

Estimating Inhomogeneous Fields Using Wireless Sensor Networks

Robert Nowak^{*}, Urbashi Mitra[†] and Rebecca Willett

February 4, 2004

Abstract

Sensor networks have emerged as a fundamentally new tool for monitoring spatial phenomena. This paper describes a theory and methodology for estimating inhomogeneous, two-dimensional fields using wireless sensor networks. Inhomogeneous fields are composed of two or more homogeneous (smoothly varying) regions separated by boundaries. The boundaries, which correspond to abrupt spatial changes in the field, are non-parametric one-dimensional curves. The sensors make noisy measurements of the field, and the goal is to obtain an accurate estimate of the field at some desired destination (typically remote from the sensor network). The presence of boundaries makes this problem especially challenging. There are two key questions: 1. Given n sensors, how accurately can the field be estimated? 2. How much energy will be consumed by the communications required to obtain an accurate estimate at the destination? Theoretical upper and lower bounds on the estimation error and energy consumption are given. A practical strategy for estimation and communication is presented. The strategy, based on a hierarchical data-handling and communication architecture, provides a near-optimal balance of accuracy and energy consumption.

^{*}R. Nowak (nowak@engr.wisc.edu) was supported by the National Science Foundation, grants CCR-0310889 and ANI-0099148, the Office of Naval Research grant N00014-00-1-0390, and the State of Texas ATP, grant 003604-0064-2001. R. Nowak and R. Willett (willett@rice.edu) are with the Departments of Electrical and Computer Engineering at the University of Wisconsin-Madison and Rice University.

[†]U. Mitra (ubli@usc.edu) was Supported by the Texas Instruments Visiting Professorship and the University of Southern California. She is with the Electrical Engineering Department at the University of Southern California.

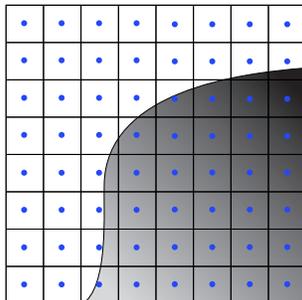


Figure 1: A wireless sensor network sampling a two-dimensional field. Dots indicate sensor locations and squares indicate the extent of each sensor's measurement.

1 Introduction

Sensor networks have emerged as a fundamentally new tool for monitoring inaccessible environments such as non-destructive evaluation of buildings and structures; contaminant tracking in the environment; habitat monitoring in the jungle; and surveillance in military zones. These *ad hoc* networks are envisioned to be a collection of embedded sensors, actuators and processors. We shall assume that communication between sensors is done in a wireless fashion. Sensor networks are distinguished from more classical networks due to strict limitations on energy consumption, the density of nodes, the simplicity of the processing power of nodes and possibly high environmental dynamics. An important problem in sensor networking applications is the estimation of spatially varying processes or fields. Consider a network sensing a field composed of two or more regions of distinct behavior (e.g., differing mean values for the sensor measurements). An example of such a field is depicted in Figure 1. In this application, the goal of the sensor network is to sense the field, construct an estimate of the field, and communicate the estimate to a desired (typically remote) destination. In contrast to prior work, we drop the assumption of homogeneity and instead consider inhomogeneous fields that are composed of two or more homogeneous (smoothly varying) regions separated by smooth boundaries. While the proposed method exhibits near-optimal theoretical properties and is practical for both homogeneous and inhomogeneous fields, we focus in this paper on inhomogeneous fields. This is a particularly challenging problem in sensor networks because boundary detection requires significant system resources. Furthermore, abrupt changes in the field are often of greatest scientific interest, and therefore we view the inhomogeneous case as especially relevant to applications.

1.1 Wireless Sensing of Inhomogeneous Fields

There are two fundamental limitations in the estimation problem. First, the accuracy of a field estimate is limited by the spatial density of sensors in the network and by the amount of noise associated with the measurement process. Second, energy constraints may limit the complexity of the estimate that is computed and ultimately communicated to a desired destination. The trade-off between accuracy and energy consumption can be characterized as follows.

Assume that n sensor nodes are arranged on an $\sqrt{n} \times \sqrt{n}$ square lattice (assuming a planar, square field). Suppose that the field being sensed consists of two homogeneous regions separated by a one-dimensional

boundary (like the case depicted in Figure 1). To be specific, let us assume that the field consists of two smoothly varying regions separated by a smooth boundary. Under these assumptions, most of the nodes will be in smooth regions away from the boundary, and only $O(\sqrt{n})$ nodes lie next to the boundary. Each sensor node makes a (noisy) measurement of the field. Specifically, each sensor makes a measurement of the field at its location which is contaminated with a zero-mean Gaussian distributed error. Assume that the noises are independent and identically distributed at each sensor. It is known that, under the assumptions on the class of fields above and the white noise model, the mean-square error (MSE) cannot, in general, decay faster than $O(n^{-\nu})$, for some $0 < \nu \leq 1$ that depends on the smoothness of the regions and the boundary [1, 2, 3]. That is, no estimator (based on centralized or distributed processing) can exceed this convergence speed-limit.

To quantify the total energy required to compute and transmit a field estimate of this accuracy, we assume a $1/r$ energy decay behavior, where r is the distance from the transmitter (in light of the $2 - d$ nature of our formulation). Other measures are possible (*e.g.*, $1/r^\alpha$ for some $\alpha \geq 1$); see [4] for examples. Note that even simple calculations require communications that consume at least $O(\sqrt{n})$ bit-meters of energy. For example, consider computing the average values between the nearest neighboring nodes in the network. This requires the transmission of $O(n)$ bits (a certain number for each node in the network) over a distance of $O(1/\sqrt{n})$ meters (the distance between nearest neighbors). Such a trivial operation is necessary for almost any imaginable field estimation process. Thus, the total energy required to compute and transmit the field estimate is at least $O(\sqrt{n})$. Note that, here and throughout the paper, we assume that local computations require a negligible amount of energy in comparison to communications. Combining this lower bound on the energy consumption with the lower bound on the MSE decay rate above yields a “best-case” trade-off between accuracy (MSE) and energy (\mathcal{E}) of the form

$$\text{MSE} \sim \frac{1}{\mathcal{E}^{2\nu}}.$$

It is important to note that this relation should not be interpreted to mean that a fixed number of sensor nodes using more energy can provide more accuracy. Rather, both the MSE and the energy consumption are functions of the number of sensor nodes, and the above relation indicates how the accuracy and energy consumption behave as the density of nodes increases. Moreover, note that the MSE cannot decay faster than $O(n^{-1})$, the parametric rate. Therefore, the very best trade-off between MSE and energy is $\text{MSE} \sim \frac{1}{\mathcal{E}^2}$, and this is generally unachievable except under very restrictive and unrealistic assumptions on the nature of the field.

