SILENCE: Distributed Adaptive Sampling for Sensor-based Autonomic Systems*

Eun Kyung Lee, Hariharasudhan Viswanathan, and Dario Pompili
NSF Center for Autonomic Computing, Department of Electrical and Computer Engineering
Rutgers University, New Brunswick, NJ
{eunkyung_lee, hari_viswanathan, pompili}@cac.rutgers.edu

ABSTRACT

Adaptive sampling and sleep scheduling can help realize the much needed resource efficiency in densely deployed autonomic sensor-based systems that monitor and reconstruct physical or environmental phenomena. This paper presents a *data-centric* approach to distributed adaptive sampling aimed at minimizing the communication and processing overhead in autonomic networked sensor-based systems. The proposed solution exploits the spatio-temporal correlation in sensed data and eliminates redundancy in transmitted data through selective representation without compromising on accuracy of reconstruction of the monitored phenomenon at a remote monitor node. In addition, the solution also exploits the same correlations for adaptive sleep scheduling aimed at saving energy in Wireless Sensor Networks (WSNs) while also providing a mechanism for ensuring connectivity to the monitor node. The data-centric joint adaptive-sampling and sleep-scheduling solution, SILENCE, has been evaluated through real experiments on a testbed monitoring temperature and humidity distribution in a rack of servers as well as through extensive simulations on TOSSIM, the TinyOS simulator.

Categories and Subject Descriptors

C.2.2 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Protocols; C.4 [**PERFORMANCE OF SYSTEMS**]: Measurement techniques

General Terms

Algorithms, Performance

Keywords

Sensor-based systems, adaptive sampling, spatial and temporal correlation, autonomic systems.

*This work was supported by the National Science Foundation under Grant No. CNS-0855091

ICAC'11, June 14–18, 2011, Karlsruhe, Germany.

1. INTRODUCTION

Cyber-Physical Systems (CPSs) are distributed, autonomic sensor and actor systems [14] that feature a tight combination of, and coordination between, the system's computational and physical elements to enable timely reaction to sensor information with an effective action. For instance, CPSs can be employed for monitoring heat and air circulation inside a datacenter to enable energy-efficient thermal management decisions such as workload distribution and cooling system optimization. Such distributed sensor-based systems are, in general, composed of heterogeneous sensor nodes that differ with respect to i) *the type of data* they can sense (e.g., temperature, humidity, vibration, air flow, images), ii) *the source of energy* for operation, and iii) *the mode of data transmission* making them constrained in terms of energy or communication cost or both, as shown in Table 1. In CPSs, as derived metrics from data are of greater interest than the raw sensed data itself, the self-managing autonomic sensing systems deployed for estimating a physical or environmental phenomenon primarily address three major issues through self-configuration, self-healing, and self-optimization. They are 1) sampling rate in space, 2) sampling rate in time, and 3) data reporting rate.

The phenomenon of interest is in general characterized by multiple manifestations. For example, temperature, humidity, and airflow rates (manifestations) are crucial for understanding thermal hotspots (the phenomenon) inside datacenters. Hence, accurate estimation of a phenomenon requires simultaneous monitoring of its multiple manifestations, which exhibit their own spatial and temporal variation characteristics. Such nodes, equipped with multiple sensors, are generally deployed in a dense fashion to ensure sensing and communication coverage because of their small sensing (or scope) and transmission range, respectively. Because of the high density of nodes, there is in general a high degree of correlation among observations of spatially proximal nodes (*spatial correlation*). Also, the degree of correlation between consecutive measurements collected at a node may vary according to the temporal variation characteristics of the manifestation (*temporal correlation*). Therefore, resource-efficient estimation of the phenomenon can be performed by exploiting the spatial and temporal correlation characteristics of its manifestations [24, 25].

In this paper, we present distributed adaptive sampling for sensor-based autonomic systems (SILENCE) for resource-efficient estimation of a phenomenon. SILENCE combines data-centric adaptive sampling with an adaptive sleep-scheduling algorithm while also providing mechanisms for ensuring connectivity in CPSs composed of a network of wireless sensors. To the best of our knowledge, this is the first work to consider such joint approach. The solution exploits the spatial and temporal correlation in the manifestations to eliminate redundancy and to reduce the cost of pro-

Table 1: Categories of sensor-based systems.

		Communication Mode	
		Wired	*Wireless*
Energy Source	*Power line*	Unconstrained	Bandwidth and connectivity constrained
	Battery	Energy constrained	Bandwidth, connectivity, and energy constrained

cessing and communicating large volumes of sensed data to a base station (or sink) for post processing. *Note that the proposed solution is relevant and applicable to any complex sensor-based system consisting of heterogeneous sensors falling under any of the four categories shown in Table 1.*

Specifically, the proposed distributed solution enables each node to decide its state (or role) and sleeping schedule independently based on *correlation* and *similarity* of its own sampled data with that of the neighboring nodes' obtained through local control messaging. This aggregated data helps a node determine whether to play the role of a *representative (*REP*)* and, consequently, to actively report data to the sink on behalf of a group of nodes; or to be an *associate (*ASSOC*)* to a REP and sleep. Putting nodes to sleep ensures that energy is not spent on packet receptions as well as sensing. This is advantageous as it saves resources (energy and bandwidth) whatever be the system category (Table 1). Two nodes are said to be sensing similar values if the difference between the means of magnitude of the manifestation k observed at the two nodes is less than a user-specified threshold e^k_{th}. Measured values of manifestation k at two nodes are said to be correlated if the correlation coefficient calculated using recent samples of data from those nodes is greater than a user-specified threshold (γ^k_{th}). The sleep duration is calculated based on the degree of temporal correlation.

The REPs also exploit the temporal correlation characteristics of their sensed data to adapt the rate of control message broadcasts and data transmissions to the sink. Furthermore, ASSOCs wake up periodically and can identify or track any variation in the spatial distribution of the manifestation over time and change their state accordingly in order to enable accurate reconstruction at the sink. The *accuracy* measure verifies whether the distributed solution follows the variation in the manifestations of the phenomenon while still satisfying user-specified thresholds.

To put the applicability of our work into context, consider the following scenario. The growing popularity of cloud computing has led to an increase in the size and number of datacenters. The operating costs are becoming extremely high with a significant portion of it being costs associated with cooling. Meanwhile, there is an increasing awareness and emphasis on green computing practices that encourage energy-efficient design, operation, and maintenance of computing infrastructure. To balance these conflicting demands, we envision the use of an intelligent, non-invasive, easily deployable, wireless network of heterogeneous sensors feeding vital information to help in the design of environment-aware and energy-efficient solutions for datacenters. For instance, consider instrumenting a large High Performance Computing (HPC) datacenter (consisting of 1000 racks and 50 blade servers in each rack) with temperature and humidity sensors on each server (50000 in total). SILENCE, when running on such a sensing infrastructure will exploit the spatial and temporal correlation in the phenomenon to eliminate redundancy and to reduce the cost of processing and communicating potentially large volumes of data (of the order of gigabits) to a monitor node.

We evaluated our scheme through real experiments on a Wireless Sensor Network (WSN) testbed of TelosB sensor motes monitoring temperature distribution in a rack of 13 servers and also through extensive simulations on TOSSIM, the TinyOS simulator. We studied the trade-off between gains in terms of energy cost savings (for sensing and communication) and loss in accuracy. Using our algorithms, we observed that we can achieve up to approximately 50% reduction in the number of nodes (REPs) transmitting data to the sink (remote processing center) while also significantly saving on energy and communication costs (approximately 30%) in our experimental and simulation scenarios. We compare our distributed heuristic approach with the benchmark provided by a centralized optimization problem. We also demonstrate the effectiveness of our solution when monitoring multiple manifestations and in high-density deployment scenarios.

The following are the contributions of our work.

