with current exact value and a new approximation to be used by subsequent queries. This process is called *consumer-initiated request and update*. Fig. 1 illustrates the data collection process.

Problem Statement: Given m user queries, our objective is to minimize sensor energy consumption in the process of answering all queries (say m). Since a sensor consumes energy even when it is not transmitting or receiving data, besides reducing the communication overhead between a sensor and the server, we also need to minimize the time a sensor is either active and/or listening even when it is not transmitting updates to the server. Assume that the probability of source- and consumer-initiated updates at each time instant are P_{su} and P_{cu} . Formally, we would like to

$$\begin{aligned}
 & \text{minimize } \bar{E} = E_{su} \cdot P_{su} + E_{cu} \cdot P_{cu} + E_{extra}. \\
 & \text{subject to (1) } a_i \leq A_i, 0 \leq i \leq m \\
 & \quad \quad \quad (2) t_i \leq D, 0 \leq i \leq m
 \end{aligned}$$

where a_i is the answer accuracy for query i and t_i is the query response time. Also note that, E_{su} is the energy required to send an update to the server, and E_{cu} is the energy required to both receive the request for the data and for transmitting the sensor value to the server. Note that E_{su} and E_{cu} are not constant. They depend upon the state at which the sensor was when the source-initiated update and consumer-initiated update occurred. Consider, for ex-

ample, a sensor that is in the sleeping state. If the sensor value changes causing it to exceed the range associated for its value at the server, it will first has to transition to the active state followed by transmitting the value to the server. Thus, the total energy spent would be the sum of energy spent to transition from the sleep state to the active state and the energy spent to transmit the update to the server. In contrast, if the value divergence occurs when the sensor is in the active state, the energy consumption would be only for transmitting the update to the server. E_{extra} is the amount of energy consumed while not receiving or transmitting any data. This is also not a constant, and it depends upon the state a sensor is in while it is free.

To achieve the objective of minimizing the energy consumption at the sensor, we need to address two issues:

- *how to maintain the database*: an optimal range needs to be maintained and adjusted for each sensor so that it reduces sensor energy consumption while still being able to meet query accuracy constraints. If the range is large, accuracy constraints of many queries will be violated resulting in expensive probes; likewise, if the range is small, sensor update would needlessly be transmitted to the server too frequently. Both cases will consume a large amount of energy. We address how to set the range such that the energy consumption is minimized in Section 3.
- *how to manage sensor state*: we need to determine sensor state transition strategies. Sensors consume power not only when sending and receiving data, but also when idling at the active and listening states. To save energy, a sensor needs to power down into a lower energy state. Powering down a sensor requires additional cost to power up when a request that needs to be processed arrives. Furthermore, it could result in increased latency for queries. In Section 4, we address optimal state transition that determines the length of sensor idling and sleeping to minimize overall energy consumption.

3. Data Precision Adjustment

In this section, how the approximation range for the sensor can be set at the server in order to minimize the energy consumption due to communication between sensors and the server. The energy cost due to communication depends upon the number of source- and consumer-initiated updates which, in turn, depends upon the range size adaptation, patterns of the changes in sensor values and query workload characteristics. Before presenting our solutions, we briefly review the approach described in [4] where the authors considered range adaptation to minimize the communication overhead between data producers and the server. Our approach builds upon some of their results.

Assuming that the communication cost incurred during a source- and consumer-initiated update is C_{su} and C_{cu} respectively, the expected cost per unit time $C = P_{su} \cdot C_{su} + P_{cu} \cdot C_{cu}$. The authors established that $P_{su} = \frac{K_1}{r^2}$ and $P_{cu} = K_2 \cdot r$, where r is the range size and K_1 and K_2 are model parameters that depend on the characteristics of source updates and queries. Therefore, $C = \frac{K_1}{r^2} \cdot C_{su} + K_2 \cdot r \cdot C_{cu}$, which is minimized when the range size $r = \sqrt[3]{\rho \cdot \frac{K_1}{K_2}}$, here $\rho = 2 \cdot \frac{C_{su}}{C_{cu}}$. At this optimal point, it can be shown that $\frac{P_{cu}}{P_{su}} = \rho$, which is a constant.

