Balanced Data Gathering in Energy-Constrained Sensor Networks⋆

Emil Falck[1], Patrik Floréen[2], Petteri Kaski[1],
Jukka Kohonen[2], and Pekka Orponen[1]

[1] Laboratory for Theoretical Computer Science, P.O. Box 5400, FI-02015 Helsinki
University of Technology, Finland. E-mail: `firstname.lastname@hut.fi`.
[2] Helsinki Institute for Information Technology/Basic Research Unit, Department of
Computer Science, P.O. Box 26, FI-00014 University of Helsinki, Finland. E-mail:
`firstname.lastname@cs.helsinki.fi`.

Abstract. We consider the problem of gathering data from a wireless multi-hop network of energy-constrained sensor nodes to a common base station. Specifically, we aim to balance the total amount of data received from the sensor network during its lifetime against a requirement of sufficient coverage for all the sensor locations surveyed. Our main contribution lies in formulating this balanced data gathering task and in studying the effects of balancing. We give an LP network flow formulation and present experimental results on optimal data routing designs also with impenetrable obstacles between the nodes. We then proceed to consider the effect of augmenting the basic sensor network with a small number of auxiliary relay nodes with less stringent energy constraints. We present an algorithm for finding approximately optimal placements for the relay nodes, given a system of basic sensor locations, and compare it with a straightforward grid arrangement of the relays.

1 Introduction

Wireless networks consisting of a large number of miniature electromechanical devices with sensor, computing and communication capabilities are rapidly becoming a reality, due to the accumulation of advances in digital electronics, wireless communications and microelectromechanical technology [1, 14, 22]. Prospective applications of such devices cover a wide range of domains [1, 7, 9, 10].

One generic type of application for sensor networks is the continuous monitoring of an extended geographic area at relatively low data rates [1, 5]. The information provided from all points of the sensor field is then gathered via multi-hop communications to a base station for further processing. We are here envisaging a scenario where environmental data are frequently and asynchronously collected

over an area, and all information is to be gathered for later postprocessing of best possible quality, including detection of possibly faulty data. This means that data aggregation [17, 18] cannot be employed.

Significant design constraints are imposed by multi-hop routing and the limited capabilities and battery power available at the sensor nodes. A number of recent papers have addressed e.g. optimal sensor placement [6, 11, 12, 16, 21] and energy-efficient routing designs and protocols [8, 13, 15, 17, 19, 20, 24] with the objective of lifetime maximization [3, 4, 17, 21].

We envisage the sensor placement to be fixed beforehand, either by an application expert according to the needs of the particular application at hand, or randomly, for example by scattering them from an airplane. For the sake of achieving a comprehensive view of the whole area to be monitored, not only should the total amount of data received at the base station be maximized, but the different sensors should be able to get through to the base station some minimum amount of data. The main contribution of this paper lies in formulating and studying this balanced data gathering problem.

The idea of incorporating a certain balancing requirement on the data gathering has also recently been proposed in [19, 20] and in [11]. In [19, 20] the authors put forth a more general model of information extraction that takes into account the nonlinear relation between transmission power and information rate. Our problem formulation can be seen as a linearized, computationally feasible version of this approach. Another difference between [19, 20] and our work pertains to the expression of the balancing, or fairness, requirement. Article [11] considers the problem of maximizing the lifetime of a sensor network, and explicates this task in terms of an integer program that counts the number of "rounds" the network is operational, assuming that each sensor sends one data packet in each round. This formulation entails a strict fairness condition among the sensors, requiring them all to send exactly equal amounts of data. We allow an adjustable trade-off between maximizing the total amount of data received at the base station and the minimum amount of data received from each sensor. Moreover our program formulation does not require integer variables.

In Sect. 2 we formulate the problem of balanced data gathering as a linear program (LP). In Sect. 3 we present some experiments on routing designs and data flows resulting from various balancing requirements. Although we do not address the issue of ideal placement of the sensor nodes, in Sect. 4 we do consider the effects, from the point of view of our balanced data gathering measure, of augmenting a given sensor network by a small number of auxiliary radio relay nodes with higher battery power levels. The locations of these relays may be chosen at will, and we present and compare two heuristics for determining good relay locations to optimize the network behavior.

