Distributed Data-Centric Adaptive Sampling for Cyber-Physical Systems

EUN KYUNG LEE, HARIHARASUDHAN VISWANATHAN, and DARIO POMPILI,
Rutgers University

A data-centric joint adaptive sampling and sleep scheduling solution, SILENCE, for autonomic sensor-based systems that monitor and reconstruct physical or environmental phenomena is proposed. Adaptive sampling and sleep scheduling can help realize the much needed resource efficiency by minimizing the communication and processing overhead in densely deployed autonomic sensor-based systems. The proposed solution exploits the spatiotemporal 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. Differently from existing adaptive sampling solutions, SILENCE employs temporal causality analysis to not only track the variation in the underlying phenomenon but also its cause and direction of propagation in the field. The causality analysis and the same correlations are then leveraged for adaptive sleep scheduling aimed at saving energy in wireless sensor networks (WSNs). SILENCE outperforms traditional adaptive sampling solutions as well as the recently proposed compressive sampling techniques. Real experiments were performed on a WSN testbed monitoring temperature and humidity distribution in a rack of servers, and the simulations were performed 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: Design, Algorithms, Performance

Additional Key Words and Phrases: Sensor networks, adaptive sampling, spatial and temporal correlation, autonomic systems, cyber-physical systems

ACM Reference Format:
Eun Kyung Lee, Hariharasudhan Viswanathan, and Dario Pompili. 2015. Distributed data-centric adaptive sampling for cyber-physical systems. ACM Trans. Autonom. Adapt. Syst. 9, 4, Article 21 (January 2015), 27 pages.
DOI: http://dx.doi.org/10.1145/2644820

21

1. INTRODUCTION

Cyber-physical systems (CPSs) are distributed, autonomic sensor and actor systems [Melodia et al. 2006] 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

This work is supported by the National Science Foundation (NSF) under grants CNS-0855091, IIP-0758566, and CNS-1117263.
Authors' addresses: E. K. Lee, H. Viswanathan, and D. Pompili, NSF Center for Cloud and Autonomic Computing (CAC), Department of Electrical and Computer Engineering, Rutgers University, New Brunswick, NJ; emails: {eunkyung_lee, hari_viswanathan, pompili}@cac.rutgers.edu.

DOI: http://dx.doi.org/10.1145/2644820

distributed sensor-based systems are, in general, composed of heterogeneous sensor nodes that differ with respect to (1) the type of data that they can sense (e.g., temperature, humidity, vibration, airflow, images), (2) the source of energy for operation, and (3) the mode of data transmission making them constrained in terms of energy or communication cost or both. 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 generally 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 ranges. Because of the high density of nodes, there is generally a high degree of correlation among observations of spatially proximal nodes (*spatial correlation*). In addition, 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 [Vuran et al. 2004; Willett et al. 2004].

In this article, 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 in CPSs composed of a network of wireless sensors. The solution exploits the spatial and temporal correlation in the manifestations to eliminate redundancy and to reduce the cost of processing and communicating large volumes of sensed data to a base station (or sink) for postprocessing. Additionally, SILENCE employs temporal causality analysis for timely reconfiguration of the sensor network in response to variations in the phenomenon. Specifically, Granger causality metric [Geweke 1982] is computed and leveraged to not only track the variation in the phenomenon being monitored but also to analyze the cause and direction of propagation of the variation itself. To the best of our knowledge, this is the first work on adaptive sampling to exploit temporal causality analysis for tracking the variation in the phenomenon and for timely reconfiguring the CPS. Another salient feature of SILENCE is that it is a generic solution in the sense that a variety of criteria can be used for REP selection, such as similarity and correlation, Hellinger distance (distance between probability distributions), histogram distance, and vector norms.

SILENCE 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' data 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 and sensing. This is advantageous, as it saves resources (energy and bandwidth). 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_{th}^{k}. 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_{th}^{k}). 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 adaptively and identify or track any variation in the spatial distribution of the manifestation over time and change their state accordingly to enable accurate reconstruction at the sink. The *accuracy* measure verifies whether the solution follows the variation in the manifestations while still satisfying user-specified thresholds.

2. OUR CONTRIBUTIONS: IN THE CONTEXT OF CPS FOR DATACENTER MANAGEMENT

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, noninvasive, 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 1,000 racks and 50 blade servers in each rack with temperature and humidity sensors on each server (50,000 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.

Another recently proposed and widely used adaptive sampling solution for accurate reconstruction of sparse signals from a few random samples is compressive sensing (CS) [Donoho 2006]. The amount of samples that has to be transmitted and stored in CS is smaller than what is deemed necessary by the Nyquist-Shannon criterion for accurate signal reconstruction. However, accurate reconstruction in CS requires *incoherent sampling* of sparse signals, which have only a few nonzero components with respect to some basis. Incoherence is jointly determined by the randomness of the sampling strategy and the sparsity of the sampled data (in the sparse basis), which usually exploits the spatial correlation of the data in the field. The knowledge of spatiotemporal characteristics in the data, however, is not incorporated in the sampling strategy. In contrast, SILENCE is data centric and exploits the spatial as well as temporal correlation in the data to elect the REPs.

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 servers and through extensive simulations on TOSSIM, the TinyOS simulator. Figure 1 shows our real experiment setup of 26 TelosB motes densely deployed on 13 servers/blades in a server rack. We focus on monitoring only one server rack in the machine room at the NSF Center for Autonomic Computing at Rutgers University. 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 the centralized optimal (but impractical) solution, a traditional non–data-centric randomized REP selection procedure for WSNs and different variations of the compressive sampling technique (with different sparse basis). We also demonstrate the effectiveness of our solution when monitoring multiple

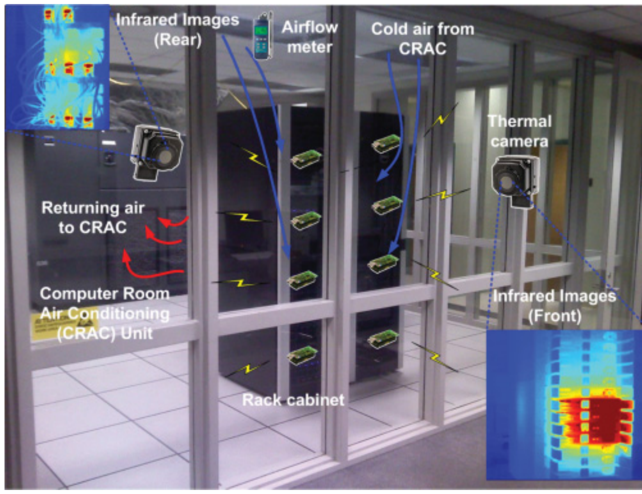

Fig. 1. A heterogeneous wireless sensor network testbed measuring multiple manifestations such as temperature, humidity, and airflow in a server room at the Center for Cloud and Autonomic Computing, Rutgers University.

manifestations and in high-density deployment scenarios. The contributions in this article build on a preliminary version of SILENCE developed in 2011 [Lee et al. 2011]. The following are the contributions of the preliminary version:

—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 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 user-defined thresholds.

In addition to the preliminary version, the following are the contributions of this article:

—We have formulated and explained in detail the centralized optimization-based approach to the problem of selecting representatives. The centralized approach, although impractical, serves as a benchmark to the performance of our proposed distributed algorithm.
—We have described and discussed another recently proposed and widely used adaptive sampling solution, CS, against which we have evaluated our solution.
—We have discussed in detail and shown through additional simulations the impact of different user-defined thresholds on the reconstruction accuracy of our approach.
—We have introduced the use of temporal causality analysis across network nodes to not only track the variation in the phenomenon in a timely manner but also to identify its cause and direction of propagation in the field of interest.

The remainder of the article is organized as follows. Section 3 highlights the contribution of this work and compares it with existing approaches. Section 4 discusses the optimal centralized approach and provides insights for the design of our distributed solution. Section 5 describes our solution for autonomic adaptive sampling. Section 6 presents the experiment and simulation setup, and discusses the performance evaluation of our proposed solution. Finally, we conclude our article in Section 7.

3. RELATED WORK

Recent works 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 [Moore et al. 2006; Abbasi et al. 2010; Banerjee et al. 2010]. In reality, thermal- or cooling-aware datacenter management schemes require information about (1) inlet and outlet fan temperature for each machine (blade or chassis), (2) CPU or core utilization for each machine, (3) computer room air conditioner (CRAC) air inlet and outlet temperatures, (4) the fan speeds (CRAC and computing unit exhaust fans), (5) power specifications of each server type, and (6) workload information (duration, start time, arrival frequency, etc.). The following are the summary of features of our solution, which differs from other solutions:

—Our solution reduces the amount of raw data required for extracting useful information in challenging real-world applications such as thermal-aware datacenter management [Lee et al. 2012b] without penalizing the effectiveness of the management decisions due to inaccurate reconstruction of the phenomenon.
—Our solution is the first work to *jointly* perform adaptive sampling for data collection and sleep scheduling for energy savings in WSNs.
—Our solution ensures data quality by employing *similarity* along with *correlation* for exploiting spatial and temporal correlation in sensor data. Previous solutions for energy saving using cluster heads/local base stations focused on how to reduce the quantity of the data transmitted to the sink but did not take the quality of the data into account.
—Our solution is integrated with a mechanism to ensure end-to-end connectivity (i.e., any REPs is guaranteed to be able to reach the sink via multihopping), which is often overlooked by the literature dealing with clustering in WSNs.

