

An Energy Conservation Method For Wireless Sensor Networks Employing a Blue Noise Spatial Sampling Technique

Mark Perillo, Zeljko Ignjatovic, and Wendi Heinzelman
Department of Electrical and Computer Engineering
University of Rochester
Rochester, NY 14627 USA
{perillo, ignjatov, wheinzel}@ece.rochester.edu

ABSTRACT

In this work, we present a method for the selection of a subset of nodes in a wireless sensor network whose application is to reconstruct the image of a (spatially) bandlimited physical value (e.g., temperature). The selection method creates a sampling pattern based on blue noise masking and guarantees a near minimal number of activated sensors for a given signal-to-noise ratio. The selection method is further enhanced to guarantee that the sensor nodes with the least residual energy are the primary candidates for deselection, while enabling a tradeoff between sensor selection optimality and balanced load distribution. Simulation results show the effectiveness of these selection methods in improving signal-to-noise ratio and reducing the necessary number of active sensors compared with simpler selection approaches.

Categories and Subject Descriptors

C.2.1 [Computer-Communications Networks]:
Network Architecture and Design—*wireless communication*

General Terms

Algorithms, Management

Keywords

wireless sensor networks, blue noise, sensor management

1. INTRODUCTION

In this work, we consider applications for wireless sensor networks where a spatially bandlimited physical phenomenon (e.g., temperature, pressure) is to be sensed and reconstructed at a central base station or local cluster head. We assume that the sensors are deployed in a manner such that their locations in the field are random (e.g., dropped

from an airplane) and that they are deployed with a higher density than what the application's accuracy requirements demand. There are several reasons why this may occur. First, signal-to-noise ratio (SNR) requirements can only be guaranteed within a reasonable likelihood if the average sensor density is high enough. Also, sensor networks may be purposely deployed with high density to achieve longer network lifetime, if it assumed that sensor activity is rotated. Finally, some applications may require resolution that is adaptive over time, necessitating a dense initial deployment of the sensors.

In any of these situations, it is sufficient to activate only a subset of the sensors, as the activation of all sensors in densely populated regions may provide minimal additional accuracy. Furthermore, sensor nodes are characterized by very low energy budgets, so the energy burden imposed by activating certain sensors may be much more significant than the improvement in cumulative sensor network data that they provide. It is important to efficiently manage the sensors so that network lifetimes of months or even years are realizable. While power consumption in the network is affected by many factors, in general, a reduction in the number of active sensors reduces overall power consumption. The sensor management problem addressed in this work is essentially to determine which sensors to deactivate so that the quality of data is sufficient enough to meet application requirements while achieving the energy efficiency that is critical in wireless sensor networks.

In this paper, we propose an algorithm for determining which sensors within the more densely covered subregions should be selected to acquire data from the environment and which nodes should remain inactive in order to conserve energy. The proposed method is especially suitable for applications where it is desirable to trade spatial resolution for sensor network longevity. Our proposed algorithm chooses sensor subsets such that the sensor positions can be mapped into the blue noise binary patterns that are used in many image processing applications [9]. The algorithm guarantees that subsets are chosen such that each subset provides a near optimal SNR for a given initial sensor distribution and a desired number of active sensors. We also enhance the algorithm so that the sensor nodes with the least residual energy are the primary candidates for deselection and present a way in which the algorithm can be distributed for practical implementation. Finally, we propose a method in

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.

which multiple subsets are chosen according to a blue noise sampling method and the scheduling of these subsets is optimized through a linear program.

The rest of this paper is organized as follows. Section 2 addresses related work. Section 3 provides a brief background on the concepts of blue noise and stochastic sampling. Section 4 presents our proposed algorithms. Section 5 provides simulation results along with analysis. Section 6 concludes the paper and suggests future work in this area.

2. RELATED WORK

The problem of node selection in wireless sensor networks has been explored to some extent and several protocols have been proposed as solutions. Some algorithms essentially produce a randomly chosen sensor subset [6]. Other protocols have been proposed to select subsets of sensors for use in coverage-preserving applications such as distributed detection [3, 16, 18, 19, 20]. The patterns of the chosen sensors resulting from these algorithms are much more evenly spread out than random subset selection. However, most do not consider load balancing in the selection of active sensors.

An alternative approach to achieving energy efficiency in overpopulated sensor networks is to scale back the amount of traffic generated in areas of high sensor density. In [5], the authors propose an approach where sensors can adjust their resolution according to the density in their neighborhood. The resolution of the individual sensors is chosen so that the accuracy of the reconstructed image of the phenomenon is approximately equal throughout the monitored region. Other distributed compression solutions include [1] and [13]. These approaches still require all sensors to remain on, which we assume to be costly despite the reduction of traffic generated and transmitted.

