

A Analysis of Strategies for Mitigating the Sensor Network Hot Spot Problem

Mark Perillo, Zhao Cheng, and Wendi Heinzelman
Department of Electrical and Computer Engineering
University of Rochester
Rochester, NY 14627
Email: {perillo,zhcheng,wheinzel}@ece.rochester.edu

Abstract

In multi-hop wireless sensor networks that are characterized by many-to-one (converge-cast) traffic patterns, problems related to energy imbalance among sensors often appear. When the transmission range is fixed for nodes throughout the network, the amount of traffic that sensors are required to forward increases dramatically as the distance to the data sink becomes smaller. Thus, sensors closest to the data sink tend to die early. Network lifetime can be improved to a limited extent by the use of a more intelligent transmission power control policy that balances the energy used in each node by requiring nodes further from the data sink to transmit over longer distances (although not directly to the data sink). Alternatively, policies such as data aggregation allow the network to operate in a more energy efficient manner. Since the deployment of an aggregator node may be significantly more expensive than the deployment of an ordinary microsensor node, there is a cost tradeoff involved in this approach. This paper provides an analysis of these policies for mitigating the sensor network hot spot problem, considering energy efficiency as well as cost efficiency.

1 Introduction

Large scale wireless sensor networks are an emerging technology that have recently gained attention for their potential use in applications such as environmental sensing and mobile target tracking. Since sensors typically operate on batteries and are thus limited in their active lifetime, the problem of designing protocols to achieve energy efficiency to extend network lifetime has become a major concern for network designers. Much attention has been given to the reduction of unnecessary energy consumption of sensor nodes in areas such as hardware design, collaborative signal processing, transmission power control policies, and all levels of the network stack. However, reducing an indi-

vidual sensor's power consumption alone may not always allow networks to realize their maximal potential lifetime. In addition, it is important to maintain a balance of power consumption in the network so that certain nodes do not die much earlier than others, leading to unmonitored areas in the network.

Previous research has shown that because of the characteristics of wireless channels, multi-hop forwarding between a data source and a data sink is often more energy efficient than direct transmission. Based on the power model of a specific sensor node platform, there exists an optimal transmission range that minimizes overall power consumption in the network. When using such a fixed transmission range in general ad hoc networks, energy consumption is fairly balanced, especially in mobile networks, since the data sources and sinks are typically assumed to be distributed throughout the area where the network is deployed. However, in sensor networks, where many applications require a many-to-one (covergecast) traffic pattern in the network, energy imbalance becomes a very important issue, as a hot spot is created around the data sink, or base station. The nodes in this hot spot are required to forward a disproportionately high amount of traffic and typically die at a very early stage. If we define the network lifetime as the time when the first subregion of the environment (or a significant portion of the environment) is left unmonitored, then the residual energy of the other sensors at this time can be seen as wasted.

Intuition leads us to believe that the hot spot problem can be solved by varying the transmission range among nodes at different distances to the base station so that energy consumption can be more evenly distributed and the lifetime of the network can be extended. However, this is only true to some extent, as energy balancing can only be achieved at the expense of using the energy resources of some nodes inefficiently. Alternatively, it may be possible to form clusters within the network and distribute the processing of sensor network data regionally. Not only does this mitigate the hot spot problem significantly, but it increases the power

efficiency of the network, as data is sent over many fewer hops on average. However, as the deployment of aggregator nodes may be more expensive than deployment of an ordinary sensor node, the optimal deployment method should be considered from a cost perspective. In this paper, we explore this tradeoff and propose cost efficient deployment solutions.

The rest of this paper is organized as follows. Section 2 explores the effectiveness of transmission power control in mitigating the hot spot problem. Section 3 investigates the effectiveness and cost efficiency of the deployment of multiple aggregator nodes as an alternative solution to the hot spot problem. Section 4 concludes the paper and suggests future work in this area.

2 Transmission Power Control

In this section, we review the transmission range distribution optimization problem, originally addressed in [4]. This problem is solved by determining how a node should distribute its outgoing data packets over multiple distances, always using the minimum transmission power necessary to send over each distance. Given the energy constraints and data generation rate of each sensor node, the lifetime of the network, which we define to be the time at which the first sensor dies, can be maximized by using this optimal distribution. In our model, we make several basic assumptions. First, we assume that the power consumption of sensor nodes is dominated by communication costs, as opposed to sensing and processing costs. We also ignore the overhead that would typically be introduced by the routing and MAC layers.