Our approach is different from traditional clustering algorithms (e.g., Younis and Fahmy [2004], Mhatre and Rosenberg [2004], and Bandyopadhyay and Coyle [2003, 2004]), which reduce global communication to prolong the lifetime of the network, or self-organization mechanisms (e.g., Sohrabi et al. [2000], Stankovic et al. [2003], Bhardwaj and Chandrakasan [2002], and Ogren et al. [2004]), which use dynamic role assignment to extend the lifetime of the network while also reducing communication cost. Such solutions, in fact, have largely overlooked 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 instead are based on constraints such as lifetime of the network, energy, and wireless link quality.

SILENCE differs from solutions that perform in-network processing of data to eliminate redundancy such as compression [Kusuma et al. 2001; Pradhan et al. 2002], data aggregation [Li et al. 2009; Krishnamachari et al. 2002], source coding [Cui et al. 2007; Pradhan and Ramchandran 2000], and routing and data compression [Scaglione and Servetto 2002; Scaglione 2003], as they require constant local communication inside a cluster/group of nodes. Instead, our solution puts the nodes to sleep 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 [Xu et al. 2008; Chachra and Marefat 2006] because it adjusts sleep duration based on the data correlation, whereas others are scheduled for increasing network lifetime.

SILENCE is also fundamentally different from the recently proposed CS technique, which can reconstruct sparse signals (on some basis) accurately using fewer random samples than what is prescribed by the Nyquist-Shannon criterion. CS has already been employed for monitoring networks [Haupt et al. 2010; Quer et al. 2009], hardware performance [Huang et al. 2012], ambient temperature, rainfall, and pollution levels [Quer et al. 2009]. CS does not incorporate the knowledge of spatial and temporal characteristics in the data (which is exploited in the selection of a sparse basis) into its sampling strategy. In contrast, SILENCE exploits the spatial as well as temporal correlation in the data (using similarity and correlation) to elect the REPs.

In SILENCE, connectivity issues after node deployment are addressed and solved by using a mechanism employing AWAKE packets to notify that an ASSOC node is no longer sleeping, but awake and ready (referred to as AWAKE – ASSOC) for relaying the packets. This is in contrast to existing solutions like clustered aggregation [Yoon and Shahabi 2007], in which all nonrepresentative nodes stay awake all the time to aid in multihop routing. In a sensor network employing SILENCE, only the minimum number of nodes required for maintaining connectivity will be awake at any point in time, thus saving energy and increasing the lifetime of the network. The scope of the connectivity in this article is to address the issues only caused by introducing our adaptive sleeping algorithm. We do not consider the case of ensuring global network connectivity (i.e., determining the minimum density of "awake" relay nodes), because it requires predeployment optimization, which in turn relies on global topology information apart from deviating from the focus of our article. There are many algorithms that deal specifically with the aforementioned problem in WSNs (e.g., connectivity map [Ha et al. 2006] and initial connectivity graph [Keshavarzian et al. 2006]), and our algorithm runs on the top of one of those.

Researchers have explored the use of spatial and temporal correlation in the measured data to determine cluster heads and to compress or duplicate data [Aggarwal et al. 2011; Zoghi and Kahaei 2009; Chen et al. 2010; Jiang et al. 2011; Liu et al. 2007]. However, SILENCE differs from these solutions in the following manner. First, SILENCE is generic in the sense that a variety of criteria can be used for REP selection. In this work, we have used similarity and correlation, whereas other metrics include Hellinger distance (distance between probability distributions), histogram distance, any vector norm, and so forth. Second, SILENCE also leverages temporal causality information during the reclustering phase for timely reconfiguration of the sensor network in response to variations in the phenomenon. This information is obtained from the Granger-causality metric, which is computed pairwise across sensor network nodes. Such analysis enables us to not only track the variation in the phenomenon being monitored but also to analyze the cause and direction of propagation of the variation itself. This way, SILENCE closes the gap between the cyber and physical worlds by incorporating information from temporal causality analysis into the REP selection procedure.

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.

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 he or she 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 greenhouse 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 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.

4. 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. This centralized optimization approach is not intended to achieve redundancy elimination using selective representation under user-specified reconstruction accuracy constraints in real-time. However, it is intended to show the complexity of the problem and motivate the need for a localized distributed solution and to serve as a benchmark for the performance of our proposed distributed solution (SILENCE) for small problem sizes. Even though the centralized approach is impractical, complicated, and nonscalable, it gives us insight to make key design decisions for our localized distributed solutions.

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 one that is reconstructed at the sink based on the information from its REP. The optimization problem 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 and the manifestation k are given, the following problem finds the optimal number and set of M REPs reporting to the sink. This optimal set is chosen such that the absolute difference between maximum error e^k_{max} in reconstruction and the error threshold e^k_{th} is minimized. The notations used in the optimization problem are listed next, where matrices are represented in bold and vectors are underlined:

—$\underline{R} = \{r_n\}_{N\times1}$, where $r_n \in \{1, 0\}$, is an $N \times 1$ vector that indicates whether a node is REP or not, and $r_n = 1$ if it is a REP node and 0 otherwise.
—$\mathbf{A} = \{a_{nm}\}_{N\times N}$, where $a_{nm} \in \{1, 0\}$, is an $N \times N$ matrix that indicates whether a node n is associated with REP node m or not, and $a_{nm} = 1$ if node n is associated with REP node m and 0 otherwise.
—$\mathbf{\Psi^k} = \{\psi^k_{ns}\}_{N\times S}$ is an $N \times S$ matrix of samples of manifestation k at each node, ψ^k_{ns} is the sth sample of manifestation k measured at sensor node n, and S is the total number of samples.
—$\underline{\Psi}^k = \{\overline{\psi}^k_n\}_{N\times1}$ is an $N \times 1$ vector of mean of sampled values (S samples) at each node n for a manifestation k.
—$\mathbf{E^k} = \{e^k_{nm}\}_{N\times N}$, where $e^k_{nm} = |\overline{\psi}^k_n - \overline{\psi}^k_m|$ is the difference between the means of sampled values (S samples) of manifestation k at nodes n and m.
—$\mathbf{C^k} = \{\gamma^k_{nm}\}_{N\times N}$, where γ^k_{nm} is the correlation between the sampled values (S samples) of manifestation k at nodes n and m.

—$\mathbf{D} = \{d_{nm}\}_{N\times N}$, where $d_{nm} \in \{1, 0\}$, is an $N \times N$ matrix that denotes whether nodes n and m are within the radio transmission range of each other or not.

$$\begin{aligned}
\textbf{Given (offline):}\quad & N, S, e^k_{th}, \gamma^k_{th}, \mathbf{\Psi}^k, \mathbf{D}\\
\textbf{Computed (online):}\quad & \mathbf{E^k}, \mathbf{C^k}\\
\textbf{Find:}\quad & \underline{R}^*, \mathbf{A}^*\\
\textbf{Minimize:}\quad & |e^k_{max} - e^k_{th}|\\
& e^k_{max} = \|\underline{e}^k\|_\infty,\\
& \underline{e}^k = \{e^k_n\}_{N\times 1} = |\mathbf{A}\cdot\underline{\Psi}^k - \underline{\Psi}^k|\\
\textbf{Subject to:}\quad &
\end{aligned}$$

$$a_{mm} = r_m, \ \forall m = 1 \ldots N; \tag{1}$$

$$\sum_{m=1}^{N} a_{nm} \cdot r_m = 1, \ \forall n = 1 \ldots N; \tag{2}$$

$$d_{nm} \geq a_{nm}, \ \forall n, m = 1 \ldots N. \tag{3}$$

The objective function has been designed to ensure that the problem does not choose all nodes in the field to be REPs to minimize the maximum reconstruction error e^k_{max}. We use the user-specified reconstruction error threshold to limit the number of REPs in the field. Constraint (1) forces a REP to be associated only with itself and also ensures that only ASSOCs are associated to REPs and not vice versa, constraint (2) forces every node that is not a REP to be associated with one and only one REP, and constraint (3) ensures that an ASSOC and its REP are within each other's radio transmission range.

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. The time complexity of the centralized problem grows exponentially with the number of REPs M and the number of nodes N. In addition, given a solution (the number M and the list of REPs $\underline{R}$), its optimality cannot be verified in polynomial time. Therefore, the centralized optimization problem is NP-hard. The centralized approach is also impractical for real-world deployment, as it assumes complete knowledge of the entire network state and sensed data at a centralized location. However, it 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, in our distributed solution 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.

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 15 is shown in Figure 2. The maximum reconstruction error falls within the threshold that 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 that the number of REPs chosen by SILENCE fluctuates between 11 and 12, which is close to the optimal solution. The difference between the centralized approach and the distributed approach can be attributed to the nonidealities that are introduced by the wireless link and to the limited sensing and communication range

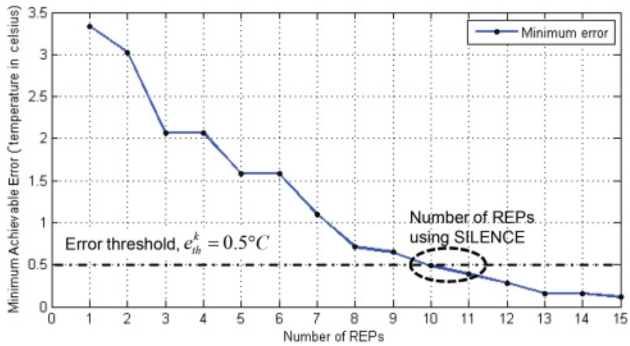

Fig. 2. The minimum of maximum reconstruction error obtained when the centralized problem is solved by varying M (number of REPs) from 1 to 15. The maximum reconstruction error falls within the threshold that we set, $e_{th}^{k} = 0.5°C$, only when the number of REPs, M, is at least 11.

of sensor nodes. Detailed comparisons between SILENCE and the centralized solution with different error thresholds are presented in Section 6.1.