3. BLUE NOISE MODEL / STOCHASTIC SAMPLING BACKGROUND

A blue noise pattern is a statistical model for describing stochastic patterns with very little frequency content below a blue noise principal frequency f_{BN} [17]. A binary blue noise mask/pattern $s_{BN}(x) \rightarrow S_{BN}(f)$ is a special case of a blue noise pattern that consists of similarly sized impulses distributed in a homogeneous manner and maintains a stochastic nature (i.e., uniform distribution of the impulses is prohibited). By distributing impulses in such a way, the resulting spectral content of the pattern is composed almost entirely of high frequency content. A binary blue noise pattern may be described by the following equations.

$$s_{BN}(x) = \sum_{x_i} \delta(x - x_i) \quad (1)$$

$$\int_{0+}^{f_{BN}} S_{BN}(f) df \ll \int_{f_{BN}}^{\infty} S_{BN}(f) df \quad (2)$$

Several algorithms have been proposed to generate binary blue noise patterns [7, 8, 9, 10, 11]. The method that is of interest to us is that which was proposed in [9] and improved in [11]. In their work, the authors propose a dart throwing method, mimicking a stochastic Poisson disc method. In this method, randomly chosen new points (impulses) are added to the point set if and only if no other points are located within a specified radius centered at the location of

the new points. A low pass spatial filter is then used to determine which points contribute the most low frequency content (the points with the highest value of the filtered sampling pattern). These points are subject to relocation within the regions with the least low frequency content. Our proposed sensor selection method uses a modification of this algorithm to create a binary blue noise pattern.

While the optimal sampling pattern for a bandlimited signal consists of samples at regular intervals (i.e., a grid pattern for images), other sampling patterns can be used with little performance dropoff. As long as the maximum spacing between sensors is small enough to meet the Nyquist sampling rate criteria, several algorithms can be used to perform ideal reconstruction, converging to a reconstructed signal with no error from the original [2, 14, 15]. In fact, when employing a sampling pattern based on a binary blue noise mask, a source signal can be reconstructed nearly as optimally as from a grid pattern. In other words, the number of sampling points necessary to perfectly reconstruct a bandlimited signal is only slightly higher when sampling with a blue noise noise pattern than when performing regular sampling at the Nyquist rate.

The maximum sample spacing imposed by the Nyquist constraint can be met using fewer sampling points when employing blue noise sampling rather than random (white noise) sampling. Even if the local sampling rate falls below the Nyquist rate in some subregions, a blue noise sampling pattern will achieve higher accuracy than a white noise sampling pattern with the same number of sampling points. To understand why this is the case, observe the stochastic sampling processes shown in Figure 1. The bandlimited signal in Figure 1(a) is sampled by the blue noise sampling pattern in Figure 1(b) as well as by the white noise sampling pattern in Figure 1(c). The spectrum of the blue noise sampled signal, shown in Figure 1(d) shows that there is very little aliasing into the signal band. This is due to the high frequency nature of the sampling pattern. However, the signal is corrupted much more by the white noise sampling, as shown in Figure 1(e), and it is apparent that more information about the original signal is lost. In reality, the reconstruction algorithms presented in [2, 14, 15] are nonlinear and not as simple as the application of a low pass filter. These algorithms can tolerate some noise leakage into the signal band, but clearly it is desirable to reduce the signal band noise as much as possible.

The blue noise sampling described here is easily extended to the two-dimensional scenario. Figure 2 illustrates a typical blue noise sampling pattern and a typical random (white noise) sampling pattern. As the figure shows, the sampling points in the blue noise pattern are more evenly spread out than those in the white noise pattern. For the same number of sampling points, a bandlimited image can be reconstructed more accurately from the blue noise sampled signal than from the randomly sampled signal.

4. BLUE NOISE SPATIAL SAMPLING IN WIRELESS SENSOR NETWORKS

In this section, we propose a method to decide which nodes in a wireless sensor network should be used to actively sense the environment and which should remain off in order to conserve energy. We consider the situation where a large number of sensor nodes are randomly deployed within

