Distributed Algorithms for Transmission Power Control in Wireless Sensor Networks

Martin Kubisch, Holger Karl, Adam Wolisz
Telecommunication Networks Group
Technische Universität Berlin
Sekr. FT 5-2, Einsteinufer 25
10587 Berlin, Germany
`kubisch|karl|wolisz@ee.tu-berlin.de`

Lizhi Charlie Zhong, Jan Rabaey
Berkeley Wireless Research Center
University of California Berkeley
2108 Allston Way, Suite 200
Berkeley, CA 94704-1302
`czhong|jan@eecs.berkeley.edu`

***Abstract*— Two algorithm for dynamically adjusting transmission power level on per node basis have been evaluated using simulative approach. Network lifetime, convergence speed as well as resulting network connectivity have been obtained for these two algorithms under a particular indoor sensor environment. The network life metrics of these two are also benchmarked against power control algorithms using global information. We show that these two algorithms outperform fixed power level assignment and is generally within a factor of two of a globally computed solution.**

I. INTRODUCTION

Sensor networks [1] — networks of tiny nodes equipped with limited sensing, computing, and radio communication capabilities — is a technological vision that is currently receiving a lot of attention from several research communities. Such sensor networks would use wireless communication to transmit their observation values to a given monitor station which would serve as a user interface.

A joint characteristic of most application scenarios is that sensors only have a limited energy supply which might not even be rechargeable, hence they have to be as energy-efficient as possible. One option is to reduce transmission power with the help of intermediate nodes as relays in place of direct communication with a remote node. While such relaying has its own disadvantages (energy is now also consumed for intermediate reception and transmission), relaying can be beneficial for improving energy efficiency (e.g.,[2]).

Yet arbitrarily reducing transmission power is not possible; at least, some direct neighbors of a sensor node must be reachable to provide the possibilities to perform relaying and to form a connected network via relaying. Therefore it is important to find distributed algorithms which determine appropriate transmission power level for every node.

We present two distributed algorithms that determine an individual transmission power level for each node of a fixed, non-mobile sensor network.[1] This computation happens in an initialization phase, and these power levels are then used for the actual data communication within the network. At the end of the initialization phase, neighborhood information is used to compute a simple shortest-path routing based on Dijkstra's algorithm [3]. One particularly interesting aspect is the fact that power assignment algorithms can compute asymmetric communication relationships. Hence, the routing protocol must be able to handle such a situation.

As a figure of merit and optimization criterion, we use the time until the first sensor node runs out of energy (all nodes start with the same fixed amount of energy). Alternatively, either time to network partition could be used or the time at which a given area is no longer covered by at least a single sensor. In order to evaluate the performance of our local algorithms, we compare the network lifetime achieved with algorithms using global knowledge.

The remainder of this paper is organized as follows. In the following section, an overview of related work is given. Section III describes both local algorithms as well as the global algorithms which serve as comparison cases. Section IV outlines the simulation setup we used to study these algorithms and presents simulation results. Finally, Section V presents conclusions and directions for future work.

II. RELATED WORK

In recent publications the problem of power control was addressed while assuming information on the angle of reception [4] or the nodes knowledge of its location [5][6]. In these publications more complicated algorithms are used, but as they use additional information a comparison can only be made using the same metric.

As there are similar approaches to solve the problem of power control [7], [8] as in this work, they differ in that ELBATT et al. needs a separate and contention-free feedback channel and uses a cellular TDMA system. KRISHNAMACHARI et al. uses an algorithm with an exponential grow in control messages in the number of nodes.

In [9] RAMANATHAN and ROSALES-HAIN provided the idea for the first algorithm used in this paper, but they used the algorithm to examine network throughput and delay in a two dimensional space with a smaller set of nodes.

---

[1] The same power is used when sending to any neighbor, irrespective of whether some neighbors are closer than others. One reason to do so is the time and hence energy expenditure that is require to set the amplifier to different power levels. Algorithmically, the use of different power levels depending on the intended receivers is easily possible.