5. 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) in which a node can be 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, whereas the ASSOCs go to sleep to save energy. The other auxiliary state in which a node can be 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.

The goal of our solution is to address the problem of redundancy elimination using selective representation under userspecified reconstruction accuracy constraints in a distributed and energy-efficient manner. To realize the goal, SILENCE relies on basic assumptions about the underlying sensor-based system of the broader CPS. First, each sensor node is aware of its position in the field. This is necessary to spatially reconstruct the data at the sink. Second, 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 such as ad hoc on-demand distance vector routing (AODV), optimized link-state routing (OLSR), and geographical routing. The proposed solution is not a best effort kind suboptimal distributed solution (such as random REPs selection). Rather, it is a localized solution that exploits the acquired knowledge about the spatiotemporal characteristics of the field and whose performance is shown to be close to that of a centralized optimal solution.

5.1. Similarity and Correlation

The key component of SILENCE—that is, the transition between roles or states of a node—depends on two important metrics: similarity and correlation. Two nodes n and

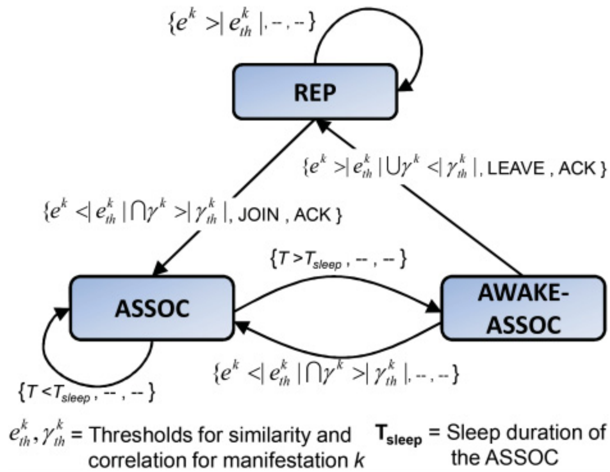

Fig. 3. The three states in which a node on the field can be at any point in time, as well as the events and messages that trigger state transitions. $\{Event, Tx, Rx\}$ represents the event and control messages that trigger a state change.

m are said to be sensing similar values if the difference between the means of their measured values of a manifestation k, $e^k_{n,m}$, 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, $\gamma^k_{n,m}$, which is calculated using S samples of data from those nodes, is greater than a user-specified threshold, γ^k_{th}. The formulas used for determining similarity and correlation are as follows:

$$e^k_{n,m} = |\overline{\psi}^k_n - \overline{\psi}^k_m|, \tag{4}$$

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

$$\gamma^k_{n,m} = \frac{\sum_{s=1}^{S}\left(\psi^k_n[s] - \overline{\psi}^k_n\right)\cdot\left(\psi^k_m[s] - \overline{\psi}^k_m\right)}{\sqrt{\sum_{s=1}^{S}\left(\psi^k_n[s] - \overline{\psi}^k_n\right)^2}\cdot\sqrt{\sum_{s=1}^{S}\left(\psi^k_m[s] - \overline{\psi}^k_m\right)^2}}. \tag{5}$$

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

5.2. State Transitions

The three states in which a node on the field can be at any point in time, as well as the events and messages that trigger state transitions, are depicted in Figure 3. Initially, after deployment, all sensor 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 and 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}, which is chosen between T^{min}_{End} and T^{max}_{End} in a uniform random manner. This randomization is purely to avoid packet collisions due to synchronization among nodes that happen to start sampling at the same time. The sampling phase (whose duration is T_{End}-T_{Start}) is followed by the *state decision phase*, whose duration is less than T^{max}_{End}-T_{End}.

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 REP. 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 frequency with which both the packets are transmitted can be 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 are 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*.

ALGORITHM 1: State decision using similarity and correlation.

$\boxed{\text{INIT\_ STATE\_ DECISION:}}$
$IS\_ASSOC = 0$, $MY\_REP = 0$
e^k_{th}, γ^k_{th} = {Threshold for similarity and correlation}
$LIST\_REP_{PA}$ = {List of potential REPs}
$\boxed{\text{STATE\_ DECISION:}}$
In $LIST\_REP_{PA}$ find $LIST\_REP^{max}_{PA}$
if SIZE ($LIST\_REP^{max}_{PA}$) == 1 **then**
    $MY\_REP = LIST\_REP^{max}_{PA}$
    $IS\_ASSOC = 1$
    Send JOIN, Receive ACK
    Sleep (T_{Sleep})
else
    In $LIST\_REP^{max}_{PA}$ find REP^{max}_{ID}
    $MY\_REP = REP^{max}_{ID}$
    $IS\_ASSOC = 1$
    Send JOIN, Receive ACK
    Sleep (T_{Sleep})
end if

With the samples in the DATA or HELLO packets that a REP receives from its neighbors, it calculates an 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 (e^k_{th}) and correlation (γ^k_{th}) 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, a node decides its future state as shown in Algorithm 1.

Initially, after deployment, all nodes in the field are REPs and are not associated with any other node (i.e., for all nodes $IS\_ASSOC = 0$ and $MY\_REP = 0$). 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. After a list of potential REPs is available ($LIST\_REP_{PA}$) in the sampling phase, each node chooses the one with the highest number of potential associates ($LIST\_REP^{max}_{PA}$) in the state decision phase. 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

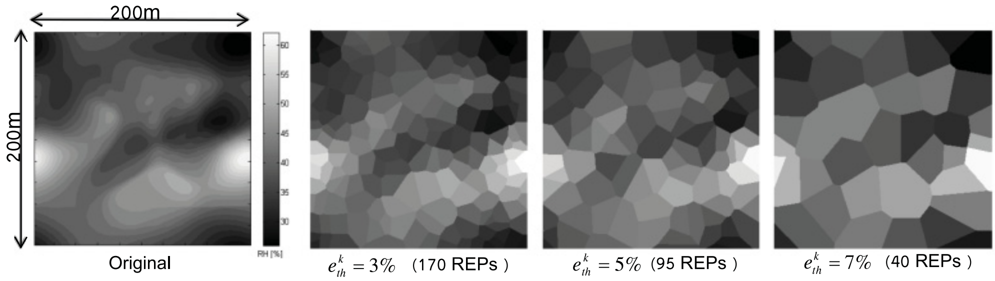

Fig. 4. Effect of user-specified thresholds on accuracy of reconstruction of the manifestation (humidity): $e^k_{th} = 3\%$, $e^k_{th} = 5\%$, and $e^k_{th} = 7\%$. The higher the error threshold, the lower the number of REPs (out of a total of 400 sensor nodes deployed over a 200m $\times$ 200m field monitoring humidity) reporting to the sink.

of avoiding undesired associations between nodes that are within each other's radio frequency 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. The long sampling phase ensures that undesired associations due to transient and erratic behaviors of sensors do not happen.

5.3. Choice of Thresholds

As mentioned earlier, SILENCE allows the user to increase (by zooming in) and decrease (by zooming out) the resolution of sensing by choosing appropriate values for the similarity (e^k_{th}) and correlation (γ^k_{th}) thresholds for a manifestation k as illustrated in Figure 4. The choice of thresholds can be based on some prior knowledge about the observed manifestations of a physical or environmental phenomenon. For example, the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) published the recommended environmental envelope (in terms of temperature and relative humidity (RH)) for IT equipment housed in datacenters. The current envelope is given by low-end temperature $18^{\circ}C$, high-end temperature $27^{\circ}C$, low-end moisture $5.5^{\circ}C$ dew point (DP), and high-end moisture 60% RH and $15^{\circ}C$ DP. The purpose of the recommended envelope is to provide guidance to datacenter operators for maintaining high equipment reliability and achieving high energy efficiency. Exceeding the recommended operating environment limits for extended time periods could result in equipment damage, which may lead to an increase in service downtime and, hence, increase in number of service-level agreement violations.

For example, when the RH exceeds 60% and remains high for extended periods, it can result in failures given the reduced conductor to conductor spacings common in many designs today. Datacenter managers can use the measured temperature, RH, and the psychrometric chart, which conveys the relationship between the two as depicted in Figure 5, to determine whether the datacenter operating point is within the recommended range. They can modify the similarity (e^k_{th}) and correlation (γ^k_{th}) thresholds to zoom in or zoom out depending on the current operating point. Figure 4 demonstrates the effect of the user-specified threshold on the reconstruction of distribution of humidity at the remote sink. The higher the error threshold, the lower the number of REPs (out of a total of 400 nodes) reporting to the sink. However, as the bound on reconstruction accuracy is increased (due to a higher e^k_{th}, from 3% to 7%), the accuracy with which the distribution of the manifestation is reconstructed (using Voronoi-based reconstruction) decreases.