It needs to be noted that our linear program based solution relies on information about all the energy costs of transmitting and receiving a unit of data, and the data rates and energy supplies of all nodes. However, knowledge of node locations as such is not required, and our model readily adjusts to obstacles and other deviations from simple radio-link models as long as the energy costs

of transmission can be determined by the nodes themselves, either by simply probing at different power levels, or using more sophisticated means such as a Received Signal Strength Indicator (RSSI) [23].

Our linear program formulation also requires all the information to be available at a single location. This assumption is realistic only if the operation time is long and the amount of control traffic small. Otherwise, the protocol must be able to decide on the basis of local information. Our results thus provide an upper bound on what is actually achievable using distributed protocols, with local and imprecise information.

2 Optimization of Balanced Data Gathering

We consider a network consisting of n sensor nodes, m relay nodes and a base station, all with predetermined locations, except for the relay nodes whose locations may be changed. For simplicity, we index the nodes so that the base station has index 1. The set of all nodes is denoted as $V = B \cup S \cup R$, where B, S, and R denote the sets consisting of the base station node, sensors, and relays, respectively.

Each node $i \in V$ has an initial energy supply of e_i units; as a special case, we set $e_1 = \infty$. The mission of the network is to gather data generated at the sensor (source) nodes to the base station (sink) node under these energy constraints, during the desired operation time T.

We assume that the sensors generate data asynchronously and in such small unit packets that the process can be modeled by assigning to each of the sensor nodes an "offered data rate" parameter s_i, $i \in S$. The energy cost of forwarding a unit of data from node i to node j is given by a parameter d_{ij} and the cost of receiving a unit of data is given by a parameter c. We also assume the transmission rates to be low enough, so that collisions and signal interference can be ignored in the model.

Our model places no restrictions on the values of the parameters d_{ij} and c. As an example, in the commonly used simple radio-link models [23], d_{ij} would be taken to be proportional to $c_t + D_{ij}^\alpha$, where c_t corresponds to the energy consumed by the transmitter electronics and D_{ij}^α corresponds to the energy consumed by the transmit amplifier to achieve an acceptable signal-to-noise ratio at the receiving node; D_{ij} is the physical distance between nodes i and j and the exponent α, $2 \leq \alpha \lesssim 4$, models the decay of the radio signal in the ambient medium. The cost c corresponds to the energy consumed by the receiver electronics.

Thus, if we introduce flow variables f_{ij} indicating the rate of data forwarded from each node i to node j, the energy constraints in the network can be expressed as $\left(\sum_j d_{ij} f_{ij} + \sum_j c f_{ji}\right) \cdot T \leq e_i$, for all $i \in V$.

Because of the energy constraints in the network, the sensors cannot usually productively achieve their full offered data rates; thus we introduce variables r_i indicating the actual "achieved data rate" at each sensor $i \in S$. One goal of the data flow design for the network is to maximize the total, or equivalently, the

average achieved data rate $(1/n)\sum_{i\in S} r_i$. However, taking this as the singular objective may lead to the "starvation" of some of the sensor nodes: typically, the average data rate objective is maximized by data flows that only forward data generated close to the sink, and do not allocate any energy towards relaying data generated at distant parts of the network.

To counterbalance this tendency, we add a "minimum achieved rate" variable ℓ, with the constraints $r_i \geq \ell$ for all $i \in S$, and try to maximize this simultaneously with the average data rate. The trade-off between these two conflicting objectives is determined by a parameter λ, $0 \leq \lambda \leq 1$, where value $\lambda = 0$ gives all weight to the average achieved rate objective, and value $\lambda = 1$ to the minimum achieved rate objective. The combined objective F_λ, see (1), will subsequently be called the "balanced data rate," or "balanced rate."