III. PROBLEM SOLUTIONS

In the following we concentrate on five different approaches. Either applicable to be used as distributed algorithm to extend the network lifetime or as compare cases to show the maximum achievable.

A. Threshold in number of neighbors

The "local mean algorithm" (LMA) works in the following way: All nodes start with the same initial transmission power ($\mathrm{TransPwr}$). Every node periodically broadcasts a life message ($\mathrm{LifeMsg}$) including its unique identity. All the other nodes, which receives such a $\mathrm{LifeMsg}$, reply with a life acknowledge message ($\mathrm{LifeAckMsg}$) including the address of the $\mathrm{LifeMsg}$ sender. In order to avoid collisions, a proper MAC layer is assumed. Before a node issues the next $\mathrm{LifeMsg}$ it counts the number of $\mathrm{LifeAckMsg}$s received ($\mathrm{NodeResp}$). If $\mathrm{NodeResp}$ is less than a minimum threshold ($\mathrm{NodeMinThresh}$), the node increases its transmission power by adding a certain amount (A_{inc}) for every missing neighbor, but having a upper bound (B_{up}).[2] If $\mathrm{NodeResp}$ is larger than a maximum threshold ($\mathrm{NodeMaxThresh}$), it decreases its transmission power by subtracting a certain amount (A_{dec}) for every supernumerary neighbor, but having a lower bound (B_{down}).[3] If $\mathrm{NodeResp}$ is between $\mathrm{NodeMinThresh}$ and $\mathrm{NodeMaxThresh}$ the node does no longer change its transmission power; it has converged. While this algorithm has a periodic nature, it is important to note that no close synchronization of nodes or global time base is required; the notion of periodic is a purely local one.

B. Threshold in mean number of neighbors

The "local mean of neighbors algorithm" (LMN) works similar to LMA except that it adds some information to the $\mathrm{LifeAckMsg}$ and it defines $\mathrm{NodeResp}$ in a different way. In addition to the address from the received $\mathrm{LifeMsg}$, the $\mathrm{LifeAckMsg}$ also contains its own number of neighbors computed in the previous cycle. The node receiving the $\mathrm{LifeAckMsg}$s calculates a mean value from its neighbors' number of neighbors — the new $\mathrm{NodeResp}$.

C. Fixed transmission power

The most simple algorithm is to assign an arbitrarily chosen transmission power level to all sensor nodes, much like it would be done at production time for sensors that do not have power control at all.

D. Global solution with equal transmission power

The Equal Transmission Power (ETP) algorithm also assigns a uniform power to all nodes, but chooses the minimal value that ensures a fully connected network.

To find the minimum transmission power, the following algorithm is used:

1) Among the node pairs that are not yet connected, choose the one with the smallest distance.
2) Set transmission power of all nodes to a value sufficient to connect these two nodes.
3) Check connectivity of the resulting network and when the network is connected, minimum power level is found.

This power value represents the optimal value for a sensor network with fixed transmission range and therewith produces symmetric communication links. This algorithm uses global information and it is not evident how to implement a corresponding local algorithm that achieves the same results.

E. Global solution with diverse transmission power

The global solution with Diverse Transmission Power (DTP) algorithm creates a connected network but does not set all transmission ranges to the same value. Instead it tries to find a minimum power level for every node individually. The algorithm works in the following way:

1) Among the node pairs that are not yet connected, choose the one with the smallest distance.
2) Set transmission power of these two nodes to a value sufficient to connect them.
3) Check connectivity of the resulting network and when the network is connected, the minimum power level is found.

This algorithm minimizes the overall transmission power consumption for the entire network, but it may result in asymmetric communication links.

Even as it is possible to construct networks where this algorithm does not find minimum power levels for all nodes, DTP vastly outperforms any other global algorithm that we have considered, therefore, we use DTP as a comparison case. Similarly to the ETP algorithm, this algorithm also uses global knowledge, and equivalent local implementations are not obvious.