This paper explores these basic trade-offs between MSE and energy consumption as functions of node density. We propose and develop field estimation algorithms based on multiscale partitioning methods. The algorithms are quite practical and map nicely onto a sensor network architecture. Moreover, we demonstrate theoretically that our methods nearly achieve the best-case MSE/energy trade-offs discussed above. The theory hinges on an application of our extension [5] of the Li-Barron bound for complexity regularized model selection [6] to bound the MSE and to bound the expected energy consumption. Since our methods (nearly) achieve the optimal trade-offs, no other schemes can be devised that will (asymptotically) perform significantly better. Simulation experiments verify the predicted theoretical performance of our methods. The algorithms and theoretical analysis in this paper expand upon our previous work [7] which considered

only piecewise constant fields.

1.2 Related Work

Due to the recent emergence of sensor network research, there is limited literature concerning field estimation via wireless networks. We are aware of several lines of investigation. First, several authors [8,9,10,11] have considered the problem of estimating a field via wireless sensor networks. In those works, rate-distortion analysis is employed to study essentially the same trade-off between accuracy and communication costs as discussed above. However, the class of fields considered by those authors differs strikingly from our investigation. Rather than consider inhomogeneous fields as we do here, they considered classes of homogeneous processes without boundaries (e.g., stationary Gaussian fields). While possibly adequate for some applications, we feel that boundaries corresponding to abrupt spatial changes in the field may occur quite frequently in many envisioned application domains. Moreover, abrupt changes in the field are likely to be the features most critical to the user. The theory and methods in [8,9,10] do not address such situations. Second, to address the important problem of boundary detection, several proposals have been made to employ classical image processing tools in sensor network algorithms. In particular, the work of [12] has led to simple and communication-efficient strategies for detecting field boundaries using wireless sensor networks. Similar, “image-processing” influenced schemes are considered in [13]. However, like most classical image edge detection methods, it is difficult to quantify the accuracy of these strategies.

We also mention the close connection between the hierarchical estimation and communication architecture employed in our methods and the data collection algorithm proposed in [14], which also exploits a hierarchy to reduce communication costs. That algorithm is based on a hierarchical compression scheme where clusterheads aggregate measurements from child nodes to pass to the next layer of hierarchy. Very similar strategies underlie our methods and are critical for achieving near-optimal trade-offs between accuracy and communication. Thus, we view our work as theoretically supporting the use of such hierarchies in sensor network architectures.

1.3 Organization

This paper is organized as follows. In Section 2 we define the problem and develop the basic algorithmic and communication structure of our field estimators. In Section 3 we study the theoretical properties of the estimators. In particular we develop bounds on the mean-squared estimation error and on the communication costs of the proposed methods. Simulation results are provided in Section 4. Section 5 draws final conclusions and discusses ongoing research. The Appendix contains key derivations and proofs.

2 Problem Formulation and Approach

Assume a two-dimensional field, denoted by f , is being sensed and that it is composed of smooth regions of Hölder- α regularity in \mathbb{R}^2 , $\alpha \in (0, 2]$, with smooth contours or *boundaries* that are Hölder- β regular in \mathbb{R}^1 , $\beta \in (0, 2]$. As special cases, if α and β are integer valued, then these cases include fields in which smooth regions are C^α (α continuous derivatives in \mathbb{R}^2) and the boundaries are C^β (curves that behave

locally like C^β one-dimensional functions). The domain of the field is assumed to be the unit square, $[0, 1]^2$. Furthermore, the range of the field is assumed to be within $[-R, R]$, for some constant $R > 0$. Extensions of our theory and methods to three-dimensional fields are straightforward, and we will focus on two-dimensional fields and sensor networks to illustrate the salient ideas.

The sensor network is a collection of n wireless nodes arranged on a uniform, square lattice in $[0, 1]^2$. Each node measures the field at its position. Denote the sensor measurements by $x \equiv \{x_{i,j}\}$, where the indices range $1 \leq i, j \leq \sqrt{n}$. The indices indicate the location of the corresponding node which is at the center of the square $S_{i,j} = [(i-1)/\sqrt{n}, i/\sqrt{n}] \times [(j-1)/\sqrt{n}, j/\sqrt{n}]$. Let $f_{i,j} = n \int_{S_{i,j}} f$, the average value of the field over the square $S_{i,j}$. The array $\{f_{i,j}\}$ can be viewed as a ‘‘pixelized’’ version of the original field. Assume that the measurement at each node has the form

$$x_{i,j} = f_{i,j} + n_{i,j}, \quad (1)$$

where $n_{i,j}$ is an observation noise that could encompass ambient noise in the environment, electronic transducer noise sources, and possibly quantization errors. The noises $\{n_{i,j}\}$ are assumed to be independently and identically Gaussian distributed with mean zero and variance σ^2 . Thus, we have $x_{i,j} \sim \mathcal{N}(f_{i,j}, \sigma^2)$, a realization of a Gaussian random variable with mean $f_{i,j}$ and variance σ^2 . The goal of the sensor network is to compute an estimate of the field f given the noisy measurements x and to transmit this estimate to a desired location (assumed to be remote from the sensor network). The basic problem is illustrated in Figure 1.

