The Precision and Energetic Cost of Snapshot Estimates in Wireless Sensor Networks

Paul G. Flikkema
Control Engineering Laboratory
Helsinki University of Technology
PL 5500, FI-02015
Finland
Tel: +358 9 451 5241
Fax: +358 9 451 5208
Email: paul.flikkema@hut.fi

Abstract— **Even for a specific application, the design space of wireless sensor networks is enormous, and traditional disciplinary boundaries are disappearing in the search for efficient integrated network architectures and protocols. There is a strong need to develop objective frameworks for the evaluation of performance and energetic cost as a function of network control at multiple levels, including the signal/data processing application, network organization, routing, MAC, and physical layers. This paper, a step in this direction, addresses the efficiency of linear estimation of a second-order random spatial field at a central server—a snapshot—in terms of the precision of the estimate and the energetic cost of computing it. We present a model that is based on the tasks of taking samples at a specific resolution and reporting them over the network to the server where the estimate is computed. It provides insight into the joint design of sampling and routing, an explicit efficiency measure for finite networks, and a performance benchmark for precision. The model, with predefined sampling locations at each resolution, implicitly requires centralized control to obtain spatially uniform sampling. We describe a decentralized reporting protocol called PROSE (Protocol for Randomized Opportunistic Sampling and Estimation) that uses localized decisions for sampling, data aggregation, and routing that are random and opportunistic. Finally, we show that the performance of PROSE compares favorably with the benchmark, and that PROSE admits simpler non-opportunistic random behavior at the cost of lower performance.**

I. Introduction

One of the most intriguing applications of wireless sensor networking technology is the monitoring of spatial phenomena. Spatial processes may consist of spatially localized events or broadly distributed processes. Either case can be captured by modeling the process as spatially correlated to a varying degree. If the dynamics are slow, but a portrait of the environment is needed relatively quickly, it is of interest to acquire a snapshot from the network. Roughly, the goal is to compute a useful estimate of the process state at a certain, recent, time at a central server. This involves three steps: sampling at the sensor nodes, reporting of the samples to the server via the network, and computing the estimate. Since in many instances energy is limited, the fundamental tradeoff is between the fidelity of the computed estimate and the energetic cost of obtaining it.

A sensible criterion for the fidelity of the estimate is its precision, defined as the inverse of a function of the error variances. In this paper we develop precise measures of the precision and energy cost for the particular case of linear estimation. The spatial field is modeled as a second order spatially-correlated random process. While the correlation of the sampled data between sensors can be arbitrary, most results are presented for the case where the correlation between samples increases with physical proximity between sensors. Sensors are assumed to be synchronized, and sample the field simultaneously.

A unique property of wireless sensor networks is that boundaries between computation and communication are blurred. Indeed, any useful computational task will involve exchange of information between sensors, so that communication is an integrated component of computation. The total energetic cost is incurred in (1) sampling and computation of the estimate and (2) in the reporting of the sample data via wireless networking. The energy required for the two tasks depends strongly on the spatio-temporal information bandwidth of the application. One application of interest is object detection, identification, and tracking via fusion of distributed images and image sequences, where the computational cost (including source coding) may in fact dominate. The opposite is true in embedded environmental sensing tasks where slowly changing scalar-valued sampled fields and correlated point measurements are of interest.

This work is motivated by the application of sampling environmental processes, including light, temperature, and soil moisture [1]. The sensors wake up at the correct time, sample the field, and then collaboratively communicate to send the samples to the server. Hence our primary goal is to characterize the precision of the snapshot estimate as a function of the energy cost. To simplify the presentation, the snapshot is considered independently of snapshots at other times. The communication cost is approximated by the number of packet transmissions rather than the number of bits. The underlying assumption is that communication of a few bits is nearly as expensive as communicating many. This approximation is expedient, but also reasonably accurate when (1) the per-sensor source rates are very low (on the order of 0.1 bits/sec or less) and (2) the communication cost is approximately constant for small packets. The latter is due to the fixed energy costs of radio warmup, packet detection and bit/frame synchronization, and radio shutdown. We assume here that each sensor freely aggregates data received from other sensors for which it functions as a relay. While we do not consider coding explicitly, it can easily exploit spatial correlations to control packet growth.

