Power-Efficient Sensor Placement and Transmission Structure for Data Gathering under Distortion Constraints

Deepak Ganesan [†] Răzvan Cristescu [††‡] Baltasar Beferull-Lozano [††‡]

[†] Department of Computer Science, University of California at Los Angeles, Los Angeles, CA 90095

[††‡] Audio-Visual Communications Laboratory (LCAV), Swiss Federal Institute of Technology (EPFL), Lausanne CH-1015, Switzerland

deepak@lecs.cs.ucla.edu, {Razvan.Cristescu, Baltasar.Beferull}@epfl.ch

ABSTRACT

We consider the joint optimization of sensor placement and transmission structure for data gathering, where a given number of nodes need to be placed in a field such that the sensed data can be reconstructed at a sink within specified distortion bounds while minimizing the energy consumed for communication. We assume that the nodes use joint entropy coding based on explicit communication between sensor nodes, and consider both maximum and average distortion bounds. The optimization is complex since it involves an interplay between the spaces of possible transmission structures given radio reachability limitations, and feasible placements satisfying distortion bounds.

We address this problem by first looking at the simplified problem of optimal placement in the one-dimensional case. An analytical solution is derived for the case when there is a simple aggregation scheme, and numerical results are provided for the cases when joint entropy encoding is used. We use the insight from our 1-D analysis to extend our results to the 2-D case, and show that our algorithm for two-dimensional placement and transmission structure provides significant power benefit over a commonly used combination of uniformly random placement and shortest path trees.

1. INTRODUCTION

Wireless sensor networks are often envisaged to comprise thousands of nodes accomplishing a sensing task. Yet, the realities of economies of scale in manufacturing and the high cost of many sensors themselves mean that these nodes are currently significantly more expensive than predicted. Therefore, typical deployed networks (e.g. habitat monitoring [1]) comprise a few hundred of nodes, each with cost of a few hundreds of dollars. While we await a future with ubiquitous cheap sensor nodes, a problem that is both immediate and necessary is to accomplish the required tasks with a limited number of resource-constrained sensor nodes.

We consider the problem of deploying a finite number of sensor nodes in a geographic area, and choosing a communication structure among the nodes on the corresponding network. A single sink is responsible for gathering the sensor data, for storage or control purposes. Since sensor nodes have limited battery power, an important goal is to minimize the total power consumption of data gathering, while keeping the sensing distortion within specified bounds. An important characteristic of typical sensor networks, that can be exploited for reducing the power consumption, is that the data measured at nodes is correlated.

Several algorithms have been proposed for energy efficient data gathering [2, 3, 4]. However, these works do not take into consideration the correlation in the data. Recent studies on the joint rate allocation and transmission structure optimization for sensor networks with correlated data can be found in [5, 6]. In particular, the result of Cristescu et al [6] is similar in that it considers optimal tree structures for data gathering that exploit correlation in the measured data. Our work adds new constraints to this problem by allowing node placement to be varied, thus introducing tough distortion constraints. This results in a more complex problem that both requires a different approach to solve and produces novel results. Node placement for optimal coverage is a well-studied (and difficult) problem (e.g. [7]). The work of [8] considers the problem of energy-efficient topology aware placement. That work does not exploit the correlation present in the data measured; also, the placement constraints considered in that work are rather event driven than related to the distortion of measurement.

A commonly used method for deploying sensor networks is the uniform random placement, since such a deployment is often the easiest and cheapest. However, we believe that there are compelling reasons for understanding the interactions between the node placement and the data and transmission structures, and the effect of these interactions on the efficiency of utilization. First, studying the impact of placements lets us understand if other easy-to-deploy configurations of sensor nodes can give us important gains in power consumption. As we show in this paper, this is likely to be the case. Second, controlled placement will be necessary for applications which have to deploy limited numbers of expensive nodes (such as seismic nodes which need high precision) and hence the location of sensors has to be optimized.

In this work, we are particularly interested in the relation between

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN'04, April 26–27, 2004, Berkeley, California, USA.

Copyright 2004 ACM 1-58113-846-6/04/0004 ...$5.00.

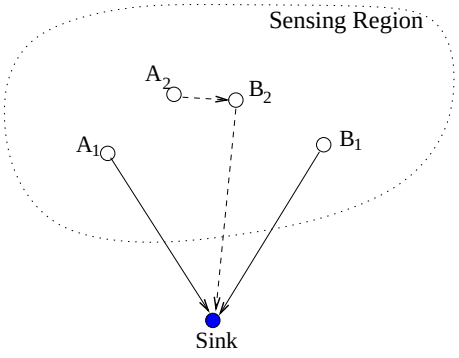

Figure 1: Two possible placement and structure configurations of nodes A and B that satisfy specified sampling distortion bounds are shown. A_1 and B_1 transmit their data over shorter cumulative distances than A_2 and B_2 respectively, but B_2 can code its data with that of A_2 to exploit their high data correlation as a result of their proximity. Thus, determining the more power-efficient configuration among many possible ones that satisfy distortion bounds is difficult due to the interplay of placement and transmission structure.

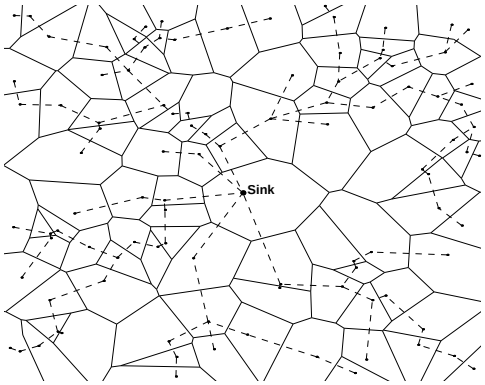

Figure 2: The Voronoi cells (solid lines) represent the distortion in each cell. The tree structure (dotted lines) represents a possible transmission structure.

the data reconstruction distortion that results from the node placement, and the power requirements of data gathering from the sensors. An instance of the problem is shown in Fig. 2 where the Voronoi cells represent the distortion, and the dotted lines describe a possible tree structure to be used for data gathering.