Table 2: Data precision adjustment proposed in [4]

Data-Precision-Adjustment(update-type, r)	
switch (update-type) {	
case source-initiated update:	
with probability $\min\{\rho, 1\}$, set $r' = r(1 + \theta)$; break ;	
case consumer-initiated update:	
with probability $\min\{\frac{1}{\rho}, 1\}$, set $r' = \frac{r}{(1+\theta)}$; break ; }	
return r' ;	

The algorithm (Table 2) exploits this observation and attempts to change the range r such that the ratio of probability of consumer-initiated update to the probability of source-initiated update can be maintained to be the constant ρ . For example, if $\rho = 1$, the algorithm attempts to ensure that the probability of consumer-initiated update is equal to the probability of source-initiated update. In case $\rho < 1$, it is desirable for source-initiated updates to be more likely than consumer-initiated updates. Thus the range is decreased on every consumer initiated update but only increased with probability ρ on source-initiated updates. Conversely, in case $\rho > 1$, the range is increased on every source-initiated update but only decreased with probability $\frac{1}{\rho}$ on consumer-initiated updates.

Table 3: Symbols used

Symbol	Meaning
r	interval size
P_{su}	probability of source-initiated update
P_{cu}	probability of consumer-initiated update
P_i	probability of being in state i ($i = a, l, s$)
T_{rx}	time needed receiving a consumer-init. request
T_{tx}	time needed sending an update
T_{ij}	transition time from state i to j ($j = a, l, s$)
PC_i	power consumption when sensor is in state i
E_{ij}	energy consumed in switching from state i to j

We note that the solution in [4] has essentially been developed for data collection in environments where data producers are not energy constrained (e.g., they could be powerful network routers) and its straightforward application is not suitable in energy constrained sensor environments. A

direct application of their solution would require that a sensor be always maintained in an active state since a server may need to access the current sensor value at any time which would result in a very high energy cost.

In the following, we present data collection strategies for sensor environments. We consider a series of sensor models based on power saving sensor states identified in the previous section. These models progressively consider more sensor states and become more complicated. For each of the above models, we discuss how to determine the ranges r_i such that the overall energy consumption is minimized while meeting the quality constraints of the query. Table 3 summarizes the symbols used in the derivation.

3.1. Always-Active Model (AA)

In this model, sensors are always active. The total normalized energy consumption is shown in Equation (1).

$$\begin{aligned} \bar{E}_{aa} &= PC_a(P_{su}T_{tx}) + PC_a(P_{cu}(T_{rx} + T_{tx})) \\ &+ PC_a[1 - P_{su}T_{tx} - P_{cu}(T_{rx} + T_{tx})] \\ &= PC_a \end{aligned} \quad (1)$$

As expected, it shows that the normalized energy consumption is equal to the power consumption at the active state. Therefore, irrespective of how the range is set, energy consumption is constant. This model serves as a baseline to study the energy savings that result in utilizing sensor states that consume less energy.

3.2. Active-Listening Model (AL)

In this model (illustrated in Fig. 2), the sensor consists of two states: active and listening. Initially the sensor is in




Figure 2: **The Active-Listening model (AL)**

the listening mode. The sensor shifts to the active state if either the sensor value diverges from the range used to represent the sensor value at the server, or if it receives a request for its current value from the server. When a sensor is in the active state, it processes all its pending requests and waits for a T_a unit of time before switching to the listening mode. The reason to wait for T_a time units in the active (higher energy) state instead of powering down to the listening (lower energy) state immediately is that switching from a lower energy state to a higher energy state is associated with a significant energy cost. From an energy perspective

it might be advantageous to wait in the higher energy state (instead of powering down) if the sensor will be required to transition back to higher energy state in the near future. Obviously, the optimal value of T_a that minimizes energy consumption depends upon the application workload and sensor value change patterns. We defer further discussion on how T_a can be set in order to minimize power consumption to Section 4. For the time being, we assume that T_a has been optimally set. With T_a fixed, we consider the problem of optimally determining the range r for the sensor that minimizes the energy consumption.