Asymptotic limits on data compression and transport capacity for the snapshot problem similar to the one considered here are presented in [2], but with the assumption that delay is permitted, so that sensors gather enough samples for efficient source coding. [3] and [4] consider the coding problem for estimation with unreliable channels for the case when the sensors are observing a common source. [5] considers bounds on multiresolution estimation performance and communication cost for deterministic fields for the ideal case of inverse-law propagation loss. In this paper, we consider a multiresolution paradigm in order to maximize spatial coverage of a spatially-correlated process; various strategies for approximation and data aggregation based on wavelet decompositions are described in [6]. Two-level coding and transmission strategies for one-dimensional networks based on the source entropy are proposed in [7]. The joint lossy coding/routing problem is shown to separate into shortest-path (least cost) routing and rate allocation problems in [8].

This paper uses the model described in Section II to capture the key signal processing parameters of *a priori* knowledge, measurement noise, inter-sensor correlation, and spatial sampling resolution. Closed-form analytical results are reported in the tractable case of uncorrelated observations, while numerical results are described in the general case of correlated samples. Section III describes the energy cost model that incorporates the communication-related parameters of inter-sensor distance and radio channel propagation loss, giving a realistic energetic cost of reporting via a wireless network. These models are idealistic in the sense that the protocols required to produce the estimates would require centralized control. We therefore introduce the concept of a reporting protocol in Section IV, where we motivate and describe PROSE, a decentralized protocol that relies only on local, independent sensor decisions. PROSE combines random and opportunistic sampling, routing, and aggregation decisions to obtain good performance and energy efficiency with the robustness of decentralized (autonomous) behavior. The performance of PROSE and a simpler, non-opportunistic variant are illustrated in Section V, followed by concluding remarks.

II. Sampling, Reporting, and Estimation Model

Consider a collection of sensors i, $i = 1, \ldots, N$ located in some environment. At the sampling instant, each sensor takes a sample x_i of a spatial process with $(N \times N)$ covariance matrix Σ_x. (This can be trivially extended to vector processes at the cost of additional notation.) The samples are corrupted by measurement noise n_i with corresponding covariance matrix Σ. The samples, collectively called the snapshot, are reported to a central server for computation of the estimate. If all samples are reported, then the resulting measurement vector at the server is $x + n$. However, we are interested in the case where a subset of the sensors take and report samples. This is modeled by an availability matrix $A = \text{diag}(a)$, where a is an N-vector such that

$$a_i = \begin{cases} 1, & \text{sample } x_i \text{ is reported;} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The resulting model becomes

$$s = Ax + n,$$

where A is the identity matrix I when all samples are reported.

The well-known linear minimum-variance estimate for this model is (see, e.g., [9])

$$\hat{x} = \Sigma_{MV} A' \Sigma^{-1} s, \quad (2)$$

with error covariance matrix

$$\Sigma_{MV} = (A' \Sigma^{-1} A + \Sigma_x^{-1})^{-1}. \quad (3)$$

The estimate can be computed only when the inverse of $A'\Sigma^{-1}A + \Sigma_x^{-1}$ exists; this is assured when Σ is positive definite [9]. The role of sample availability is modeled by the structure of A, which is the identity matrix with element $A(i,i)$ zero when the sample from sensor i is not available to the server.If $A = I$ and no prior information on x is available (Σ_x^{-1} is the zero matrix), then $\Sigma_{MV} = \Sigma$.