More specifically, N nodes need to be deployed over a finite geographic region $\mathcal{A}$, whose two-dimensional area is A. Each of these nodes takes samples from a three dimensional random field $X(u, v, t)$, where (u, v) is the spatial location and t is the time-axis. Each sensor transmits periodically its sensed data, through multi-hop routing, to a sink located at the center of $\mathcal{A}$. The sink can reconstruct the sensor field in the region within specified maximum and average distortion bounds (D_{max} and D_{avg} respectively). We assume that sampling in the time domain is sufficiently high (above Nyquist frequency) for the sink to fully reconstruct the time-axis. Thus, we only need to consider the distortion in reconstructing snapshots along the spatial axes, $X(u, v)$.

To enable data gathering, the sensor nodes build a routing tree rooted at the sink, and transmit data along this tree. Note that there are situations where tree structures are not optimal, but for the sake of simplicity we will limit our study to data gathering trees. The data gathering procedure is periodic and originates at the leaves, proceeding iteratively towards the sink through multihop forwarding. At each iteration, a junction node receives data from its children, decodes the received data, jointly codes the decoded data with its own data, and forwards the encoded data to its parent on the tree.

We use a simple but relevant energy related cost function to study the interplay between placement (which determines locations of nodes) and transmission structure (which determines how sources on the tree are connected). Namely, the total cost of data gathering over the tree structure described above can be written as

$$\sum_{i=1}^{N} Rate(i) \times CommunicationCost(i) \qquad (1)$$

where $Rate(i)$ is the total amount of data transmitted by node i, and $CommunicationCost(i)$ is the per-bit transmission cost from node i to its parent on the tree. Since the data at nodes is correlated, the rate at node i, $Rate(i)$ depends on the particular set of sources that are in the sub-tree rooted at i, and on their locations. Similarly, the transmission cost per-bit, $CommunicationCost(i)$, depends on the identity of the particular node used as next hop to sink, and on its position. Thus, the total power consumption depends on both the placement and the transmission structure.

An illustration of this interplay is seen in Fig. 1, where even in a simple example with two nodes, it is not easy to determine the most power-efficient configuration. In terms of transmission distance, configuration 1 is better since A_1 and B_1 transmit each over smaller distance than A_2 and B_2, respectively. However, in configuration 2, A_2 and B_2 are closer to each other than nodes A_1 and B_1, hence they are likely to be stronger correlated. This correlation can be exploited by properly choosing the transmission structure ($A_2 \rightarrow B_2 \rightarrow sink$). This can significantly reduce the amount of data that node B_2 needs to generate, by using side information from node A in coding its data, thus reducing the total power consumption. In general, if the data is correlated, the shortest path tree (SPT) is not necessarily optimal as transmission structure [6].

The optimal solution involves searching through the spaces of all possible configurations that satisfy distortion bounds, and all possible transmission structures that are feasible given radio reachability limitations. While this combined optimization has not been considered in prior work, it is known that even a subset of our problem is NP-hard [5, 6] (namely, the transmission structure optimization for a given placement).

The rest of this paper is structured as follows. In Section 2, we formulate precisely our problem, and describe the sensing, communication, aggregation and data reconstruction models. In Section 3, we consider the one-dimensional variant of the problem, provide analytical solutions for a simplified aggregation model and analyze the joint encoding case numerically in detail. In Section 4, we extend the solutions from the one-dimensional case to the two-dimensional case and show that it out-performs typical random placement approaches. Finally, in Section 5, we discuss various other interesting performance aspects of our solution, for instance, even though we optimize only for total power, our solution provides large benefits in terms of per-node energy utilization.

To the best of our knowledge, this is the first study of the interaction between node placement under distortion constraints, transmission structure optimization, and rate allocation in the context of sensor

networks that measure correlated data.

2. PROBLEM FORMULATION

We assume that we are given N nodes, that need to be placed in a two dimensional region $\mathcal{A}$, of area A. For simplicity, we assume that the region is circular, and has a radius L. The placement of nodes, $P = \{(x_i, y_i) \in \mathcal{A}, 1 \leq i \leq N\}$ is constrained by two distortion metrics, the maximum distortion D_{max}, defined as the maximum acceptable distortion at any point in $\mathcal{A}$, and the average distortion D_{avg}, defined as the distortion per unit area over $\mathcal{A}$.

2.1 Sensing Model

A frequently used sensing model is the Gaussian random field [9, 6]. This model is suitable for analysis and can provide the essential intuition to solve the problem in practice. We assume that the field is a continuous-space two dimensional stationary random field $X(u, v)$, where u and v represent the geographic co-ordinates of points in the region $\mathcal{A}$. Without loss of generality, we assume that the random field has zero mean, that is, $\mu_X = E[X(u, v)] = 0 \quad \forall\, u, v$. We make simple assumptions about the nature of the random field, and make no assumptions on whether this field is band-limited or not. Let $\mathbf{R}_X(\tau_x, \tau_y)$ denote the covariance function associated to the random field $X(u, v)$. The correlation between two points (u_i, v_i) (position of node i), and (u_j, v_j) (position of node j) is given by $r_X(i, j) = \mathbf{R}_X(u_j - u_i, v_j - v_i)$. The correlation functions are $\mathbf{R}_X(\tau_x, \tau_y) = e^{-a\sqrt{\tau_x^2+\tau_y^2}}$, or in terms of Euclidean distance, $\mathbf{R}_X(d) = e^{-ad}$. Thus, $r_X(i, j) = e^{-ad_{ij}}$, where d_{ij} is the Euclidean distance between nodes i and j. Such a model is typical of spatially varying data and is used widely in spatial statistics [10].