The sensor energy consumption under this model is shown in Equation (2).

$$\begin{aligned} \bar{E}_{al} &= [P_l(E_{la} + PC_aT_{tx}) + P_a(PC_aT_{tx})]P_{su} \\ &+ [P_l(PC_lT_{rx} + E_{la} + PC_aT_{tx}) \\ &+ P_a(PC_a(T_{rx} + T_{tx}))]P_{cu} + (PC_aP_a + PC_lP_l) \\ &\cdot [1 - (P_l(T_{la} + T_{tx}) + P_aT_{tx})P_{su} \\ &- (P_l(T_{rx} + T_{la} + T_{tx}) + P_a(T_{rx} + T_{tx}))P_{cu}] \end{aligned} \quad (2)$$

The energy consumption depends upon the probabilities P_a and P_l of the sensor being in the active and listening state. We next show how these probabilities can be expressed in terms of the probability of source and consumer initiated updates. If $T_a = 0$, the sensor state transition matrix capturing the state transition probabilities is as follows:

$$\mathbf{P} = \left(\begin{array}{c|cc} & \text{listening} & \text{active} \\ \hline \text{listening} & 1 - P_a & P_a = P_{su} + P_{cu} \\ \hline \text{active} & 1 - P_{aa} & P_{aa} = P_{su} + P_{cu} \end{array} \right)$$

The long-term probability that the system will be in each state can be obtained by computing the steady state vector of the Markov Chain. Therefore, we get

$$P_l = 1 - P_{su} - P_{cu} \text{ and } P_a = P_{su} + P_{cu}.$$

As mentioned before, $P_{su} = \frac{K_1}{r^2}$ and $P_{cu} = K_2 \cdot r$ [4]. To find the minimum $\bar{E}_{al}$, we can find the root of the derivative $\frac{d\bar{E}_{al}}{dr}$, and we get $r^* = \sqrt[3]{\frac{2K_1}{K_2}}$. At this optimal point, $\frac{P_{cu}}{P_{su}} = \frac{K_2}{K_1} \cdot r^{*3}$. Thus, energy consumption is minimized when the ratio $\frac{P_{cu}}{P_{su}} = 2$ which is a constant. Following similar derivation for the case when $T_a > 0$, we observe the same conclusion. Since the data precision adjustment algorithm shown in Table 2 maintains the ratio of P_{cu} over P_{su} to be a constant, it can be used in conjunction with the AL model to minimize energy.

3.3. AS model

In the AS model (described in Figure 3) the sensor toggles between the sleeping and active modes. Initially, the sensor is in the sleep state. Similar to the AL model, sensor shifts to the active state if the sensor value diverges from the range used to represent the sensor value at the

server. Since a sensor in the sleep state cannot receive request from the server, it periodically wakes up on a timeout if it has been sleeping uninterrupted for T_s time units. Such a timeout based transition is necessary in order to meet the quality requirements of queries that would have resulted in the consumer-initiated update at the sensor. The sensor, on switching to the active state sends its current value to the server. Note that this update can be used by the server to answer those queries that would have resulted in consumer-initiated requests while the sensor was in the sleep state. The sensor remains in the active state while there are requests for its value. After it has handled all the requests, it switches to the sleeping state after waiting for T_a time units without handling any requests. We next discuss the energy consumption for the AS model.




Figure 3: **The Active-Sleeping model (AS)**

In the AS model, total energy consumption(Equation (3)) consists of (a) energy consumed by source-initiated updates. Besides the energy spent in transmitting source-initiated updates, there is energy involved in transitioning from sleeping to active if the sensor is sleeping when the source-initiated update is due; (b) energy consumed by transition from the sleeping state to the active state and the associated value updates when a sensor wakes up due to time-out; (c) energy consumed by consumer-initiated updates; and (d) energy consumed while sleeping or being active without receiving or transmitting.

