Experimental Evaluation of Lifetime Bounds for Wireless Sensor Networks

Hartmut Ritter
and Jochen Schiller
Freie Universität Berlin
Berlin, Germany
Email: {hritter, schiller}@inf.fu-berlin.de

Thiemo Voigt
and Adam Dunkels
and Juan Alonso
Swedish Institute of Computer Science
Stockholm, Sweden
Email: {thiemo, adam, alonso}@sics.se

Abstract— **In this paper we present a method for experimental lifetime measurements of sensor networks. Despite the importance of experimental validation, none of the lifetime models proposed so far has been validated experimentally. One of the reasons for the absence of practical validations might be the long lifetime of batteries which make the validation of the proposed models non-trivial and time consuming. Our solution enables validation of lifetime models within a reasonable amount of time. We also use our method to validate a simple mathematical model that provides bounds on the lifetime of sensor networks.**

I. INTRODUCTION

Wireless sensor networks have received a lot of attention from the research community during the last years. While most of the work has focused on more theoretical issues, there are also some deployments such as the sensor network on Great Duck Island [18]. Deploying real sensor networks is a challenging task both because of the dependency on environmental factors and because of the computation and memory constraints of the sensor nodes. A further challenge are the power constraints of the nodes in particular for unattended sensor networks. Most sensor nodes are battery-driven and once the battery of a node has depleted the node becomes useless. In multi-hop networks the connectivity of the network breaks if critical nodes run out of battery. Therefore, the lifetime of networks is finite. Hence, in order to estimate if a network can fulfill its task, it is important to be able to predict the longevity of the network before deploying it.

To predict the longevity of a sensor network, mathematical models providing upper bounds on the lifetime of sensor networks have been developed [3], [4], [7], [8], [13]. However, to the best of our knowledge, none of these models has been validated experimentally using real sensor nodes. One of the reasons for this might be the long lifetime of the batteries which makes the comparison of theory and practice of the models a tedious task that might take weeks or months. Towards this end, we present a hardware methodology using special capacitors with very large capacities, so-called GoldCaps, that enables short-term experiments with a duration no longer than a few hours.

Using GoldCaps on the ESB nodes developed at FU Berlin [2], we measure the longevity of some small sensor networks and compare these results to the results of a mathematical model developed by some of us [1]. This model computes bounds on the lifetime of continuous sensor networks. It concentrates on the energy consumption of routings since communication is the most expensive sensor node activity in terms of energy, generally far more expensive than sensing and computing [17]. Our results show that the measured lifetime of the sensor networks agrees well with the lifetime predictions by the mathematical model.

The main contribution of this paper is a hardware approach that enables experimental evaluation of lifetime boundaries of WSNs. We demonstrate the feasibility of the hardware approach by validating a simple mathematical model for lifetime bounds of sensor networks presented earlier by Alonso et al. [1].

The remainder of this paper is outlined as follows: First, we present our methodology for experimental evaluation of lifetime boundaries for sensor networks in the next section. In Section III we present the mathematical model for bounding the lifetime of sensor networks. The following section compares the longevity results of the mathematical model to a real world estimation using the methodology presented in Section II. After discussing some related work in Section V we conclude with some final remarks and future work.

II. A METHOD FOR MEASURING THE LIFETIME OF SENSOR NETWORKS

While calculating the lifetime of sensor networks using the data sheets of all involved parts is at least in theory feasible, only measurements make sure that all influences and side effects are taken into account. However, in the special case of sensor networks where nodes are small and distributed, instrumentation of lifetime measurements of all nodes is not trivial. Simply equipping all nodes with replaceable (primary) batteries or rechargeable (secondary) batteries and measuring the lifetime of each node will give only a very rough picture. The reasons are twofold: First, the capacities even of freshly charged secondary batteries vary heavily, especially if some are new and some were used during many charging cycles. Second, it is really impractical to use batteries or rechargeable batteries as with energy-saving sensor nodes, lifetimes can be weeks to months or even years. For short term experiments and comparisons this is problematic, as for technical reasons the remaining capacity of a battery cannot be measured with sufficient precision.