The power consumption model that we use is such that the amount of energy to transmit a bit can be represented as $E_{bit,tx} = E_{elec} + \epsilon d^\alpha$, and the amount of energy to receive a bit can be represented as $E_{bit,rx} = E_{elec}$, where E_{elec} represents the electronics energy, ϵ is determined by the transmitter amplifier's efficiency and channel conditions, and α represents the path loss exponent [3]. It has been shown that using our power model, there exists an ideal energy efficient transmission range wireless networks in the absence of the convergecast-specific hot spot problem equal to $d^* = \sqrt[\alpha]{\frac{2E_{elec}}{(\alpha-1)\epsilon}}$ [1, 2]. In all simulations, we used values of $E_{elec} = 50 \text{ nJ/bit}$ and $\epsilon = 100 \text{ pJ/bit/m}^2$, resulting in a value of $d^* = 32\text{m}$.

In this section, we provide simulation results in order to observe the optimal transmission range distributions for two-dimensional sensor network deployments of uniform sensor distribution. The initial energy and data generation rate of all nodes in the network were arbitrarily set to 1 J and 1 bit/second. In these simulations, we compared the lifetimes obtained using the optimal distribution, using the




Figure 1. Optimal transmission range distribution for sensor fields with radii of 150m (a) and 250m (b).

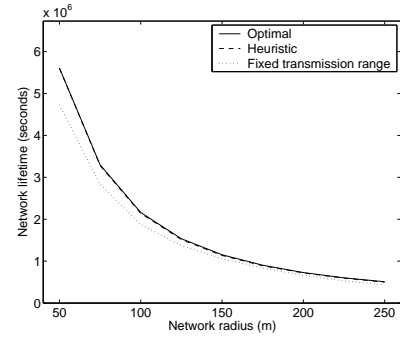


Figure 2. Comparison of lifetime using different transmission power control policies.

optimal fixed network-wide transmission range, and using a heuristic routing scheme. In this heuristic scheme, routing costs are assigned to sensors to be the inverse of their residual energy. Link costs are set equal to a weighted sum of the energy consumed by the transmitting node and the receiving node, as given by Equation 1.

$$C_{link}(s_i, s_j) = \frac{1}{E_{res}(s_i)} E_{tx,bit}(s_i, s_j) + \frac{1}{E_{res}(s_j)} E_{rx,bit}(s_i, s_j) \quad (1)$$

Finally, minimum cost routes are found through Dijkstra's algorithm.

The optimal transmission range distributions are shown in Figure 1 for network radii of 50m, 150m, and 250m. We have observed similar results in randomly deployed two-dimensional sensor networks. However, we have omitted these results due to space limitations. Figure 2 shows network lifetime as a function of network radius for two-dimensional networks using the optimal transmission range distribution, using the optimal fixed transmission range, and using the heuristic routing scheme.




Figure 3. Lifetime vs. percentage of the total energy consumed in the network.

2.1 Energy Inefficiency

While transmission range distribution optimization is somewhat effective in extending network lifetime compared to schemes using a fixed transmission range, this improvement is limited because of the energy inefficiency forced on the sensors farthest from the data sink in order to evenly distribute the energy load among the nodes. In fact, in order to achieve near-optimal network lifetimes, it is only necessary to use a fraction of the energy available in the network. Figure 3 shows how network lifetime increases with the total energy used in the network when the energy consumption of each individual sensor node is limited at $1 J$. If energy is to be allocated among the nodes in any way in order to maximize network lifetime (i.e., if only the total energy consumption, but not individual energy consumption, is limited), network lifetime should increase proportionally with the energy consumed. This lifetime is illustrated by the dotted line in the figure. With the individual energy consumption constraints imposed by our assumption of uniform node distribution and equal energy allocation, however, the obtainable lifetime, illustrated by the solid line in Figure 3, is found to be only a fraction of this. The shape of the energy-lifetime curve implies that when all of the network energy is completely used, marginal network lifetime improvement is minimal and energy is being used inefficiently.