- We propose a decentralized data-centric adaptive sampling scheme (SILENCE) that elects appropriate representatives for selective data reporting to the sink while maintaining user-specified reconstruction accuracy.
- We also propose to combine adaptive sampling with a data- and communication-centric sleep scheduling to bridge the gap between data-centric distributed sensing and connectivity issues in a sparsely connected WSN.
- We allow the user to steer the performance of the sensor-based system (in terms of the number of REPs reporting to the sink and the accuracy in reconstruction of the manifestations) through specification of two key thresholds.

The remainder of the paper is organized as follows. Sect. 2 highlights the contribution of this work and compares it with existing approaches. Sect. 3 discusses the optimal centralized approach and provides insights for the design of our distributed solution. Sect. 4 describes our scheme for autonomic adaptive sampling in detail. Sect. 5 presents the experiment and simulation setup and discusses the performance evaluation of our proposed solution. Finally, we conclude our paper in Sect. 6.

2. OUR CONTRIBUTION

Recent work on energy-efficient thermal management of datacenters assume that the information required to make thermal-aware decisions (such as cooling system optimization and/or workload redistribution) is readily available without considering the communication and computation overhead involved in the collection and processing of huge amounts of raw sensor measurements [16, 2, 3]. In reality, thermal- or cooling-aware datacenter management schemes require information about i) inlet and outlet fan temperature for each machine (blade or chassis), ii) CPU or core utilization for each machine, iii) Computer Room Air Conditioner (CRAC) air inlet and outlet temperature, iv) the fan speeds (CRAC and computing unit exhaust fans), v) power specifications of each computing machine type, and vi) workload information (duration, start time, arrival frequency, etc).

A significant contribution of our work is to reduce amount of information required for challenging real-world applications such as thermal-aware datacenter management without penalizing the effectiveness of the management decisions due to inaccurate reconstruction of the phenomenon. Ours is also the first work to jointly consider adaptive sampling for collecting data and sleep-scheduling for energy savings while taking connectivity issues in WSNs into account. Previous solutions for energy saving using

cluster-heads/local base stations focus on how to reduce the quantity of the data transmitted to the sink, but do not take the quality of the data in account. However, our solution ensures data quality by employing similarity along with correlation for exploiting spatial and temporal correlation in sensor data. The number of REPs selected and their spatial distribution is determined by the reconstruction accuracy that the user wants eventually. This solution is integrated with a mechanism to ensure end-to-end connectivity (i.e., any REPs to the sink) that are directly/indirectly overlooked in many literatures related to clustering algorithms in WSNs.

Our approach is different from traditional *clustering algorithms* (e.g., [27, 15, 5, 4]), which reduce global communication to prolong the lifetime of the network, or *self-organization mechanisms* (e.g., [22, 23, 6, 17]), which use dynamic role assignment to extend the lifetime of the network while also reducing communication cost. Such solutions, in fact, have largely overlooked autonomic integration between self-configuration and data interpretation. In such schemes, the group leaders (cluster heads/ local base stations) are selected not based on the quality of the data that the user needs but based on the constraints such as lifetime of the network, energy, and wireless link quality.

Also, SILENCE differs from solutions that perform in-network processing of data to eliminate redundancy such as *compression* (e.g., [11, 18]), *data aggregation* (e.g., [13, 10]), *source coding* (e.g., [8, 19]), and *routing and data compression* (e.g., [21, 20]) as they require constant local communication inside a cluster/group of nodes. Instead, our solution puts the nodes to sleep in order to reduce local communication and to save energy, while ensuring user-specified levels of accuracy in data reconstruction. This approach is different from previous sleep-scheduling algorithms [26, 7] because it adjusts sleep duration based on the data correlation while others are scheduled for increasing network lifetime. In SILENCE, connectivity issues are addressed and solved by using a mechanism employing AWAKE packets to notify that the node is no longer sleeping, but awake and ready for relaying the packets.

An interesting feature of SILENCE is the use of similarity along with temporal correlation for determining the spatial correlation in sensor data. We achieve resource efficiency (reduction of costs for processing and global communication of sensor data) by allowing only a subset of nodes (REPs) to send meaningful data to the sink while the rest of the nodes (ASSOCs) sleep. This representation of a group of ASSOCs by a single REP is possible only due to the use of similarity along with correlation. If only correlation were considered, a REP would end up representing nodes that experience only a correlated trend in variation of the manifestation with that of its own and not similar values.

Interestingly, SILENCE also allows the user to steer the performance of the sensor-based system (in terms of the number of REPs reporting to the sink and the accuracy in reconstruction of the manifestations) through specification of two important thresholds for similarity (e^k_{th}) and correlation (γ^k_{th}) for each manifestation k associated with a phenomenon. The user or administrator of the networked system can decide the granularity at which s/he needs data from the sensing system. For instance, in datacenters, sensed temperature and humidity values are crucial as they convey the operating environment of servers that may be handling sensitive and crucial data. Their values directly reflect on the performance of the machines and, hence, they should be monitored at a very fine granularity. Another scenario could be a networked sensing system monitoring a green house botanic garden where larger thresholds could be set as it may be sufficient to obtain measurements at a coarse granularity. SILENCE is easily implementable and scalable to large and dynamic sensor-based systems. The highly decentralized nature of our solution is demonstrated using a simulated scenario in which a localized change in the spatial distribution of the manifestation is identified and dealt with locally without involving nodes that are not spatially proximal and unconcerned with that change.

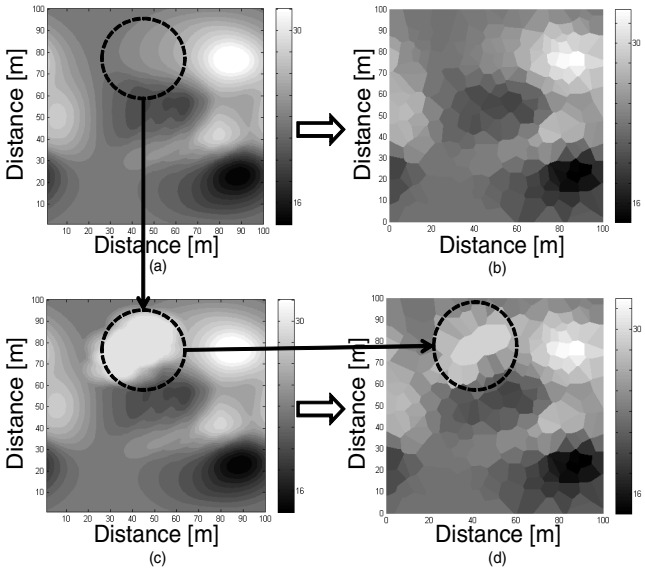

Figure 1: (a) Spatial distribution of a manifestation (temperature) in a 2D field; (b) Reconstructed field with data reported only by 210 REPs out of 400 nodes; (c) Variation in spatial distribution of the manifestation (area marked by dashed circle); (d) Reconstruction of the modified field with data from 218 REPs after some AWAKE − ASSOCs change state to become REPs.

The following simulated scenario illustrates the reconstruction of the spatial distribution of a manifestation at the sink. The distribution of temperature in a 100×100 m^2 field and its remote reconstruction (at the sink) with data obtained from the nodes of a sensor-based system that employs SILENCE are shown in Fig. 1. The nodes determine their roles in a distributed fashion through the exchange of control messages. Figure 1(a) shows the actual spatial distribution of temperature in the field and Fig. 1(b) shows the reconstruction at the sink with the data obtained from 210 REPs out of 400 nodes deployed. When there is a change in the spatial distribution of data over time in the same field, as shown in Fig. 1(c) (marked by the dashed circle), some ASSOCs capture the variation and change their state to help the sink reflect the same in its reconstruction, as depicted in Fig. 1(d). Note that SILENCE neither assumes any pre-deployment optimization of nodes nor a priori knowledge of spatio-temporal correlation of the manifestations.

3. CENTRALIZED APPROACH

In this section, we discuss the REP selection problem of SILENCE for a deployment of N nodes in a 2D field as a centralized optimization problem. The motivation for this exercise is twofold: firstly, even though optimal centralized solution is impractical, complicated, and non-scalable, it gives us insight to make key design decisions for our heuristic distributed solutions; secondly, the performance of the optimal centralized approach serves as a benchmark for comparison.