Our hierarchical approach to estimation and communication is as follows. The sensor measurements can be viewed as sampling the field over a partition of n sub-squares of sidelength $\frac{1}{\sqrt{n}}$, as shown in Figure 1. In principle, this initial partition can be generated by a *recursive dyadic partition* (RDP) as follows. First divide the domain into four sub-squares of equal size. Repeat this process again on each sub-square. Repeat this $1/2 \log_2 n = J$ times. This gives rise to a *complete* RDP of resolution $\frac{1}{\sqrt{n}}$ (the rectangular partition of the sensing domain shown in Figure 1). The RDP process can be represented with a quadtree structure. The quadtree can be pruned back to produce an RDP with non-uniform resolution.

Let \mathcal{P}_n denote the set of all RDPs, including the initial complete RDP and all possible prunings. For each RDP $P \in \mathcal{P}_n$, there is an associated quadtree structure (generally of non-uniform depth corresponding to the non-uniform resolution of most RDPs). The leafs of each quadtree represent dyadic (sidelength equal to a negative power of 2) square regions of the associated partition. For a given RDP and quadtree, each sensor node belongs to a certain dyadic square. We consider these squares ‘‘clusters’’ and assume that one of the nodes in each square serves as a ‘‘clusterhead,’’ which will assimilate information from the other nodes in the square. Notice that if one considers all RDPs in \mathcal{P}_n , then each sensor node actually belongs to a nested hierarchy of $1/2 \log_2 n$ dyadic squares of sidelengths $\frac{1}{\sqrt{n}}, \frac{2}{\sqrt{n}}, \frac{4}{\sqrt{n}}, \dots, 1$, respectively. Thus, we have a hierarchy of clusters and clusterheads.

Consider a certain RDP $P \in \mathcal{P}_n$. Define the estimator of the field as follows. On each square of the partition, we consider two possible estimators of the field in that region. One estimator is simply the least squares fitting of a constant on the corresponding square, in other words the mean of the measurements in the square. A discrete collection of such models is generated by quantizing the dynamic range $[-R, R]$ of the field to $n^{1/4}$ levels (in other words using $1/4 \log n$ bits to represent the mean value). This choice for

the quantization is due to the fact that ultimately the MSE of the estimator is $O(n^{-1/2})$ (as shown in the next section) and this choice puts the squared quantization error at that level. The second estimator is a least squares fitting of a piecewise linear model known as a *platelet* [15]. A platelet is formed by splitting a dyadic square into two wedge-shape regions (defined by a line connecting two points on the boundary of the square) and fitting a planar surface to the field on each wedge. Each platelet consists of six parameters, which describe the two surfaces, and a pair of points on the boundary of the square, which describe the wedge. A discrete collection of platelets are generated by discretizing the surface parameters to $1/3 \log n$ bits each and selecting the boundary points from a finite collection of points spaced evenly $\delta = n^{-2/3}$ units apart along the boundary of the square. Again, these choices of quantization are made so that the squared quantization error matches the overall MSE, which is shown to be $O(n^{-2/3})$ in the next section. The platelet estimator is determined by selecting the element of this collection that minimizes the sum of squared errors between it and the data. In summary, the two types of estimators employed at each leaf of the RDP (denoted by P) are:

- Estimator 1, denoted by $\hat{f}^h(P)$: least squares fit of a constant to the data in each square.
- Estimator 2, denoted by $\hat{f}^p(P)$: least squares fit of a platelet to the data in each square.

Estimator 1 is very similar to a Haar wavelet estimator, and therefore we will call it the “Haar estimator.” Estimator 2 will be called the “platelet estimator.” The remaining issue is the selection of the pruned RDP. The complete (unpruned) RDP will, of course, provide a perfect fit to the data (i.e., $\hat{f}^h(P) = \hat{f}^p(P) = x$). This choice of partition is undesirable for two reasons. First, the data are noisy and fitting constants or platelets over larger squares will affect an averaging that reduces the noise. Second, the unprocessed data x will require the maximum amount of energy to transmit to the destination. It turns out that properly penalizing partitions based on the number of squares (or leafs in the corresponding tree) provides both a more accurate estimate of the field and a compressed representation of the data more suitable for transmission to the destination.

The *complexity penalized* estimator is defined as follows. Let $f^m(P)$ denote a model of the field (based either on constant ($m = h$) or platelet ($m = p$) models on each square of P). The empirical measure of performance is the sum-of-squared errors between $f^m(P)$ and the data $x = \{x_{i,j}\}$.

$$R(f^m(P), x) = \sum_{i,j=1}^{\sqrt{n}} (f_{i,j}^m(P) - x_{i,j})^2. \quad (2)$$

For a fixed partition, P , the choice of $f^m(P)$ that minimizes $R(f^m(P), x)$ is simply given by the least squares fits on each square, as discussed above. For $m = h$ or $m = p$, define the complexity penalized estimator

$$\hat{f}_n^m = \arg \min_{\substack{f^m(P) : \\ P \in \mathcal{P}_n}} R(f^m(P), x) + 2\sigma^2 g^m(n) |P|, \quad (3)$$

where $|P|$ denotes the number of (square or wedge-shaped) terminal regions in the partition P , and $g^m(n)$ is a certain monotonically increasing function of n that discourages unnecessarily high resolution partitions (appropriate choices of $g^m(n)$ will be discussed in the next section). It is well known that the optimization

in (3) can be solved using a bottom-up tree pruning algorithm [16, 2, 15]. This is possible because both the sum-of-squared errors and the penalty are additive functions, and therefore the squared error plus penalty cost can be separated into terms associated with each individual square of the candidate partition.

The hierarchy of clusterheads facilitates this process in the sensor network. First consider the case for the Haar estimator \hat{f}^h . At each level of the hierarchy, the clusterhead receives the means from the four clusterheads below it, and compares the total cost of these estimates (sum-of-squared errors plus penalties) with the total cost of the mean estimate on the larger square associated with the clusterhead. Implementation of platelet approximations is more challenging in a decentralized sensor network structure. This is due to the fact that the optimization involved requires that all data be passed from the bottom up to the root of the tree structure [15]. More specifically, all the data in each square are needed to select the best wedge split. This may require a prohibitive amount of communication in large networks. Therefore, we advocate the following approximate strategy. Rather than passing all data from finer scales to coarser scales, we will instead pass only the optimized platelet fits from below. For example, each parent node will receive the best platelet fits from its four children nodes, and select an optimal platelet fit for itself based on these, rather than on the complete data from the four children. In the simulation section, it is demonstrated that this approximate strategy works about as well as the optimal one. The key advantage of the approximate strategy is that now only a handful of summary statistics need to be passed up the tree structure, and the communication requirements for \hat{f}^p are on the same order as those for \hat{f}^h .