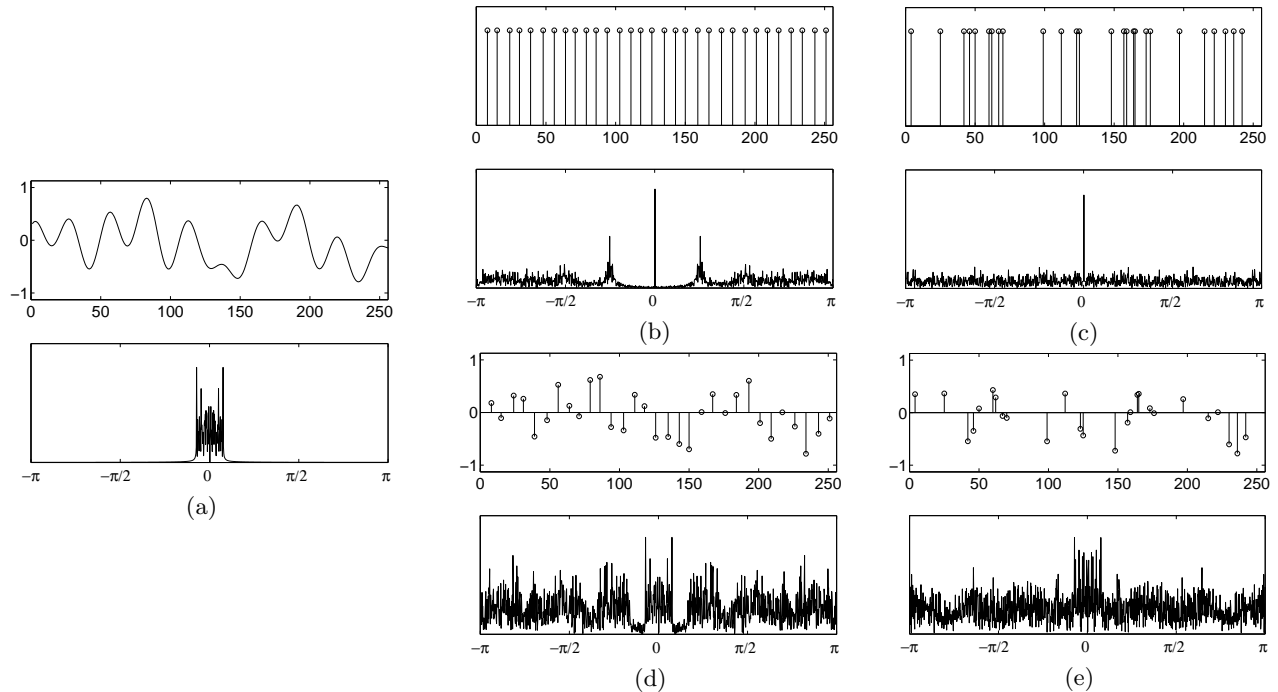


Figure 1: Stochastic sampling process. The original signal in (a) is sampled with a blue noise sampling pattern (b) and a white noise sampling pattern (c). The resulting sampled signals are shown in (d) and (e), respectively.




Figure 2: Blue noise (a) and white noise (b) sampling patterns.

the region to be monitored such that in all subregions of the area, it is likely that the density of the sensor nodes is more than necessary to meet the SNR requirements of the application. In such situations, accuracy of the reconstructed phenomenon can be traded for energy efficiency. In other words, a smaller subset of sensor nodes can be used to observe the area. Although there may not always be a direct correlation between power consumption in the network and the number of active sensors, reducing the number of active sensors generally leads to better energy efficiency.

Since the objective of our application is to reconstruct a spatially bandlimited signal (e.g., a temperature field), we can consider the cumulative sensor data to be a stochastically sampled signal, sampled at the active sensor locations, and apply stochastic sampling theory to our selection algorithm. Specifically, since the random deployment of the sensors limits us such that we must stochastically (rather

than regularly) sample the phenomenon, we would like to sample with a blue noise pattern so that the accuracy of the reconstructed signal is as high as possible for a given number of sensors. Thus, the sensor selection algorithm was designed using the intuition of the algorithm for creating blue noise patterns presented in [11].

We assume a grid overlaid on the region where the sensor network is deployed with resolution high enough that each sensor's location can be precisely mapped to the nearest grid point and each point on the grid is associated with no more than one sensor. If a grid point has a sensor mapped to it, the grid point is assigned a value of 1; otherwise, it is assigned a value of 0. The resulting binary pattern $s_{initial}[i, j]$ should have white noise spectral characteristics because of our assumption about the random deployment of sensors.

In our proposed method, a low pass filter relaxation algorithm is iteratively applied to the current sampling pattern/point set $s_{current}[i, j]$, which is initially set to the initial sensor pattern $s_{initial}[i, j]$. The characteristics of the low pass spatial filter $h_{BN}[i, j]$, which is a symmetric filter based on a one-dimensional impulse response $h_{BN}(d)$, should depend on the characteristics of the phenomenon that is being observed. More specifically, the coefficients of the low pass filter are determined such that the frequency content of the observed variable falls within the filter's pass band. Meanwhile, the selection of the order of the filter is essentially a tradeoff between the desired performance of the algorithm and computational cost. The algorithm filters the sampling pattern to create $f[i, j]$, such that $f = s_{current} * h_{BN}$. Next, the algorithm finds the maximum value of $f[i, j]$ at the grid points where any currently active sensors reside, deactivates the corresponding sensor, and updates $s_{current}$ by zeroing

the grid point that the sensor is mapped to. Essentially, this results in the removal of a sensor that contributes significant low spatial frequency content. These steps are carried out iteratively until the maximum filter output drops below a predetermined threshold or the number of remaining active sensor nodes drops below a certain value. The resulting sampling pattern $s_{BN}[i, j]$ is shown to have blue noise spectral characteristics [4].