IV. SIMULATION RESULTS

A. Investigation scenario

In order to evaluate and compare these algorithms, we simulated the energy consumption and resulting network lifetime for a particular indoor sensor network scenario. The simulator used for this task was written using the OMNeT++ [10] simulation tool.

The physical layout consists of four rooms connected by a hallway as shown in Figure 1. All rooms are 3.5 m high, the grey bars indicate doors (1 m wide), the black dots show the position of two monitor nodes, which are positioned 1.2 m above the ground. Walls are assumed to be infinitesimally thin, constituting no obstacle for radio communication. The path loss coefficient is set to 2. Based on this layout, 32 different placements of sensor nodes were generated by placing 318 nodes randomly on the walls, ceilings, and floors (using a uniform node distribution), with the exception of the doors, which are assumed to reach up to the ceiling.

---

[2] Formally: $\mathrm{TransPwr} \longleftarrow \min\{B_{\mathrm{up}} \cdot \mathrm{TransPwr}, A_{\mathrm{inc}} \cdot (\mathrm{NodeMinThresh} - \mathrm{NodeResp}) \cdot \mathrm{TransPwr}\}$.

[3] Formally: $\mathrm{TransPwr} \longleftarrow \max\{B_{\mathrm{down}} \cdot \mathrm{TransPwr}, A_{\mathrm{dec}} \cdot (1 - (\mathrm{NodeResp} - \mathrm{NodeMaxThresh}) \cdot \mathrm{TransPwr}\}$.

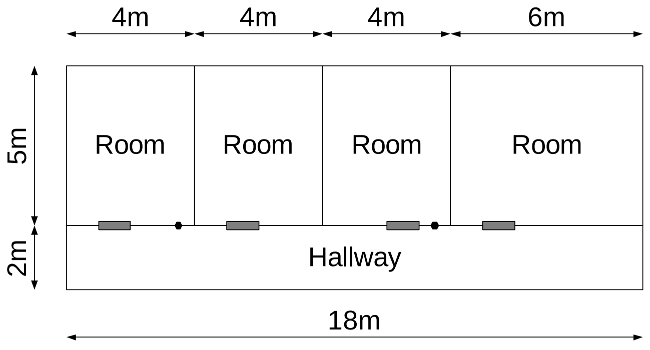

Fig. 1. Physical room layout. Grey bars represent doors, black circles represent monitor nodes.

For each of these placements, every algorithm computes the transmission power level assignments for each sensor node, and afterwards the routing tables are computed for all nodes. For the ETP and DTP algorithms, this results in a single configuration of power level assignments.The final transmission power level for the local algorithms depends on the initial transmission power to be used by each node. The initial power value is varied linearly in 56 steps from a transmission range of 25 mm to 1.4 m. For each of these initial power levels, the initialization phase is computed using 100 cycles.

The local algorithms use a transmission power increase value (A_{inc}) of 10% with an upper bound (B_{up}) of two times the old transmission power and a decay value (A_{dec}) of 2% with an lower bound (B_{down}) of half the old value (with higher values the implementation tends to oscillate). Based upon these configurations, the data communication phase is simulated and the results for all of these varied initial conditions are averaged to form a single value.

Note that the energy consumption during this initialization phase is not taken into account. This is due to the fact that distribution of routing information is unlikely to happen this way in a real sensor network. Information needed is also available in other, sensor-network-related protocols [11], thus it can be incorporated and energy consumption is negligible.

After this initialization, the data communication phase is simulated. The traffic in the network consists of request/replies initiated by the monitor. These requests are directed at a randomly chosen sensor node (the monitor is assumed to have sufficient information about the sensors). The nodes receive the request and answer with an (arbitrary) reply value. These requests are sent every half a second, alternating between the two monitor nodes.

These simulation runs result in a total of 32 network lifetime samples for the global algorithms and 32×56 samples for the local algorithms; network lifetimes, confidence levels for the average lifetime and comparisons are shown in Section IV-C. But first, the question of convergence speed of the local algorithms as well as whether these algorithms actually reach a fully connected network is interesting.