$$\begin{aligned} \bar{E}_{as} = & [(E_{sa} + PC_a T_{tx})P_s + PC_a T_{tx} P_a]P_{su} \\ & + (P_{sa} - P_{su}P_s)(E_{sa} + PC_a T_{tx}) \\ & + [PC_a(T_{rx} + T_{tx})]P_a P_{cu} + (PC_a P_a + PC_s P_s) \\ & \cdot [1 - ((T_{sa} + T_{tx})P_s + T_{tx}P_a)P_{su} \\ & + (P_{sa} - P_{su}P_s)(T_{sa} + T_{tx}) + (T_{rx} + T_{tx})P_{cu}P_a] \end{aligned} \quad (3)$$

We next derive the optimal setting of the range r_i for the sensor that minimizes the energy consumption under the assumption that the T_a has already been set. If $T_a = 0$, the sensor switches to the sleeping state as soon as there is no requests waiting. Therefore,

$$P_{sa} = P_{su} \left[\frac{1}{T_s P_{su}} \right] = \frac{1}{T_s} + \alpha P_{su}, 0 \leq \alpha < 1.$$

$$P_{aa} = P_{su} + P_{cu}$$

The sensor state transition matrix is as follows:

$$\mathbf{P} = \begin{pmatrix} & \text{sleeping} & \text{active} \\ \text{sleeping} & 1 - P_{sa} & P_{sa} \\ \text{active} & 1 - P_{aa} & P_{aa} \end{pmatrix}$$

Long term probabilities of the sensor being in the sleep- ing and active states are as follows:

$$P_s = \frac{1 - P_{su} - P_{cu}}{\frac{1}{T_s} + \alpha P_{su} + 1 - P_{su} - P_{cu}} \text{ and } P_a = \frac{\frac{1}{T_s} + \alpha P_{su}}{\frac{1}{T_s} + \alpha P_{su} + 1 - P_{su} - P_{cu}}.$$

If $T_a > 0$, the sensor stays active for a period of time so that bursty update requests can be processed without state switching. We can derive state transition probabilities as follows. The probability of switching from sleeping to active:

$$P_{sa} = P_{su} \left[\frac{1}{T_s P_{su}} \right] = \frac{1}{T_s} + \alpha P_{su}, 0 \leq \alpha < 1$$

The probability of switching from active to sleeping:

$$\begin{aligned} P_{as} &= (P_{su} + P_{cu}) \left(\left[\frac{1}{T_a (P_{su} + P_{cu})} \right] - 1 \right) \\ &= \frac{1}{T_a} - \beta (P_{su} + P_{cu}), 0 < \beta \leq 1 \end{aligned}$$

The sensor state transition matrix is as follows:

$$\mathbf{P} = \begin{pmatrix} & \text{sleeping} & \text{active} \\ \text{sleeping} & 1 - P_{sa} & P_{sa} \\ \text{active} & P_{as} & 1 - P_{as} \end{pmatrix}$$

Long term state probabilities are:

$$P_s = \frac{\frac{1}{T_s} - \beta (P_{su} + P_{cu})}{\frac{1}{T_s} - \beta (P_{su} + P_{cu}) + \frac{1}{T_s} + \alpha P_{su}} \text{ and}$$

$$P_a = \frac{\frac{1}{T_s} + \alpha P_{su}}{\frac{1}{T_s} - \beta (P_{su} + P_{cu}) + \frac{1}{T_s} + \alpha P_{su}}.$$

For both $T_a = 0$ and $T_a > 0$, by applying these probability formula into Equation 3, the total energy consumption can be expressed as a ratio of two complex polynomials of range size r (see [2] for detail). Since is not possible to express the ratio P_{cu} to P_{su} in terms of other parameters, the basic strategy for range setting described in Table 2 can not be used. Instead, we need to monitor parameters K_1 , K_2 , α and β at runtime. For this purpose, the following information for the sliding window of last k updates is maintained: (1) the number of sensor state transitions (N_{sa} and N_a) of the last k updates; and (2) the number of source- or consumer-initiated updates (N_{su} or N_{cu}) of the last k updates. Using this information, the values of K_1 , K_2 , α and β is estimated. For example, K_1 is set to be $P_{su} \cdot r^2$, where P_{su} is estimated as the number of source-initiated updates (N_{su}) divided by T , where T is the time period of the current window. The parameter K_2 can be estimated similarly. Given these parameter values, we find the roots of $\frac{dE_{as}}{dr}$ and compare the energy values at the roots to determine the value of r that minimizes the energy consumption $\bar{E}_{as}$. Since the computation is too complex to be performed at the resource-constrained sensors, it is done at the server. Note that since the value of α and β depends upon the number of sensor state transitions during the window, the sensors determine the values of α and β and piggy-back these values for the last k updates with the k^{th} update.