The power consumption incurred depends on two factors: the distance from a transmitting node to its parent, and the aggregate amount of data transmitted over that distance.

2.2 Communication Model

For the transmission power, we use a standard transmission model that assumes that the power per bit for transmission over a wireless link is a function of the distance between the transmitter and receiver. We assume that there is an underlying transmission scheduling protocol (such as SMAC [11]) that schedules transmissions over a tree to avoids collisions. For the scope of this paper, we will ignore the protocol overhead resulting from creating the schedules, assuming that this overhead is small in comparison with the data size.

If the distance between the transmitter and receiver is d, then the power is $E \propto d^{\kappa}$ where κ is called the path-loss exponent (typically $2 \leq \kappa \leq 4$ [12].[1] In addition, each node has a maximum power at which it can transmit, which places a limit on the maximum transmission range, $d^{\kappa} \leq E_{max}$. Our model is simplistic to keep the optimization manageable for this paper. In practice, two additional factors need to be considered: (a) radios often have non-isotropic propagation, and, (b) radios adjust power levels in discrete steps rather than at arbitrarily fine granularity.

The radio communication constraint overlaps with the maximum distortion constraint that we will discuss shortly. Both these constraints limit the maximum separation between nodes in the network. Therefore, we do not consider this constraint explicitly in the rest of this paper. We will consider the impact of communication bounds in greater depth in an extended report.

[1] In this work, we do not consider reception overhead, which would increase the cost of communication per-hop, and can only improve our results.

2.3 Aggregation Model

For our aggregation model, we assume that each node performs complete joint entropy coding of all the data coming from its corresponding sub-tree. We assume that each node quantizes its samples with an independent quantizer and all quantizers use the same quantization interval. In order to express the amount of data in bits resulting from first quantizing independently at each node and then performing joint entropy at a given node of all that data in its corresponding sub-tree, we will use the differential entropy as described in [13]. The differential entropy of a k-dimensional multivariate normal distribution $\mathcal{N}_k(\mu_X, \mathbf{R}_X)$ is:

$$h(\mathcal{N}_k(\mu_X, \mathbf{R}_X)) = \frac{1}{2}\log(2\pi e)^k \det(\mathbf{R})$$

We approximate the joint discrete entropy associated to the quantized samples by assuming a high-resolution uniform scalar quantization with step-size Δ for all the nodes. Thus, as $\Delta \to 0$, the distortion at each node $\to \Delta^2/12$ ([13]). Also, due to the fine quantization, we assume that noise is uncorrelated with the signal.

$$H(\mathcal{N}_k(\mu_X, \mathbf{R}_X)) \approx h(\mathcal{N}_k(\mu_X, \mathbf{R}_X)) - k\log\Delta$$

For any node i let T_i represent the set of nodes in the sub-tree rooted at node i. Node i receives encoded data from its children, first decodes it, and then jointly compresses it together with its own data (quantized samples). The total data rate sent from node i is approximately:

$$H(\mathcal{N}_k(0, \mathbf{R}[T_i])) = \frac{1}{2}\log(2\pi e)^{|T_i|}\det(\mathbf{R}[T_i]) - |T_i|\log\Delta \quad (2)$$

where $\mathbf{R}[T_i]$ denotes the covariance matrix associated with the nodes in the sub-tree T_i, and $|T_i|$ represents the number of nodes of the sub-tree.

2.4 Data Reconstruction Model

We now describe the reconstruction procedure used by the sink to reconstruct the entire continuous-space sensor field given the encoded data from a discrete set of sample points at the positions of the N nodes. The sink periodically receives quantized values $\hat{X}(u_1, v_1), \hat{X}(u_2, v_2)..\hat{X}(u_N, v_N)$ from the N sensing nodes placed at points $(u_1, v_1), (u_2, v_2)..(u_N, v_N)$ respectively. In general, given these N quantized values, an interpolation procedure will result in a reconstruction, $\hat{X}(u, v)$, that gives the samples at any location (u, v) in the region $\mathcal{A}$. In this work, we use a nearest-neighbor reconstruction procedure, which, although very simple, helps us understand the complex interactions in our problem and focus on the power minimization issue. In future work, we plan to improve our results using better interpolation models.[2]

Let V_i be the Voronoi cell corresponding to the sensor node i located at position (u_i, v_i) (which is the centroid of V_i). Then:

$$\begin{aligned}\hat{X}(u, v) &= \hat{X}(u_i, v_i) \text{ iff } ||(u, v) - (u_i, v_i)|| \\ &\leq ||(u, v) - (u_j, v_j)||, (\forall) j \neq i\end{aligned}$$

Given that the sink uses a nearest neighbor reconstruction procedure, we formulate the coverage and distortion constraints. The

[2] Note that optimal reconstruction is difficult, because of the different issues of aliasing, non-uniform quantization, etc.

first constraint is a coverage constraint, which ensures that the set of Voronoi cells covers the region $\mathcal{A}$:

$$\bigcup_{i=1,..,N} V_i = A \tag{3}$$

To evaluate the maximum distortion, we use the fact that for an isotropic radially decreasing correlation model, the maximum distortion points are located along the boundaries of the Voronoi cells. By definition, all Voronoi cells are convex because of the property of minimum distance decoding, therefore, the furthest points in each Voronoi cell are the corners of the cell. The maximum distortion constraint controls the distance from these furthest points in each cell to the centroid of the corresponding Voronoi cell. Thus, for any point (u, v) in region $\mathcal{A}$, the distortion of reconstruction when it is assigned the same value as the nearest sampled point (u_i, v_i) is:

$$\begin{aligned} Distortion(u,v) &= MSE(u,v) \\ &= E[(\hat{X}(u_i,v_i) - X(u,v))^2] \leq D_{max} \end{aligned}$$

with the i sensor being closest to (u, v). Note that the error is computed between the quantized version of the closest sample given by $\hat{X}(u, v)$ which is received at the sink, and the actual unquantized random variable, $X(u, v)$.
Using the correlation model in Section 2.1, we obtain the distortion between the unquantized random variables at the two points:

$$\begin{aligned} MSE(u,v) &= E[(\hat{X}(u,v) - X(u,v))^2] \\ &= E[(\hat{X}(u_i,v_i) - X(u,v))^2] \\ &= E[(X(u_i,v_i) + n_Q(u_i,v_i))^2] + E[X(u,v)^2] \\ &\quad -2E[(X(u_i,v_i) + n_Q(u_i,v_i))X(u,v)] \\ &= E[X(u_i,v_i)^2] + E[X(u,v)^2] \\ &\quad -2E[X(u_i,v_i)X(u,v)] \\ &= \sigma^2 + \sigma^2 - 2\sigma^2 e^{-ad_{ij}} \\ &= 2\sigma^2(1 - e^{-ad_{ij}}) \leq D_{max} \end{aligned} \tag{4}$$

where n_Q is the quantization noise between the quantized random variable and the original random variable, which is assumed to be small due to fine quantization. As expected, the above equation shows that the mean square error (MSE) is a concave and monotonically increasing function of the distance between the location (u, v) and the closest sample point. Therefore, the maximum distortion bounds the distance between any point in the region $\mathcal{A}$, and the nearest sample point. Thus, for the particular model that we consider, the maximum allowed distance R_{max}, is:

$$R_{max} = -\frac{1}{a}\log(1 - \frac{D_{max}}{2\sigma^2}) \tag{5}$$

The average distortion constraint, defined as the mean square error in data reconstruction over the entire region $MSE(\mathcal{A})$, can be computed by integrating $MSE(u, v)$ over $\mathcal{A}$.

$$\begin{aligned} AvgDistortion(\mathcal{A}) &= MSE(\mathcal{A}) \\ &= \frac{1}{A}\int_{\mathcal{A}} MSE(u,v)dudv \leq D_{avg} \end{aligned} \tag{6}$$

where $MSE(u, v)$ is calculated as shown in (4). Then, for N sensors, each with corresponding Voronoi cell V_i,

$$\begin{aligned} MSE(\mathcal{A}) &= \frac{1}{A}\sum_{i=1}^{N} MSE(V_i) \\ &= \frac{1}{A}\sum_{i=1}^{N}\int_{V_i} MSE(u,v)dudv \\ &= \frac{1}{A}\sum_{i=1}^{N}\int_{V_i} E[(\hat{X}(u,v) - X(u,v))]dudv \\ &\leq D_{avg} \end{aligned} \tag{7}$$

2.5 Objective

We state now formally our objective to minimize the total power cost, given the constraints described so far. Namely, our problem is to find a placement P of nodes, $|P| = N$, and a tree ST rooted in the sink, that spans the nodes in P, such that to

$$\text{Minimize (1)} \quad \text{under constraints (3), (4), (6).} \tag{8}$$

3. OPTIMAL PLACEMENT AND STRUCTURE IN THE ONE-DIMENSIONAL CASE

Although our final goal is to solve the problem in the two-dimensional case, the one-dimensional case is significantly more tractable since, as we will show shortly, the transmission structure optimization is trivial, and it is possible to understand the placement problem in isolation.
We adapt the problem statement in Section 2 for the one-dimensional instance. In this case, let $X(s)$, $0 \leq s \leq L$, represent the measured random field along a one-dimensional line of length L, as shown in Fig. 3. Since all transmission terminates at the sink, instances of the problem where the sink is between nodes (not at the corner of a line) can be split into two independent optimizations, one each for nodes on either side of the sink with the sink at the corner. Further, due to symmetry, it does not matter which end of the line the sink is placed. Therefore, the N nodes $1, 2, .., N$ are placed in sequence along the line with the sink at the left corner such that node 1 is closest to the sink and node N is the furthest. We denote the distance between node i and node $i - 1$ by r_i; $\{r_i\}_{i=1}^{N}$ represents the set of unknowns in the one-dimensional optimization.

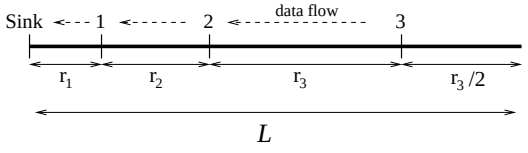

Figure 3: One-dimensional node placement.

3.1 One-Dimensional Constraints

The coverage and distortion constraints can be readily formulated similar to (3), (4) and (6). The Voronoi cells for the one-dimensional problem are the mid-points between adjacent pairs of sample points (nodes) as shown in Fig. 4. Without loss of generality, we use a boundary extension of the Voronoi cell for the last node, r_N. Thus, the coverage constraint can be rewritten as:

$$\text{Coverage Constraint: } \sum_{i=1}^{N} r_i + r_N/2 = L. \tag{9}$$

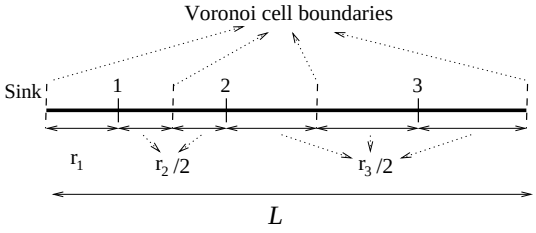

Figure 4: Voronoi cells and Maximum distortion distances for linear placement.