4.1 Incorporation of an Energy Cost

One of the most important design goals for wireless sensor networks is energy efficiency, which must be met in order to allow networks to operate unattended for extended periods of time. Deselecting nodes according to an algorithm such as the one described in the previous section clearly reduces overall power consumption in the network. In addition to this, it is also important to rotate activity among nodes and to specifically avoid the use of nodes with little remaining lifetime. Consider a network that chooses active sensors based on a method similar to the one described in this section, where chosen sensors operate until their energy supply becomes completely depleted. Toward the end of the network's lifetime, the set of available sensors to choose from will be more sparsely located and more sensors may be needed to achieve the same SNR as in the early stages of the network. In order to avoid this scenario, we would like to extend the time before any of the sensors in the network deplete their energy supply and die.

Here, we propose a modification to the blue noise sensor selection algorithm in which we incorporate energy costs. In this modified selection method, the low pass filter output is combined with an energy cost so that

$$Cost(S_i) = f[x(S_i), y(S_i)]^\alpha \times Cost_{energy}(S_i)^{(1-\alpha)} \quad (3)$$

where α ($0 \leq \alpha \leq 1$) represents a tuning parameter that allows the system designer to balance a tradeoff between ideality of the sensor pattern and balanced load distribution and $Cost_{energy}(S_i)$ represents a monotonically decreasing function of S_i 's residual energy. In essence, we are combining a redundancy cost with an energy cost. A typical energy cost assignment might be as simple as the inverse of a node's residual energy. During each step, the sensor with the highest overall cost is deselected. These steps are repeated iteratively, as in the original algorithm, until the active subset consists of the desired number of sensors or guarantees a sufficient SNR. Following the selection algorithm, the selected subset of nodes is used to monitor the environment and send data to the base station for a round of arbitrary length. Following the completion of a round, the selection algorithm is repeated with updated energy information from the sensors.

4.2 Distribution of the Algorithm

As wireless sensor networks are expected to grow to orders of thousands of nodes, it is desirable to distribute organization algorithms. Advantages of distributed algorithms include load distribution and fault tolerance, among others. Conveniently, the nature of our proposed sensor selection algorithm makes it easily distributable. The deactivation of a sensor is affected only by sensors with distances of less than $\frac{N \times T_r}{2}$ from the sensor, where N represents the order of the filter used to create the blue noise sampling pattern

and T_r represents the arbitrary resolution of the image grid to which the sensors are mapped. In fact, there is not necessarily any need for the concept of a discrete grid, as long as each sensor S_i is able to translate distance to its neighbors S_n to the appropriate values of the filter's impulse response $h_{BN}(dist(S_i, S_n))$.

As long as sensors are synchronized and begin the sensor selection algorithm simultaneously, they may set a backoff timer according to

$$B(S_i) = W - K \times f(x(S_i), y(S_i)) \quad (4)$$

$$f(x, y) = s_{initial}(x, y) * h_{BN}(x, y) \quad (5)$$

where W represents the maximum backoff window value, and K is chosen according to the expected node density and the filter used to create the blue noise sampling pattern. Using this approach, sensors whose locations correspond to high values of f will have low backoff timers, as they should have highest priority for deactivation. When sensor S_i 's timer expires, it broadcasts a beacon to its neighbors, informing them of its intended deactivation. When the neighboring nodes S_n receive this deactivation beacon from S_i , they adjust their calculated filter output by subtracting $h_{BN}(dist(S_n, S_i))$ and increase their backoff time as

$$B(S_i) = B(S_i)' + K \times h_{BN}(dist(S_n, S_i)) \quad (6)$$

where $B(S_i)'$ represents the value of $B(S_i)$ immediately before the beacon is received. It is convenient that sensors can only increase their backoff time as a result of receiving another node's beacon, meaning that the nodes must only synchronize once during the algorithm. The algorithm should terminate once the density of the nodes approaches a threshold based on the desired SNR. In terms of the distributed algorithm, this means that once a node resets its backoff value past a certain threshold B_{max} , it should withhold its beacon, stop listening for beacons, and assume that it will remain active for the next round. The energy cost modification could easily be incorporated into this distributed version of the protocol by making minor changes to Equations 4, 5, and 6.

An illustration of how the algorithm works (using sensors deployed in a single dimension for clarity) is shown in Figure 3. Figure 3(a) shows the initial values of f and the corresponding timer values at each of the sensor nodes. Since sensor S_2 has the highest value of f , it is the first to send a beacon. Consequently, the other nodes update their values of f and their backoff timers, as shown in Figure 3(b). S_2 's beacon causes sensors S_1 and S_3 to set their timers beyond B_{max} , and the algorithm terminates at these nodes. At this point, S_5 has the highest value of f . After its backoff timer expires, S_5 sends a beacon and again, each of the other active sensors updates f and their backoff timers accordingly, as shown in Figure 3(c). At this time, all of the remaining sensors have increased their backoff timers beyond B_{max} and remain active. It should be noted that as long as all sensors are synchronized, the distributed version of the algorithm should converge to the exact same solution as the centralized version.