The energy required for the above process is the sum of two terms: (1) the in-network communication required to obtain the estimate at the final clusterhead; (2) the out-of-network cost of transmitting the estimate to the destination. Note that the in-network communication only requires the transmission of data and/or sufficient statistics over short distances within the sensor network, while out-of-network communication is typically over a much greater distance and may dominate the overall cost of communication. The energy requirements will be discussed in detail in the next section.

3 Theoretical Performance

In the introduction, we described a fundamental trade-off between accuracy (MSE) and energy (\mathcal{E}) of the form

$$\text{MSE} \sim \frac{1}{\mathcal{E}^{2\nu}} .$$

In this section, we demonstrate that the hierarchical methods proposed above nearly achieve this optimal trade-off for the cases of $\nu = 1/2$ and $\nu = 2/3$, which correspond to two different levels of assumed smoothness in the field and boundaries. The theory hinges on an application of our extension [5] of the Li-Barron bound for complexity regularized model selection [6] to upper bound the MSE of the estimators and to bound the expected energy consumed by communications to compute the estimates. Since our methods (nearly) achieve the optimal trade-offs, no other schemes can be devised that will (asymptotically) perform significantly better.

Assume that the two-dimensional field denoted by f is composed of smooth regions of Hölder- α regularity in \mathbb{R}^2 , $\alpha \in (0, 2]$, with smooth boundaries that are Hölder- β regular in \mathbb{R}^1 , $\beta \in (0, 2]$. We will focus

on two classes:

- $C1$: Hölder-1 regular smooth regions separated by Hölder-1 regular boundaries.
- $C2$: Hölder-2 regular smooth regions separated by Hölder-2 regular boundaries.

The class $C2$ is a subset of $C1$ and places stronger assumptions on the smoothness than class $C1$, and therefore one expects better performance (in terms of MSE and energy consumption) is possible in class $C2$, which turns out to be true. Stronger smoothness assumptions are, of course, possible; say Hölder- α , $\alpha > 2$. However, the algorithms we proposed above are not capable of taking advantage of such additional assumptions; polynomial approximations to the surfaces and boundaries would be necessary for smoother functions, and we do not pursue such possibilities in this paper. We will show that our algorithms are nearly optimal for classes $C1$ and $C2$, achieving the trade-offs mentioned in the paragraph above. If one were to consider even smoother conditions, then the best trade-off for which one could aim is

$$\text{MSE} \sim \frac{1}{\mathcal{E}^2} ,$$

since the MSE cannot decay at a rate faster than $O(n^{-1})$, the parametric rate, and the communication energy must be at least $O(\sqrt{n})$ bit-meters, as discussed in the introduction.

3.1 MSE Performance

The performance of the proposed estimators can be studied in terms of (minimax) lower bounds and upper bounds on the MSEs. The minimax error is the error incurred by the best possible (Haar or platelet) estimator on the hardest possible field in class $C1$ or $C2$. The mean square error is defined as

$$\text{MSE}(\hat{f}^m) = \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\hat{f}^m(i, j) - f(i, j) \right)^2 \right] . \quad (4)$$

First consider the $C1$ class. Taking $g^h(n) \approx (1/4) \log n$, let \tilde{f} denote any estimator based on the data x . Then

$$\begin{aligned} & O(n^{-1/2}) \\ & \leq \inf_{\tilde{f}} \sup_{f' \in C1} \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\tilde{f}(i, j) - f'(i, j) \right)^2 \right] \\ & \leq \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\hat{f}^h(i, j) - f(i, j) \right)^2 \right] \\ & \leq O \left(\left(\frac{\log n}{n} \right)^{1/2} \right) . \end{aligned} \quad (5)$$

Similarly, for class $C2$ take $g^p(n) \approx (4/3) \log n$. Then

$$\begin{aligned}
& O(n^{-2/3}) \\
& \leq \inf_{\tilde{f}} \sup_{f' \in C2} \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\tilde{f}(i,j) - f'(i,j) \right)^2 \right] \\
& \leq \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\hat{f}^p(i,j) - f(i,j) \right)^2 \right] \\
& \leq O \left(\left(\frac{\log n}{n} \right)^{2/3} \right). \tag{6}
\end{aligned}$$

The minimax lower bounds follow from fairly standard ‘‘hypercubical’’ type arguments that re-cast the field estimation problem into a simpler series of independent hypothesis tests, for which the error rates are easy to determine [1, 2].

The upper bounds are based on a combination of approximation theory (from the field of wavelets and harmonic analysis) coupled with upper bounds on the variance of estimators based on such approximations. We begin by recalling a fundamental upper bound on expected error of complexity penalized estimators. This particular bound was originally developed for mixture density modeling [6], and we later extended it to more general settings [5]. Here we state a specialized version of the bound, tailored to the estimator proposed in (3).