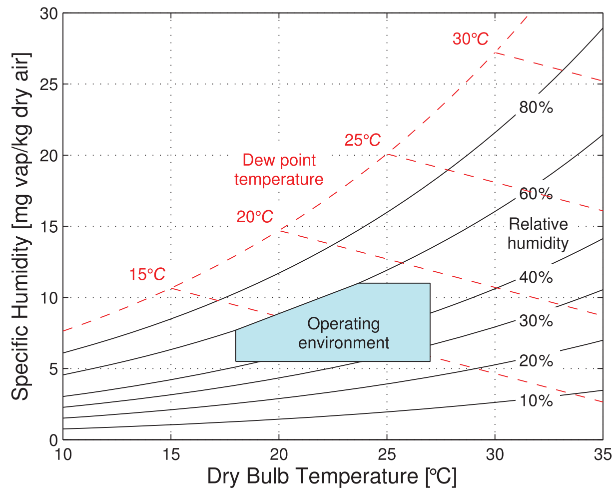

Fig. 5. Psychrometric chart showing ASHRAE's recommended environmental envelope (operating conditions) for IT equipment housed in datacenters. The current envelope is given by low-end temperature 18°C, high-end temperature 27°C, low-end moisture 5.5°C DP, and high-end moisture 60% RH and 15°C DP.

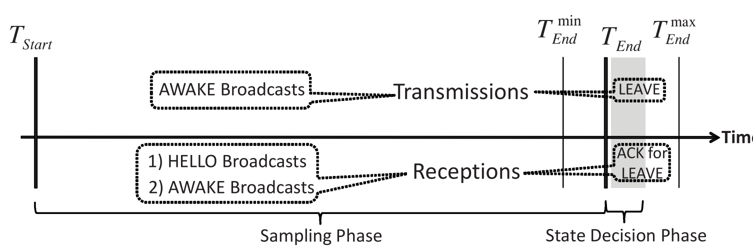

Fig. 6. The activity of an AWAKE – ASSOC (an ASSOC after its sleep duration). In the sampling phase, an AWAKE – ASSOC listens to HELLO broadcasts from all neighboring REPs, including its own. At the end of the sampling phase, it verifies similarity and correlation of its own sampled data and that of its REPs received recently to decide whether a state change is necessary.

5.4. 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—that is, $T_{Sleep} = \phi(e^1 \ldots e^K, \gamma^1 \ldots \gamma^K) \cdot T_{Sleep}^{min}$, where K is the maximum number of manifestations. To ensure connectivity of all active nodes (REPs) with the sink at all times, we enforced an upper bound on the sleep duration of all ASSOCs. In other words, every node (REPs or AWAKE – ASSOCs) will have at least one other active node within its radio range and in the positive advance set toward the sink. A generic node belongs to the positive advance set with respect to a sender node if it is closer to the destination (in this case, the sink) than the sender. The activity of an AWAKE – ASSOC (an ASSOC after its sleep duration) is shown in Figure 6. In the sampling phase, an AWAKE – ASSOC listens to HELLO broadcasts from all 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 REPs received recently. If each of the thresholds e_{th}^{k} and γ_{th}^{k} 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 relay opportunistically packets for multihop 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 multihopping. 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. The AWAKE packet contains the duration of the availability ($T_{End} - T_{Current}$) and the necessary routing information (e.g., link quality information in OLSR, location information in geographical routing) to help forward data packets to the sink. 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).

5.5. Temporal Granger Causality Analysis

In most applications, reconstruction of the underlying phenomenon at different points in time (snapshots) is not sufficient. The ultimate goal of any CPS is to understand the source of spatial and temporal variation in the phenomenon and the direction in which it propagates. This will allow the CPS to exploit sensor nodes that are in a better position to efficiently track the phenomenon. We focus on detecting the magnitude of influence the direction of propagation of variation in a phenomenon by applying real-time, in situ temporal Granger (G)-causality analysis. Note that G-causality provides a much more stringent condition on causation than just observing high correlation with some lag–lead relationship. SILENCE aims at analyzing in real time the *direction of propagation* and *magnitude* of phenomenon changes (i.e., temperature, humidity) between nodes in a cluster. Specifically, G-causality is computed and leveraged to not only track the variation in the phenomenon being monitored but also to analyze the cause and direction of propagation of the variation itself.

Granger-causality: A time series $\mathbf{x} = \{x_1, x_2, \ldots, x_t, \ldots\}$ is said to G-cause another time series $\mathbf{y}$ if including information about the past of $\mathbf{x}$ significantly increases the prediction accuracy of the current value y_t of $\mathbf{y}$ in comparison to predicting it based *only* on the past values of $\mathbf{y}$ alone. G-causality was initially introduced in Geweke [1982], where the authors implemented it using two vector autoregressive (AR) models; the first, called *restricted model*,

$$x_t = \sum_{j=1}^{P} a_j x_{t-j} + \delta_{1t}, \quad y_t = \sum_{j=1}^{P} a_j y_{t-j} + \gamma_{1t}, \tag{6}$$

calculates how much two time series, $\mathbf{x}$ and $\mathbf{y}$, can be "explained" by their own past (x_{t-j} and y_{t-j}, with $j = 1, 2, \ldots$), resulting in residual error variances $\Delta_1 = var(\delta_{1t})$ and $\Gamma_1 = var(\gamma_{1t})$ (the model order is represented by P, which specifies how many previous time points are taken into account, and the length of the time series by T, with $P < T$).

In the second, called *unrestricted model*,

$$x_t = \sum_{j=1}^{P} a_j x_{t-j} + \sum_{j=1}^{P} b_j y_{t-j} + \delta_{2t},$$
$$y_t = \sum_{j=1}^{P} a_j y_{i-j} + \sum_{j=1}^{P} b_j x_{t-j} + \gamma_{2t}, \tag{7}$$

the prediction is based on the time series' own past and the past of the other time series. This results in residual error variances $\Delta_2 = var(\delta_{2t})$ and $\Gamma_2 = var(\gamma_{2t})$. The linear influence from $\mathbf{x}$ to $\mathbf{y}$, $\mathcal{F}_{\mathbf{x}\to\mathbf{y}}$, and from $\mathbf{y}$ to $\mathbf{x}$, $\mathcal{F}_{\mathbf{y}\to\mathbf{x}}$, can now be calculated as the ratio between the variances of the residual error—that is,

$$\mathcal{F}_{\mathbf{x}\to\mathbf{y}} = \ln\frac{var(\gamma_{1t})}{var(\gamma_{2t})} = \ln\frac{\Gamma_1}{\Gamma_2},\ \mathcal{F}_{\mathbf{y}\to\mathbf{x}} = \ln\frac{var(\delta_{1t})}{var(\delta_{2t})} = \ln\frac{\Delta_1}{\Delta_2}. \tag{8}$$

A reduction in error variance when including the past of another time series results in a larger $\mathcal{F}$-ratio. The difference G-causality—that is, $\mathcal{F}_{\mathbf{x}\to\mathbf{y}} - \mathcal{F}_{\mathbf{y}\to\mathbf{x}}$—was calculated to assess the dominant direction of information flow.

Selection of time lag: Selecting the time lag is an important problem in G-causality. The estimation of AR models requires as a parameter the number of time lags P to include (i.e., the model order). Too few lags can lead to a poor representation of the data, whereas too many of them can lead to problems in the model estimation [Seth 2010]. Two criteria have been introduced in the literature, namely the Akaike information criterion (AIC) [Akaike 1974] and the Bayesian information criterion (BIC) [Schwarz 1978] to estimate the model order. For n variables, we have

$$AIC(P) = \ln(|\mathbf{\Sigma_2}|) + \frac{2Pn^2}{T},$$
$$BIC(P) = \ln(|\mathbf{\Sigma_2}|) + \frac{\ln(T)Pn^2}{T}. \tag{9}$$

where $\mathbf{\Sigma_2}$ is the noise covariance matrix of the unrestricted model and $|\cdot|$ indicates the determinant of a matrix. AIC is calculated for a set of model orders, and the order that gives the minimum value of AIC is selected as the model order of the AR model to determine G-causality between two time series.

In SILENCE, we have employed G-causality to elect/change REP within a cluster as a timely response to change in the underlying phenomenon. We present an example temporal G-causality analysis on a small cluster of 11 nodes numbered 1 through 11. A closer look at the metric, *causality flow*, which is represented using a bar graph Figure 7(a), a network flow diagram Figure 7(b), and in matrix form Figure 7(c), teaches us the following about the cluster. Any change in the phenomenon is first experienced by nodes 9 and 10 before any other node in the cluster. This is reflected in a high causality flow score in the bar diagram as shown in Figure 7(a). The network flow diagram (Figure 7(b)) and the matrix representation (Figure 7(c)) also indicate that nodes 9 and 10 *cause* all other nodes. When nodes with high causality flow scores are elected as REP, the reconstruction error decreases significantly. Therefore, a good choice of REP in this cluster would be node 9 or 10.