The implementation of this distributed algorithm requires very little computation in the sensor. For each neighbor, the nodes must initially translate location to distance, requiring

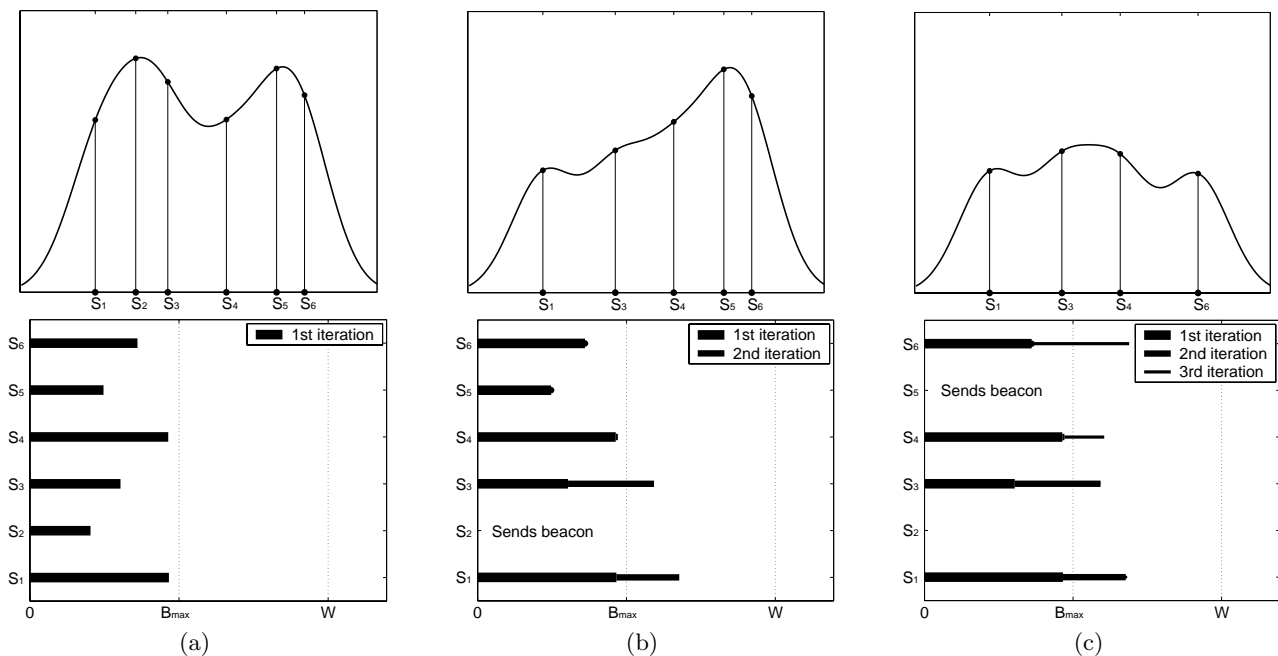


Figure 3: Distribution of the sensor selection algorithm. Figures (a)-(c) show the value of f and the current backoff timer values for the sensors at three time instances.

a few operations. Following this, only a few lookups, additions, and possibly subtractions (following the reception of a deactivation beacon) are necessary per neighbor. A full-blown filtering operation is not necessary since most of the image grid being filtered is zero-valued. For the version of the algorithm that considers energy consumption, another multiplication is necessary for each neighbor. The energy overhead for implementing this algorithm is clearly dominated by the transmission of the short beacon messages, and since we can expect the number of each node's neighbors to be bounded to the order of tens, we do not expect the overhead to be very significant. Furthermore, the processing capabilities required by the algorithm are very reasonable.

4.3 Lifetime Optimization

Here, we propose another energy efficient sensor selection approach. In this approach, active sensor subsets are created in a similar manner as in the original blue noise sensor selection algorithm. However, rather than deterministically deselecting nodes at locations with the highest filter output, nodes are deselected with a weighted probability proportional to this value. Having introduced this factor of randomness, many active subsets, each providing the desired blue noise sampling characteristics, can be found. It is possible to schedule the use of each of these subsets so that the total lifetime of the monitoring application is maximized [12]. This optimization can be performed through a simple linear programming approach, which accounts for initial energy at individual sensors. Unfortunately, such an optimized schedule is not as robust or as easily distributed as the previously described distributed algorithm. For example, if a sensor failure occurs, the use of all sensor subsets of which it was a member may no longer provide the SNR requirements and should no longer be used. At this point, the optimal schedule would need to be recalculated and all

previous sensor activity might be considered misguided in hindsight. Similar consequences might stem from the use of mobile sensors. Also, such an optimal schedule would not be able to take advantage of additional sensor deployment. In light of these observations, this method should only be used in static topology networks where one sensor has the computational power necessary to solve large linear programs.