Let \mathcal{F}_n^m denote the set of all possible models of the field. This set contains either Haar approximations ($m = h$) or platelet approximations ($m = p$), as described in Section 2. Recall that the parameters of the models are quantized, so the set \mathcal{F}_n^m consists of a finite number of models. Let \tilde{f} denote a generic element of \mathcal{F}_n^m , and recall that the true field is denoted by f and our estimators are denoted by \hat{f}^m . Assume that $g^m(n)$ satisfies the summability condition (Kraft inequality)

$$\sum_{\tilde{f} \in \mathcal{F}_n^m} e^{-g^m(n)|\tilde{f}|} \leq 1, \tag{7}$$

where $|\tilde{f}|$ denotes the number of (square or wedge-shaped) leaves in the partition associated with \tilde{f} . It is shown in the Appendix that $g^h(n) \approx 1/4 \log n$ and $g^p(n) \approx 4/3 \log n$ satisfy (7). Then, we can apply Theorem 7 in [5] to upper bound the MSE of the estimator \hat{f}_n^m as follows:

$$\begin{aligned}
& \text{MSE}(\hat{f}^m) \\
& = \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\hat{f}^m(i,j) - f(i,j) \right)^2 \right] \\
& \leq \min_{\tilde{f} \in \mathcal{F}_n^m} \left\{ \frac{2}{n} \sum_{i,j=1}^{\sqrt{n}} \left(\tilde{f}_{i,j} - f_{i,j} \right)^2 + \frac{8\sigma^2 g^m(n)}{n} |\tilde{f}| \right\} \tag{8}
\end{aligned}$$

The upper bound involves two terms. The first term, $2 \sum_{i,j=1}^{\sqrt{n}} \left(\tilde{f}_{i,j} - f_{i,j} \right)^2$, is a bound on the bias or approximation error. The second term, $8\sigma^2 g^m(n) |\tilde{f}|$, is a bound on the variance or estimation error. The

bias term, which measures the squared error between the best possible model in \mathcal{F}_n^m and the true field, can be easily bounded. Under the assumption that the underlying field belongs to class $C1$ or $C2$, well known results from approximation theory [2, 3, 15] can be applied. In particular, if $|\tilde{f}| = M$ (corresponding to a partition with M squares or wedges), then for class $C1$, using either the Haar or platelet estimator, the bias term is upper bounded by $O(M^{-1})$ and for class $C2$, using the platelet estimator, it is upper bounded by $O(M^{-2})$. Then by equating the bias and variance terms and solving for M in terms of n leads to the MSE upper bounds stated above. Also note that the Haar estimator has the same MSE bound for both classes $C1$ and $C2$ because it cannot exploit the additional field smoothness in $C2$. Further details are given in the Appendix.

3.2 Expected Communication Requirements

Hierarchical networks of sensors, such as those described above, allow field estimates to be calculated and transmitted with a significant reduction in energy. In this section we will analyze the energy consumption associated with the field estimation methods described above. We assume that a certain received energy level, E_0 , is required for reliable communication, and that the transmitted energy decays like $1/r$, where r is the distance from the transmitter. Therefore, the energy required to reliably send 1 bit over r meters grows linearly in r . Other models for energy decay can be considered and incorporated into our framework [4], but we will not pursue this further here.

Without the hierarchical structure of the Haar and platelet estimation methods, one might naively transmit each sensor's measurements to a remote processor at a cost of $\mathcal{E} \approx kn$ bit-meters, assuming that a constant number of bits, k , is transmitted per measurement. On the other hand, we argued in the introduction that $\mathcal{E} \geq O(\sqrt{n})$. Here we show that the Haar and platelet estimators can be constructed and communicated with $O(\sqrt{n})$ bit-meters of expected communication.

The communication requirements can be split into two components. The “in-network” cost of computing the optimal estimator, and the “out-of-network” cost of transmitting the estimate to a remote destination. The out-of-network cost is proportional to the number of leafs in the partition associated with the optimal estimator. The size of the partition associated with the optimal estimator can be deduced from the MSE bound as follows. Suppose the optimal partition consists of M (square or wedge-shaped) regions. Then the variance is at least $O(M/n)$ since we have $O(M)$ free parameters to estimate ($O(1)$ parameters per region). Recall that the MSE is less than or equal to $O\left(\left(\frac{\log n}{n}\right)^\nu\right)$, $\nu = 1/2$ for the Haar estimator, and $\nu = 2/3$ for the platelet estimator. Since the variance must be less than the MSE, it follows that $E[M] \leq O(n^{1-\nu} \log^\nu n)$. That is, for the Haar estimator $E[M] \leq O((n \log n)^{1/2})$ and for the platelet estimator $E[M] \leq O((n \log^2 n)^{1/3})$.

In the case of the Haar estimator, \hat{f}^h , it turns out that the expected in-network communication costs are also proportional to the size of the partition associated with the optimal estimator. This is easy to see as follows. At the first stage of the hierarchy, $O(n)$ values must be transmitted over a distance of $O(n^{-1/2})$ meters (each group of four nodes send their measurements to the nearest cluster head, $O(n^{-1/2})$ meters away), requiring $O(\sqrt{n})$ bit-meters of energy. At each subsequent stage in the hierarchy, we expect at most $O((n \log n)^{1/2})$ squares will remain. Therefore, ignoring the log factor, $O(\sqrt{n})$ values must be transmitted

to clusterheads which are $O(2^j/\sqrt{n})$, $j = 1, \dots, 1/2 \log n$, meters away. Thus, in total the energy required for computing \hat{f}^h is $O(\sqrt{n})$ bit-meters.

The expected in-network costs for the platelet estimator, \hat{f}^p , are higher. As pointed out in Section 2, the optimization involved requires that all data be passed from the bottom up to the root of the tree structure [15]. This leads to an in-network energy requirement of $O(n)$ bit-meters. However, the approximate strategy, based on passing only the optimized platelet fits up through the hierarchy, has the same energy requirements as the Haar estimator. In this case, the first stage requires $O(n^{1/2})$ bit-meters and subsequent stages involve passing $O(\sqrt{n})$ parameter values to clusterheads.

In conclusion, the expected in-network energy cost of $O(n^{1/2})$ and out-of-network energy cost of less than $O(n^{1/2})$ bit-meters combined with the upper and lower bounds on the MSEs produce the stated trade-offs at the beginning of this section.