Different sensors may submit different types of data. At each unit of time, the average amount of data transmitted from one sensor might be one bit, and from another sensor ten bits; however, the one bit may be equally valuable for the application as the other sensor's ten bits. As a generalization, we can assign weights w_i to the data rates from different sensors according to their importance. A natural choice is $w_i = 1/s_i$, which normalizes the data rates of all sensors to the interval $[0,1]$, and expresses the idea that an equal proportion of each sensor's offered data should be transmitted. For simplicity, we have used equal offered data rates and equal weights in our experiments.

Our model can now be formulated as the following linear program, which can then be solved using standard techniques. Note that a linear programming approach is taken also in [11, 17, 21, 24].

$$\max \quad F_\lambda := (1-\lambda)\frac{1}{n}\sum_{i\in S} w_i r_i + \lambda\ell \tag{1}$$

$$\text{s.t.} \quad \sum_{j\in V} f_{1j} = 0,$$

$$\sum_{j\in V} f_{ij} = r_i + \sum_{j\in V} f_{ji}, \qquad i \in S,$$

$$\sum_{j\in V} f_{ij} = \sum_{j\in V} f_{ji}, \qquad i \in R,$$

$$\sum_{j\in V} T d_{ij} f_{ij} + \sum_{j\in V} T c f_{ji} \leq e_i, \quad i \in V,$$

$$r_i \leq s_i, \qquad i \in S,$$

$$w_i r_i \geq \ell, \qquad i \in S,$$

$$f_{ij} \geq 0, \qquad i,j \in V,$$

$$f_{ii} = 0, \qquad i \in V.$$

A flow matrix f_{ij}, obtained as a solution to this linear program, can easily be used to route approximately r_i unit-size data packets from each source node $i \in S$ to the sink node 1, assuming that all the r_i and f_{ij} values are large. At

each node k, simply forward the first $\lfloor f_{k1} \rfloor$ packets to node 1 (the sink), the next $\lfloor f_{k2} \rfloor$ packets to node 2, the next $\lfloor f_{k3} \rfloor$ packets to node 3, and so on. A somewhat more elegant solution is to randomize the routing strategy, so that each incoming packet at node i is forwarded to node j with probability $f_{ij} / \sum_k f_{ik}$.

3 Experimental Results

As already mentioned, we do not address the problem of optimal sensor placement, but take the sensor locations as given. Since our focus is on studying the effect of the balancing factor λ on the resulting data flows and sensor data rates in the network, we choose in most of our experiments to place the sensor nodes in a regular grid. This eliminates the coincidental effects arising in, e.g., a random node placement from the variations in the distances between the nodes. However, we illustrate also the case of random node placement in Sect. 3.2.

3.1 Node Placement in a Regular Grid

In our first experiments, we place 100 sensors in a 10×10 grid in a square of dimensions 1 km $\times$ 1 km and the base station at the middle of one of the sides of the square. All sensors have an energy constraint of 20 J and an offered amount of data of 100 Mbits during the operation of the network. If the time of operation is $T = 10^6$ seconds, this translates to an average offered data rate of 100 bit/s. Transmission and reception costs are computed as $d_{ij} = 100$ nJ/bit $+$ 0.01 nJ/bit/m$^2 \cdot D_{ij}^2$ and $c = 100$ nJ/bit. These values are comparable to those in [4, 15].

The effects of the balancing factor λ on the resulting flows and rates in the network are illustrated in Figs. 1, 2 and 3. In these examples, the offered rates are relatively high, making the network heavily energy-constrained: the achieved data rates are limited more by the network's ability to transport data than by the sensors' ability to generate it.

If no balancing is required ($\lambda = 0$), the objective is simply to maximize the average data rate from all sensors. Since the sensors nearest to the base station can provide a large contribution to the average by sending at relatively low energy cost, the optimum solution indeed allocates most of the network's resources to this goal. Very little data is received from the most distant sensors. As higher balancing factors are used ($\lambda = 0.5$, $\lambda = 1$), the distant nodes get a bigger share of the network's transport resources – and accordingly, the area is more evenly covered by observations. This comes at a cost of reducing the average rate.

Depending on the characteristics of the network, the effect of balancing can be quite large. With no balancing the achieved rates are over 12 bit/s on the average, but below 2 bit/s for the most distant nodes. With full balancing, all sensors get to transmit at an equal rate of about 7.4 bit/s.