5. SIMULATIONS

We simulated a network of sensors randomly deployed within a $128m \times 128m$ region, of which a given number were activated to monitor a bandlimited phenomenon whose average signal power was normalized to unity. We compared our proposed blue noise sensor selection approach with the random selection of sensors (essentially, a white noise sampling approach) and a grid-based approach. In the grid-based approach, a hexagonal grid was overlaid onto the region where the sensors were deployed and the closest sensor to each grid point was selected for activation. Following the selection of nodes, the original data image (e.g., temperature field) was reconstructed from the sensor samples using the Voronoi reconstruction algorithm [2] and the mean square error was calculated.

In our first simulations, there were 250 sensors deployed to sense a signal that was bandlimited to $\frac{1}{20}m^{-1}$. We measured the mean square error of the reconstructed signal using each approach as we increased the number of sensors activated. Figure 4 shows that, as expected, for all methods the mean square error decreases as the number of activated sensors increases. However, the blue noise sampling method performs slightly better than the grid-based sampling approach and much better than the random (white noise) sampling approach. At first, it might seem surprising that the grid-based approach does not perform the best. However, this

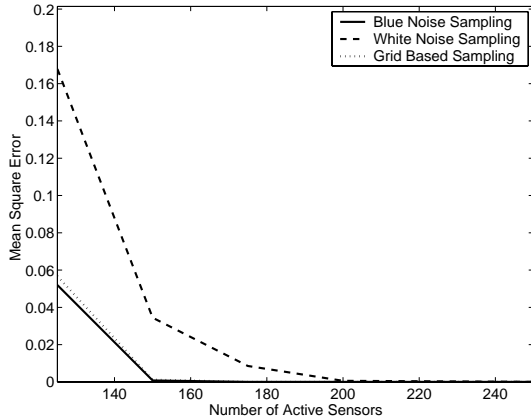


Figure 4: Mean square error of reconstructed signal using the blue noise, grid-based, and random selection methods for a phenomenon bandlimited to $\frac{1}{20}m^{-1}$.

method is only expected to perform very well when the number of selected active sensors is much less than the number of deployed sensors. Otherwise, the sampling pattern created from the grid-based method may not resemble a grid at all, due to the random initial placement of the sensors. The blue noise selection algorithm also holds the advantages of being able to easily distribute the algorithm.

In Figure 4, the error goes to zero at approximately 150 sensors using the blue noise and grid-based approaches and approximately 200 sensors using the random sampling approach, meaning that at these points, the Nyquist sampling rate criteria is being met in all subregions. Even when the sensor deployment is not sufficient to meet this criteria, a subsampling approach can be beneficial if the deselected sensors are chosen correctly. In Figure 5, we plot the results from a scenario similar to the one above, but with a phenomenon that is bandlimited to $\frac{1}{15}m^{-1}$. In this case, even when all 250 sensors are activated, the mean square error of the reconstructed signal does not go to zero. However, with an intelligent selection scheme such as the blue noise selection algorithm, a limited number of sensors may be deactivated without much further loss in signal quality. Figure 5 shows that with the blue noise selection algorithm, the deselection of 20 sensors has almost no effect on signal quality, and even the deselection of 40 sensors does not greatly affect signal quality.

Next, we ran some simple simulations to show the benefit of including an energy cost. In these simulations, 400 nodes were deployed, of which 150 were selected for activation. We used the inverse of a node’s residual energy as its energy cost, so that

$$Cost(S_i) = f(x(S_i), y(S_i))^\alpha \times \left(\frac{1}{E(S_i)}\right)^{(1-\alpha)} \quad (7)$$

where $E(S_i)$ represents the residual energy of sensor S_i , which was initially distributed uniformly between 2 J and 10 J in these simulations. We varied α to observe the tradeoff between ideality of the sampling pattern and appropriate distribution of the energy load. For a large value of α , the algorithm is essentially unchanged from the original version,

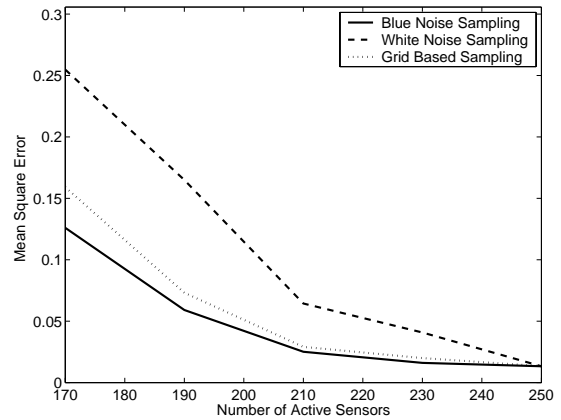


Figure 5: Mean square error of reconstructed signal using the blue noise, grid-based, and random selection methods for a phenomenon bandlimited to $\frac{1}{15}m^{-1}$.