4 Simulation Experiments

The theoretical properties of the proposed multiscale estimators are supported in simulation experiments. We first simulated observations of a sample field for $n = 16^2$, 32^2 , and 256^2 , and compared the performance of the Haar and platelet estimators. The noise variance is $\sigma^2 = 1$, and all estimates were computed using the theoretical penalties developed in Section A. These results were compared with the approximate platelet estimator, which as discussed above has a communication cost comparable to the Haar estimator. Figure 2 displays the Haar estimates for the three different network sizes. The partitions displayed in the third column show that the size of the partition scales with the square root of the number of sensors, as predicted. The estimates in Figure 3 show that the platelet method results in smoother estimates with more sharply defined edges on significantly small partitions. Again, the partition sizes grow with the cube root of the size of the network, as predicted by the theory above. The partition images also demonstrate the impact of the wedge fits on accurate boundary estimation. Compare these partitions with those generated by the Haar analysis. Because of the fine partition near the boundaries in the Haar estimate, estimates of the values of the field in those regions will have a large variance. The wedge-shaped blocks, in contrast, are generally larger, and so estimates of those regions are made using more sensors, thus lowering the estimator variance. Figure 4 shows that approximate platelet estimator, calculated using only the estimates from the next finer scale instead of all the observations, vary only slightly from the optimal platelet estimator in Figure 3. Thus, the platelet approximations should be an effective tool even when in-network communication costs prohibit extensive data passing.

Finally, Figure 5 shows Haar, platelet, and approximate platelet estimates for a network with $n = 64^2$ sensors. The estimates in Column 1 were calculated from data with a noise variance of 1, while the estimates in Column 2 were calculated from data with a noise variance of 10. As predicted, weighting the penalization by the variance of the noise produces good estimates at different SNRs.

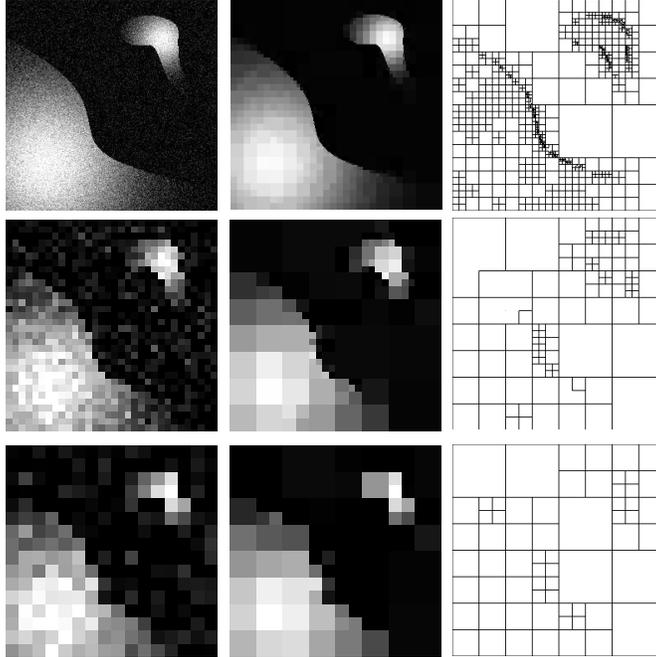


Figure 2: Effect of sensor network density (resolution) on field estimation with the Haar estimator. Column 1 is the noisy set of measurements, Column 2 is the estimated field, and Column 3 is the associated partition. In Row 1, 65536 sensors are used for form a partition with 889 regions, resulting in an MSE of $9.85e - 2$. In Row 2, 1024 sensors are used to form a partition with 127 regions, resulting in an MSE of $4.86e - 1$. In Row 3, 256 sensors are used to form a partition with 64 regions, resulting in an MSE of $7.26e - 1$.

5 Conclusions

In this work, we have proposed a method for boundary estimation in sensor networks. The boundary estimate is determined via complexity regularization of a hierarchical tree-based estimation method. We demonstrated theoretically that our method nearly achieves the optimal trade-off $\text{MSE} \sim 1/\mathcal{E}^{2\nu}$, which shows that no other scheme can be devised that will (asymptotically) perform significantly better. Simulation experiments agreed very well with the theoretical predictions.

A Estimator penalization

We penalize the platelet estimates according to the length of a prefix code which can uniquely describe each possible platelet estimate (*i.e.*, codes which satisfy the Kraft inequality). These codelengths are proportional to the size the partition associated with each model, and thus penalization leads to estimates that favor smaller partitions. In particular,

$$g^p(n) \equiv 2/3 \log_e 2 + 4/3 \log_e n. \quad (9)$$

To see how this is derived, consider constructing a unique code for every $\tilde{f} \in \mathcal{F}_n^p$. If \tilde{f} has k leaves, then k can be decomposed into $3m + 1$ square-shaped leaves plus some number of wedge-shaped leaves, where m

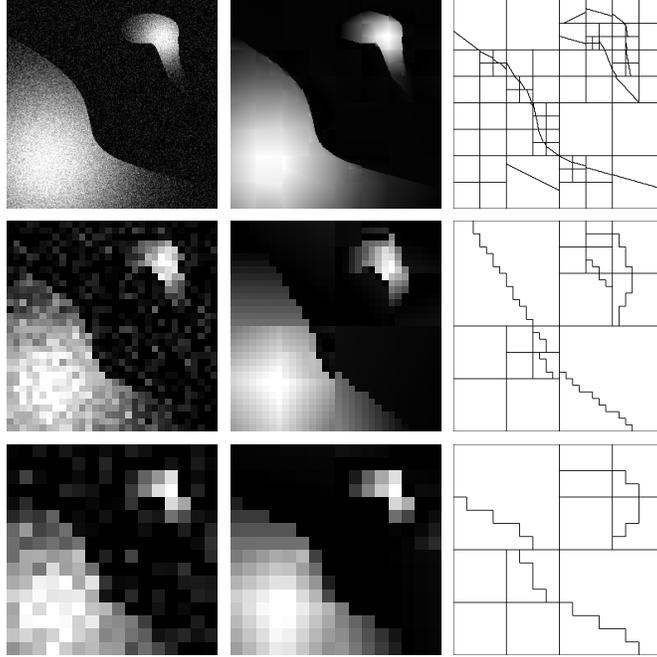


Figure 3: *Effect of sensor network density (resolution) on field estimation with platelets. Column 1 is the noisy set of measurements, Column 2 is the estimated boundary, and Column 3 is the associated partition. In Row 1, 65536 sensors are used for form a partition with 105 regions, resulting in an MSE of $3.64e - 2$. In Row 2, 1024 sensors are used to form a partition with 25 regions, resulting in an MSE of $3.33e - 1$. In Row 3, 256 sensors are used to form a partition with 16 regions, resulting in an MSE of $3.19e - 1$.*