The maximum error e^k_{max} in reconstruction observed over all N nodes in the field should be minimized to find the optimal set of

ASSOC and REP nodes. Here, the error in reconstruction is defined as the absolute difference between the actual value measured by a node on the field and the value that is reconstructed at the sink based on the information from its REP. The optimization problem should aims at finding the *optimal set* of REPs that minimize this maximum error in reconstruction to a value below the threshold specified by the user. When the number of nodes, N, the number of REPs, M, and the manifestation k are given, the problem should find the optimal set of M REPs that would minimize the maximum error e^k_{max} in reconstruction. The objective of the optimization problem should be to minimize the maximum error (given by the infinity norm of a vector of error values) in reconstruction of manifestation k, when only M out of the available N nodes report their measured values to the sink. Constraints to the problem should force the total number of REPs to be equal to M, should force a REP to be associated only with itself and also ensure that only ASSOCs are associated to REPs and not vice-versa, and finally, should force every node that is not a REP to be associated with one and only one REP.

As our objective is to find the minimum number of REPs, M, and the best possible set of those from the total number of nodes, N, so that the error threshold set by the user is satisfied, we should solve $N \cdot \binom{N}{M}$ of these problems for every snapshot by fixing M in each problem. The complexity of the optimization problem increases as the number of nodes N increases and is impractical for real-world deployment. However, the centralized problem provides us with insights for devising distributed mechanisms to select the best set of REP nodes while minimizing the error in reconstruction of the phenomenon. To achieve the dual objective of energy efficiency and minimization of reconstruction error, we follow a divide and conquer approach to split the problem of finding the best set of REPs into a number of localized optimization problems. In this strategy, measured values are locally exchanged between nodes and REPs are elected through a distributed election procedure.

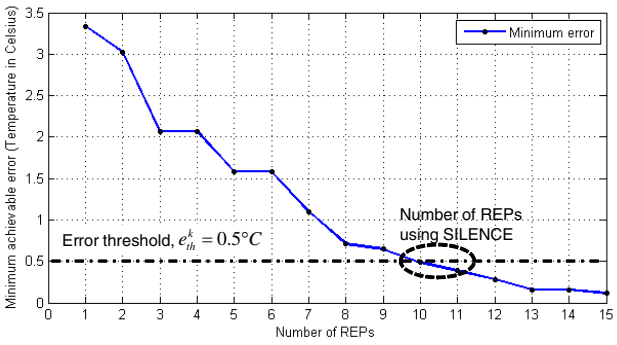

Figure 2: Error in reconstruction for different number of REPs.

We also performed comparisons between the centralized approach and SILENCE on a 16 node sensor network deployed in a 50×50 m^2 field. The minimum of maximum reconstruction error obtained when the centralized problem is solved by varying M from 1 to 16 is shown in Fig. 2. The maximum reconstruction error falls within the threshold we set $e^k_{th} = 0.5\,^\circ C$ only when the number of REPs, M, is at least 11. From our simulation in TOSSIM for the same deployment, field and error threshold, we observed the number of REPs chosen by SILENCE fluctuates between 11 and 12, which is close to the optimal solution. Difference in the solutions of the centralized approach and the distributed approach can be attributed to the non-idealities that are introduced by the wireless link and limited sensing and communication range of sensor nodes. Detailed comparisons between SILENCE and the centralized solution with different error thresholds are presented in Sect. 5.1.

4. PROPOSED SOLUTION

SILENCE enables each node in autonomic sensor-based systems to determine its role or state dynamically and independently (*self-configure*) through the exchange of control messages. The three states (two primary and one auxiliary) that a node can be in are REP, ASSOC, and AWAKE – ASSOC. Primarily, a node can either be a REP or an ASSOC; REPs periodically transmit sampled data to the sink along with the list of their ASSOCs so that the sink can accurately reconstruct the spatial distribution of the manifestations, while the ASSOCs go to sleep to save energy. The other auxiliary state that a node can be in is a special case of the primary ASSOC state: an AWAKE – ASSOC is an ASSOC that has woken up to verify if its newly sampled data is still correlated and similar to that of its REP and to aid in relaying packets to the sink. It changes its state to become a REP and reports to the sink if there is a significant variation in the spatial distribution of data over time (*self-heal*). *Thus, the solution allows only an appropriate (small) subset of nodes to send meaningful data to the sink and strives to incur only lower-than-acceptable loss of accuracy in the reconstruction of the phenomenon.*

To realize the adaptive sampling goals, SILENCE relies on very basic assumptions about the underlying sensor-based system of the broader CPS. Firstly, each sensor node i is aware of its position $\mathbf{p}_i$ in the field in the field. This is necessary to spatially reconstruct the data at the sink. Secondly, we assume that REPs in the field can communicate with the sink over multiple hops, if required, through the use of an appropriate routing communication protocol, i.e., geographical routing.

4.1 Similarity and Correlation

The key component of SILENCE, i.e., transition between roles or states of a node is based on i) events and ii) control messages, depends on two important metrics: 1) *similarity* and 2) *correlation*. Two nodes n and m are said to be sensing similar values if the difference between the means of their measured values of a manifestation k, e^k_{nm}, is less than a user-specified threshold, e^k_{th}. Measured values of manifestation k at n and m are said to be correlated if the correlation coefficient, γ^k_{nm}, which is calculated using S samples of data from those nodes, is greater than a user-specified threshold, γ^k_{th}. The formulae used for determining similarity and correlation are as follows,

$$e^k_{n,m} = |\overline{\psi}^k_n - \overline{\psi}^k_m|, \tag{1}$$

where $\overline{\psi}^k_n$ and $\overline{\psi}^k_m$ are the means of S samples of the manifestation k at nodes n and m, respectively. The correlation coefficient $\gamma^k_{n,m}$ is given by,

$$\gamma^k_{n,m} = \frac{\sum_{s=1}^{S}[\psi^k_{ns} - \overline{\psi}^k_n] \cdot [\psi^k_{ms} - \overline{\psi}^k_m]}{\sqrt{\sum_{s=1}^{S}[\psi^k_{ns} - \overline{\psi}^k_n]^2} \cdot \sqrt{\sum_{s=1}^{S}[\psi^k_{ms} - \overline{\psi}^k_m]^2}}. \tag{2}$$

If n and m are sensing similar and correlated values, then n is a potential-ASSOC of m and viceversa. Actual-ASSOCs are finalized after control message exchanges.

4.2 State Transitions

The three states that a node on the field can be in at any point in time as well as the events and messages that trigger state transitions are depicted in Fig. 3. Initially, after deployment, all the sensor

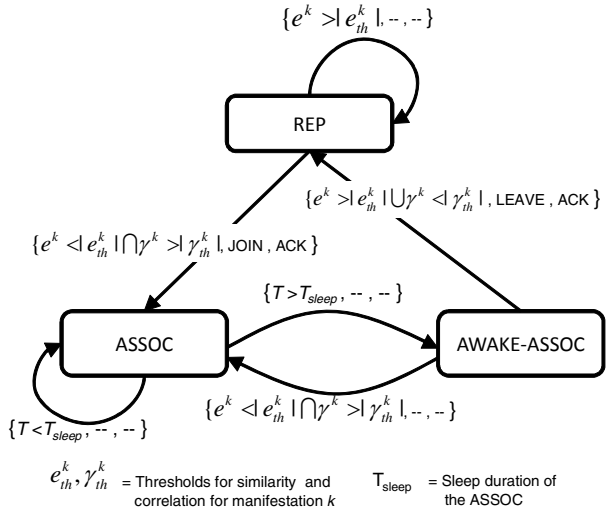

Figure 3: Illustration of states and state transitions ($\{Event, Tx, Rx\}$ represents the event and control messages that trigger a state change).