To gain insight on the effect of A on the error covariance matrix, we consider the case where the *a priori* data uncertainty and measurement noises are component-wise uncorrelated and uniform so that $\Sigma_x = \sigma_x^2 I$ and $\Sigma = \sigma^2 I$. Then

$$\Sigma_{MV}(i,j) = \begin{cases} \left(\frac{1}{\sigma_x^2} + \frac{1}{\sigma^2} \right)^{-1}, & i = j, A(i,i) = 1; \\ \sigma_x^2, & i = j, A(i,i) = 0; \\ 0, & i \neq j. \end{cases} \quad (4)$$

Since the data uncertainty and the noise variances are uniform over all sensors, we can define a scalar summary measure, the normalized average spatial sampling rate $\theta = K/N$, $0 \leq \theta \leq 1$ where $K = \text{tr}(A)$ is the number of sensors that sample and report to the server. Then the mean per-sensor error variance is

$$\bar{\sigma}_{MV} = \frac{1}{N} \text{tr} \left(\Sigma_{MV} \right) = \sigma_x^2 \left(1 - \theta \frac{\sigma_x^2}{\sigma_x^2 + \sigma^2} \right). \quad (5)$$

It is clear that the average performance increases linearly with the normalized average sampling rate as observed by the server. However, since the data is assumed uncorrelated, the per-sensor estimates fall into two classes: poor for non-reporting sensors and relatively good for the reporting sensors.

In the following, we borrow from statistics to define the overall precision of the estimate

$$\Phi = \left[\text{tr} \left(\Sigma_{MV} \right) \right]^{-1}, \quad (6)$$

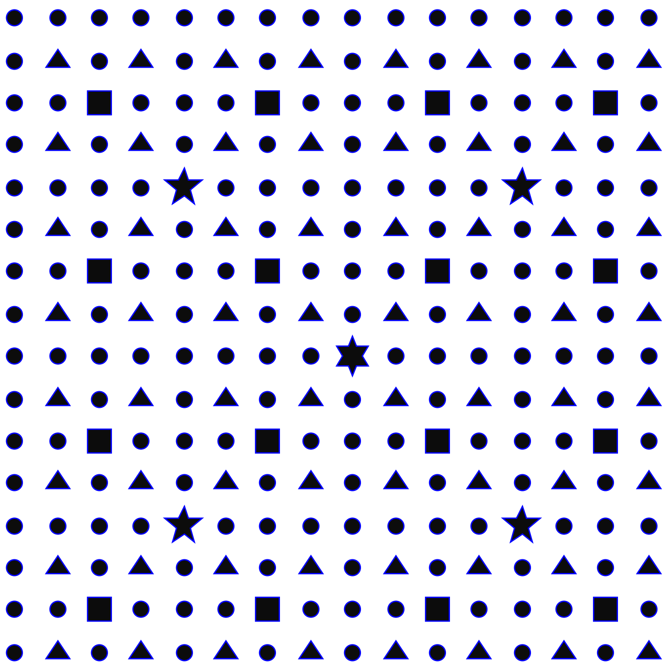

Fig. 1. Two-dimensional multiresolution grid of 256 sensors. Sensor nodes are represented by polygons, with fewer sides as the spatial sampling rate increases. The 6-pointed star denotes the central server. At level $L = \log_4 256$, all nodes sample the field.