As described in (4), the maximum distortion constrain restricts the maximum distance from any node to the edge of the nearest neighbor cell. Hence,

$$r_1 \leq R_{max} \qquad \frac{r_i}{2} \leq R_{max}\ ; \quad (\forall) \quad 2 \leq i \leq N \tag{10}$$

Similar to (7), the average distortion (defined as the distortion per unit length) is given by:

$$\begin{aligned} MSE(L) &= \frac{1}{L}\int_L MSE(u)du \\ &= \frac{1}{L}\sum_1^N \int_{V_i} E\left[(\hat{X}(u) - X(u))\right] du \\ &= \frac{1}{L}\left[\int_0^{r_1} 2\sigma^2(1-e^{-ad}) + 2\sum_2^N \int_0^{\frac{r_i}{2}} 2\sigma^2(1-e^{-ad}) \right. \\ &\qquad \left. + \int_0^{\frac{r_n}{2}} 2\sigma^2(1-e^{-ad})\right] \\ &= 2 - \frac{2}{La}\left(2N - e^{-ar_1} - 2\sum_2^N e^{-ar_i/2} - e^{-ar_N/2}\right) \end{aligned} \tag{11}$$

We analyze the joint placement-structure optimization in two steps. First, we show that the optimal structure is simple shortest path routing. We then proceed to optimizing the placement for different choices of aggregation functions.

3.2 Optimal Structure is Shortest Path

Proposition 1: In a one-dimensional sensor network where there is a single sink and joint entropy coding is used at each hop, shortest path communication is optimal in terms of minimizing total energy. We only provide a brief outline of the proof of this proposition due to lack of space. If node i transmits its data to a node j where $j > i$ (i.e. j is further from the sink than i), the data from j must be eventually routed through i to minimize power consumption. This results from the fact that power per-bit increases super-linearly with distance (since κ is between 2 and 4), hence it is always better to multi-hop through as many intermediate hops as available [14]. Thus, to minimize power consumption, the aggregate data from j must be routed through i. On the other hand, joint entropy coding is a monotonically increasing function of the number of sources, hence, if node i transmits its data to a node j where $j > i$, the jointly coded data at j is larger than the amount of data when i did not transmit to j, and this consumes more power to transmit to the sink.

3.3 Optimizing Placement for 1-d transmission

Given that the transmission structure is shortest path forwarding from nodes towards the sink, the placement problem can be reformulated as:

$$\begin{aligned} \{r_i\}_{i=1}^N &= \arg\min_{r_i} \sum_{i=1}^N H(X_i, X_{i+1}, .., X_N) r_i^\kappa \\ &= \arg\min_{r_i} \sum_{i=1}^N \left[\frac{1}{2}\log(2\pi e)^{(N-i+1)} \det(\mathbf{R}[T_i]) \right. \\ &\qquad \left. -(N-i+1)\log\Delta\right] r_i^\kappa \end{aligned} \tag{12}$$

under coverage (9) and distortion constraints (10),(11). In the above equation, T_i (the sub-tree rooted in i) is the set of nodes $i, i+1, .., N$ as in (2).
It is hard to solve (12) analytically for an arbitrary correlation structure. We thus first obtain a closed-form solution for a simplified scenario, where we assume zero correlation between data sampled at different nodes. In this case, $\mathbf{R}[T_i]$ is diagonal $(\forall) i \in \{1..N\}$.

3.3.1 Analytical Solution for Independent Data at Nodes

In this case, the optimization in (12) reduces to minimizing $\sum_1^N \left[(N-i+1)(\frac{1}{2}\log(2\pi e) - \log\Delta)\right] r_i^\kappa$ since $\det(\mathbf{R}[T_i])$ is unity when nodes are uncorrelated and the variance of the random process is one. Dropping the constant scaling factor that does not impact the minimization, we get:

$$\{r_i\}_{i=1}^N = \arg\min_{r_i} \sum_{i=1}^N (N-i+1) r_i^\kappa \tag{13}$$

As an example, suppose in Fig. 3 the samples at nodes 1, 2 and 3 are uncorrelated, and each of the nodes has one unit of data to transmit to the sink. In this case, node 3 transmits one unit of data, node 2 transmits two units (its own unit + one forwarded unit), and node 1 transmits three units. Let us now see the impact of each constraint in the above optimization.
First, we consider the optimization in (13) when only the coverage constraint (9) is active. Using a Lagrangian multiplier, we obtain

$$\{r_i\}_{i=1}^N = \arg\min_{r_i} \sum_{i=1}^N (N-i+1) r_i^\kappa - \lambda\left[(\sum_{i=1}^N r_i) + r_N/2\right]$$

We solve the above Lagrangian optimization using partial derivatives:

$$\begin{aligned} r_i &= \left[\frac{\lambda}{\kappa(N-i+1)}\right]^{-(\kappa-1)}, \quad (\forall)\ 2 \leq i \leq N; \\ r_N &= \left[\frac{3\lambda}{2\kappa}\right]^{-(\kappa-1)} \\ where\ \lambda &= \left[\frac{L}{\sum_1^{N-1}(\frac{1}{\kappa(N-i+1)})^{-(\kappa-1)} + \frac{3}{2}(\frac{3}{2\kappa})^{-(\kappa-1)}}\right]^{\kappa-1} \end{aligned}$$

Fig. 5 shows the optimal $\{r_i\}_{i=1}^N$ in a placement with $N = 15$ nodes over a line of length $L = 200$, for the case of quadratic and cubic path-loss exponents, $\kappa = 2$ and 3. As expected, nodes

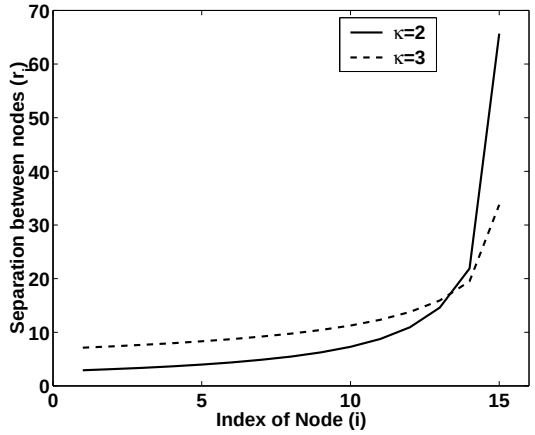

Figure 5: Placement for different values of pathloss exponent (κ). As κ increases, placements become more uniform, since communication power dominates differences in aggregated data size.