During the data communication phase, energy is consumed for both transmitting and receiving data packets as well as for idle phases. The power consumption during the idle phase is 0.1 μW, for receiving 0.5 mW, for sending 1 mW;[4] a sensor node's initial energy supply is 100 J. At an assumed bit rate of 10 kbit/s, these values correspond to 1 μJ and 0.5 μJ to send and receive a bit, respectively. The transmission power levels are set by the above algorithms such that transmission errors only happen with negligible probability (more precisely, a node is only considered to be reachable if the SNR is at least -90 dBm), hence transmission errors are not considered.

B. Convergence and connectivity

The local algorithms LMA and LMN use a number of iterations before settling down to particular transmission power levels. An individual node has converged when its number of neighbors is between $\mathrm{NodeMinThresh}$ and $\mathrm{NodeMaxThresh}$. The values for these thresholds were determined using some preliminary experiments with the global algorithm described in Section III-E: using this algorithm, the number of neighbors of most nodes is between four and seven; these values were therefore used as thresholds for the local algorithms.

A desirable property of such an algorithm would be that all nodes converge very quickly. Evidently, the initial transmission power (equivalent to the initial range of transmission) plays an important role for these algorithms. Figure 2 shows that most nodes converge within a small number of cycles and some few nodes take a medium number of cycles.

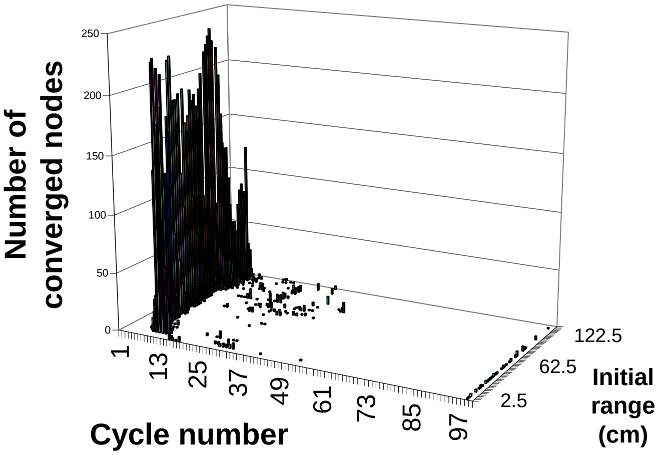

Fig. 2. Histogram of time to convergence (in cycles) depending on intitial transmission range (in cm)

For certain initial power levels (particularly, at small initial power), it also happens that a very small number

---

[4] The ratio of this values might be surprising, but the assumption is that a sending node can wake up nodes in its vicinity by using a low-bandwidth signaling channel that is easy to demodulate even for a low-power receiver, e.g. the ”Frisbee” model in [12]

of nodes do not converge at all. As an example, consider a case where a single node is located far away from a cluster of nodes that quickly form a connected group.

While a larger value of NodeMaxThresh increases the likelihood of a connected network, it also results in larger transmission power levels.

Figure 3 shows the percentage of nodes that were not reachable from the monitor stations for the local algorithms from Sections III-A and III-B. The largest value is 1.65% for LMA and 0.04% for LMN, and on the average, 0.003% are not connected considering the mean number of neighbors and 0.37% when only taking into account the number of direct neighbors.[5] These results suggest that LMN creates a much stronger connected network than LMA. Hence, distributing and taking into account the number of neighbors a node has, results in a much stronger connected network.