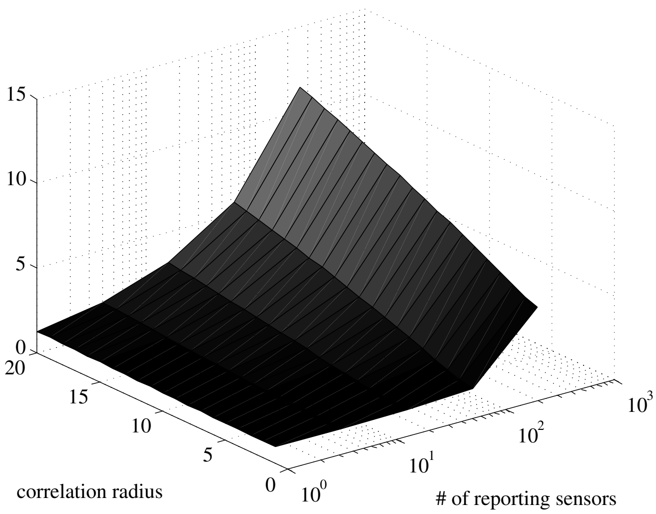

Fig. 2. Precision (inverse sum-variance) for uniform multiresolution sampling (as in Figure 1) as a function of correlation radius (for a minimum correlation value of 0.1) and number of reporting sensors; here the prior data uncertainty and measurement noise are $\sigma_x = 1$ and $\sigma = 1/4$ respectively.

which will generally be a non-decreasing function of the inter-sensor data correlation and the spatial sampling rate.

The general case of correlated data is easily handled by numerical evaluation of (3). Let $d(i,j)$ be the distance between two sensors i and j. We model the covariances using the covariance matrix

$$\Sigma_x(i,j) = \sigma_x^2 e^{-\frac{d(i,j)}{\alpha}}, \tag{7}$$

where typically $d(i,j)$ is the Euclidean distance and α is a constant that defines the *correlation radius*, i.e., the distance from a particular sensor that defines a minimum correlation value; for example, $d(i,j) = 5\alpha$ implies almost complete decorrelation (8% correlation) at unit distance.

To this point we have made no assumptions about the physical placement of the sensors. In the remainder of this section and in the next, we introduce more structure to provide insight into sampling and reporting. We assume that the N sensors are deployed in a two-dimensional $\sqrt{N} \times \sqrt{N}$ lattice or grid with distance d between neighboring nodes in each axis. Additional structure is gained by assuming that $N = 4^L$ so that spatial subsampling is easily achieved by sampling at 4^l sensors where $l = 0, \ldots, L$. Hence the spatial sampling rate is 4^{l-L}, and $l = L$ is the base resolution level with complete sampling and reporting. This is shown in Figure 1 for $N = 256$ sensors. The depicted subsampling, related to quadtree representations used in image coding, assumes that correlation increases with physical proximity.

The average precision improves as the correlation radius increases and as the number of sampling nodes increases, as shown in Figure 2 for $N = 256$. While the latter effect is clear from (5), the former effect shows that the correlation between sensor measurements is exploited by the estimator to improve the overall precision.

III. Reporting Cost and Efficiency

We assume that the energetic cost of communication is determined by the distance r between two nodes and the propagation loss exponent β. In typical wireless channels, $3 \leq \beta \leq 5$. If P_T is the transmitted power and P_R is the received power, they are related by $P_R = P_T r^{-\beta}$. Thus the cost scales with distance according to r^β.

We consider two reporting scenarios. In both, we assume full aggregation of the data along the path to the server so that each node transmits one packet per snapshot. While the packet length will increase somewhat, this growth can be controlled with straightforward compression techniques [10]. The first, called *maximum-sleep*, is suggested by the multi-resolution sampling structure: we assume that only only the sensors that sample communicate, so that the maximum number of nodes can remain entirely asleep. This could model the situation where the sensors are deployed at a particular spatial resolution. To conserve energy, each sensor at level l sends its information to the nearest sensor at level $l-1$. If sampling is performed at level l, $0 \leq l < \log_4 N$, the incremental reporting cost is

$$\gamma_l = 4^l \left(\frac{\sqrt{N}}{2^{l+1}} \sqrt{2} d \right)^\beta .$$

Since at the highest level $l = L$ each remaining node is within distance d of a node already reporting, we have