An alternative approach would be to limit the data rates of the nearby nodes from above (e.g. as in [19]), in order to prevent them from sending unnecessarily
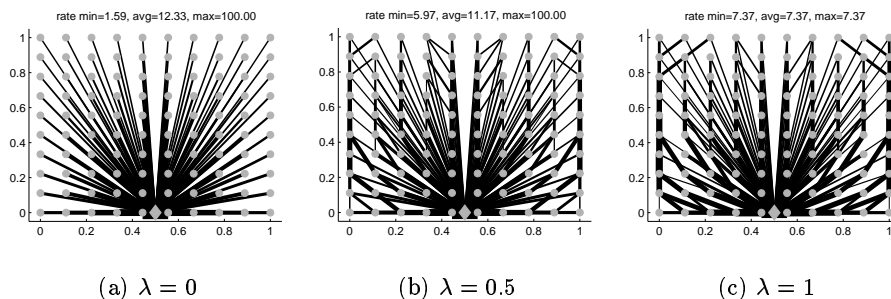
(a) $\lambda = 0$ (b) $\lambda = 0.5$ (c) $\lambda = 1$

Fig. 1. Sensor network with 100 sensor nodes placed in a grid: Optimal flow solutions for different values of the balancing factor λ. Line widths are proportional to the square root of the flow.
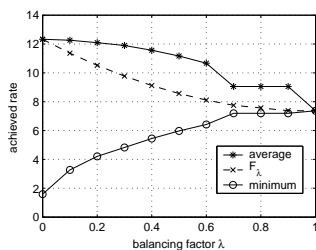

Fig. 2. Average, balanced and minimum data rates (bit/s), as a function of the balancing factor.
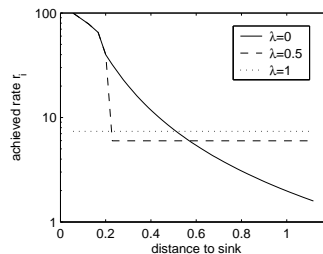

Fig. 3. Achieved data rates r_i for each individual sensor (logarithmic scale).

large amounts of data, and to save energy for the data from the more distant nodes. However, our approach of using a lower limit more directly models the intuitive goal of gathering enough data equally from all parts of the area to be monitored.

3.2 Random Node Placement

To illustrate random node placement, we made a second set of experiments; see Fig. 4. Network parameters are otherwise the same as before, but the 100 nodes were placed randomly in the square area, using uniform distribution. The effects of balancing are very similar to those described above for the regular grid.

3.3 Addition of Obstacles

Our LP formulation readily allows additional constraints to be included. For instance, we could limit the transmission powers of the nodes or the channel capacities of the links.
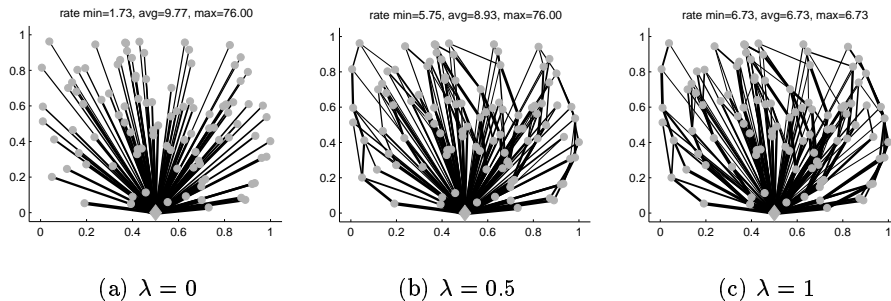
(a) $\lambda = 0$ (b) $\lambda = 0.5$ (c) $\lambda = 1$

Fig. 4. Sensor network with 100 randomly placed sensor nodes: Optimal flow solutions for different values of the balancing factor λ.

(a) Large, rectangular obstacles, $\lambda = 0.5$.

(b) Many small random obstructions: 50% of links obstructed, $\lambda = 0.5$.