Figure 6: Selection optimality versus load balancing tradeoff imposed by the assignment of an energy cost.

while a small value of α means that good load distribution has become the more critical goal. This tradeoff is illustrated in Figure 6. For large values of α , a very low mean square error is achieved, but many nodes with little residual energy are selected. The average residual energy of the selected nodes is the mean of the energy distribution — 6 J. As we decrease the value of α , the average residual energy of the selected nodes increases, as desired, while there is initially a very small rise in mean square error. As the value of α decreases further, it can be seen that mean square error begins to increase more rapidly. A typical application may want to operate somewhere near the knee of the curve, simultaneously attaining the goals of accuracy of the reconstructed data image and good load balancing.

Finally, we ran simulations to observe the lifetime achievable through the optimization of the sensor schedule, as described in Section 4.3. We compared the optimal schedule’s lifetime with that of a blind selection method, in which we select the sensor subset based on the unaltered blue noise se-

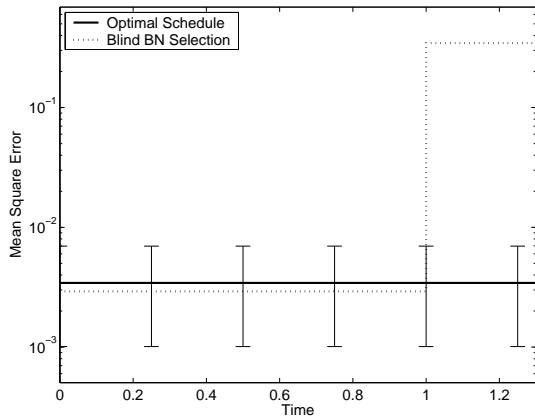


Figure 7: Mean square error of optimized sensor schedule (mean and standard deviation) compared with mean square error using the blind selection approach.

lection algorithm, and iteratively reselect the subset once a sensor in the subset dies. In these simulations, we deployed 100 sensors in a $128m \times 128m$ field and chose active subsets of 50 nodes to observe a phenomenon that was bandlimited to $\frac{1}{35}m^{-1}$. For the optimized schedule, we chose 100 sensor subsets of 50 nodes each according to the blue noise sampling pattern algorithm (with the necessary modifications to add a factor of randomness to the selection algorithm). Nodes were each given a lifetime of 1 (units are arbitrary).

In Figure 7, we show the average mean square error of the reconstructed signal as a function of time using the optimized sensor schedule and using the blind selection algorithm, averaged over 20 trials. Since the order in which the subsets are used in the optimized schedule is arbitrary, we show the average and standard deviation of these subsets instead of a time plot. The blind selection scheme initially chooses a sensor subset that reconstructs the data image with a low mean square error, as Figure 8(a) shows for a single trial. However, once the energy of the nodes in the first subset is used up (at time 1), there are only 50 sensors remaining from which to choose the next set of 50. The resulting sensor subset’s sampling pattern, shown in Figure 8(b), does not have the desired blue noise properties. This is reflected in the high mean square error in the late stages of network operation. On the other hand, the optimization procedure uses many high quality sensor subsets, a sampling of which is shown in Figure 9. The optimal scheduling of these sets allows the network to perform well over a longer period of time, for an average length of 1.43.

To observe the performance of the optimization program when the energy among the deployed nodes is nonuniform, we simulated a network in which the nodes’ initial energy was randomly distributed so that sensor lifetime ranged from 3 to 10 time units, taking only integer values in order to simplify the simulations. Again, we found 100 sensor subsets from which we calculated the optimal schedule and compared with the blind approach. The results are shown in Figure 10. The optimized schedule allows the network to operate with a low mean square error for an average time length of 6.3, while the average mean square error of the




Figure 8: Initial (a) and subsequent (b) sets chosen through blind blue-noise sensor selection method.




Figure 9: Sampling of sensor subsets used in the optimal sensor scheduling method.

blind approach reaches an unacceptable level by this time. Of course, the amount by which lifetime can be extended is affected by factors such as how many subsets we run the optimization for as well as how much randomness we add to the selection algorithm, which allows sensor pattern ideality to be traded for sensor subset diversity.

6. CONCLUSIONS

We have presented an algorithm for the selection of active sensors in a wireless sensor network whose application is to reconstruct the data image of spatially bandlimited physical phenomenon such as temperature. Through simulation results, we have shown that significant improvement in overall accuracy of the reconstructed signal can be achieved by this algorithm compared with other simple approaches. We have also shown that by incorporating an energy cost into the algorithm, one can trade accuracy of the reconstructed data image for well balanced load distribution among the sensors in the network. Finally, if multiple sensor subsets are chosen through a more randomized version of the blue noise selection algorithm, the subsets can be scheduled in such a way that the network can produce a higher quality signal over a longer period of time compared with a simple approach of iteratively selecting a new sensor subset whenever a currently active sensor dies.