$$\gamma_L = \left(N - \sum_{l=1}^{L-1} 4^l \right) d^\beta .$$

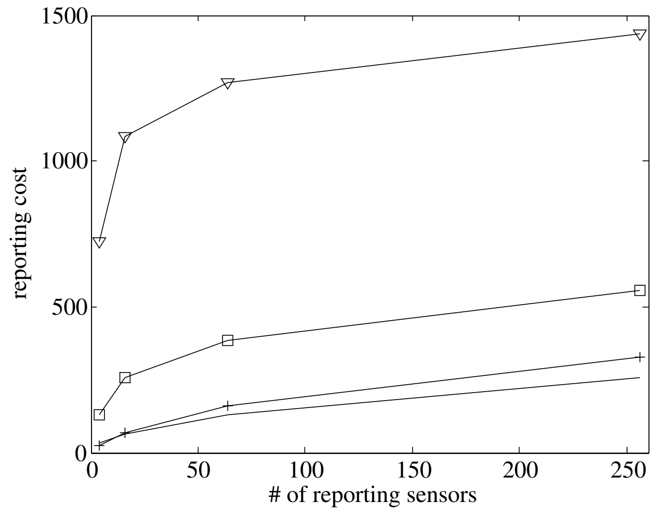

Fig. 3. Reporting cost for multiresolution sampling and reporting. Results are shown for multihop (no symbols) and maximum-sleep reporting for $\beta = 1$ (+), $\beta = 2$ (□), and $\beta = 3$ (▽).

Hence the total reporting cost for sampling at level l is

$$\Gamma_l = \sum_{m=1}^{l} \gamma_m \quad (8)$$

In the second *multihop* reporting scenario, all sensors are available to relay the samples to the server; in this case many sensors serve only as relays for the information. We assume that a multihop path with the shortest number of single (distance d) hops to the server is always used. Let the number of nodes reporting samples be K, where $K = 4^l$ in keeping with the multiresolution approach (Figure 1). Then a convenient and reasonably tight upper bound on the reporting cost is

$$\psi(K, N) = \sqrt{KN} d^{\beta}. \quad (9)$$

This is an upper bound since a slightly lower cost can be achieved by routing and aggregating all samples along one axis in the direction of the server followed by routing and aggregation along two paths toward the server along the other axis. While (9) can be proved by a careful counting argument, a more intuitive approach is to note that when K nodes sample (at level $\log_4 K$), the average number of hops per sensor is $\sqrt{N/K}$. Note that $\psi(N, N) = Nd^{\beta}$, or that each sensor must perform on average one and only one transmit/receive transaction in full reporting.

The relative costs of the two approaches for any network size and propagation loss exponent can be found from (8) and (9). For realistic values of β (i.e., except perhaps in the case of very dense networks where the propagation is near-field), the longer hops (between levels) used in maximum-sleep reporting are expensive, as shown in Figure 3. Note that the cost for multi-hop reporting is independent of β since $d = 1$.

A measure of the efficiency of sampling and multi-hop reporting as a function of resolution and correlation radius is the ratio of the precision to the cost,

$$\eta(K, N) = \kappa \frac{\Phi}{\psi(K, N)}$$

where Φ (from (6)) is understood to be a function of the data uncertainty, measurement noise, sampling rate and correlation radius, and the constant κ is determined primarily by the efficiency of the transmit power amplifier and the modulation technique.

The efficiency strongly depends on the measurement noise power and the data uncertainty. Considering the case when the data are uncorrelated, when $\sigma_x^2 \gg \sigma^2$, we find from (5) that the mean per-sensor precision is

$$\bar{\phi}_{MV} = (\bar{\sigma}_{MV})^{-1} \approx \frac{1}{(1-\theta)\sigma_x^2};$$

increasing the sampling rate yields large benefits. Here, the per-sensor energy usage is $\sqrt{KN}/N = \sqrt{\theta}$, which is ultimately dominated by the inverse-law precision. When measurement noise dominates, more measurements have little benefit, so precision grows much more slowly (and efficiency suffers). Indeed, in many scenarios, the efficiency is a *decreasing* function of the sampling rate, since the precision grows too slowly relative to the cost.