At the same time the server monitors K_1 and K_2 ; upon receiving the k^{th} update, the server computes the new optimal range which is transmitted to the sensor.

3.4. ALS Model

In this model (illustrated in Fig. 4), the sensor is initially in the sleeping state. It switches to the active state when a source-initiated update occurs or when it has been sleeping for T_s time units without interruption. When it is in the active state, it processes all the waiting requests. After it has been free in the active state for T_a time units, it goes to the listening state. Once in the listening state, any source initiated update or consumer initiated update will trigger the sensor to go to the active state; otherwise, if it is idling for T_l time units, it goes to sleep. The sensor energy consumption




Figure 4: The Active-Listening-Sleeping model

tion is shown in Equation (4).

$$\begin{aligned}
 E_{als}^- &= (P_s E_{sa} + P_l E_{la} + P C_a T_{tx}) P_{su} \\
 &+ (P_{sa} - P_{su} P_s)(E_{sa} + P C_a T_{tx}) \\
 &+ [P_l (P C_l T_{rx} + E_{la} + P C_a T_{tx}) \\
 &+ P_a P C_a (T_{rx} + T_{tx})] P_{cu} \\
 &+ (P C_a P_a + P C_l P_l + P C_s P_s) \\
 &\cdot [1 - (P_s T_{sa} + P_l T_{la} + T_{tx}) P_{su} \\
 &- (P_{sa} - P_{su} P_s)(T_{sa} + T_{tx}) \\
 &- ((P_l + P_a)(T_{rx} + T_{tx}) + P_l T_{la}) P_{cu}]
 \end{aligned} \quad (4)$$

Similar to the AS model, optimal range size can be set based on the parameters monitored at runtime. The detailed derivation, though conceptually simple, is quite complex, so we refer the interested readers to [2]. We can apply the same data collection approach as in the AS model.

4. Adaptive Sensor State Transition

In the various sensor models discussed above, transitions among states besides being triggered by the sensor value diverging from its representation at the server, also occur due to timeouts. In AL, AS and ALS model, T_a needs to be determined. We use the AS model to show how an optimal

value of T_a can be derived. The development for the ALS model is similar in nature and interested readers are referred to [2] for details.

Waking up a sensor at the sleeping state requires additional energy and latency, so it is not obvious that putting the sensor to sleep immediately after it finishes the requests at hand is the most energy efficient choice. Depending upon the request arrival rate, the sleeping period could be so short that powering up costs are greater than the energy saved in that state. On the other hand, waiting too long to power down may not achieve the best energy reductions possible. Thus, a careful selection of T_a is important. Intuitively, if updates (either initiated by source or consumer) are not bursty, it is better to set T_a to be zero; otherwise, the sensor should remain active for a while before going to sleep, so that more requests can be answered in time and frequent state switching can be avoided. Hence, a good understanding of source- and consumer-initiated update patterns will help in determining the optimal active time.

Let us assume that $f(t)$ is the probability of receiving any type of requests at any time instant t . Let $\pi(t)$ be the probability of being silent for t time units, i.e., there are no requests before t until a request arrives at time t . Since any incoming request means the end of the silent period, $\pi(t) = f(t)$. If we assume that requests (either source or consumer initiated update requests are uniformly distributed in interval $(0, T_a + T_s]$ (since we know that at the end of T_s , there must be a timeout update request), then $\pi(t) = f(t) = \frac{1}{T_a + T_s}$. In this case, the expected energy consumption for a single silent period will be

$$\begin{aligned}
 E &= \int_0^{T_a} \pi(t) P C_a t dt \\
 &+ \int_{T_a}^{T_a + T_s} \pi(t) [P C_a T_a + P C_s (t - T_a) + E_{sa}] dt \\
 &= \frac{P C_a T_a^2 + 2 P C_a T_s T_a + (P C_s T_s^2 + 2 E_{sa} T_s)}{2(T_a + T_s)}
 \end{aligned}$$

Since $T_a \geq 0$ and $E > 0$, E is non-decreasing and is minimal when $T_a = 0$. i.e., the sensor should go to sleep immediately after it finishes all the requests at hand.