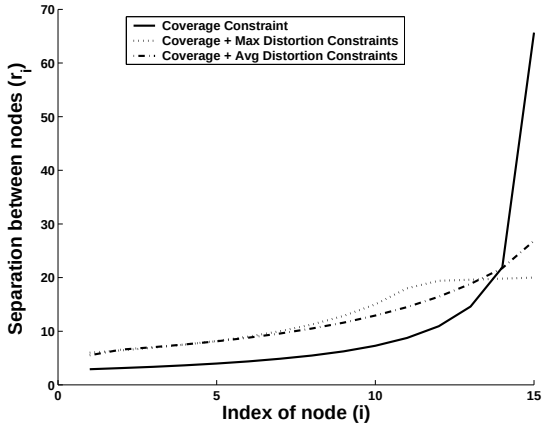

Figure 6: Impact of distortion constraints. The R_{max} constraint places a ceiling on the maximum r_i whereas the D_{avg} constraint equalizes r_is to reduce the average distortion.

further from the sink transmit smaller amounts of data over longer distance than closer ones, which need to transmit larger amounts of aggregated data. However, the impact of increasing data load for nodes closer to the sink is balanced by the effect of the path-loss exponent κ, since communication power increases super-linearly with distance. Thus, the optimal choice reflects a balance between these two opposing factors. As the path-loss exponent is increased from $\kappa = 2$ to $\kappa = 3$, the communication overhead dominates, hence, spacing between nodes becomes more uniform

The maximum and average distortion bounds impact on placement in significantly different ways as shown in Fig. 6. The maximum distortion constraint places a ceiling on the maximum separation between nodes, whereas the average distortion constraint reduces the mean error by making cells more equally sized.

3.3.2 Performance of Joint Coding Case

For the analysis of the correlated data case, we use MATLAB optimization toolbox for numerically finding the optimal solution of (12). We use the model in 2.1 with parameter $a = 0.001$ (high correlation) in this study. For this model, the placement of nodes for the joint coding case has essentially the same behavior as that for the zero correlation case. This behavior results from the continuous and slowly decaying nature of our correlation model.

We evaluate the performance of our optimal one-dimensional placement by comparing its power consumption to that of a commonly considered uniform placement. We show the relative performance for two values (a relatively small one and a relatively large one, given our model) for each variable, as shown in Table 1. For a reasonable choice of distortion and network parameters, we see that even for low number of nodes we get a factor of 1.2 to 2 benefit over the uniform spacing case. For a larger network, these gains increase. Among other parameters, increasing R_{max} increases the flexibility in placement, hence it makes possible to further minimize power consumption. Increasing L for fixed R_{max} has the opposite effect, since the feasible placement region reduces. Increasing the number of nodes N (for fixed R_{max}), increases both the correlation between nodes (hence the aggregation benefit), and reduces the average per-hop distance for multihop transmission, thus, power consumption reduces.

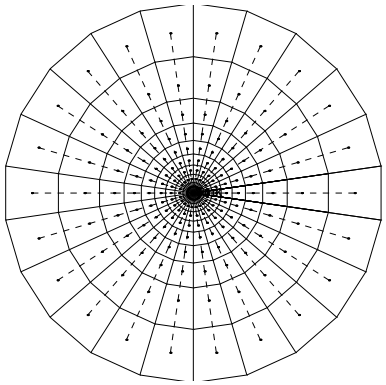

Figure 7: The Voronoi cells (solid) and transmission structure (dashed) for a wheel placement, for data-gathering at a sink located at the center of the circular region $\mathcal{A}$.

4. TWO-DIMENSIONAL PLACEMENT AND STRUCTURE

While the one-dimensional problem instance can be well-understood since node placement optimization separates from transmission structure optimization, the two-dimensional case is significantly more complex. Even for a fixed location of nodes, the problem of optimizing the transmission structure for power efficient data gathering is NP-complete. If the aggregation function at nodes is concave and dependent on the number of nodes that relay via that node, then the optimization problem includes the Steiner tree problem; moreover, the problem is NP-complete also when the aggregation function is known [5, 15]. If the aggregation function at a node depends also on the transmission structure among the nodes that relay via that nodes, then, even for very simple settings, the problem remains NP-complete [6]. We expect that the arbitrary positioning of nodes and the distortion constraints can only increase the complexity of the joint placement and transmission optimization problem, thus we conjecture that our problem is NP-complete too. Our goal in this section is, therefore, to provide good approximation algorithms based on intuition from the one-dimensional solution.

Pathloss Exp (κ)	R_{max}		N		L		D_{avg}	
	Small (10)	Large (20)	Small (10)	Large (20)	Small(100)	Large(300)	Small (0.05)	Large(0.15)
$\kappa = 2$	1.3	1.8	1.3	1.8	1.9	1.3	1.5	1.6
$\kappa = 3$	1.3	1.5	1.2	1.6	1.5	1.3	1.4	1.5

Table 1: Gain of optimal placement over uniform placement for different settings.

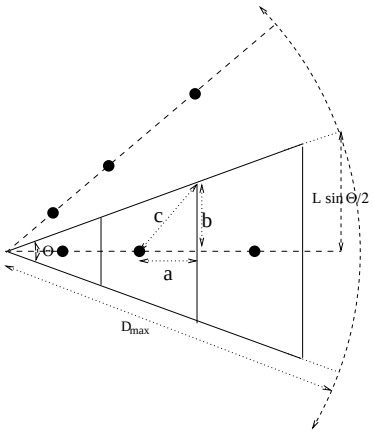

Figure 8: Deriving the R_{max} constraint in 2d placement. For any of the Voronoi cells, the furthest points are the corners of the cell.