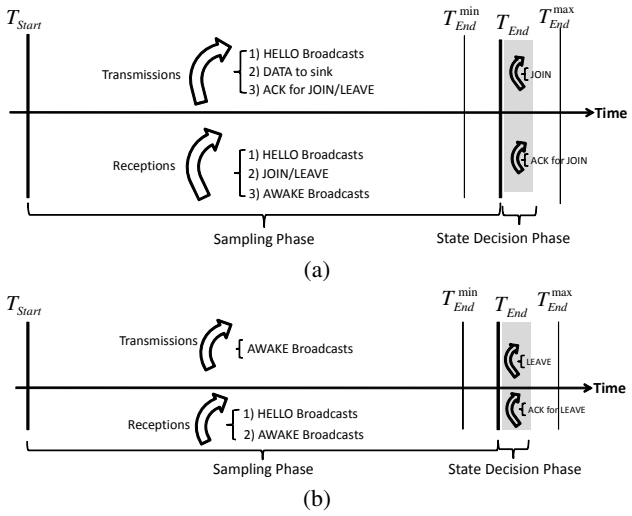

Figure 4: (a) Incoming and outgoing messages at a REP **node over time; (b) Incoming and outgoing messages at a** AWAKE – ASSOC **over time.**

nodes in the field are in the REP state. Each REP periodically samples every manifestation k of the phenomenon and transmits it to the sink if there is connectivity or advertises the measured values to its neighbors through HELLO broadcasts in an independently-determined sampling duration. For each node, the Sampling Phase starts at T_{Start} and ends at T_{End}, randomly chosen between T_{End}^{min} and T_{End}^{max}, as shown in Fig. 4(a).

HELLO packet from a node contains time series values of the manifestations (sampled data) observed at that node, additional information about the number of neighboring nodes that can be potential-ASSOCs, and the number of neighboring nodes that are actual-ASSOCs of that REPs. In the meantime, a REP sends DATA packets to the sink periodically and also receives a number of HELLO broadcasts from its neighboring REPs. The structure of HELLO and DATA packets are shown in Fig. 5. The frequency with which both the packets are transmitted is adaptively determined from the rate of change of the manifestations over time. This is evident from the way the time series measurements of each manifestation is conveyed in the packets as <Value, Duration> pairs: the slower the rate of change of manifestations at a sensor node, the lower the frequency of transmission of both types of packets. This component of the solution is called adaptive data reporting.

Figures 4(a) and (b) show the two different phases, namely, the Sampling Phase (whose duration is T_{End}-T_{Start}) and the State Decision Phase (whose duration is less than T_{End}^{max}-T_{End}) that a REP and ASSOC can be in. The curved arrow pointing towards the time axis represents incoming messages/packets and the curved arrow pointing away from the time axis represents outgoing packets. With the samples in the DATA or HELLO packets that each REP receives from its neighbors, a REP calculates its updated list of potential-ASSOCs and actual-ASSOCs, and appends this information in its future HELLO broadcasts. At T_{End}, the REP stops sampling and makes a decision whether to stay in the same state or transition to be an ASSOC to another REP. This decision is entirely based on similarity and correlation between its own sampled data and the data of neighboring nodes obtained from received HELLO packets. With the most recently updated list of potential-ASSOCs based on similarity and correlation, a node decides on its future state as shown in Algorithm 1. A REP switches to the ASSOC state only if it finds a suitable REP that satisfies both the similarity and correlation thresholds and has a higher number of potential-ASSOCs than itself.

Algorithm 1 State decision using similarity and correlation.

INIT_ STATE_ DECISION:
$IS\_ASSOC = 0, MY\_REP = 0$
e_{th}^k, γ_{th}^k = {User-specified threshold for similarity and correlation}
$LIST\_REP_{PA}$ = {List of potential REPs}

STATE_ DECISION:
In $LIST\_REP_{PA}$ find $LIST\_REP_{PA}^{max}$
if SIZE ($LIST\_REP_{PA}^{max}$) == 1 **then**
  $MY\_REP = LIST\_REP_{PA}^{max}$
  $IS\_ASSOC = 1$
  Send JOIN, Receive ACK
  Sleep (T_{Sleep})
else
  In $LIST\_REP_{PA}^{max}$ find REP_{ID}^{max}
  $MY\_REP = REP_{ID}^{max}$
  $IS\_ASSOC = 1$
  Send JOIN, Receive ACK
  Sleep (T_{Sleep})
end if

After a list of potential REPs is available ($\text{LIST\_REP}_{\text{PA}}$) in the Sampling Phase, each node chooses the one with the highest number of potential associates ($\text{LIST\_REP}_{\text{PA}}^{\max}$) in the State Decision Phase (Fig. 4(a)). Node IDs are used to break the deadlock – a situation where two REPs have similar number of potential associates and one has to be chosen for association – if there is one. The transition is complete with the exchange of JOIN and ACK packets that contain sleeping duration T_{Sleep} based on the correlation. After the state transition, the new ASSOC node goes to sleep during T_{Sleep}. Conversely, if the REP chooses to continue in the same state, the whole cycle starting from T_{Start} repeats itself. The system is capable of avoiding undesired associations between nodes that are within each other's RF range and that are accidentally measuring similar values. This ability of the system can be attributed to the intentionally long Sampling Phase employed at every REP (initially all nodes are REPs) and AWAKE-ASSOC as shown in Figs. 4(a) and (b). The long Sampling Phase ensures that undesired associations due to transient and erratic behaviors of sensors do not happen.

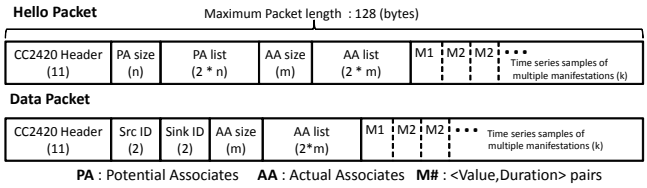

Figure 5: Structures of HELLO and DATA packets.

4.3 Sleep Scheduling and Connectivity

In the case of WSNs, ensuring connectivity among the REP nodes and the sink/base station is as important as saving energy and wireless communication bandwidth. In SILENCE, the sleep duration T_{Sleep} for an ASSOC can be fixed or can be adaptively adjusted based on the degree of similarity and correlation of the ASSOC and REPs data, i.e., $T_{Sleep} = \phi(e^1 \dots e^K, \gamma^1 \dots \gamma^K) \cdot T^{min}_{Sleep}$, where K is the maximum number of manifestations. The activity of an AWAKE – ASSOC (an ASSOC after its sleep duration) is shown in Fig. 4(b). In the Sampling Phase, an AWAKE – ASSOC listens to HELLO broadcasts from all the neighboring REPs including its own. It does not transmit any HELLO packets. At the end of the Sampling Phase, it verifies similarity and correlation of its own sampled data and that of its REP's received recently. *If each of the thresholds e^k_{th} and γ^k_{th} for all the manifestations are satisfied, the node switches back to* ASSOC *state and sleeps again.* If any of the threshold is not satisfied, the AWAKE – ASSOC switches to REP state after exchanging LEAVE and ACK packets with its old REP.

Even though REPs are always 'ON' for relaying packets from other nodes, reliable connectivity from every REP to the sink node cannot be ensured in WSNs. This is because SILENCE has elected the REPs in a data-centric manner. Hence, AWAKE – ASSOCs should be utilized to opportunistically relay packets for multi-hop communication. However, connectivity between AWAKE – ASSOCs cannot always be ensured because they may not have information about the next node to route the packet toward the sink. Hence, AWAKE – ASSOCs use an AWAKE broadcast packet for announcing availability for multi hopping. The best next hop (ASSOC/REP) with respect to delay or distance could be selected by the REP because it has the sleeping duration and location information of its ASSOCs exchanged during association. AWAKE packet contains duration of the availability ($T_{End} - T_{Current}$) and location information for geographical routing. Neighboring AWAKE – ASSOCs and REPs overhear this packet and choose the next best hop depending on the available information and the routing strategies (e.g., minimum delay, minimum number of hops, minimum energy).