is the number of dyadic splits in the partition and the number of wedge-shaped leafs is bounded above by $3m + 1$. The total number of k leafs, then, is bounded above by $2(3m + 1)$. The quadtree representing the RDP of the data can be encoded by representing each of the $4m + 1 < 2k/3$ nodes with a 0 for an internal node and a 1 for a leaf node. This can easily be verified with an inductive argument. Each of these dyadic squares can then be split into two wedge-shaped blocks. We can allow as few as $n^{2/3}$ possible wedgelet splits per dyadic block without affecting the error decay rate. The splits can then be encoded if we allot $1/3 \log_2 n$ bits to each wedge-shaped block. Finally, the three coefficients of a planar fit on one of the wedge-shaped blocks can be uniformly quantized to one of $n^{1/3}$ levels without changing the error decay rate, and so all three coefficients can be encoded with a total of $\log_2 n$ bits. Thus the necessary codelength for an estimate with k leafs in the partition requires $k(2/3 + 4/3 \log_2 n)$ bits.

From the Kraft inequality, we know that the existence of this uniquely decodable scheme guarantees that

$$\sum_{\tilde{f} \in \mathcal{F}_n^p} 2^{-k(2/3 + 4/3 \log_2 n)} \leq 1.$$

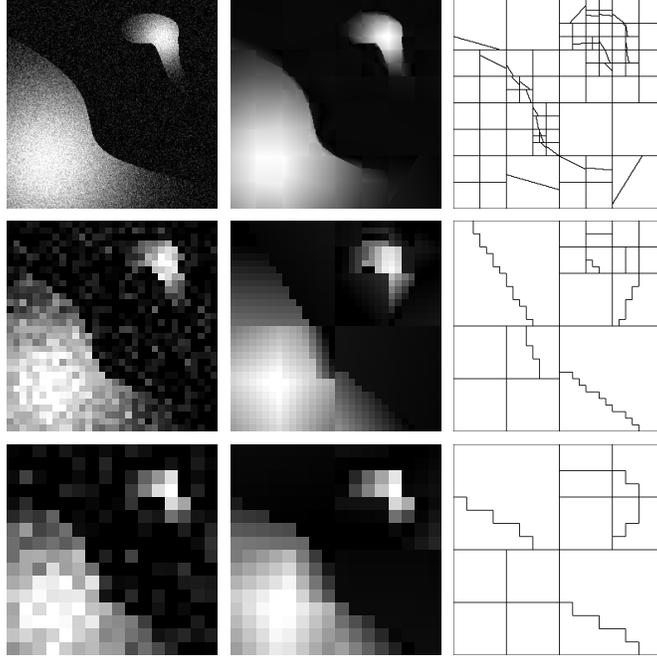


Figure 4: *Effect of sensor network density (resolution) on field estimation with platelets without passing all the data from fine to coarse levels. Column 1 is the noisy set of measurements, Column 2 is the estimated boundary, and Column 3 is the associated partition. In Row 1, 65536 sensors are used for form a partition with 107 regions, resulting in an MSE of $7.84e - 2$. In Row 2, 1024 sensors are used to form a partition with 23 regions, resulting in an MSE of $4.70e - 1$. In Row 3, 256 sensors are used to form a partition with 15 regions, resulting in an MSE of $3.71e - 1$.*

Therefore, if $g^p(n) = k(2/3 \log_e 2 + 4/3 \log_e n)$, then

$$\begin{aligned}
 \sum_{\tilde{f} \in \mathcal{F}_n^p} e^{-g^p(n)} &= \sum_{\tilde{f} \in \mathcal{F}_n^p} e^{-\log_e(2)k(2/3+4/3 \log_2 n)} \\
 &= \sum_{\tilde{f} \in \mathcal{F}_n^p} 2^{-k(2/3+4/3 \log_2 n)} \\
 &\leq 1,
 \end{aligned}$$

as desired.

By a similar argument, Haar estimates should be penalized as

$$g^h(n) \equiv 4/3 \log_e 2 + 1/4 \log_e n. \quad (10)$$

The analysis follows as in the platelet case, except that (1) the wedgelet locations do not need to be encoded, (2) only one Haar coefficient is needed in place of the three platelet coefficients, and (3) $n^{1/4}$ quantization levels are required to guarantee near-optimal error decay rates. This means that the necessary codelength for an estimate with k leafs in the partition requires $k(4/3 + 1/4 \log_2 n)$ bits, which results in the above penalization.

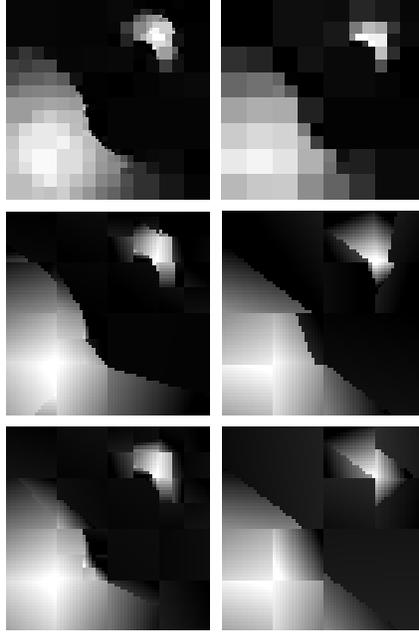


Figure 5: *Effect of noise variance on Haar and platelet estimates. Estimates in the first column were made using noisy data with variance 1. Estimates in the second column were made using noisy data with variance 10. In Row 1, Haar estimates with 4096 sensors resulted in MSEs of $2.78e - 1$ (for $\sigma^2 = 1$) and $1.03e + 0$ (for $\sigma^2 = 10$). In Row 2, platelet estimates with 4096 sensors resulted in MSEs of $1.73e - 1$ (for $\sigma^2 = 1$) and $5.18e - 1$ (for $\sigma^2 = 10$). In Row 3, approximate platelet estimates with 4096 sensors resulted in MSEs of $2.34e - 1$ (for $\sigma^2 = 1$) and $9.75e - 1$ (for $\sigma^2 = 10$).*