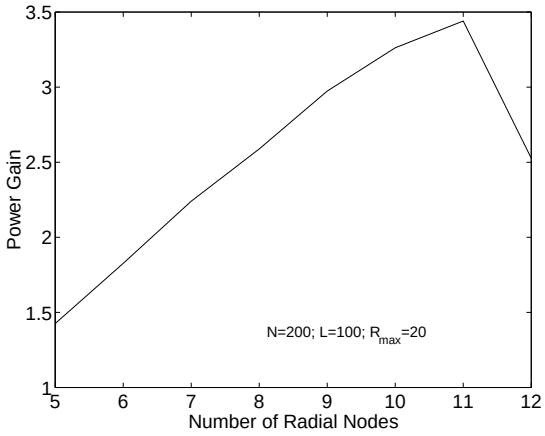

Figure 9: As the number of nodes per-spoke (n_{radial}) is increased, gains initially improve but eventually $\bar{R}_{max}$ is too constrained, hence gains reduce.

4.1 Placement Strategy

Our two-dimensional placement strategy replicates the linear placement along a wheel structure as shown in Fig. 7. The wheel comprises n_{spoke} spokes where each spoke has n_{radial} nodes placed along it. Each node transmits data using shortest path forwarding along the spoke on which it is placed. Note that the shortest path might not be always optimal for explicit communication coding [6], however we restrict to such simple gathering trees, which can be constructed distributedly in polynomial time. We study in more detail the placement problem; the study of alternative transmission structures is subject of further work.

Fig. 7 shows both the transmission structure and the Voronoi cells for such a placement. Besides being analytically tractable, the wheel structure captures the essential behavior that we would want from an efficient two-dimensional placement and transmission structure. The network is dense closer to the sink where the data load is higher, and sparse further away from the sink.

While the two-dimensional placement is simple once we have decided that the placement and transmission structure is along a wheel, many questions remain to be solved. How do we place the n_{radial} nodes along each spoke such that the distortion bounds are not violated over the entire two-dimensional area, $\mathcal{A}$? Given N nodes, how many nodes, n_{radial}, do we place along each spoke and how many spokes, n_{spoke}, do we place angularly over the wheel? A lookahead into the results in Fig. 9 suggests that performance gains are not only sensitive to the choices of n_{spoke} and n_{radial}, but there is a non-obvious choice of these parameters that provides maximum power benefit.

Note that, even if we propose a deterministic placement strategy, our approach provides meaningful insight into the design of random radial distributions for nodes placement in an arbitrary area in practical scenarios.

4.1.1 Maximum distortion bound

We will translate the maximum distortion constraint for the two-dimensional case, to a one-dimensional bound that we can solve using the technique described in Section 3. Consider the Voronoi cells for a single spoke as shown in Fig. 8. As discussed in Section 2.4, the maximum distortion corresponds to the distortion at the point that is furthest from its nearest sampled point. By definition, such a point should lie on the Delaunay triangulation of the sample points. Since each Voronoi cell is convex, this point lies on one of the corners of the cell. Due to radial symmetry, the Voronoi cells for nodes on each spoke are identical, hence, it suffices to consider the maximum distortion bound on cells corresponding to any one of the spokes.

Consider the triangle with sides a, b and c that is formed between any sampling point and one of the corners of its Voronoi cell (see Fig. 8). The maximum distortion constraint is satisfied if $c \leq R_{max}$ (5). Since the radius of the region $\mathcal{A}$ is L, the angle between any two spokes is $\theta = \left(\frac{2\pi}{n_{spoke}}\right)$, and $b \leq L \sin\left(\frac{\theta}{2}\right)$, then it follows that a sufficient condition for $c \leq R_{max}$ is $a \leq \sqrt{R_{max}^2 - L^2 \sin^2\left(\frac{\theta}{2}\right)}$. Thus, in order to have the two-dimensional maximum distortion distance of R_{max} satisfied, it suffices to place nodes along each spoke such that the one-dimensional distortion distance along each line is bounded by:

$$\bar{R}_{max} = \sqrt{R_{max}^2 - L^2 \sin^2\left(\frac{\pi}{n_{spoke}}\right)} \qquad (14)$$

Due to the place limitations, we do not consider the 2D average distortion case from this paper; we will analyze it in a more extended publication shortly.

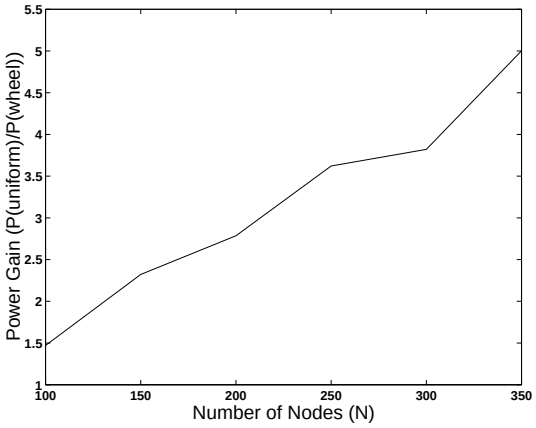

Figure 10: As the number of nodes placed in the network is increased, the gains increase.

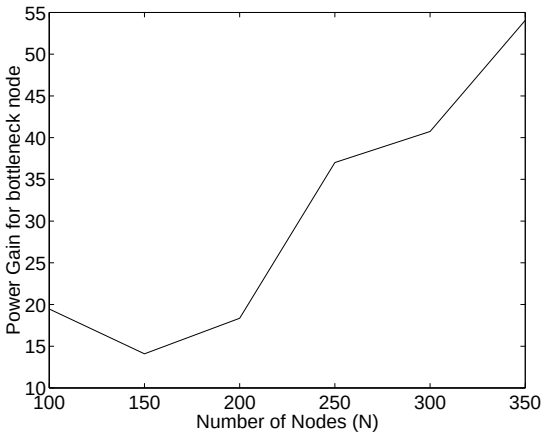

Figure 11: The power gains at the bottleneck node are very large, with 50 times improvement for large N.