5.6. Toy Example

We present a *toy example* to help understand REP selection in SILENCE better. Figure 8 shows the node ID and number of potential-ASSOCs for every node. Solid circles represent the nodes that have similar and correlated values in time. Dotted

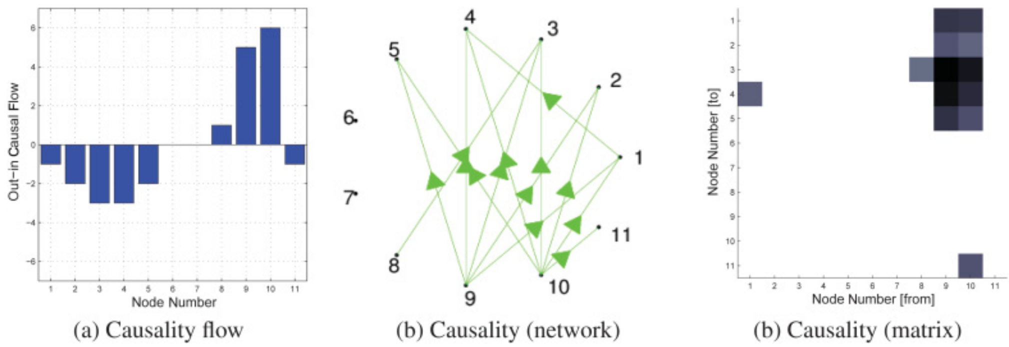

Fig. 7. Temporal G-causality analysis on a cluster of 11 nodes with 1 REP and 10 ASSOCs. This analysis reveals causal sources and sinks. Part (a) shows the causality flow (sign and magnitude), whereas parts (b) and (c) show the causality information in network form and matrix form, respectively.

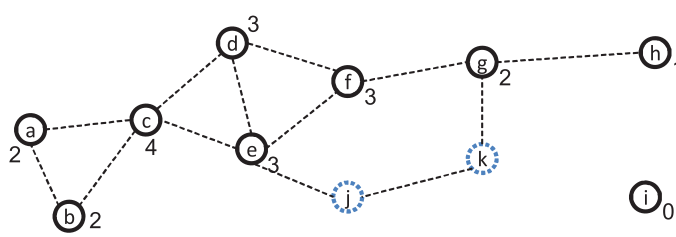

Fig. 8. A toy example illustrating the REP selection mechanism. Lowercase letters represent node IDs, whereas the numbers that follow represent the number of potential-ASSOCs. Nodes with the highest number of potential-ASSOCs in every one-hop neighborhood are elected REP. Here, c, f, i, j, and k are REPs.

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, which senses and communicates its data. Based on its 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 if it has 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, nodes j and k) will continue to be REPs. Nodes c and f will decide to continue to be REPs, as they know that they have the highest number of potential ASSOCs based on the information about their neighborhood. Nodes 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 that 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 *chaining*. The length of the aforementioned window is fixed in our solution. However, the window length can be adapted based on the relationship between the nodes' sensing scope and radio range. If the sensing scope is small compared to the radio range, then the window can be large to allow for chaining to avoid redundancy in data reporting. However, if the sensing scope is small compared to the radio range, then chaining is undesirable because it may affect global connectivity between the sparsely distributed REPs and the sink.

5.7. Competing Approach: Compressive Sensing

In this section, we explain CS, another recently proposed and widely used adaptive sampling solution against which we will evaluate SILENCE. CS has gained popularity as an effective tool for accurate reconstruction of sparse signals from few random samples. The number of samples used for reconstruction can be well below the number prescribed by the Nyquist-Shannon sampling theorem for perfect reconstruction of sampled signals. Measured features of the field (i.e., temperature, humidity) are not naturally sparse. Hence, reconstruction of the distribution of a feature $\mathbf{f}$ using a few of its random samples $\mathbf{b}$ requires knowledge of the following: (1) a linear transformation $\mathbf{\Psi}$ under which the feature $\mathbf{f}$ to be estimated becomes sparse (say, as $\mathbf{x}$)—that is, $\mathbf{f} = \mathbf{\Psi} \cdot \mathbf{x}$, and (2) a measurement matrix $\mathbf{\Phi}$ that describes how the random samples $\mathbf{b}$ have been obtained from the sensors on the field (either as pure random samples or linear combinations of samples)—that is, $\mathbf{b} = \mathbf{\Phi} \cdot \mathbf{f}$. With the knowledge of the basis and the measurement matrices, an estimate of the sparse signal $\mathbf{x}$ (which is the original signal $\mathbf{f}$ under transformation $\mathbf{\Psi}$) can be obtained by using L_1 minimization. The L_1 minimization problem is given by

$$
\begin{aligned}
\textbf{Given:}\quad & \mathbf{\Phi}, \mathbf{\Psi}, \mathbf{b};\\
\textbf{Find:}\quad & \mathbf{f}\\
\textbf{Minimize:}\quad & \|\mathbf{x}\|_{L_1};\\
\textbf{Subject to:}\quad & \mathbf{A} \cdot \mathbf{x} = \mathbf{\Phi} \cdot \mathbf{\Psi} \cdot \mathbf{x} = \mathbf{b} = \mathbf{\Phi} \cdot \mathbf{f}.
\end{aligned}
$$

The solution to this LP can be used to reconstruct the nonsparse signal $\mathbf{f}$. According to CS theory, the number of random samples $M = |\mathbf{b}|$ required to reconstruct the signal $\mathbf{f}$ with a very low error requires the knowledge of the sparsity K of $\mathbf{x}$, which is $\mathbf{f}$ under some linear transformation $\mathbf{\Psi}$. The sparsity of $\mathbf{x}$ depends on the spatiotemporal correlation of the feature, and as this can change from time to time, optimization of M, which is determined by Φ, in real time is a challenging task. However, if M is fixed, the transformation $\mathbf{\Psi}$ needs to be adapted according to the spatiotemporal variations to render $\mathbf{x}$ highly sparse. Adapting the transformation $\mathbf{\Psi}$ in real time is again nontrivial.

In our case, the data from sensors have a degree of correlation in time and similarity in space, especially when they come from closely located sensors. Suppose that $\mathbf{f}$ is a vector that contains temperature measurements of datacenters. Using multiple realizations of $\mathbf{f}$, corresponding to different times, we need to find the $\mathbf{\Psi}$ that makes $\mathbf{x}$ sparse. We adopt several bases that take and do not take into account spatial correlation to emphasize how choosing the right basis is important for CS and compare their reconstruction error with SILENCE. For the basis that does not take into account the spatial correlation, we choose discrete cosine transform (DCT) because it is a simple yet powerful transform with which insignificant high-frequency components can be discarded to get a sparse representation. For the basis that takes into account spatial correlation, we choose HorzVer-diff basis [Quer et al. 2009]. To obtain a representation in the HorzVer-diff basis, the input signal $\mathbf{x}$ (a matrix) is subjected to (1) pairwise subtraction of the elements along the columns of $\mathbf{x}$ and then to (2) pairwise subtraction of the elements of the resulting matrix along its rows.

We compared the performance of different sampling and reconstruction approaches (with mean reconstruction error as the metric) under the following simulated scenario: reconstruction of the spatial distribution of temperature in a 100×100 m^2 field at a remote sink with data selectively sampled from 400 sensors deployed in a uniform random manner. Figure 9 shows the mean error of reconstruction achieved by the following approaches: (1) Random-Voronoi, uniform random sampling and Voronoi-based reconstruction; (2) CS (DCT), compressive sensing using the DCT basis; (3) CS (HorzVer-diff), compressive sensing using the HorzVer-diff basis; and (4) SILENCE.

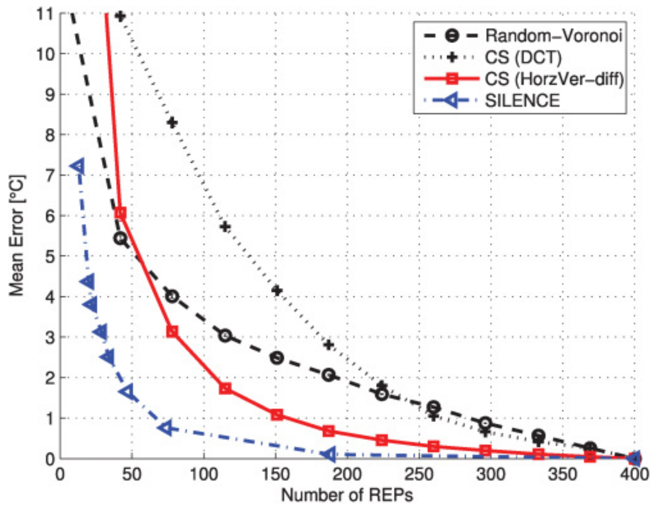

Fig. 9. Mean error in reconstruction achieved by different sampling and reconstruction approaches: (1) Random-Voronoi, uniform random sampling and Voronoi-based reconstruction; (2) CS (DCT), compressive sensing using the DCT basis; (3) CS (HorzVer-diff), compressive sensing using the HorzVer-diff basis; and (4) SILENCE.

We also tried two other bases for the CS approach: Wavelet and VerHorz-diff bases. However, we have only presented the performance when DCT and HorzVer-diff bases are employed, as they represent the worst and best, respectively, among the four bases. It is intuitive that the mean reconstruction error decreases as the number of samples (or REPs) increases for all approaches. SILENCE outperforms all other approaches showing low reconstruction error. We also observed that CS (HorzVer-diff) shows lower reconstruction error than CS (DCT), implying that the selection of sparse basis is crucial in using the CS approach.