B Rate of MSE Decay

First consider the case for the Haar estimator, \hat{f}^h . Consider a complete RDP with M^2 squares of sidelength $1/M$. It is known that if the boundary is a Lipschitz function, or more generally has a box counting dimension of 1 (the case for both classes $C1$ and $C2$), then the boundary passes through $\ell \leq CM$ of the squares, for some constant $C > 0$ [2, 15]. Furthermore, there exists a pruned RDP with at most $C'M$ leaves, where $C' = 8(C + 2)$, that includes the above ℓ squares of sidelength $1/M$ that contain the boundary [2, 15].

Now consider the upper bound (8), which as stated earlier follows as from an application of Theorem 7 in [5].

$$\begin{aligned}
& \frac{1}{n} \sum_{i,j=1}^{\sqrt{n}} E \left[\left(\hat{f}_{i,j}^h - f_{i,j} \right)^2 \right] \\
& \leq \min_{\tilde{f} \in \mathcal{F}_n^h} \frac{1}{n} \left\{ 2 \sum_{i,j=1}^{\sqrt{n}} \left(\tilde{f}_{i,j} - f_{i,j} \right)^2 + 8\sigma^2 g^h(n) |\tilde{f}| \right\} \\
& \leq \min_{\tilde{f} \in \mathcal{F}_n^h} 2 \int_{[0,1]^2} (\tilde{f} - f)^2 + 8\sigma^2 \frac{1}{4} \frac{\log n}{n} C' M
\end{aligned}$$

where the discretized squared error is bounded by the corresponding continuous counterpart. The squared

error $\int_{[0,1]^2} (\tilde{f} - f)^2 \leq \frac{K_1}{M} + \frac{K_2}{\sqrt{n}}$, where the first term is due to the error between the $1/M$ resolution partition along the boundary, and the $1/\sqrt{n}$ term is due to the quantization error overall. Thus, the MSE behaves like

$$\text{MSE} \leq O(1/M) + O(1/\sqrt{n}) + O\left(M \frac{\log n}{n}\right).$$

Taking $M = \sqrt{\frac{n}{\log n}}$ produces the desired result: $\text{MSE} \leq O(\sqrt{\log n/n})$.

The platelet case is handled in an analogous manner. The only significant change is that the squared error $\int_{[0,1]^2} (\tilde{f} - f)^2 \leq \frac{K_1}{M^2} + \frac{K_2}{n^{2/3}}$, where the first term is due to the error between the plate approximation in the smooth regions and the wedge approximation to the boundary, and the $n^{-2/3}$ term is due to the quantization error overall. Taking $M = \sqrt[3]{\frac{n}{\log n}}$ produces the desired result: $\text{MSE} \leq O((\log n/n)^{2/3})$.

References

- [1] A. P. Korostelev and A. B. Tsybakov. *Minimax theory of image reconstruction*. Springer-Verlag, New York, 1993.
- [2] D. Donoho. Wedgelets: Nearly minimax estimation of edges. *Ann. Statist.*, 27:859 – 897, 1999.
- [3] E. J. Candes and D. L. Donoho. Curvelets - a surprisingly effective nonadaptive representation for objects with edges. In *Curves and Surfaces*, 1999. L.L. Schumaker et al. (eds).
- [4] R. Willett, A. Martin, and R. Nowak. Backcasting: A new approach to energy conservation in sensor networks. submitted to the third international symposium of Information Processing in Sensor Networks (IPSN'04), April 2004.
- [5] E. Kolaczyk and R. Nowak. Multiscale likelihood analysis and complexity penalized estimation. *Annals of Statistics* (to appear). Also available at www.ece.wisc.edu/~nowak/pubs.html, 2002.
- [6] Q. Li and A. Barron. Mixture density estimation. In S.A. Solla, T.K. Leen, and K.-R. Müller, editors, *Advances in Neural Information Processing Systems 12*. MIT Press, 2000.
- [7] R. Nowak and U. Mitra. Boundary estimation in sensor networks: Theory and methods. In *Information Processing in Sensor Networks*, pages 80–95, Springer, 2003.
- [8] A. Scaglione and S. Servetto. On the interdependence of routing and data compression in multi-hop sensor networks. In *Mobicom*, 2002.
- [9] D. Marco, E. Duarte-Melo, M. Liu, and D. Neuhoff. On the many-to-one transport capacity of dense wireless sensor networks and the compressibility of its data. In *Information Processing in Sensor Networks*, pages 1–16, Springer, 2003.
- [10] S. D. Servetto. Sensing lena—massively distributed compression of sensor images. In *Proceedings IEEE ICIP'03*, September 2003. to appear.

- [11] S.S. Pradhan and K. Ramchandran. Distributed source coding using syndromes (discus). *IEEE Transactions on Information Theory*, 49(3), 2003.
- [12] K. Chintalapudi and R. Govindan. Localized Edge Detection in Sensor Fields. *Ad-hoc Networks Journal*, 2003. to appear.
- [13] Divya Devaguptapu and Bhaskar Krishnamachari. Applications of localized image processing techniques in wireless sensor networks. In *Proceedings SPIE Aerosense '03*, April 2003.
- [14] Deepak Ganesan, Ben Greenstein, Denis Perelyubskiy, Deborah Estrin, and John Heidemann. An evaluation of multi-resolution search and storage in resource-constrained sensor networks. In *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*., 2003. to appear.
- [15] R. Willett and R. Nowak. Platelets: a multiscale approach for recovering edges and surfaces in photon-limited medical imaging. *IEEE Transactions on Medical Imaging*, 22(3), 2003.
- [16] L. Breiman, J. Friedman, R. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, Belmont, CA, 1983.