IV. Random Opportunistic Sampling and Estimation

The estimation performance results of Section II provide guidelines for deployment density as a function of the spatial statistics, and the previous section describes the reporting cost and efficiency for multi-resolution sampling. While the multiresolution structure provides optimum coverage of the field at a given sampling density, the implied control algorithms and protocols do not take into account available network resources, including energy supplies and sampling locations. They also do not suggest simple algorithms for distributed control and its advantages of robustness. Therefore it is of great interest to consider *reporting protocols* with good performance that that balance resource consumption and allow graceful degradation in the event of failures.

Any reporting protocol must encompass the tasks of (i) determining when each sensor should sample, and (ii) how the sensors communicate the samples to the server. Given that many sensor networks are ad hoc due to mobility and/or communication channel dynamics, the protocol should be based on local decisions with a minimum of centralized control. It should also minimize energy costs, allow for spatially distributed samples at varying spatial resolutions, and provide for a desired distribution of energy consumption over the sensors.

One approach to these challenges is random sampling: in the simplest case, each sensor flips a coin (this can be implemented by sampling sensor radio noise) to determine at every sampling epoch whether to take a sample. The probability of sampling, λ, is set to achieve a desired average spatial resolution. Here we assume that each sensor's decision is independent of all others to maximize robustness. The only global parameter is the sampling probability λ, which can be easily distributed as part of network route discovery and/or maintenance.

The performance of random sampling is found by recalling the availability matrix $A = \text{diag}(a)$, with a a $\{0,1\}$-valued N-vector. Hence any randomly-sampled snapshot can be modeled by treating a as a random vector; the assumption of independent decisions by the sensors implies that the elements $a_i = Y_i$ of a are independent binary random variables such that

$$P(Y_i = y_i) = \begin{cases} \lambda, & y_i = 1; \\ 1-\lambda, & y_i = 0. \end{cases} \tag{10}$$

This determines A so that the error covariance matrix can be found from (3).

The samples must be reported to the server; since we have demonstrated that multihop reporting is the most efficient for typical wireless channels, we know that sensors on a path between a sampling sensor and the server must communicate anyway. This suggests a joint sampling and reporting protocol that merges the two tasks, which we call PROSE (Protocol for Randomized Opportunistic Sampling and Estimation). PROSE is based on the observation that the energetic cost of sampling may be several orders of magnitude lower than the cost of reporting. Hence PROSE is based on a two-phase determination of sampling and reporting. In the first phase, the outcome of Y_i determines whether to take a sample and *originate* a report to the server. The second phase determines only whether a sample is taken; that is, the outcome X_i is conditionally dependent on Y_i as follows:

$$P(X_i = x_i | Y_i = y_i) = \begin{cases} \mu, & x_i = 1, y_i = 0; \\ 1-\mu, & x_i = 0, y_i = 0; \\ 1, & x_i = 1, y_i = 1; \\ 0, & x_i = 0, y_i = 1. \end{cases} \tag{11}$$

Here, Y_i is the first (sample and originate) decision; if this decision is negative, then the final decision X_i is made based on a second coin flip. This can be interpreted as follows: a sensor makes a random decision whether to become an *originating* node that both samples and originates a report. If this decision is negative, then the sensor makes another random decision on whether to sample or not; however, it cannot originate a report regardless of the outcome of the second decision. On the other hand, it can opportunistically aggregate its sample into a packet that it aggregates from received packets and forwards, so this is a "sample and opportunistically aggregate" decision. Hence μ is the probability that a sensor will be opportunistic given that it decides not to sample and originate. Since any node may be called on to forward a report via a reception and transmission, the second decision also determines whether the node will opportunistically aggregate its sample into the packet for the ultimate use of the estimation computation at the server.