6. 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 simulator. We also compared its performance against the centralized optimal REP selection approach. 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. As 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 the results, are detailed in the following discussion.

6.1. Real Experiments

The temperature sensors on the motes measure the external temperature and are used to estimate heat generation and distribution [Lee et al. 2012a] during the operation of the servers on our testbed shown in Figure 1. 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. Figure 10 visualizes the strong spatial and temporal correlation in measured temperature data (single snapshot from sensors placed close to the outlet

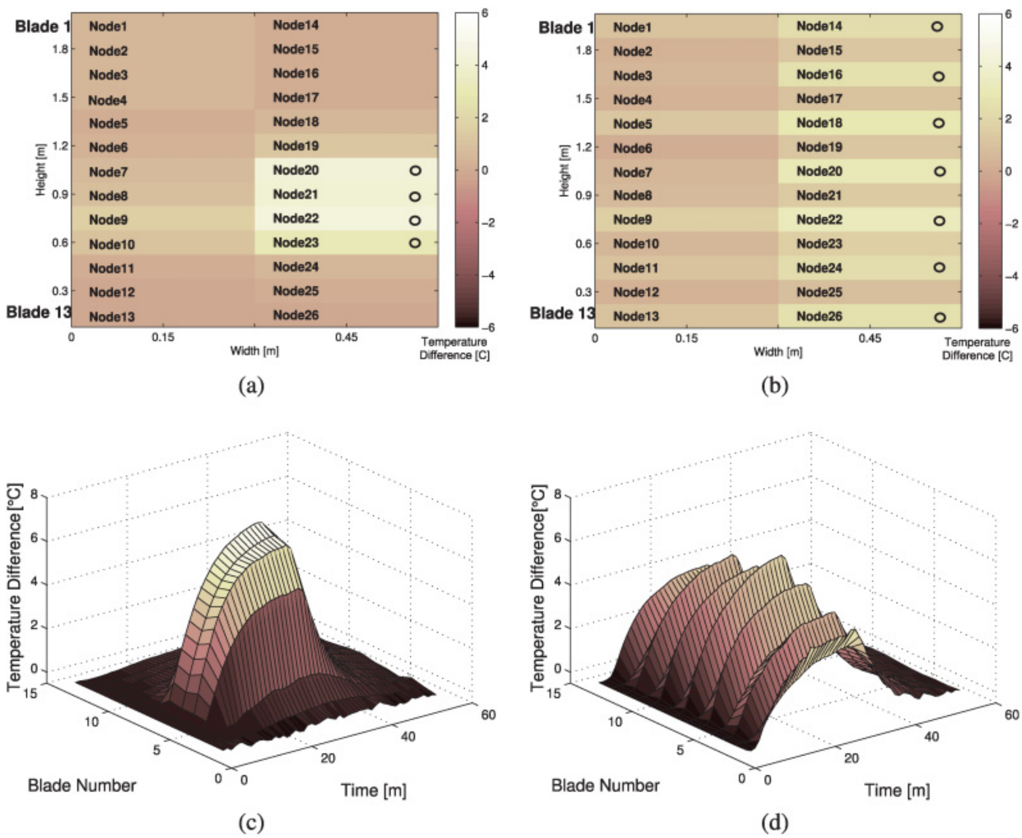

Fig. 10. Experiment results. Spatial distribution ((a) and (b)) and temporal distribution ((c) and (d)) of measured data (temperature) among the 26 sensor nodes deployed on a rack of 13 blade servers. One sensor is placed right in front of the outlet fan, and the other is placed farther away at the other end.

fans) among the 26 sensor nodes based on the workload distribution. The blades are numbered 1 to 13 from top to bottom.

In Figure 10(a), it can be seen that when blades 7 through 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 Figure 10(b), where the odd-numbered blades are operational. The depicted spatial correlation was observed at a particular snapshot where the thermal impact of running blades are high. We observe significant spatial correlation in server temperature. Figure 10(c) and (d) 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 Figure 10(a) and (b). 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 and correlation threshold to $e_{th}^{temp} = 0.5°C$ and $\gamma_{th}^{temp} = 0.75$, respectively.

Figure 11(a) shows the number of REPs transmitting data to the sink as the manifestation changes over time and the reconstruction accuracy—that is, the number of nodes

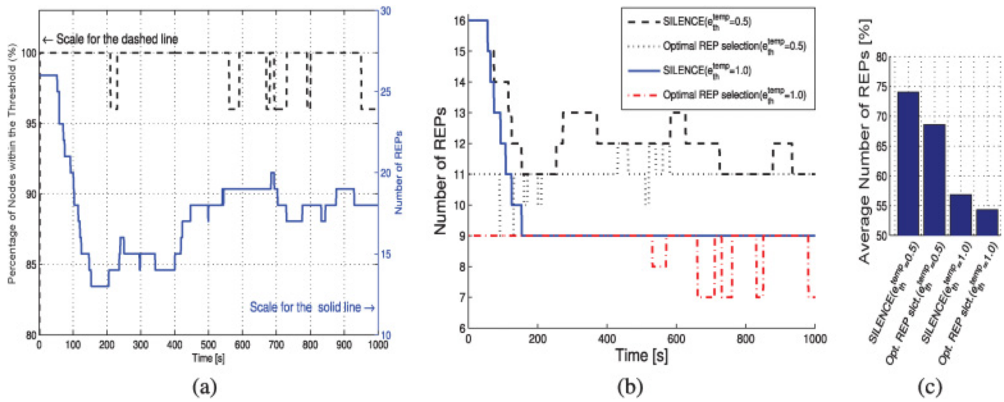

Fig. 11. Experiment results. (a) Number of REPs and percentage of nodes within the error threshold specified by the user (here, $e_{th}^{k} = 0.5°C$ for temperature) over time. Comparison of performance of SILENCE with centralized optimal REP selection in terms of number of REPs in time (b) and in terms of average percentage REPs in long-duration experiment (24 hours) (c).

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 nodes 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 centralized optimal REP selection for different error thresholds ($e_{th}^{temp} = 0.5°C$ and $1.0°C$). Figure 11(b) shows that the number of REP nodes selected by SILENCE is comparable to the optimal value for different error thresholds. Figure 11(c) shows the average percentage of REP nodes in the network in long duration (24 hours). It implies that even in the long-duration experiment, the performance of SILENCE is close to that of optimal value selection for different error thresholds.

6.2. Simulations

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; therefore, we performed additional simulations on the TinyOS simulator, TOSSIM. For simulations, it is key that the right models are used for the phenomenon. The spatiotemporal 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). This scaling of spatiotemporal correlation was done to show how our solution is tied to the variation in the manifestations of the observed phenomenon (data-centric nature). For a slowly varying field, 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

Table I. Parameters of the Model Used for Simulations

Deployment Parameters	
Terrain dimension	200×200 m^2
Topology	Uniform random
Number of nodes	400
Channel Parameters	
Path loss exponent	3.3
Shadowing standard deviation	5.5dB
Reference distance (D_0)	1 m
Power decay for D_0	–30dB
Radio Parameters	
Transmission power	–1dBm
White noise standard deviation	4dB
Radio noise floor	–105dBm
Hardware variance (for asymmetric links)	
Covariance matrix	$\mathbf{S} = [s_{11}s_{12}; s_{21}s_{22}]$
s_{11} (variance of noise floor)	3.7
s_{12} (covariance between S1 and S2)	–3.3
s_{21} (same as s_{12})	–3.3
s_{22} (variance of output power)	6.0

of performance on slow-, medium-, and fast-varying fields is intended to convey the idea that the proposed solution is not limited to a slow-varying band-limited phenomenon. The deployment, channel, and radio parameters used in our simulations are based on Zuniga [2010] and are listed in Table I.

The channel model used for simulations is anisotropic and asymmetric, as pointed out in Zuniga [2010]. Link asymmetry has both dynamic and static components. The dynamic component is due to thermal noise, which leads to a dynamic variation of a node's noise floor readings at runtime. This dynamic variation is usually modeled as a Gaussian random variable with mean 0 and a standard deviation (we use 4dB). The static component is caused by hardware variance (i.e., variance in the output power and baseline noise floor across nodes). Hardware variance is modeled as a multidimensional Gaussian process, where a covariance matrix captures the variances of the output power, the noise floor, and their correlation. The elements of the matrix $[s_{11}s_{12}; s_{21}s_{22}]$ represent the following: s_{11}, the variance around the mean value of baseline noise floor of the radio; s_{22}, variance of output power; and s_{12} and s_{21}, covariance between s_{11} and s_{22}. The higher the values of s_{11} and s_{22}, the greater the asymmetry in the link.

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 energy cost for communication takes into account the energy spent on both *control* (HELLO, JOIN, LEAVE, ACK, and AWAKE packets) and *data traffic* (DATA packets). The energy cost is calculated as follows:

$$E = E_{elec} + E_{comm}; \quad E_{comm} = V \cdot I \cdot \frac{L}{R}, \tag{10}$$

where E_{elec} [J] is the energy consumed by the electronic circuit, E_{comm} [J] is the energy consumed for communication, V [V] is the battery voltage, I [A] is the current, L [Byte] is the packet size, and R [Byte/s] 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 [Heinzelman et al. 2000] referred to as LEACH-gen.