In the future, we would like to develop the details and per-

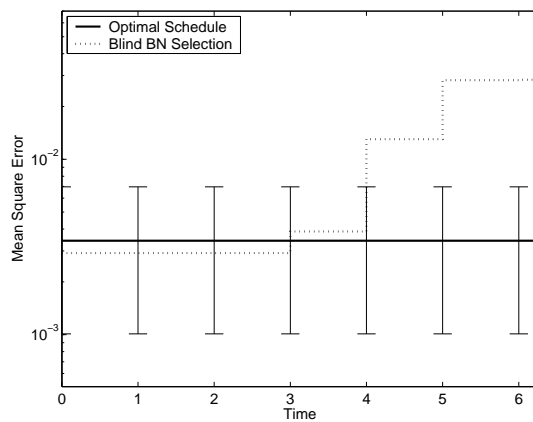


Figure 10: Mean square error of optimized sensor schedule (mean and standard deviation) compared with mean square error using the blind selection approach, for randomly distributed initial energy.

form simulations of the distributed version of this algorithm using a round based approach and observe the achievable lifetime of the network as well as the degradation of SNR over time. Also, we have worked under the assumption that all sensors send traffic directly to a base station or local cluster head over a single hop. Thus, the energy cost assignment is based solely on a sensor's own residual energy since that sensor is the only one affected by its activation. This is not the case in multihop networks, and for such networks, it may be more appropriate to incorporate a cumulative path cost. We wish to explore this issue in the future as well.

7. REFERENCES

- [1] J. Chou and D. Petrovic. A Distributed and Adaptive Signal Processing Approach to Reducing Energy Consumption in Sensor Networks. In *Proceedings of the Twenty Second International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2003.
- [2] H. G. Feichtinger and K. H. Grochenig. Theory and Practice of Irregular Sampling. In J. Benedetto and M. Frazier, editors, *Wavelets: Mathematics and Applications*, pages 305–363. CRC Press, 1994.
- [3] H. Gupta, S. Das, and Q. Gu. Connected Sensor Cover: Self-Organization of Sensor Networks for Efficient Query Execution. In *Proceedings of the Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2003.
- [4] S. Hiller, O. Deussen, and A. Keller. Tiled Blue Noise Samples. In *Proceedings of Vision, Modeling, and Visualization*, 2001.
- [5] P. Ishwar, A. Kumar, and K. Ramchandran. Distributed Sampling for Dense Sensor Networks: A Bit-Conservation Principle. In *Proceedings of the Second International Workshop on Information Processing in Sensor Networks (IPSN)*, 2003.
- [6] R. Iyer and L. Kleinrock. QoS Control For Sensor Networks. In *Proceedings of the IEEE International Conference on Communications*, 2003.
- [7] S. Lloyd. Least Square Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, March 1982.
- [8] D. Mitchell. Spectrally Optimal Sampling for Distribution Ray Tracing. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*, 1991.
- [9] T. Mitsa and K. J. Parker. Digital Halftoning Using a Blue-Noise Mask. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, 1991.
- [10] T. Mitsa and K. J. Parker. Digital Halftoning Technique Using a Blue-Noise Mask. *Journal of the Optical Society of America A*, 9(11):1920–1929, November 1992.
- [11] K. Parker, T. Mitsa, and R. Ulichney. A New Algorithm for Manipulating the Power Spectrum of Halftone Patterns. In *Proceedings of SPSE's 7th International Congress on Non-Impact Printing*, 1991.
- [12] M. Perillo and W. Heitzelman. Simple Approaches for Providing Application QoS Through Intelligent Sensor Management. *Elsevier Ad Hoc Networks Journal*, 1(2-3):235–246, September 2003.
- [13] S. S. Pradhan, J. Kusuma, and K. Ramchandran. Distributed Compression in a Dense Microsensor Network. *IEEE Signal Processing Magazine*, 19(2):51–60, March 2002.
- [14] K. D. Sauer and J. P. Allebach. Iterative Reconstruction of Bandlimited Images from Nonuniformly Spaced Samples. *IEEE Transactions on Circuits and Systems*, 34(12):1497–1506, December 1987.
- [15] R. Stasinski and J. Konrad. Improved POCS-based image reconstruction from irregularly-spaced samples. In *Proceedings of the XI European Signal Processing Conference*, 2002.
- [16] D. Tian and N. Georganas. A Node Scheduling Scheme for Energy Conservation in Large Wireless Sensor Networks. *Wireless Communications and Mobile Computing Journal*, 3(2):271–290, March 2003.
- [17] R. A. Ulichney. Dithering with Blue Noise. *Proceedings of the IEEE*, 76(1):56–79, January 1988.
- [18] X. Wang, G. Xing, Y. Zhang, C. Lu, R. Pless, and C. Gill. Integrated Coverage and Connectivity Configuration in Wireless Sensor Networks. In *Proceedings of Sensys*, 2003.
- [19] T. Yan, T. He, and J. A. Stankovic. Differentiated Surveillance for Sensor Networks. In *Proceedings of Sensys*, 2003.
- [20] F. Ye, G. Zhong, J. Cheng, S. Lu, and L. Zhang. PEAS: A Robust Energy Conserving Protocol for Long-lived Sensor Networks. In *Proceedings of the Twenty Third International Conference on Distributed Computing Systems*, 2003.