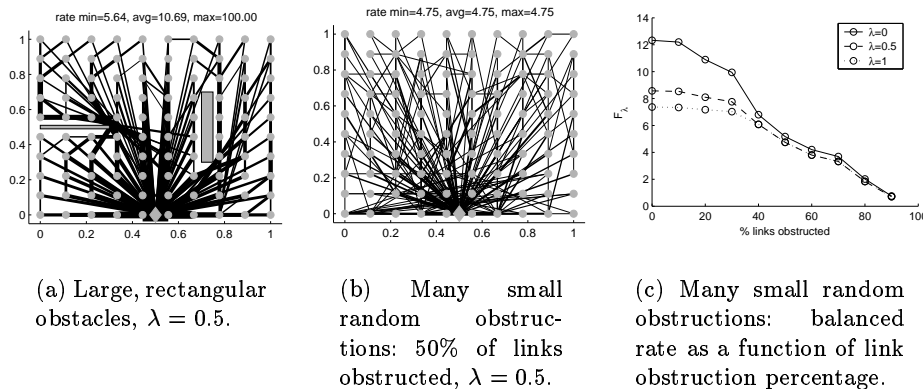(c) Many small random obstructions: balanced rate as a function of link obstruction percentage.

Fig. 5. Experiments with networks with nonuniform signal propagation.

Since our model allows arbitrary energy costs D_{ij}, it is not restricted to idealized transmission conditions like the popular unit disk model. We can study situations where there are large impenetrable obstacles within the area, as illustrated in Fig. 5(a). This is done by assigning an infinite energy cost to any link that intersects an obstacle.

In Figs. 5(b) and 5(c) we illustrate the effects of a different kind of nonuniform propagation. Instead of large obstacles, a random subset of all links is made unavailable by assigning them an infinite cost. This corresponds to small obstructions scattered throughout the area.

As expected, the achieved balanced rate F_λ decreases as more links are obstructed. However, the performance of the network degrades gracefully. As seen in the curve for $\lambda = 1$, even with half of all links randomly obstructed, the simu-

lated network is able to transport a minimum rate of 4.7 bit/s from every sensor node, or about 65 % of the minimum rate of 7.4 in the unobstructed network.

4 The Effect of Relay Nodes

The performance of the sensor network can be improved by augmenting the network by a number of auxiliary relay nodes. Unlike sensor nodes, whose locations are assumed to be predetermined, the locations of the relay nodes may be chosen to optimize the network performance. The relay nodes do not generate data themselves, but are solely used for forwarding data to other nodes in the network. Furthermore, the relay nodes may have considerably higher initial energy supplies than the sensor nodes.

In this section we consider the effect of relay nodes on network performance, and present and compare two simple techniques for determining good relay node locations.

4.1 Relay Node Placement Methods: Grid and Incremental Optimization

In the case of a square area, a straightforward method to place $m = k^2$ relay nodes is to position them in a regular $k \times k$ grid inside the square; see Fig. 7(a).

For a more sophisticated approach, one notes that in order to find an optimal placement for a set of relay nodes within a given sensor network, the locations of all the relay nodes should be considered at the same time. We, however, try to approximate the optimal solution by placing relay nodes into the network one at a time.

The algorithm performs a multidimensional search [2] in the following manner. Given a starting point y, a suitable direction d is first determined, and then the flow problem is optimized in this direction by performing a line search. Thereafter, a new direction d' is chosen and, again, the flow problem is optimized starting from the previous optimum in the direction d'. The process is repeated until a good enough solution is found, or the algorithm converges to a (possibly local) optimum.

In this case the optimizable quantity is the balanced data gathering measure F_λ, which can be calculated directly with the full LP model from Sect. 2, although requiring a large number of function evaluations. The optimal objective function value for the LP model with a given balancing factor λ and the considered relay node in location y is denoted $F_\lambda(y)$.

For our algorithm we have chosen as the starting point y^1 the center of mass of the differences of the offered data rates s_i and achieved data rates r_i:

$$y^1 = \left(\frac{\sum x^i{}_1 (s_i - r_i)}{\sum (s_i - r_i)}, \frac{\sum x^i{}_2 (s_i - r_i)}{\sum (s_i - r_i)} \right),$$

where $x^i = (x^i{}_1, x^i{}_2)$ are the coordinates of the sensor nodes in S.