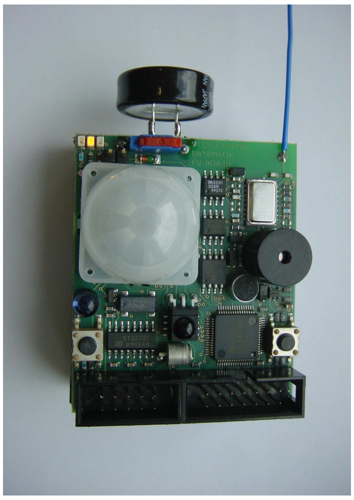

Fig. 1. Sensor board with GoldCap attached on top

A solution of this problem is the usage of so-called GoldCaps (sometimes also called SuperCaps). GoldCaps are a special kind of capacitors that come with very high capacity, e.g. 1 Farad (1F, as opposed to some µF or nF typical for standard capacitors). These devices can be charged very quickly and power a sensor node for a reasonable amount of time. In our case, using some simple energy saving mechanisms we can run a node sending data packets every 10 seconds for about 70 minutes powered only by a single 1F GoldCap. Such a node with attached GoldCap can be seen in Figure 1. Obviously, the node's lifetime heavily depends on the energy consumption of the node and the energy saving schemes. In general, the time t that a sensor node can be powered with a GoldCap of capacity C follows

$$t = C(U_{max} - U_{min})/I, \tag{1}$$

where U_{max} is the max voltage of GoldCap, U_{min} is the threshold voltage (no operation below U_{min} possible), and I the current consumption of the sensor node. For example, with a capacity of 1F, U_{max} = 4V, U_{min} = 2V and I = 250µA in low-power mode this results in 1h 13min operation.

To validate the reliability of our solution, we experimentally verified the equality of the lifetime of a number of GoldCap-equipped sensor nodes. The sensor nodes were running with the radio turned on, all sensors enabled, and were continuously transmitting data over the RS232 port. We defined that a node was alive as long as it was able to transmit data over the serial link. Our experiments showed that the difference between the shortest and longest lifetime over a number of runs with different sensors was only slightly above 5%.

The main advantage of this hardware methodology is that it enables short measurement runs. The proposed methodology is useful beyond lifetime analysis of sensor networks. It enables experimental evaluation and comparison of the energy-efficiency of all kind of communication protocols used in sensor networks, ranging from modulation and MAC layer schemes up to implementation properties of the application layer.

Furthermore, using GoldCaps makes sense far beyond experimental setups. Together with solar cells they allow completely autonomous operation of sensor nodes, as the GoldCap can be charged during daylight time and provide sufficient energy during the night. With some more advanced power saving and routing mechanisms [20] we can run nodes in our lab during the night completely out of the GoldCap. Examples for commercial use of GoldCaps on sensor nodes are the nodes of Enocean [11].

III. A MATHEMATICAL MODEL OF THE ENERGY CONSUMPTION OF ROUTINGS

In this section we will summarize the important aspects of a mathematical model for the energy consumption of routings and lifetime boundaries of sensor networks presented earlier by Alonso et al. [1]. The original paper provides both upper and lower bounds of the energy consumptions of routings. In this paper, we concentrate on the lower bounds of the energy consumption of routings and provide an upper bound of the lifetime of a sensor network. For more details and the formal proofs, see [1].

A. *The considered network type*

Our mathematical model considers continuous sensor networks [19] in which "the sensors communicate their data continuously at a pre-specified rate". An example of such a network is the network at Great Duck Island [18]. In such a network sensor nodes read sensor values, send them in a multi-hop fashion to a base station and go to sleep until the next iteration. Sending values in this procedure includes forwarding packets originated from other hosts, except for leaf nodes that only transmit their sensor value. This process is iterated until nodes run out of energy and connectivity is broken.

B. *Energy consumption during one iteration*

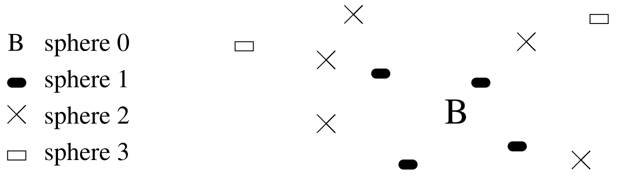

Fig. 2. A sensor network partitioned in spheres

One of the key ideas of the model is to partition the set of all sensor nodes V into subsets S_0, ..., S_n satisfying $V = S_0 \cup S_1 \cup .. \cup S_n$, $S_i \cap S_j = \emptyset$ for all $i \neq j$ and no S_i is empty. S_i is the set of nodes that can be reached from the base station B in i hops (thus $S_0 = \{B\}$), but not less than i hops. We call S_i the *sphere* of radius i around S_0. Figure 2 shows an example network. Note that the current version of our model assumes that all nodes transmit at the same constant power. Further our model implies that no data aggregation is done and gathered data is transmitted unchanged to the base station. This means, that each node in sphere S_n, the sphere containing the leaf nodes, transmits exactly one packet in each iteration. A node in sphere S_{n-1} transmits the packets it receives from leaf nodes in sphere S_n plus one packet with its own sensor value.