Consider the simplest case where all sensors sample, but only those whose first decision is positive originate a report. This is selected by choosing $\mu = 1$. Then any node on a path from an originating node can aggregate its sample into the packet before forwarding at a very low additional cost. In the ideal case where the energetic costs of sampling and aggregation are zero, the additional information is free. In general, the network-wide cost of sampling is proportional to

$$n \cdot P(X_i = 1) = n \cdot (\lambda + \mu(1-\lambda)).$$

Reducing μ decreases the network-wide cost of snapshot sampling and aggregation, and it can be selected to balance the benefit and cost of the additional precision resulting from opportunistic behavior. Notice that the joint distribution of X_i, Y_i is completely determined by (10,11). The two-phase protocol captures relevant deterministic special cases: if $\mu = 0$ then only random sampling is performed and opportunistic behavior is disabled. On the other hand, if $\lambda = 1$, then every sensor samples, guaranteeing maximum precision.

The final aspect of PROSE is the forwarding of samples to the server using randomized link selection. We assume here that a set of shortest-path routes is available via an underlying routing layer. The basic idea is that the sensor chooses a one-hop link at random from a set of links that form the first hop in shortest-path routes. This can be achieved using a gradient-based algorithm, where each sensor maintains a list of links to each of a set of sensors that are one hop closer to the server. (The list can be established and maintained as long as at least one path to the server exists by periodically updating the distance to the server via exchanges of distance information with neighbors.) These links are each thus the first in a shortest (or minimum-cost) path. If the set consists of one link, then the packet is forwarded; otherwise, the sensor chooses a link at random and forwards the packet. The same decision process is used by every node that forwards a packet. The randomized link selection has two consequences: first, the energy cost is be distributed more evenly. It also increases route diversity, thus increasing the potential number of sensors that can opportunistically aggregate their samples.

Numerous embellishments to the protocol are possible that could exploit more characteristics of the problem or give a greater degree of local control. For example, it might be useful for nodes to have different propensities for opportunism, so that sensors nearest the server do not have a higher level of opportunistic participation in the estimate. Finally, sensors could be endowed with the ability to autonomously control their behavior, perhaps in collaboration with neighbors [11]; in this case, λ and μ would become local variables λ_i, μ_i.

V. Performance of Randomized Snapshot Protocols

The performance of PROSE is random: each estimate is computed from a random number of samples from sensors in randomly chosen locations. To assess its precision performance, it can be seen that error covariance can be computed from (1,3) with a now a random vector. In the non-opportunistic ($\mu = 0$) case, the number of non-zero entries in a is binomially distributed with parameters (N, λ). In this paper,

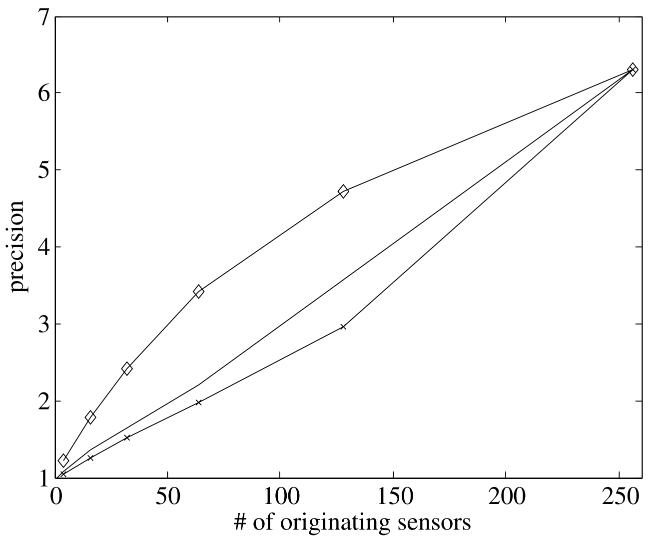

Fig. 4. Relative precision for PROSE for $\mu = 1$ ($\diamond$) and $\mu = 0$ ($\times$). Precision for multiresolution sampling is shown as a solid line for comparison. The correlation radius is 5 (for a minimum correlation of 0.1); prior data uncertainty and measurement noise are the same as in Figure 1.

the general case is evaluated using Monte-Carlo evaluation of (1,3).