<div style="border:1px solid">

Initialization Step Choose the number of iterations M and let $x^1 = (x^1{}_1, x^1{}_2)$ be the location of the sink node.

Find the present optimal achieved data rates r_i by solving the corresponding linear program.

Choose initial location for the new relay node as $y^1 = (\frac{\sum x^i{}_1(s_i - r_i)}{\sum s_i - r_i}, \frac{\sum x^i{}_2(s_i - r_i)}{\sum s_i - r_i})$, let j=1, and go to the main step.

Main Step Repeat M times.

1. Let $d^j = (y^j - x^1)$, let μ^j be a value that maximizes $F_\lambda(y^j + \mu d^j)$, and let $y^{j'} = y^j + \mu^j d^j$. Go to step 2.
2. Let $d^{j'} \perp d^j$, let $\mu^{j'}$ be a value that maximizes $F_\lambda(y^{j'} + \mu^{j'} d_j{}')$, and let $y^{j+1} = y^{j'} + \mu^{j'} d_j{}'$.

   Increment j by one.

</div>

Fig. 6. The incremental relay node placement algorithm for placing one relay node.

The idea is to place a new relay node initially in a region of the network where the achieved data rates r_i are small compared to the offered rates s_i. It is reasonable to think that the ideal location of the node would be, at least with high probability, somewhere between this region and the sink. Therefore the first search line is chosen in direction of the sink node. This idea is extended for determining the remaining search directions as the algorithm proceeds. The search directions are chosen pairwise: in the direction of the sink node and orthogonal to it. Line searches can, in principle, be performed by almost any standard one-dimensional search method, the main limiting factors being the complexity and possible roughness of the objective function F_λ. Fig. 6 summarizes our algorithm for finding a good location for a relay node.

4.2 Experimental Results with Relay Nodes

The objective of the experiments was to analyze both the impact of relay nodes on the balanced data gathering measure and the performance of our incremental relay node placement algorithm. We used the obstacle-free sensor network given in Sect. 3 with balancing factor $\lambda = 0.5$. Relay nodes were assigned 100 times the energy of the sensor nodes (i.e., 2 kJ). The performance of our incremental algorithm was compared to the placement in a regular grid. The incremental algorithm was run using a uniform line search with 20 equidistant line points for each direction and with three different direction pairs ($M = 3$). An example of $m = 9$ relay nodes placed with the incremental algorithm is presented in Fig. 7(b). As can be seen from the figure, the relay nodes form a routing backbone and the sensor nodes exhibit a clustering behavior around the relay nodes.

The results for different numbers of relay nodes are shown in Fig. 7(c). A clear improvement in network performance can be seen with an increasing number of relay nodes, even with relatively simple relay node placement schemes. If the cost of a relay node is not considerably higher than the cost of a sensor node, augmenting sensor networks having tight energy constraints by relay nodes is worthwhile.
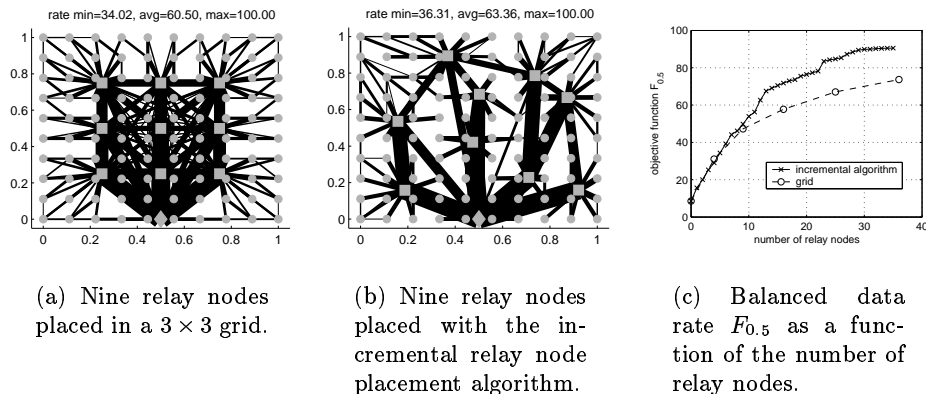
rate min=34.02, avg=60.50, max=100.00

rate min=36.31, avg=63.36, max=100.00

(a) Nine relay nodes placed in a 3×3 grid.