2.2 Networks with Limited Transmission Ranges

Finally, we observe the effect of setting a maximum transmission range on network lifetime. Figure 4 illustrates the obtainable lifetime as a function of network radius for various maximum transmission ranges. Limiting the maximum transmission range severely affects network lifetime, especially for large network radii. In addition to the penalty in network lifetime incurred from the uniform deployment constraint (as illustrated in Figure 3), an additional penalty of is incurred from limiting the transmission range. Since all transmissions in a network with a well-balanced traffic

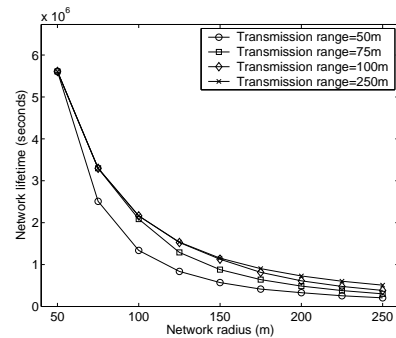


Figure 4. Network lifetime as a function of network radius for limited transmission ranges.

load are very likely to be around $32m$, this penalty can also be attributed to the convergecast traffic model. Clearly, solutions that balance the traffic load more efficiently must be considered. This gives further motivation for the work presented in the following section, in which more energy efficient and cost efficient solutions are considered.

3 Clustering Approach

In the network deployment plans outlined in the previous section, only one node serves as the data sink for the entire network at one time, even if multiple aggregator-capable nodes have been deployed. In this section, we consider a clustering approach in which multiple aggregator-capable nodes are deployed and each sink collects data from only part of the sensor network for the entire network lifetime. Such clustering schemes have been proposed for wireless sensor networks in [3, 8]. Previous work in this area deals primarily with homogeneous networks, in which any of the deployed nodes is capable of acting as cluster head. In this section, we consider heterogeneous networks, where nodes equipped with the capability of acting as cluster head (e.g., those with larger batteries, more processing power and memory, and possibly a second radio to link back to a central base station) are significantly more expensive than ordinary microsensors. In our model, a sensor may send its traffic to whichever cluster head it chooses (typically, but not necessarily, the closest cluster head).

We must first define a general metric to describe the efficiency of a network deployment plan. For this purpose, we use the normalized network lifetime \tilde{L} , which measures how many total bits can be transported on the network per unit of energy. For a given network scenario and power model, linear programming allows us to find the optimal lifetime L^{opt} that can be achieved. Lifetime can be arbitrarily increased by simply increasing the energy density in the

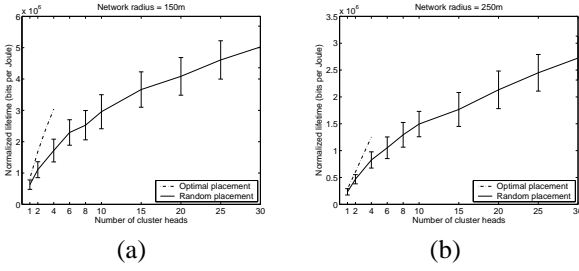


Figure 5. Normalized lifetime vs. number of cluster heads deployed for networks with a radius of 150m (a) and 250m (b).

network (either by scaling up the deployed sensor density or the average energy per sensor). Also, since we assume that protocols that manage the amount of traffic sent (e.g., [7, 6, 5]) may be used so that the density of active sensors does not necessarily correspond to the density of deployed sensors, lifetime can similarly be increased by decreasing the required active sensor density. Similarly, lifetime can be increased by reducing the bit rate among active sensors. To account for these factors, the normalized network lifetime \tilde{L} is given as

$$\tilde{L} = L^{opt} \left(\frac{R_a \lambda_a}{\lambda_e} \right) \quad (2)$$

where λ_a represents the density of active sensors, R_a represents the average bit rate among active sensors, λ_e represents the energy density of the network (we assume uniform distribution of energy), and L^{opt} is the maximum lifetime achievable with the given parameters.

We ran simulation to observe the energy efficiency of the clustering approach. The energy model used was the same as the previous section and each node's maximum transmission range was set at 75m. The relationship between the normalized lifetime (given in bits per Joule) and the number of cluster heads that are deployed is shown in Figures 5(a) and 5(b) for a 150m radius network and a 250m radius network, respectively. In both figures, the normalized lifetime is given for optimal cluster head placement as well as random placement. As expected, when more cluster heads are deployed, the network lifetime improves significantly in both cases. It is obvious that for a smaller network, fewer number of cluster heads are enough to solve the "hot spot" problem, and the figures verify this.