4.4 Toy Example

We present a *toy example* to help understand REP selection in SILENCE better. Figure 6 shows node ID and number of potential-ASSOCs for every node. Solid circles represents the nodes that have similar and correlated values in time. Dotted lines between the circles represent whether a communication link exists between the nodes or not. Once the network is deployed, every node is a REP that senses and communicates its data. Based on their own data and on the information from the neighbors, each node computes its potential number of ASSOCs and advertises that it could be a REP should it have at least one potential associate. Nodes with no potential-ASSOCs due to absence of communication with others (here node i) or due to dissimilarity in data (here node j and k), will continue to be REPs. Nodes c and f will decide to continue to be REPs as they know they have the highest number of potential ASSOCs based on the information about their neighborhood. Nodes

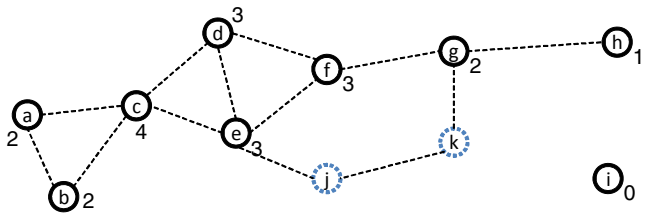

Figure 6: Example of REP selection (small-lettered alphabet represents node IDs, while the numbers below represent the number of potential-ASSOCs).

a, b, d, and e will decide to be ASSOCs of node c and will send a JOIN message to c. Once they receive an ACK from c they switch states. Similarly, node h will send a $JOIN$ message to the node g, which will send a $JOIN$ request to node f. An important feature to note in this scheme is that $JOIN$ and ACK transactions happen in a window of time and an ASSOC does not go to sleep immediately after an ACK. This allows h to join g, which may have already decided to be an ASSOC to f. In that case, g accepts the request from h and will also notify both f and h about this situation.

5. PERFORMANCE EVALUATION

The performance of SILENCE, our adaptive sampling scheme for autonomic sensor-based systems, was evaluated through real experiments using a WSN composed of TelosB motes and repeatable simulations on TOSSIM, the TinyOS 2.x simulator. We also compared its performance against the centralized optimal REP selection procedure. As the time complexity of the combinatorial problem is a limiting factor, we could only compare the performance for a small deployment of nodes. We perform real experiments to demonstrate that SILENCE is easily implementable and realizable in practice. Since large-scale experiments could not be conducted, we use simulations to study the performance of our algorithm in large-scale and high-density deployment scenarios. The experiment and simulation setup as well as their results are detailed in the following.

5.1 Real Experiments

Figure 7 shows our real experiment setup of 26 TelosB motes densely deployed on 13 servers/blades in a server rack. The blades

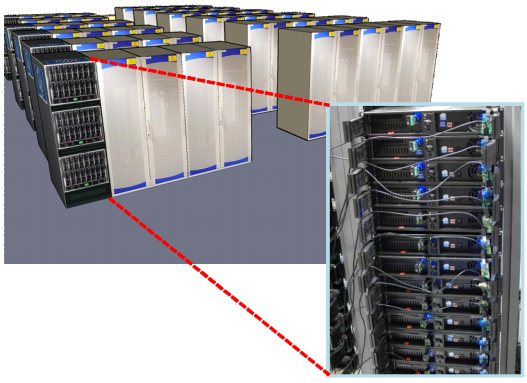

Figure 7: Experiment setup - 26 telosB motes (deployed in front of the blades) measuring multiple manifestations such as temperature, humidity, and luminescence.

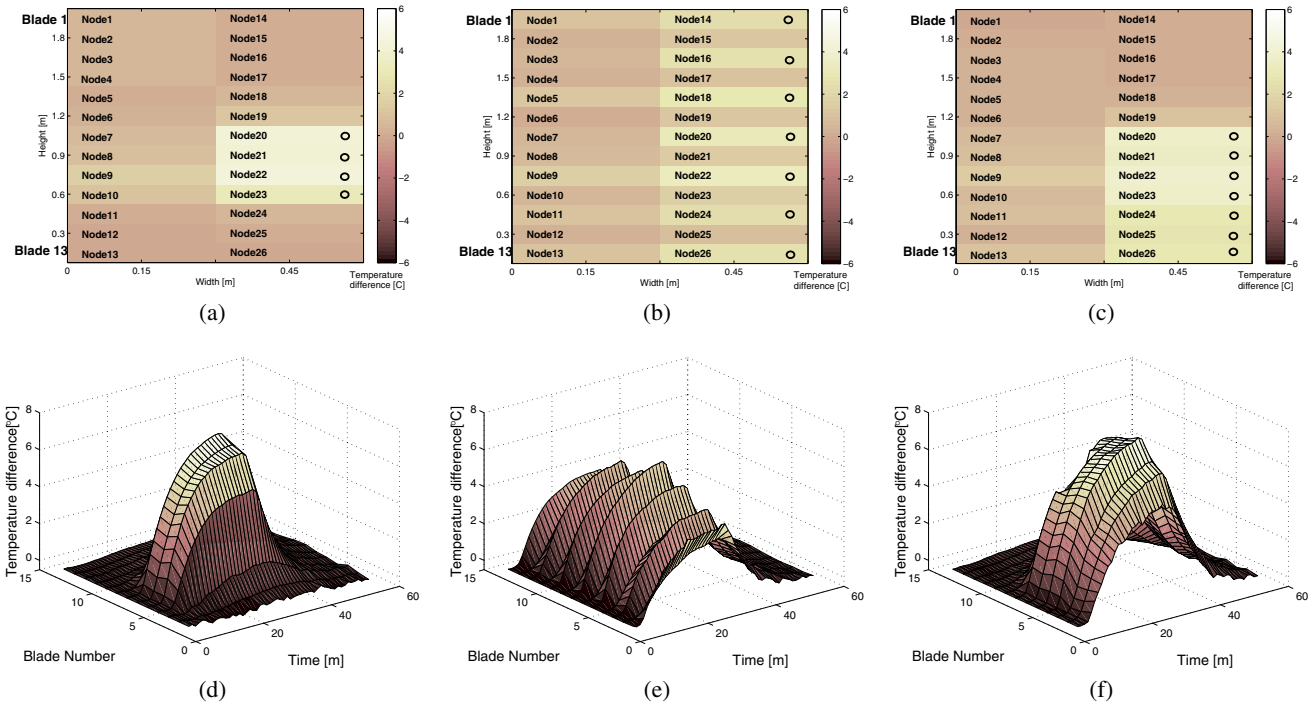

Figure 8: Experiment results - Spatial correlation (a, b, and c) and temporal correlation (d, e, and f) in measured data (temperature) among the 26 nodes.

are shown in the context of a large datacenter though we focus on monitoring only one server rack in the machine room at the NSF Center for Autonomic Computing at Rutgers University. The temperature sensors on the motes measure the external temperature and are used to estimate heat generation and distribution during the operation of these servers [12]. Two sensor motes are deployed at the front of each server blade, one right in front of the outlet fan and one farther away at the other end. Figures 8(a)-(f) visualize the strong spatial and temporal correlation in measured temperature data (single snapshot from sensors placed close to the outlet fans) among the 26 sensor nodes based on the workload distribution. The blades are numbered 1 to 13 from top to bottom.

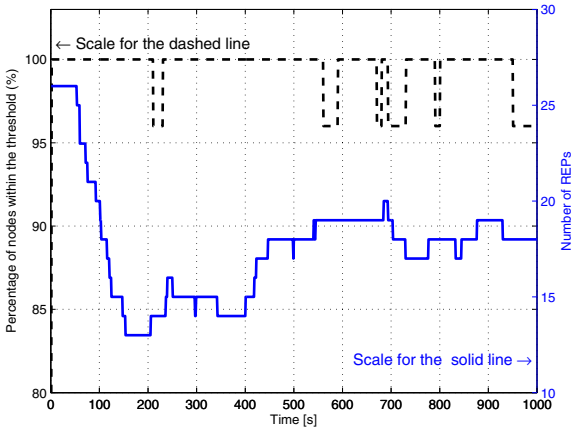

Figure 9: Experiment results - Number of REPs and percentage of nodes within the error threshold specified by the user (here $e_{th}^{k} = 0.5\,^{\circ}C$ for temperature) over time.