(b) Nine relay nodes placed with the incremental relay node placement algorithm.

(c) Balanced data rate $F_{0.5}$ as a function of the number of relay nodes.

Fig. 7. Experiments with 100 sensor nodes and relay nodes.

The incremental placement algorithm performs somewhat better than the grid placement algorithm, but is more demanding computationally. Unlike the incremental relay node placement algorithm, the straightforward grid placement of the relay nodes cannot be expected to perform as well for sensor networks where sensor nodes are placed arbitrarily, or where the area to be covered is irregular or partially obstructed.

5   Conclusions

We have considered the problem of energy-efficient data gathering in sensor networks, with special emphasis on the goal of balancing the average volume of data collected against sufficient coverage of the monitored area. We have formulated a linear programming model of the task of finding optimal routes for the data produced at the sensor nodes, given a balancing requirement in terms of a balancing factor $\lambda \in [0, 1]$.

Experiments with the model show that for reasonable values of the balancing factor, a significant increase in coverage is achieved, without any great decrease in the average amount of data gathered per node. In the examples considered in Sect. 3.1, for $\lambda = 0.5$, the minimum amount of data collected from any node was increased fourfold from the case $\lambda = 0$, with only about 10% decrease in the total volume of data gathered at the base station.

We have also considered the effects of augmenting a given network of sensors by a small number of auxiliary, freely positionable relay nodes with relatively high initial battery power levels. In the examples considered in Sect. 4.2, already an additional four relay nodes allocated in a network of 100 sensor nodes yielded a more than threefold increase in the value of the balanced data gathering objective function $F_{0.5}$; with nine relay nodes a fivefold increase was achieved. These improvements were obtained by a simple grid placement of the relay nodes; for

larger numbers of relay nodes better results were achieved by an incremental relay placement heuristic, applying techniques of multidimensional line search.

In our experiments with obstacles, sensor networks were seen to be fairly robust against even a fairly high number of obstructions. This was achieved through the use of global optimization at a central location, where information about all link costs in the network was available. It remains to be studied how closely this global optimum can be approximated by distributed algorithms that have access to local information only. The effects of possible node faults during the operation of the network are also a topic for further research.

Acknowledgement

References

1. I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, E. Cayirci, "Wireless sensor networks: a survey." *Computer Networks 38* (2002), 393–422.
2. M. S. Bazaraa, H. D. Sherali, C. M. Shetty, *Nonlinear Programming: Theory and Algorithms.* John Wiley & Sons, New York NY, 1993.
3. M. Bhardwaj, A. Chandrakasan, "Bounding the lifetime of sensor networks via optimal role assignment." *Proc. 21st Ann. Joint Conf. of the IEEE Computer and Communications Societies (InfoCom'02, New York NY, June 2002), Vol. 3*, 1587–1596. IEEE, New York NY, 2002.
4. M. Bhardwaj, T. Garnett, A. Chandrakasan, "Upper bounds on the lifetime of sensor networks." *Proc. IEEE Intl. Conf. on Communications (ICC'01, Helsinki, Finland, June 2001), Vol. 3*, 785–790. IEEE, New York NY, 2001.
5. A. Cerpa, J. Elson, D. Estrin, L. Girod, M. Hamilton, J. Zhao, "Habitat monitoring: application driver for wireless communication technology." *Proc. 2001 ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean (San José, Costa Rica, April 2001)*, 20–41. ACM Press, New York NY, 2001.
6. K. Chakrabarty, S. S. Iyengar, H. Qi, E. Cho, "Grid coverage for surveillance and target location in distributed sensor networks." *IEEE Transactions on Computers 51* (2002), 1448–1453.
7. C.-Y. Chong, S. P. Kumar, "Sensor networks: evolution, opportunities, and challenges." In [22], pp. 1247–1256.
8. M. Chu, H. Haussecker, F. Zhao, "Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks." *Intl. J. on High Performance Computing Applications 16* (2002), 293–313.
9. *Communications of the ACM 43* (2000) 5 (May). Special section on Embedding the Internet.
10. *Communications of the ACM 45* (2002) 12 (December). Special section on Issues and Challenges in Ubiquitous Computing.
11. K. Dasgupta, M. Kukreja, K. Kalpakis, "Topology-aware placement and role assignment for energy-efficient information gathering in sensor networks." *Proc. 8th IEEE Symp. on Computers and Communication (ISCC'03, Kemer-Antalya, Turkey, July 2003)*, 341–348. IEEE, New York NY, 2003.