Our implementation of LEACH-gen does not employ the localized election mechanism for REPs 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 intracluster communication 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 intracluster 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 intracluster 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, the 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, Figure 12 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%, and 25% reduction in the number of nodes transmitting data to the sink (REPs) for slow, moderate, and fast variation rates for the field as shown in parts (a), (b), and (c) of Figure 12, respectively. The percentage of nodes for which the reconstruction accuracy is satisfied is shown as well. 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 12(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. Hence, we infer that SILENCE allows for timely reaction to fast changes in the field.

Figure 13 shows the performance of SILENCE for the three different scenarios mentioned earlier. Figure 13(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. The energy cost for SILENCE is more than the one incurred by a sensor system employing LEACH-gen and CS (HorzVer-diff), which select REPs in a uniform-random manner. This can be attributed to the control overhead incurred by SILENCE. However, the average error in reconstruction is higher for LEACH-gen and CS (HorzVer-diff) compared to SILENCE, as shown in Figure 13(b). This is because SILENCE selects REPs based on the manifestations (data), whereas LEACH-gen and CS (HorzVer-diff) randomly choose REPs and do not adapt the number of REPs to the data variation. The inability of LEACH-gen and CS (HorzVer-diff) to adapt is evident in Figure 13(c). CS (HorzVer-diff) shows lower reconstruction error than LEACH-gen, as CS (HorzVer-diff) uses L_1 optimization to minimize

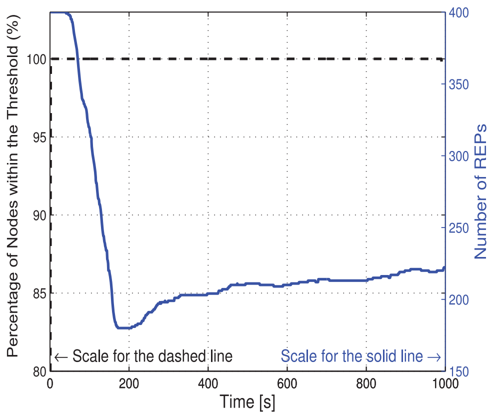

(a) Slowly-varying field

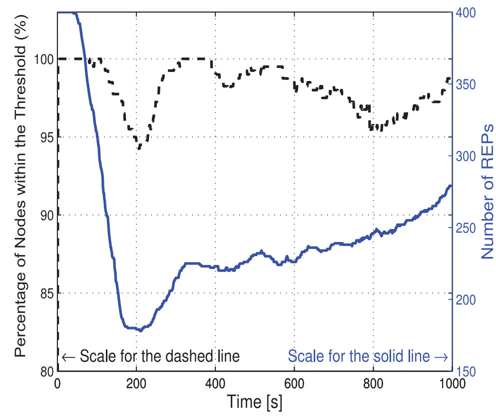

(b) Moderately-varying field

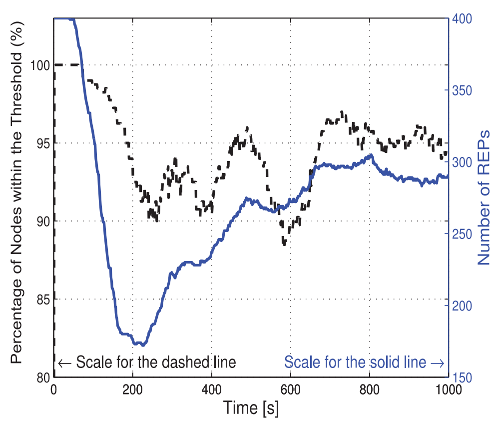

(c) Fast-varying field

Fig. 12. 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 slow-varying field (a), moderate-varying field (b), and fast-varying field (c).

reconstruction error. The extra control overhead incurred by SILENCE (manifesting itself as a marginal increase in energy expenditure as shown in Figure 13(a) for a slow-varying field) is justified by the significant reduction in reconstruction error compared to LEACH-gen and CS (HorzVer-diff), as shown in Figure 12(c).

From Figure 13, 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.

6.3. 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: (1) two nodes n and m

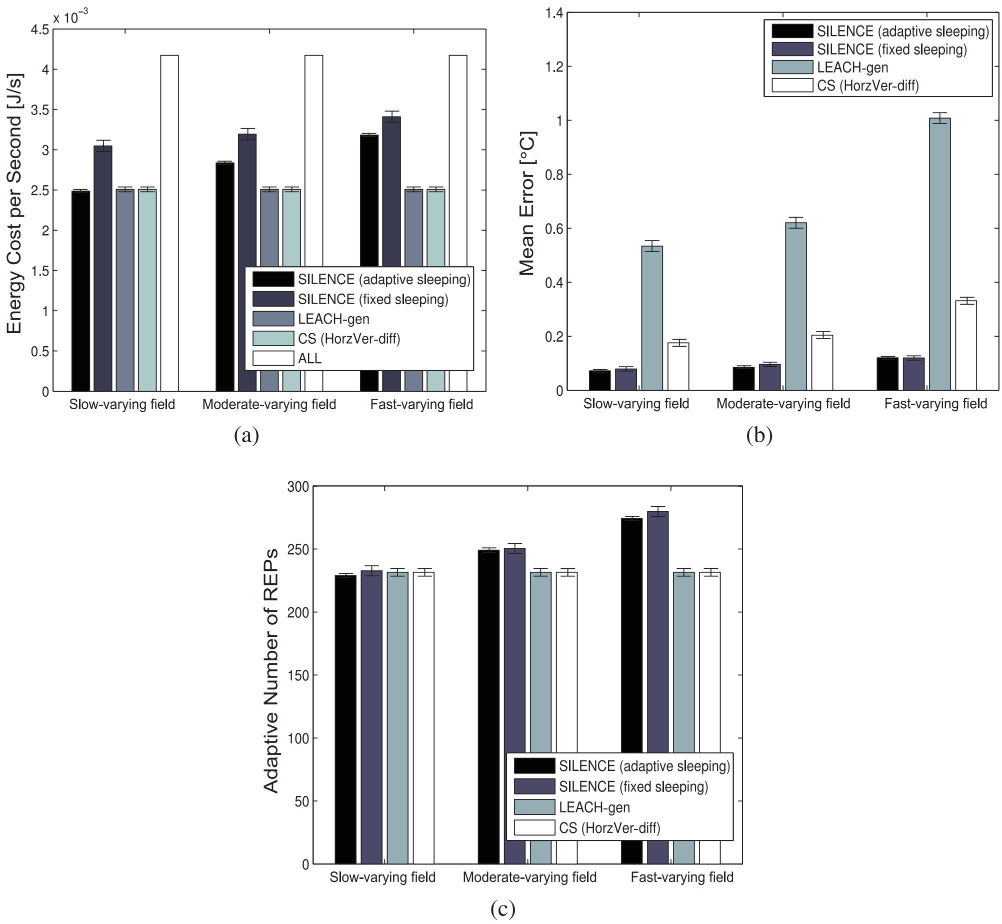

Fig. 13. Simulation results. (a) Average energy cost incurred by different adaptive sampling schemes. (b) Average error of different adaptive sampling schemes for different rates of variation in manifestations. (c) Average number of REPs of different adaptive sampling schemes.

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 (2) a different set of REPs are selected individually for each manifestation k to transmit data to the sink.

Figure 14 shows the performance of SILENCE when we follow the first strategy to monitor and reconstruct two manifestations, temperature and humidity, simultaneously. The model for humidity is again based on the actual data observed in our measurements and experiments on the testbed setup. Figure 14(a) shows a snapshot of the spatial distribution of two manifestations (above), temperature and humidity, and their reconstruction (below) based on the data obtained from the REP nodes selected using the first strategy. From Figure 14(b), one can clearly observe that approximately 40% of the nodes are sufficient to reconstruct two manifestations with 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 they are considered separately may be the same.

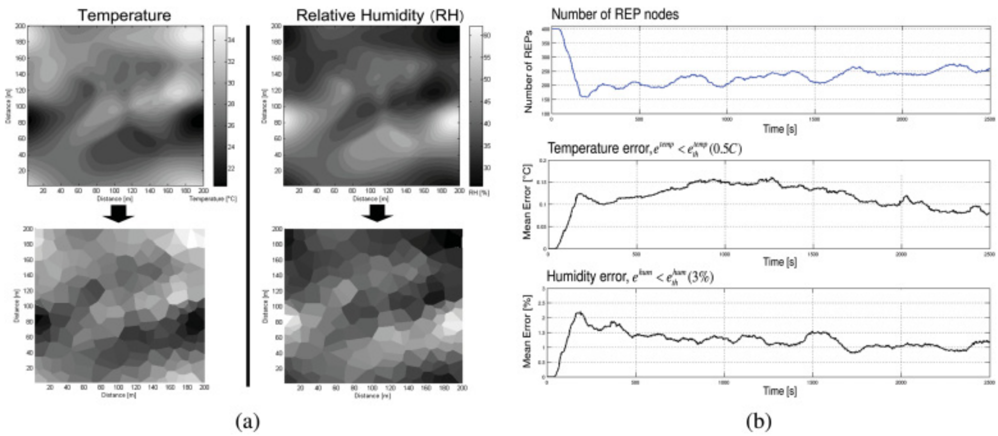

Fig. 14. Simulation results. (a) Actual distribution and reconstruction of temperature and humidity. (b) Performance of SILENCE scheme with multiple manifestations when the following strategy is used: two nodes n and m are said to be potential-ASSOCs of each other only when they sense similar and correlated values for *all* manifestations, and REPs are selected accordingly.