In Fig. 8(a) it can be seen that when blades 7-10 from the top are operational (marked with solid circles) there is a significant difference in temperature between the operational and idle server blades. Moreover, due to the spatial proximity of blades 6 and 11 to the operational ones they also experience a higher temperature compared to the other idle ones. Similar spatial correlation can be observed in the scenarios depicted in Figs. 8 (b), and (c), where the odd numbered blades and blades 7-13 are are operational respectively. The dipicted spatial correlation was observed at a particular snapshot where the thermal impacts of running blades are high. We observe significant spatial correlations between server temperature.

Figures 8 (d), (e), and (f) show the temporal correlation of the measured data among the 13 sensor nodes (in front of outlet fans of each blade) over time for the operational scenarios depicted in Figures 8 (a), (b) and (c). We run SILENCE on this sensor network to verify whether it effectively exploits the spatial and temporal correlation in the observed manifestation (temperature) to elect only an appropriate number of REPs for reporting data to the sink. We set the error threshold $e_{th}^{temp} = 0.5\,^{\circ}C$ and correlation threshold $\gamma_{th}^{temp} = 0.75$. Figure 9 shows the number of REPs transmitting data to the sink as the manifestation changes over time and the reconstruction accuracy i.e., number of nodes within the error threshold specified by the user. It can be observed that as the thresholds are relaxed, the number of REPs decreases further for the same phenomenon as only the minimum number of REPs required for a specified reconstruction accuracy are selected. The savings in terms of communication cost is evident as a fewer number of nodes are reporting data to the sink compared to the base case (when all nodes transmit).

To verify whether SILENCE achieves its objective of selecting the appropriate number of REPs, we compare its result with the outcome of a centralized optimal REP selection for different error thresholds ($e_{th}^{temp} = 0.5\,^{\circ}C\,and\,1.0\,^{\circ}C$). Figure 10 shows that the

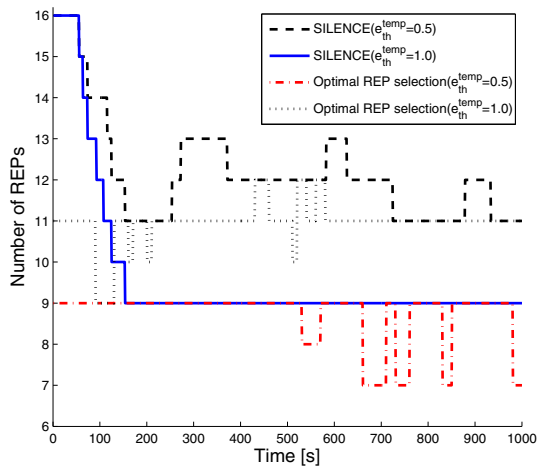

Figure 10: Experiment results - Comparison of performance of SILENCE with centralized optimal REP selection.

number of representative nodes selected by SILENCE is comparable to the optimal value for different error thresholds.

5.2 Simulations

As the size of our testbed was a limiting factor in the study of the performance of SILENCE under high-density and large scale deployment scenarios, we performed additional simulations on the TinyOS simulator, TOSSIM. For simulations, it is key that the right models are used for the phenomenon. The spatio-temporal distribution of manifestations considered in simulations (temperature and humidity) is modeled based on the characteristics of actual data measured in the testbed. We assume that we have 400 motes measuring temperature and humidity data, which has the same range (minimum to maximum) as the real observed testbed data and spatial correlation scaled in distance for simplification. The temporal correlation was also scaled in time to obtain three different fields that vary at different rates (*slow*, *moderate*, and *fast*). We set the rate of variation of manifestations between two significant values (differing by at least 5%) to be the same rate as the one observed in actual measured data. For a moderately varying field we modified the rate of variation to 5 times the original (or slow field) and for a fast varying field we made it 10 times the original. The manifestations' rate of variation in space and time are uncontrollable in real experiments, which also serves as a motivation for our simulation study. The study of performance on slow-, medium, and fast-varying fields is intended to convey the idea that the proposed solution is not limited to slow-varying band-limited phenomenon. The deployment, channel, and radio parameters used in the simulation are based on [1] and are listed in Table 2.

The two metrics used to measure the performance of our solution are: 1) *the energy cost* (energy spent in Joules per second in the sensor-based system) due to sensing and communication, and 2) *the accuracy of reconstruction* of the manifestations (percentage of nodes with reconstructed values lying within the error thresholds specified by the user). The amount of data injected by our solution into the network includes all *control* (HELLO, JOIN, LEAVE, ACK and AWAKE – ASSOC packets) and *data traffic* (DATA packets). The energy cost is calculated as follows.

$$E = E_{elec} + E_{comm}; \quad E_{comm} = V \cdot I \cdot \frac{L}{R}, \tag{3}$$

where E_{elec} [J] is the energy consumed by the electronic circuit, E_{comm} [J] energy consumed for communication, V is the battery voltage (Volts), I is the current (Ampere), L is the packet size, and R is the radio transmission/reception rate. SILENCE using adaptive sleep scheduling is compared against SILENCE employing fixed sleep scheduling and a generalized LEACH-like solution [9] referred to as LEACH-gen.

Our implementation of LEACH-gen, does not employ the localized election mechanism for REP selection. The nodes randomly decide to be REPs or ASSOCs based on a probability distribution. We considered two options for modification LEACH for a fair comparison with SILENCE. Option (1) was to mimic original LEACH in which the ASSOCs do not go to sleep. In fact, in original LEACH, there is constant intra-cluster communication so to perform some signal processing on the data, which is then transmitted to the sink by the REP. Comparison of SILENCE with such an approach would be unfair because even though the reconstruction error for LEACH can be comparable to SILENCE (if appropriate data processing is done at REPs), the energy consumption due to permanently awake ASSOCs and network congestion due to high overhead for constant intra-cluster communication will be very high for LEACH. Option (2) was to mimic LEACH only with respect to the REP selection process (based on a probability distribution) and then put the ASSOCs to sleep. This way, there is no intra-cluster data exchange after REP selection and no data processing at the REPs just as in SILENCE. Our intention was to maintain the focus only on the REP selection process, method for elimination of redundancy in reported data, and their effect on reconstruction of the phenomenon. Hence, we chose option (2).

For slow, moderate, and fast varying fields, Fig. 11 shows the number of REPs transmitting sensed data of different manifestations (temperature and humidity values) to the sink at different points in time and the corresponding percentage of nodes with value estimates (reconstruction) within the error thresholds for the two manifestations specified by the user. Even with a stringent target reconstruction accuracy of $e_{th}^{temp} = 0.5\,^{\circ}C$ and $e_{th}^{hum} = 3\%$, SILENCE achieves up to 50%, 40%, 25% reduction in the number of nodes transmitting data to the sink (REPs) for slow, moderate, and fast variation rates for the field, respectively (Fig. 11). The

Table 2: Parameters of the model used for simulations

Physical deployment parameters	
Terrain dimension	200x200 m^2
Topology	Uniform random
Number for nodes	400
Channel parameters	
Path loss exponent	3.3
Shadowing standard deviation	5.5 dB
Reference distance (D_0)	1 m
Path loss at reference distance	-30 dBm
Radio Parameters	
Transmission power	-1 dBm
White gaussian noise	4 dB
Radio noise floor	-105 dBm
Hardware Variance (For highly asymmetric links)	
Covariance matrix	S = [S11 S12 ; S21 S22]
S11 (variance of noise floor)	3.7
S12 (covariance btw S1 and S1)	-3.3
S21 (same as S12)	-3.3
S22 (variance of output power)	6.0