While $T_a = 0$ is optimal if request inter-arrival pattern follows uniform distribution, in practice, this assumption is rarely true, and the problem of finding $\pi(t)$ remains. Our approach is to learn $\pi(t)$ at runtime and adaptively select T_a accordingly. The basic idea is as follows: we choose a window size w in advance. The algorithm keeps track of the last w idle period lengths and summarizes this information in a histogram. Periodically, the histogram is used to generate a new T_a .

The set of all possible inter-arrival period lengths $(0, T_a + T_s)$ is partitioned into n intervals, where n is the number of bins in the histogram. Let t_i be the left endpoint of the i^{th} interval. The i^{th} bin has a counter which in-

indicates the number of idle periods among the last w idle periods whose length fall in the range $[t_i, t_{i+1})$. The bins are numbered from 0 to $n - 1$ and $t_i = 0, t_n = T_a + T_s$.

The counter for bin i is denoted by c_i . The threshold for changing states is selected among n possibilities: $t_0, \dots, t_{n-1}$. We estimate the distribution π by the distribution which generates an idle period of length t_i with probability c_i/w for $\forall i \in \{0, \dots, n-1\}$. $\sum_{i=0}^{n-1} c_i = w$. Thus T_a is chosen to be the value t_m that minimizes the energy consumption as follows:

$$\min_{t_m} \left\{ \sum_{j=1}^{m-1} \frac{c_j}{w} PC_a t_j + \sum_{j=m}^n \frac{c_j}{w} [PC_a t_m + PC_s(t_j - t_m) + E_{sa}] \right\}$$

Similar derivation can be done to obtain T_a for AL/ALS models and T_l for the ALS model. (See details in [2]).

5. Performance Evaluation

The objective of our simulation is to compare the performance of various sensor state models (AA, AL, AS and ALS) for quality-aware data collection in terms of energy consumption and average query response time. We also study the impact of T_a value, as well as range size adjustment on the system performance.

We built a simulator in C, consisting of a server, a database and a number of sensors. User queries are posed at the server which then returns their results. User queries arrival times at the server are Poisson distributed with mean inter-arrival time set at 2 seconds. Each query is accompanied by an accuracy constraint specifying the maximum acceptable width of the result. The accuracy constraint are sampled uniformly from the range $[0, 40]$.

Table 4: Parameters used in the simulation

Symbol	value	symbol	value
PC_a	14.88mW	E_{sl}	0.025μJ
PC_l	12.50mW	E_{la}	0.014μJ
PC_s	0.016mW	E_{sa}	0.119μJ
T_{sl}	4μs	T_{rx}	0.05ms
T_{la}	12μs	T_{tx}	1.67ms
T_{sa}	16μs		

The sensor-related parameters (Table 4) were obtained from the specification of a typical sensor node [5]. Each sensor holds one exact numeric value, and the database holds all the interval approximations. Sensor values are picked randomly and uniformly from the range $[-150, 150]$; they perform a random walk in one dimension: every second, the values either increases or decreases by an amount sampled uniformly from $[0.5, 1.5]$.

Experimental Results: Figure 5 shows sensor energy consumption and query response time of the four sensor models (AA, AL, AS and ALS). Not surprisingly, the energy consumption of the AA model is the highest, and its query response time is the lowest. This is the model where no energy is saved; sensors are always active, thus any consumer-initiated requests can be detected and processed immediately. As shown in Table 4, listening state consumes similar amount of power to active state, thus the AL model does not decrease energy consumption to a great extent. However, most of the time, the sensor is in the listening state when most requests arrive. It switches to the active state so that it can actually send out the updates. This power-up process takes time, which explains why query response time under the AL model is higher than the AA model. Models that incorporate sleeping state reduce energy consumption significantly. However, this comes at the price of higher query response time, since it takes more time for a sensor to switch from the sleeping state to the active state than from the listening state to the active state. Given that decreasing sensor energy consumption is our objective, the AS model outperforms the other models significantly due to its low energy cost. We, therefore, restrict the remainder of the performance study to the AS model.