7. CONCLUSIONS

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. With regard to multiple manifestations, we are exploring more sophisticated strategies, such as ones that associate weights to each manifestation (while determining similarity and correlation). This will enable SILENCE to jointly handle uncorrelated manifestations (e.g., humidity and luminescence).

REFERENCES

Z. Abbasi, G. Varsamopoulos, and S. K. S. Gupta. 2010. Thermal aware server provisioning and workload distribution for Internet data centers. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing (HPDC'10)*. 130–131.

C. Aggarwal, A. Bar-Noy, and S. Shamoun. 2011. On sensor selection in linked information networks. In *Proceedings of the 2011 International Conference on Distributed Computing in Sensor Systems (DCOSS'11)*. 1–8.

H. Akaike. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19, 6, 716–723.

S. Bandyopadhyay and E. Coyle. 2003. An energy-efficient hierarchical clustering algorithm for wireless sensor networks. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'03)*. 1713–1723.

S. Bandyopadhyay and E. J. Coyle. 2004. Minimizing communication costs in hierarchically clustered networks of wireless sensors. *Computer Networks* 44, 1, 1–16.

A. Banerjee, T. Mukherjee, G. Varsamopoulos, and S. K. S. Gupta. 2010. Cooling-aware and thermal-aware workload placement for green HPC data centers. In *Proceedings of the 2010 International Green Computing Conference (IGCC'10)*. 245–256.

M. Bhardwaj and A. P. Chandrakasan. 2002. Bounding the lifetime of sensor networks via optimal role assignments. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'02)*. 1587–1596.

S. Chachra and M. Marefat. 2006. Distributed algorithms for sleep scheduling in wireless sensor networks. In *Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA'06)*. 3101–3107.

Z. Chen, S. Yang, L. Li, and Z. Xie. 2010. A clustering approximation mechanism based on data spatial correlation in wireless sensor networks. In *Proceedings of the Wireless Telecommunications Symposium (WTS'10)*. 1–7.

T. Cui, L. Chen, T. Ho, S. H. Low, and L. L. H. Andrew. 2007. Opportunistic source coding for data gathering in wireless sensor networks. In *Proceedings of the International Conference on Mobile Adhoc and Sensor Systems (MASS'07)*. 1–11.

D. L. Donoho. 2006. Compressed sensing. *IEEE Transactions on Information Theory* 52, 4, 1289–1306.

J. Geweke. 1982. Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association* 77, 378, 304–313.

R. W. Ha, P. Ho, X. S. Shen, and J. Zhang. 2006. Sleep scheduling for wireless sensor networks via network flow model. *Computer Communications* 29, 13–14, 2469–2481.

J. Haupt, W. U. Bajwa, M. Rabbat, and R. Nowak. 2010. Compressed sensing and network monitoring. *Next Wave* 18, 3, 16–25.

W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. 2000. Energy-efficient communication protocol for wireless microsensor networks. In *Proceedings of the 33rd Annual Hawaii International Conference on System Science (HICSS'00)*. 8020.

T. Huang, N. Kandasamy, and H. Sethu. 2012. Evaluating compressive sampling strategies for performance monitoring of data centers. In *Proceedings of the 2012 IEEE Network Operations and Management Symposium (NOMS'12)*. 655–658.

H. Jiang, S. Jin, and C. Wang. 2011. Prediction or not? An energy-efficient framework for clustering-based data collection in wireless sensor networks. *IEEE Transactions on Parallel and Distributed Systems* 22, 6, 1064–1071.

A. Keshavarzian, H. Lee, and L. Venkatraman. 2006. Wakeup scheduling in wireless sensor networks. In *Proceedings of the 7th International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'06)*. 322–333.

B. Krishnamachari, D. Estrin, and S. Wicker. 2002. The impact of data aggregation in wireless sensor networks. In *Proceedings of the 22nd International Conference on Distributed Computing Systems*. 575–578.

J. Kusuma, L. Doherty, and K. Ramchandran. 2001. Distributed compression for sensor networks. In *Proceedings of the 2001 International Conference on Image Processing (ICIP'01)*. 82–85.

E. K. Lee, I. Kulkarni, D. Pompili, and M. Parashar. 2012a. Proactive thermal management in green datacenter. *Journal of Supercomputing* 60, 2, 165–195.

E. K. Lee, H. Viswanathan, and D. Pompili. 2011. SILENCE: Distributed adaptive sampling for sensor-based autonomic systems. In *Proceedings of the 8th ACM International Conference on Autonomic Computing (ICAC'11)*. 61–70.

E. K. Lee, H. Viswanathan, and D. Pompili. 2012b. VMAP: Proactive thermal-aware virtual machine allocation in HPC cloud datacenters. In *Proceedings of the 19th International Conference on High Performance Computing (HiPC'12)*. 1–10.

X. Li, X. Xu, S. Wang, S. Tang, G. Dai, J. Zhao, and Y. Qi. 2009. Efficient data aggregation in multi-hop wireless sensor networks under physical interference model. In *Proceedings of the IEEE 6th International Conference on Mobile Adhoc and Sensor Systems (MASS'09)*. 353–362.

C. Liu, K. Wu, and J. Pei. 2007. An energy-efficient data collection framework for wireless sensor networks by exploiting spatiotemporal correlation. *IEEE Transactions on Parallel Distributed Systems* 18, 7, 1010–1023.

T. Melodia, D. Pompili, and I. F. Akyildiz. 2006. A communication architecture for mobile wireless sensor and actor networks. In *Proceedings of the 2006 3rd Annual IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON'06)*. 109–118.

V. Mhatre and C. Rosenberg. 2004. Design guidelines for wireless sensor networks: Communication, clustering, and aggregation. *Ad Hoc Networks Journal* 2, 1, 45–63.

J. Moore, J. S. Chase, and P. Ranganathan. 2006. Weatherman: Automated, online and predictive thermal mapping and management for data centers. In *Proceedings of the IEEE International Conference on Autonomic Computing (ICAC'06)*. 155–164.

P. Ogren, E. Fiorelli, and N. E. Leonard. 2004. Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment. *IEEE Transactions on Automatic Control* 49, 8, 1292–1302.

S. S. Pradhan, J. Kusuma, and K. Ramchandran. 2002. Distributed compression in a dense microsensor network. *IEEE Signal Processing Magazine* 19, 2, 51–60.

S. S. Pradhan and K. Ramchandran. 2000. Distributed source coding: Symmetric rates and applications to sensor network. In *Proceedings of the Data Compression Conference (DCC'00)*. 363–372.

G. Quer, R. Masiero, D. Munaretto, M. Rossi, J. Widmer, and M. Zorzi. 2009. On the interplay between routing and signal representation for compressive sensing in wireless sensor networks. In *Proceedings of the Information Theory and Applications Workshop (ITA'09)*. 206–215.

A. Scaglione. 2003. Routing and data compression in sensor networks: Stochastic models for sensor data that guarantee scalability. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT'03)*.

A. Scaglione and S. D. Servetto. 2002. On the interdependence of routing and data compression in multi-hop sensor networks. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MobiCom'02)*. 140–147.

G. Schwarz. 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 2, 461–464.

A. K. Seth. 2010. A MATLAB toolbox for Granger causal connectivity analysis. *Journal of Neuroscience Methods* 186, 2, 22–26.

K. Sohrabi, J. Gao, V. Ailawadhi, and G. J. Pottie. 2000. Protocols for self-organization of a wireless sensor network. *IEEE Personal Communications* 7, 1, 16–27.

J. A. Stankovic, T. F. Abdelzaher, C. Lu, L. Sha, and J. Hou. 2003. Real-time communication and coordination in embedded sensor networks. *Proceedings of the IEEE* 91, 7, 1002–1022.

M. C. Vuran, O. B. Akan, and I. F. Akyildiz. 2004. Spatio-temporal correlation: Theory and applications for wireless sensor networks. *Computer Networks* 45, 3, 245–259.

R. Willett, A. Martin, and R. Nowak. 2004. Backcasting: Adaptive sampling for sensor networks. In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN'04)*. 124–133.

X. Xu, Y. Hu, W. Liu, and J. Bi. 2008. Data-coverage sleep scheduling in wireless sensor networks. In *Proceedings of the 7th International Conference on Grid and Cooperative Computing (GCC'08)*. 342–348.

S. Yoon and C. Shahabi. 2007. The clustered aggregation (CAG) technique leveraging spatial and temporal correlations in wireless sensor networks. *ACM Transactions on Sensor Networks* 3, 1, Article No. 3.

O. Younis and S. Fahmy. 2004. Distributed clustering in ad-hoc sensor networks: A hybrid, energy-efficient approach. In *Proceedings of the 23rd Conference of the IEEE Communications Society (INFOCOM'04)*.

M. R. Zoghi and M. H. Kahaei. 2009. Efficient sensor selection based on spatial correlation in wireless sensor networks. In *Proceedings of the 14th International CSI Computer Conference (CSICC'09)*. 627–632.

M. Zuniga. 2010. Building a Network Topology for TOSSIM. Retrieved October 30, 2014, from http://www.tinyos.net/tinyos-2.x/doc/html/tutorial/usc-topologies.html.

Received March 2013; revised June 2014; accepted July 2014