We can find the number of sensors $N_s(L, N_c)$ that need to be deployed to achieve a lifetime L when N_c cluster heads are deployed, as we did in the previous section when we considered the necessary number of sensors when N_l data sink locations are used. For the scenario in which we

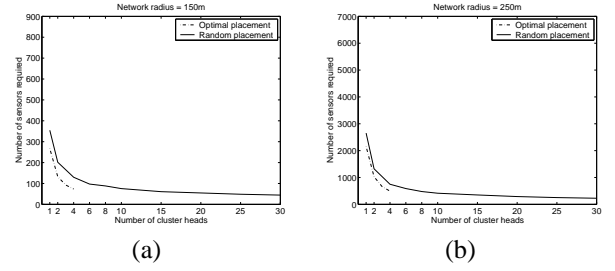


Figure 6. Number of sensors required vs. number of data sinks deployed for networks with a radius of 150m (a) and 250m (b).

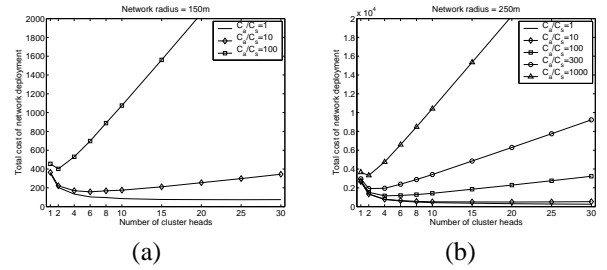


Figure 7. Network deployment cost vs. number of cluster heads deployed for networks with a radius of 150m (a) and 250m (b).

plan a sensor network to operate for one year with active sensors sending data at 1 bit per second and activated at a density of 0.0001 sensors/m², the required number of sensors are shown in Figures 6(a) and 6(b), respectively.

If the cost of a cluster head node is C_a , the total cost of deploying a heterogeneous clustering network is

$$C(L, N_c) = C_a N_c + C_s N_s(L, N_c) \quad (3)$$

This cost is plotted in Figures 7(a) and 7(b) for a 150m radius network and a 250m radius network, respectively. As expected, the most cost-efficient number of cluster heads deployed becomes larger as the price of a cluster head becomes smaller. Also, for the same cost of a cluster head, more cluster heads should be deployed in the larger network since the "hot spot" problem is worst in large networks.

4 Conclusions

We have studied the optimal transmission range distribution that allows the lifetime of sensor networks to be maximized. Based on this model, we revealed the upper bound

of the lifetime of several typical scenarios and demonstrated the inability to make good use of the energy of nodes furthest from the base station, even when utilizing the optimal distribution. Thus, varying the transmission power of individual nodes cannot alone solve the hot spot problem. In addition to transmission power control, we have considered the use of a clustering hierarchy, where heterogeneous sensors are deployed, some of which can act as a data aggregator/compressor. When analyzing the use of each strategy, we also considered the necessary extra costs incurred and show how the network configuration can be optimized for cost efficiency in each case.

References

- [1] Y. Chen, E. Siner, and S. Wicker. On selection of optimal transmission power for ad hoc networks. In *Proceedings of the Thirty-Sixth Hawaii International Conference on System Sciences (HICSS-36)*, 2003.
- [2] J. Deng, Y. Han, P. Chen, and P. Varshney. Optimum transmission range for wireless ad hoc networks. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2004.
- [3] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. An application-specific protocol architecture for wireless microsensor networks. *IEEE Transactions on Wireless Communications*, 1(4):660–670, Oct. 2002.
- [4] M. Perillo, Z. Cheng, and W. Heinzelman. On the problem of unbalanced load distribution in wireless sensor networks. In *Proceedings of the Global Telecommunications Conference (GLOBECOM) Workshop on Wireless Ad Hoc and Sensor Networks*, 2004.
- [5] M. Perillo and W. Heinzelman. DAPR: A protocol for wireless sensor networks utilizing an application-based routing cost. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2004.
- [6] 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 the First ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2003.
- [7] 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 (ICDCS)*, 2003.
- [8] O. Younis and S. Fahmy. Distributed clustering in ad-hoc sensor networks: A hybrid, energy-efficient approach. In *Proceedings of the Twenty-Third Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2004.