12. S. S. Dhillon, K. Chakrabarty, S. S. Iyengar, "Sensor placement for grid coverage under imprecise detections." *Proc. 5th ISIF Intl. Conf. on Information Fusion (FUSION'02, Annapolis MD, July 2002)*, 1581–1587. International Society for Information Fusion, Chatillon, France, 2002.

13. D. Estrin, R. Govindan, J. Heidemann, S. Kumar, "Next century challenges: scalable coordination in sensor networks." *ACM/IEEE Intl. Conf. on Mobile Computing and Networks (MobiCom'99, Seattle WA, August 1999)*, 263–270. ACM Press, New York NY, 1999.

14. D. Estrin et al., *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers.* The National Academies Press, Washington DC, 2001.

15. W. R. Heinzelman, A. Chandrakasan, H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks." *Proc. 33rd Hawaii Intl. Conf. on System Sciences (HICSS'00, Maui HW, January 2000)*, 8020–8029. IEEE Computer Society, Los Alamitos CA, 2000.

16. A. Howard, M. J. Matarić, S. Sukhatme, "Mobile sensor network deployment using potential fields: a distributed, scalable solution to the area coverage problem." In H. Asama, T. Arai, T. Fukuda, T. Hasegawa (Eds.), *Distributed Autonomous Robotic Systems 5.* Springer-Verlag, Berlin, 2002.

17. K. Kalpakis, K. Dasgupta, P. Namjoshi, "Efficient algorithms for maximum lifetime data gathering and aggregation in wireless sensor networks." *Computer Networks 42* (2003), 697–716.

18. B. Krishnamachari, D. Estrin, S. Wicker, "The impact of data aggregation in wireless sensor networks." *Intl. Workshop on Distributed Event-Based Systems (DEBS'02, Vienna, Austria, July 2002)*, 575–578. IEEE, New York NY, 2002.

19. B. Krishnamachari, F. Ordóñez, "Analysis of energy-efficient, fair routing in wireless sensor networks through non-linear optimization." *Proc. IEEE 58th Vehicular Technology Conference (VTC 2003-Fall, Orlando FL, October 2003)*, Vol. 5, 2844–2848. IEEE Computer Society, Los Alamitos CA, 2003.

20. F. Ordóñez and B. Krishnamachari, "Optimal information extraction in energy-limited wireless sensor networks." To appear in *IEEE JSAC, special issue on Fundamental Performance Limits of Wireless Sensor Networks*, 2004.

21. J. Pan, Y. T. Hou, L. Cai, Y. Shi, S. X. Shen, "Topology control for wireless sensor networks." *Proc. 9th Ann. Intl. Conf. on Mobile Computing and Networking (MobiCom'03)*, 286–299. ACM Press, New York NY, 2003.

22. *Proceedings of the IEEE 91* (2003) 8 (August). Special Issue on Sensor Networks and Applications.

23. T. S. Rappaport, *Wireless Communications: Principles & Practice.* Prentice Hall, Upper Saddle River NJ, 1996.

24. N. Sadagopan, B. Krishnamachari, "Maximizing data extraction in energy-limited sensor networks." *Proc. 23rd Conf. of the IEEE Computer and Communications Society (InfoCom'04, Hong Kong, March 2004)*, to appear.