Figure 6 compares system performance under fixed T_a with system performance using adaptive T_a . Since $T_a = 0$ was shown to be optimal when request arrival follows uniform distribution, we compare the adaptive approach to the approach that fixes T_a to be 0. The results show that adaptive T_a saves energy by half and also decreases query response time. When requests are bursty, it saves energy and shortens query waiting time by remaining active for a certain time period. Adapting T_a to user query patterns and sensor value changes performs better than fixing its value.

Figure 7 depicts the impact of range size by showing four different cases: (a) $r = 0$: the database stores single instantaneous values instead of intervals; (b) set r to be the average accuracy constraint: on an average, queries can be satisfied by stored values; (c) adaptive r as shown in our approaches: the optimal r is found periodically to minimize the energy consumption; (d) a large r . When r is 0, all queries can be answered by just retrieving values from the database, so query response time is minimized; but each change in sensor value needs to be reported to the server, which consumes a large amount of energy. When r is set to be very large, most source value changes will not exceed current range, so the likelihood of source-initiated updates is low. However, the coarse data representation is not sufficient for most of the queries, hence a number of consumer-initiated updates will occur. As a result, the average query response time is very high. Figure 7 shows that our adaptive approach significantly outperforms other approaches.

Performance Summary: Performance studies indicate




Figure 5: System performance comparison of proposed sensor models




Figure 6: Impact of T_a adaptation on system performance




Figure 7: Impact of range size adaptation on system performance

that the AS model consumes the least amount of sensor energy; our proposed strategies of intelligent sensor state transition reduce energy consumption to a great extent; optimized range size adjustment works effectively with corresponding sensor models and saves more energy than static range or storing exact values.

6. Related Work and Concluding Remarks

Energy efficiency is one of the major concerns in sensor networks. To prolong system lifetime of sensor networks, various approaches have been devised to exploit low duty-cycle operation or the cooperation among sensor nodes. In contrast, our energy saving approaches take into consideration application-level information to optimize sensor energy consumption.

Since many real-world applications can tolerate data imprecision at varying levels, the error tolerance of applications can be exploited to reduce energy consumption during sensor data collection. In this paper, we have studied energy efficient data collection mechanisms for distributed sensor environment that explores the tradeoff between sensor data accuracy and energy consumption. Both theoretical analysis and experimental results validated the effectiveness of

our approaches. Dealing with aggregate continuous queries with the same objective and constraints is a direct extension of this work.

References

- [1] I. Akyildiz, W. Su, Y. Sankarasubramanian, and E. Cayirci. Wireless sensor networks: A survey. *Computer Networks*, 38(4), March 2002.
- [2] Q. Han, S. Mehrotra, and N. Venkatasubramanian. Energy efficient data collection in distributed sensor environments. Technical report, University of California-Irvine, 2003. <http://www.ics.uci.edu/~qhan/techrep/sensor-techrep.pdf>.
- [3] I. Lazaridis and S. Mehrotra. Capturing sensor generated time series with quality guarantees. In *IEEE ICDE*, 2003.
- [4] C. Olston, B. T. Loo, and J. Widom. Adaptive precision setting for cached approximate values. In *ACM SIGMOD*, 2001.
- [5] RT Monolithics Inc., <http://www.rfm.com/>. *ASH Transceiver TR1000 Data Sheet*.
- [6] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. Srivastava. Optimizing sensor networks in the energy-latency-density design space. *IEEE Tran. Mobile Computing*, 1(1), January-March 2002.
- [7] A. Sinha and A. Chandrakasan. Dynamic power management in wireless sensor networks. *IEEE Design and Test of Computers*, 18(2), March/April 2001.