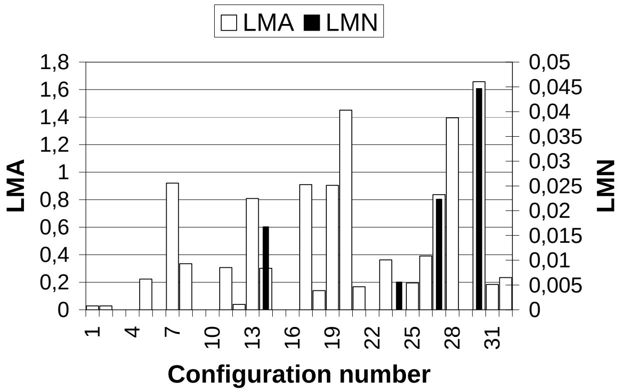

Fig. 3. Percentage of nodes not connected for different network configurations

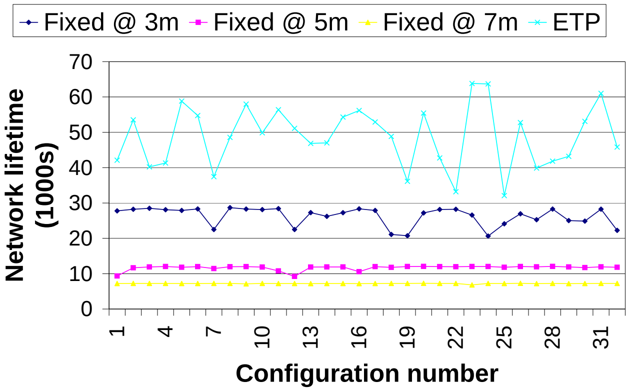

Fig. 4. Network lifetimes for different configurations for fixed transmission power and equal transmission power assignments

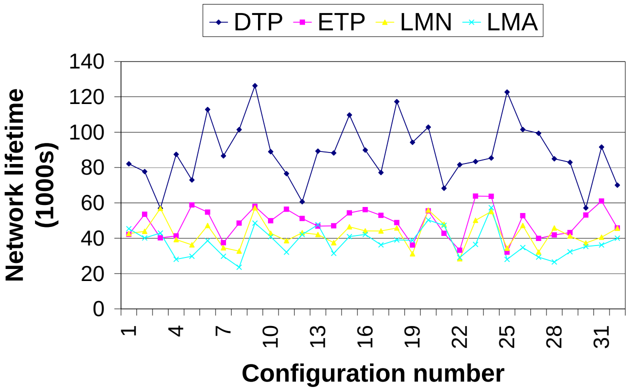

Fig. 5. Network lifetimes (in thousands of seconds) for different configurations for global and local algorithms

C. *Network lifetime*

The most interesting performance metric for such a sensor network is the network lifetime: the longer every single node is capable to communicate, the better the transmission power levels were chosen. In order to give a rough idea of possible network lifetimes achievable in our configuration, Figure 4 shows the results for the simple fixed assignments of transmission powers as well as the case for the global ETP algorithm. As expected, ETP outperforms the fixed value algorithms and achieves an average network lifetime of about 48800 seconds, while the fixed network algorithms' achievements are considerably below that. It is also interesting to see that the lifetime of the network depends heavily on the actual network configuration; differences are up to 50%.

More interesting is the comparison between the global and the local algorithms in Figure 5.

It comes as no surprise that DTP vastly outperforms all other algorithms with its global knowledge. As a general impression, the lifetime achieved by DTP up to 209% longer than that obtained by LMN and even higher for LMA. On average, DTP achieves network lifetimes that are about twice as long as LMA, LMN, and ETP. Figure 5 also suggests that the local algorithms and the ETP algorithm perform quite similarly.

Figure 6 shows the confidence intervals.

Applying a simple graphical interpretation, we can infer that the ETP outperforms both local algorithms. This comes by no surprise as ETP uses global information, but it is important to note that LMN does not only create a much stronger connected network, it also performs in mean by 14% better than LMA.

V. CONCLUSIONS

We can state that using heuristics which consider the number of neighbors a node has, result in a sufficiently connected network, provide improvements in network lifetime and is in the range of symmetric algorithms using perfect knowledge. While they are not able to outperform sophisticated algorithms, they perform usually within a factor of two of their lifetime.

---

[5] Note that these numbers are averaged over different initial transmission power levels, hence percentages do not correspond to an integer number of nodes.