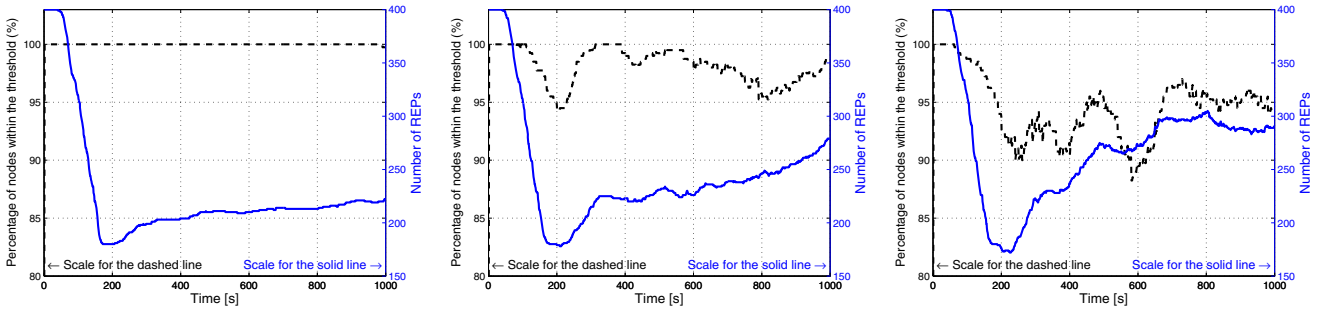

(a) Slowly varying field (b) Moderately varying field (c) Fast varying field.

Figure 11: Simulation results - Number of REPs over time and percentage of nodes within error threshold set by the user (here $e_{th}^k = 0.5\,^\circ C$ for temperature) for (a) a slowly varying field, (b) moderately varying field, and (c) fast varying field.

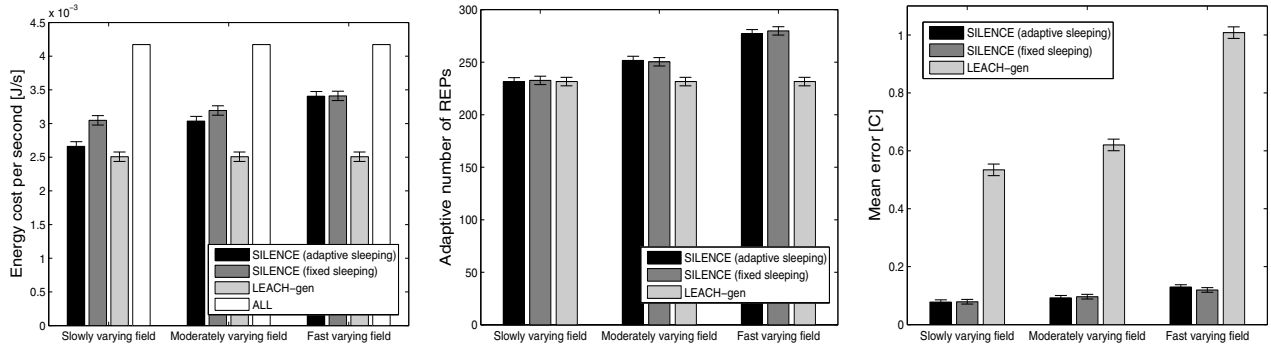

(a) (b) (c)

Figure 12: Simulation results - (a) Average energy cost incurred by different adaptive sampling schemes; (b) Average number of REPs of different adaptive sampling schemes; (c) Average error of different adaptive sampling schemes for different rates of variation in manifestations.

percentage of nodes for which the reconstruction accuracy is satisfied is also shown. The reconstruction error fluctuates and the threshold is violated for only up to $\sim$10% of the nodes in the fast varying scenario thus providing an insight into the limit on our solution's performance. The transient and steady state of the networked system when employing SILENCE can be clearly identified in the graphs. As the average of the error alone does not provide the complete picture, we have also shown the percentage of nodes that are within the error threshold (specified by the user) throughout the duration of the experiment/simulation. Figure 11(c) clearly shows that even for a fast varying field, which is 10 times faster than the slow-varying field, the number of nodes within the threshold is not lower than 90% in the worst case. Also, the average error (over all nodes) in Fig. 12(c) for a fast varying field is less than the threshold specified by the user. Hence, we infer that SILENCE allows for timely reaction to fast changes in the field.

Figure 12(a) shows the energy cost (network energy expenditure in Joules per second) incurred by SILENCE with fixed sleep scheduling, SILENCE with adaptive sleep scheduling, and LEACH-gen for the three different scenarios mentioned earlier. The energy cost for SILENCE is more than the one incurred by a sensor-system employing LEACH-gen. This can be attributed to the control overhead incurred by SILENCE. The graphs depicting control overhead is not shown here as it follows the same trend as Fig. 12(a) and due to space constraints. The average error in reconstruction, however, is higher for LEACH-gen compared to SILENCE as shown in 12(c). This is because of SILENCE's sensitivity to variation in the manifestations of the phenomenon unlike LEACH-gen (a random REP selection procedure), which only uses information about the average number of REPs and does not adapt. LEACH-gen's inability to adapt is evident in Fig. 12(b). The extra overhead incurred by SILENCE is justified by the significant reduction in reconstruction error compared to LEACH-gen as shown in Fig. 12(c). From Figs. 12 (a), (b), and (c) one can clearly infer that SILENCE captures the variation in the manifestations and self-heals by adjusting the number of (REPs) using marginally additional control overhead in exchange for higher reconstruction accuracy at the sink. The control overhead to select the best set of REPs is negligible compared to the case when all nodes in the WSN are transmitting sensed data to the sink.

Multiple Manifestations: We also performed simulations to verify the performance of SILENCE when multiple manifestations (temperature and humidity) are taken into consideration. Several strategies can be adopted to adapt SILENCE for sensor networks monitoring more than one manifestation. Following are two possible strategies: i) two nodes n and m are said to be potential-ASSOCs of each other only when they sense similar and correlated values for all the manifestations, i.e., only when e_{th}^k and γ_{th}^k are satisfied for all $k = 1, \ldots, K$ and REPs are selected accordingly, and ii) a different set of REPs are selected individually for each manifestation k to transmit data to the sink.

Figure 13 shows the performance of our solution when we follow the first strategy considering two manifestations, temperature and humidity. The model for humidity is again based on the actual data observed in our measurements and experiments on the testbed setup. It can be seen that roughly 40% of the nodes are suf-

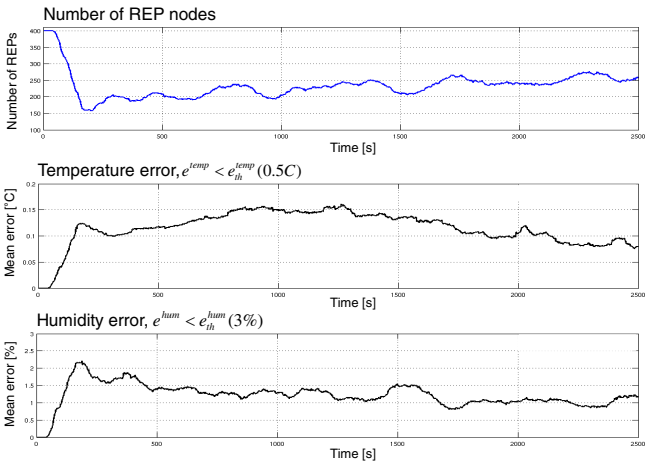

Figure 13: Simulation results - Performance of SILENCE scheme with multiple manifestations.

ficient to reconstruct two manifestations, temperature and humidity at a very high reconstruction accuracy, by leveraging similarity and correlation of sensor values (of both manifestations) in proximal sensor nodes. This is because the two manifestations are by nature highly correlated and the choice of REPs even when the they are considered separately may be the same. However, more sophisticated strategies such as ones that associate weights to each manifestation while determining similarity and correlation have to be adopted when the monitored manifestations are highly uncorrelated (e.g., humidity and luminescence).