We compare the performance of PROSE with its non-opportunistic special case and with the multiresolution (i.e., uniformly spaced) sampling. First note that all three cases will have the same performance when the target number of samples is N; here $\lambda = 1$ so that all sensors sample and report in all three cases. When $\mu > 0$, performance depends on the random routes chosen from each of the originating sensors. As indicated in Figure 4, the average performance when $\mu = 0$ is degraded relative multiresolution sampling; this is due to the fact that for the randomized protocols the abscissa is actually the mean of the above binomial distribution, and that the samples are not uniformly distributed; the latter results a spatial sampling loss. The performance gain of PROSE over multiresolution sampling and reporting is due simply to its ability to aggregate samples along the routes from originating nodes to the server. We note that a performance gain similar to that of PROSE could also be obtained by multiresolution sampling with opportunistic sampling on random routes to the server, but an implementation would lack the advantages of a decentralized random protocol, including robustness. In fact, multiresolution opportunistic sampling can be considered a special case of PROSE where an originating sensor q belonging (based on its location) to a pre-selected spatially subsampled set Q has $\lambda_q = 1$, $q \in Q$, with $\lambda_i = 0$ for all other sensors.

VI. Conclusion

This paper considers the energetic cost of precision in the computation of snapshot estimates of correlated random fields at a centralized server (base station, sink) in wireless sensor networks. The performance and efficiency of linear estimation based on uniform multiresolution sampling for two routing algorithms is determined primarily as a benchmark for any means of implementing the snapshot. We also describe a protocol called PROSE for the problem that encompasses sampling, routing, and data aggregation. It randomizes each of these tasks to achieve good performance, robustness, and spatial distribution of energy consumption. A spatially subsampled grid model is used to make the results intuitively digestible, capturing the essence of the problem—subsampling of sensors embedded in a field whose correlation is a continuous function of physical proximity—in a convenient way. More importantly, PROSE is independent of the physical network topology, and is applicable to subsampling of correlated fields for a large spectrum of detection, estimation, and model inference problems.

References

[1] Z. Yang *et al.*, "WiSARDNet: a system solution for high performance in situ environmental monitoring," in *Proc. 2nd Int'l Workshop on Networked Sensor Systems (INSS 2005)*, 2005.
[2] D. Marco *et al.*, "On the many-to-one transport capacity of a dense wireless sensor network and the compressibility of its data," in *Proc. IPSN '03*, 2003.
[3] P. Ishwar *et al.*, "On rate-constrained distributed estimation in unreliable sensor networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, pp. 765–775, Apr. 2005.
[4] K. Liu, H. E. Gamal, and A. Sayeed, "On optimal parametric field estimation in sensor networks," in *Proc. IEEE Statistical Signal Processing Workshop*, 2005.
[5] R. Willett, A. Martin, and R. Nowak, "Backcasting: Adaptive sampling for sensor networks," in *Proc. IPSN '04*, 2004.
[6] D. Ganesan, D. Estrin, and J. Heidemann, "Dimensions: why do we need a new data handling architecture for sensor networks?" *ACM SIGCOMM Computer Communication Review*, vol. 33, pp. 143–148, Jan. 2003.
[7] T. ElBatt, "On the cooperation strategies for dense sensor networks," in *Proc. IEEE International Conference on Communications (ICC '04)*, 2004.
[8] R. Cristescu and B. Beferull-Lozano, "Lossy network correlated data gathering with high-resolution coding," in *Proc. IPSN '05*, 2005.
[9] D. G. Luenberger, *Optimization by Vector Space Methods*. Wiley, 1969.
[10] P. Flikkema, B. West, and B. Ruggeri, "Network-aware lossless source coding of spatio-temporally correlated sensor data," in *Proc. IEEE CAS Workshop on Wireless Communication and Networking*, 2003.
[11] J. Kay and J. Frolik, "Quality of service analysis and control for wireless sensor networks," in *Proc. 1st IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS 2004)*, 2004.