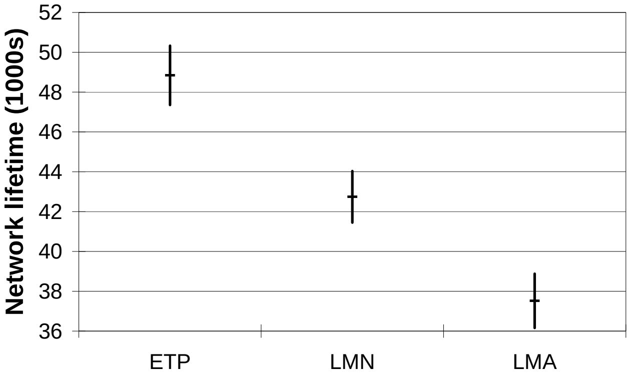

Fig. 6. Confidence intervals of network lifetimes for local algorithms and ETP; 95% confidence level

Additionally, these algorithms are structured similarly to other mechanisms that are conjectured to be deployed in sensor networks, e.g., locationing mechanisms. It would hence be possible to integrate these algorithms and to amortize their joint resource consumption.

A number of interesting questions remain for future work, e.g., to use the number of neighbors in the announcement messages and to weight this information against the transmission power with which it was sent or how a non-reliable MAC influences the algorithms. We intend to investigate these areas in the near future.

REFERENCES

[1] D. Estrin, R. Govindan, J. S. Heidemann, and S. Kumar, “Next century challenges: Scalable coordination in sensor networks,” in *Proc. 5th Ann. Intl. Conf. on Mobile Computing and Networking*. Seattle, WA: ACM, Aug. 1999, pp. 263–270.

[2] J. P. Monks, J.-P. Ebert, A. Wolisz, and W. mei W. Hwu, “A study of the energy saving and capacity improvement potential of power control in multi-hop wireless networks,” in *Workshop on Wireless Local Networks, Tampa, Florida, USA, also Conf. of Local Computer Networks (LCN)*, Nov. 2001.

[3] E. Dijkstra, “A note on two problems in connection with graphs,” *Numerical Mathematics*, vol. 1, pp. 269–271, 1959.

[4] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang, “Distributed topology control for power efficient operation in multihop wireless ad hoc networks,” in *vol. 3*. Anchorage: IEEE, Apr. 2001.

[5] V. Rodoplu and T. H.-Y. Meng, “Minimum energy mobile wireless networks,” *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 8, pp. 1333–1344, Aug. 1999.

[6] C. Savarese, J. M. Rabaey, and J. Beutel, “Locationing in distributed ad-hoc wireless sensor networks,” in *Proc. International Conference on Acoustics, Speech, and Signal Processing*, 2001.

[7] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao, “Power management for throughput enhancement in wireless ad-hoc networks,” in *ICC 2000*. New Orleans, LA: IEEE, June 2000.

[8] B. Krishnamachari, R. Bejar, and S. B. Wicker, “Distributed constraint satisfaction and the bounds on resource allocation in wireless networks,” in *Sixth International Symposium on Communications Theory and Application*. Ambleside, UK: ISCTA, July 2001.

[9] R. Ramanathan and R. Rosales-Hain, “Topology control of multihop wireless networks using transmit power adjustement,” in *Proc. IEEE Infocom*, Tel-Aviv, Israel, Mar. 2000.

[10] A. Varga, *OMNeT++: Discrete Event Simulation System*, Technical University of Budapest, Faculty of Electrical Engineering and Informatics, Mar. 2001, http://www.hit.bme.hu/phd/vargaa/omnetpp.htm.

[11] J. Li, J. Jannotti, D. De Couto, D. Karger, and R. Morris, “A scalable location service for geographic ad-hoc routing,” in *Proc. of the 6th ACM International Conference on Mobile Computing and Networking (MobiCom 20000)*, Aug. 2000, pp. 120–130.

[12] A. Cerpa, J. Elson, D. Estrin, L. Girod, M. Hamilton, and J. Zhao, “Habitat monitoring: Application driver for wireless communications technology,” in *Proc. ACM SIGCOMM Workshop on Data Communications*, Latin America and the Caribbean, Apr. 2001.