6. CONCLUSIONS AND FUTURE WORK

We designed, developed, and implemented a distributed adaptive sampling solution, SILENCE, to reduce redundancy in raw data through selective representation without compromising on accuracy of reconstruction of the phenomenon at the sink. We used similarity and correlation in the sensed data and on the fly optimized the number of representatives reporting to the sink in a distributed manner in both space and time domains. SILENCE was evaluated through experiments on a testbed of sensors monitoring temperature distribution in a rack of servers and through extensive simulations on TOSSIM, the TinyOS simulator. The results obtained through experiments and simulations are encouraging and provide insights into the performance gains that can be achieved by our autonomic adaptive sampling solution in terms of energy efficiency, reduction in communication overhead, and most importantly reconstruction accuracy.

Our current testbed is small and limited to 26 sensor motes. We are building a larger wireless sensor network testbed at our research center and large-scale experiments will be performed once it is available for use. Our current and future research efforts are targeted towards investigating distributed mechanisms to adapt powerful signal processing techniques like compressive sensing for incorporation into SILENCE for efficient joint autonomic adaptive sampling of multiple correlated and uncorrelated manifestations.

7. REFERENCES

[1] Building a network topology for tossim. http://www.tinyos.net/tinyos-2.x/doc/html/tutorial/usc-topologies.html.

[2] Z. Abbasi, G. Varsamopoulos, and S. K. S. Gupta. Thermal aware server provisioning and workload distribution for internet data centers. In *Proc. of Intl. Symp. on High Performance Distributed Computing (HPDC)*, Chicago, IL, June 2010.

[3] Ayan Banerjee and Tridib Mukherjee and Georgios Varsamopoulos and Sandeep K. S. Gupta. Cooling-Aware and Thermal-Aware Workload Placement for Green HPC Data Centers. In *Proc. of Intl. Green Computing Conf. (IGCC)*, Chicago, IL, Aug. 2010.

[4] S. Bandyopadhyay and E. Coyle. An Energy-Efficient Hierarchical Clustering Algorithm for Wireless Sensor Networks. In *Proc. of IEEE Conf. on Computer Communications (INFOCOM)*, San Francisco, CA, Apr. 2003.

[5] S. Bandyopadhyay and E. J. Coyle. Minimizing Communication Costs in Hierarchically Clustered Networks of Wireless Sensors. *Computer Networks*, 44(1):1–16, Jan. 2004.

[6] M. Bhardwaj and A. P. Chandrakasan. Bounding the Lifetime of Sensor Networks Via Optimal Role Assignments. In *Proc. of IEEE Conf. on Computer Communications (INFOCOM)*, New York, NY, June 2002.

[7] S. Chachra and M. Marefat. Distributed Algorithms for Sleep Scheduling in Wireless sensor Networks. In *Proc. of Intl. Conf. on Robotics and Automation (ICRA)*, Orlando, FL, May 2006.

[8] T. Cui, L. Chen, T. Ho, S. Low, and L. Andrew. Opportunistic Source Coding for Data Gathering in Wireless Sensor Networks. In *Proc. of Mobile Adhoc and Sensor Systems (MASS)*, Pisa, Italy, Oct. 2007.

[9] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-Efficient Communication Protocol for Wireless Microsensor Networks. In *Proc. of Hawaii Intl. Conf. On System Science (HICSS)*, Maui, HI, Jan. 2000.

[10] B. Krishnamachari, D. Estrin, and S. Wicker. The Impact of Data Aggregation in Wireless Sensor Networks. In *Proc. of IEEE Intl. Conf. on Distributed Event-Based Systems (DEBS)*, Vienna, Austria, July 2002.

[11] J. Kusuma, L. Doherty, and K. Ramchandran. Distributed Compression for Sensor Networks. In *Proc. of IEEE Intl. Conf. on Image Processing (ICIP)*, Thessaloniki, Greece, Oct. 2001.

[12] E. K. Lee, I. Kulkarni, D. Pompili, and M. Parashar. Proactive Thermal Management in Green Datacenter. *The Journal of Supercomputing (Springer)*, June 2010.

[13] X.-Y. Li, X. Xu, S. Wang, S. Tang, G. Dai, J. Zhao, and Y. Qi. Efficient Data Aggregation in Multi-hop Wireless Sensor Networks under Physical Interference Model. In *Proc. of Intl. Conf. Mobile Adhoc and Sensor Systems (MASS)*, Macau (S.A.R.), China, Oct. 2009.

[14] T. Melodia, D. Pompili, and I. F. Akyildiz. A Communication Architecture for Mobile Wireless Sensor and Actor Networks. In *Proc. of IEEE Conf. on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, Reston, VA, Sept. 2006.

[15] V. Mhatre and C. Rosenberg. Design Guidelines for Wireless Sensor Networks: Communication, Clustering and Aggregation. *Ad Hoc Networks Journal*, 2(1):45–63, Jan. 2004.

[16] J. Moore, J. S. Chase, and P. Ranganathan. Weatherman: Automated, Online and Predictive Thermal Mapping and Management for Data Centers. In *Proc. IEEE Conf. on Autonomic Computing (ICAC)*, pages 155–164, Dublin, Ireland, June 2006.

[17] P. Ogren, E. Fiorelli, and N. E. Leonard. Cooperative Control of Mobile Sensor Networks: Adaptive Gradient Climbing in a Distributed Environment. *IEEE Transactions on Automatic Control*, 49(8):1292–1302, Aug. 2004.

[18] S. S. Pradhan, J. Kusuma, and K. Ramchandran. Distributed Compression in a Dense Microsensor Network. *IEEE Signal Processing Magazine*, 19(2):51–60, Mar. 2002.

[19] S. S. Pradhan and K. Ramchandran. Distributed Source Coding: Symmetric Rates and Applications to Sensor Network. In *Proc. of Data Compression Conf. (DCC)*, Snowbird, UT, Mar. 2000.

[20] A. Scaglione. Routing and Data Compression in Sensor Networks: Stochastic Models for Sensor Data that Guarantee Scalability. In *Proc. IEEE Intl. Symp. on Information Theory (ISIT)*, Yokohama, Japan, July 2003.

[21] A. Scaglione and S. D. Servetto. On the Interdependence of Routing and Data Compression in Multi-Hop Sensor Networks. In *Proc. of Intl. Conf. on Mobile Computing and Networking (MobiCom)*, Atlanta, GA, Sept. 2002.

[22] K. Sohrabi, J. Gao, V. Ailawadhi, and G. J. Pottie. Protocols for Self-Organization of a Wireless Sensor Network. *IEEE Personal Communications*, 7(1):16–27, Oct. 2000.

[23] J. A. Stankovic, T. F. Abdelzaher, C. Lu, L. Sha, and J. Hou. Real-time Communication and Coordination in Embedded Sensor Networks. *Proc. of the IEEE*, 91(7):1002–1022, July 2003.

[24] M. C. Vuran, O. B. Akan, and I. F. Akyildiz. Spatio-temporal Correlation: Theory and Applications for Wireless Sensor Networks. *Computer Networks (Elsevier Science)*, 45(3):245–259, June 2004.

[25] R. Willett, A. Martin, and R. Nowak. Backcasting: Adaptive Sampling for Sensor Networks. In *Proc. of the Intl. Conf. on Information Processing in Sensor Networks (IPSN)*, Berkeley, CA, Apr. 2004.

[26] X. Xu, Y.-H. Hu, W. Liu, and J. Bi. Data-Coverage Sleep Scheduling in Wireless Sensor Networks. In *Proc. of Intl. Conf. on Grid and Cooperative Computing (GCC)*, Shenzhen, China, Oct. 2008.

[27] O. Younis and S. Fahmy. Distributed Clustering in Ad-hoc Sensor Networks: A Hybrid, Energy-Efficient Approach. In *Proc. of IEEE Conf. on Computer Communications (INFOCOM)*, Hong Kong (S.A.R.), China, Mar. 2004.