4.2 Choosing the Number of Spokes and Nodes per Spoke

The tradeoff involved in finding the optimal choice of n_{spoke} and n_{radial} can be understood from (14). From the one-dimensional analysis, we know that separately increasing either r_{radial} or $\bar{R}_{max}$, and keeping the other constant, can reduce power consumption. However, in this case, the two parameters have opposite effects on each other. For instance, if r_{radial} is increased, n_{spoke} decreases and, therefore, so does $\bar{R}_{max}$ (14). This interaction is illustrated in Fig. 9.

How do we obtain a good choice of n_{radial} and n_{spoke}? While an exact solution for determining the optimal choice of wheel placement parameters is hard to find, an intuitive and effective approximation is to find the placement that maximizes $n_{spoke} \times \bar{R}_{max}$. Our approximation algorithm performs within 10% of the optimal (computed through exhaustive search) for configurations that we have tested.

Fig. 10 plots the improvement of an efficient wheel placement (as determined by the above metric) over the uniformly random placement in the circular area $\mathcal{A}$, where the transmission structure is a shortest path tree. When the number of nodes in the network, N is low, the performance gains are low, since the distortion bounds provide less flexibility to optimize placement. As N increases, these gains increase up to a factor of 5. Thus, not only does our wheel placement consistently out-perform a random placement with shortest path trees, it can potentially provide an order of magnitude improvement if sufficient flexibility is allowed in terms of number of nodes and distortion bounds.

5. CONCLUSIONS AND FURTHER WORK

To summarize, we have formulated an optimization problem that considers jointly node placement, transmission structure and data structure in a data gathering sensor network, in terms of an energy-related cost function. We have studied in detail the 1-D case, namely we provided a closed-form solution for the node placement when the data is independent, and outlined the methodology to solve numerically the case when aggregation at nodes results in joint entropy coding, for arbitrary correlation structures. We used our insights from the 1-D setting to propose an approximation algorithm that places nodes in a radial "wheel" structure in the 2-D case. We show that significant power gains can be obtained with such a node placement scheme over commonly used uniformly random placements.

We have dealt with only a subset of issues raised by the above formulation in this paper. We now briefly mention some extensions to our study, and some insights that demonstrate other substantial benefits of our placement.

Optimizing bottleneck energy consumption: Our problem formulation optimizes for the *total* power consumption in data-gathering, but in a practical scenario, metrics such as network lifetime are likely to be as important. A commonly used metric for network lifetime is the time at which the first node dies, in other words, what is the power consumption of the bottleneck node. For instance, in a typical sensor network, the bottleneck node is the one that is closest to the sink, since it forwards a large amount of traffic. Fig. 11 shows that the power consumption of the bottleneck node in the optimal placement is two orders of magnitude lower than in the uniform random placement. Thus, huge gains in network lifetime can be expected from our optimized placement. In future work, we seek to incorporate this metric explicitly into our optimization.

Extension to Slepian-Wolf coding: While our work considers the explicit communication case, where nodes have to receive data to be able to code jointly, Slepian-Wolf type of coding [13] enables coding of data at nodes without explicit communication, assuming that knowledge about the joint statistics are available at each node. The Slepian-Wolf instance of our problem can be easily solved by applying a result from [6], which shows that the optimal transmission structure is always an SPT, and the optimal rate allocation can be found by solving an LP optimization problem.

Robustness to Node Failures and Inaccuracies in Correlation Model: Our work makes some simplifying assumptions regarding node stability, radio connectivity and the availability and stability of a correlation model. In practice, all these parameters are subject to large variability due to environmental factors. For instance, it is well-known that nodes fail and radios are unpredictable. Thus, a shortest path-based radial transmission structure may not always be the most efficient due to packet-loss or failure of intermediate hops. A routing strategy such as Directed Diffusion ([4]) will re-route data through a different spoke, changing the aggregation benefit, a factor that we have not considered in this work. In addition, our work optimizes for a well-known correlation model, but a desirable feature in practice may be easy deployment under varying condi-

tions. In future work, we intend to evaluate the robustness of our design to these variables.

6. REFERENCES

[1] M. Hamilton. James San Jacinto Mountains Reserve.

[2] S. Lindsey, C.S Raghavendra, and K. Sivalingam. Data gathering in sensor networks using the energy*delay metric. In *IPDPS*, 2001.

[3] W.R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocols for wireless microsensor networks. In *HICSS*, January 2000.

[4] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *ACM/IEEE Mobicom*, Boston, MA, 2000.

[5] A. Goel and D. Estrin. Simultaneous optimization for concave costs: single sink aggregation or single source buy-at-bulk. In *ACM-SIAM SODA*, 2003.

[6] R. Cristescu, B. Beferull-Lozano, and M. Vetterli. On network correlated data gathering. In *INFOCOM*, 2004. to appear.

[7] S. Eidenbenz. Approximation algorithms for terrain guarding. *IPL*, 2002.

[8] K. Dasgupta, M. Kukreja, and K. Kalpakis. Topology–aware placement and role assignment for energy-efficient information gathering in sensor networks. In *IEEE ISCC*, 2003.

[9] D. Marco, E.J. Duarte-Melo, M. Liu, and D.L. Neuhoff. On the many-to-one transport capacity of a dense wireless sensor network and the compressibility of its data. In *IPSN*, 2003.

[10] N.A.C. Cressie. *Spatial Statistics*. John Wiley and Sons, 1991.

[11] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *IEEE Infocom*, June 2002.

[12] T. S. Rappaport. *Wireless Communications, Principles and Practice*. Prentice Hall, 1996.

[13] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 1991.

[14] G.J. Pottie and W.J. Kaiser. Wireless integrated network sensors. *Communications of the ACM*, 43(5):51–58, May 2000.

[15] A. Meyerson, K. Munagala, and S. Plotkin. Cost-distance: Two-metric network design. In *IEEE FOCS*, 2000.