Corresponding to the spheres S, we introduce *balls* of radius i denoted B_i, with $B_i = S_0 \cup .. \cup S_i$. Further, we introduce $s_i = |S_i|$, $b_i = |B_i|$, $N = |V|$, r as the energy consumption for receiving one packet and t as the energy required to transmit one packet. Using these definitions, we set

$$m_i = \frac{N - b_i}{s_i} r + \frac{N - b_i + s_i}{s_i} t. \qquad (2)$$

In the equation above, $N - b_i$ is the total number of nodes outside B_i and thus the total number of packets that the set of nodes in sphere S_i receive in each iteration. The nodes in S_i must forward all the packets they receive from the outer spheres plus s_i packets containing the sensor values of the nodes in S_i. The best the routing algorithm can do is to balance the energy consumption for receiving and transmitting packets across all the nodes in S_i, therefore the denominator s_i. Thus, m_i provides a lower bound on the energy consumption (for receiving and transmitting packets) for the node in S_i that consumes the most energy of all nodes in this sphere during one iteration.

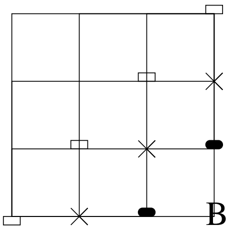

Fig. 3. A sensor network where the nodes one hop from the base station die first

For many sensor networks, $\max\{m_1, ..., m_n\}$ will be equal to m_1, i.e. the node that consumes most energy during one iteration is one hop away from the base station. An example of such a network is shown in Figure 3. In this network the nodes that are one hop away from the base station consume most energy during one iteration. In this case, we call S_1 the bottleneck sphere.

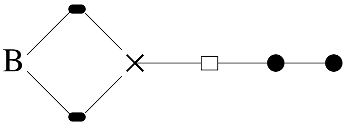

Fig. 4. A sensor network where the node two hops from base station dies first

But there exist networks where $\max\{m_1, ..., m_n\}$ is not equal to m_1, see for example Figure 4. In this example, $m_1 = 2r + 3t$ while $m_2 = 3r + 4t$. The routing achieving this m_2 distributes the packets sent by the node in S_2 evenly among the nodes in S_1. Hence $\max\{m_1, ..., m_n\} = m_2$ meaning that S_2 is the bottleneck sphere.

C. Bounding network lifetime

For most networks, the energy consumption m^T of the nodes in the bottleneck sphere for T iterations is $T \max\{m_1, ..., m_n\} = Tm_i$. In these cases, the traffic can be balanced evenly between the nodes in the bottleneck sphere. Hence, all nodes in the bottleneck sphere run out of energy during the same iteration, breaking connectivity. An example of such a network is the one in Figure 3. Note that this assumes that the routing is optimal since as discussed above m_i is a lower bound.

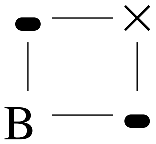

Fig. 5. A sensor network where $m_1^T \neq Tm_1$

Figure 5 presents a network where $m_1^T \neq Tm_1$. In this network, $m_1 = \frac{1}{2}r + \frac{3}{2}t$ and $m_2 = t$. However, the packet from the node in sphere S_2 cannot be split into two small packets. Thus, when the number of iterations T is odd, then one of the nodes in S_1 must forward one more packet than the other. Hence, $m_1^T \neq Tm_1$, namely $m_1^T = \frac{T+1}{2}r + \frac{3T+1}{2}t > \frac{T}{2}r + \frac{3T}{2}t = Tm_1$.

Suppose each node has the exact same amount of energy EE. Then, from the discussions above, it is obvious that the maximum number of iterations T_{max} a sensor network can perform before running out of energy under the given assumptions is bounded by the following expression:

$$T_{max} \leq \frac{EE}{\max\{m_1, ..., m_n\}}$$

This means, that whatever routing we use, the sensor network cannot perform more than T_{max} iterations before connectivity breaks.

D. Discussion

We have chosen to validate this model since it is simple but still quite general. The model and thus its results can be applied regardless of topology, routing algorithm or radio energy model. While we have chosen the mathematical model in its simplest form, it is also easy to extend the model to, for example, include the energy consumption for taking sensor readings. To include sensor readings, Equation 2 just needs another term $s_i g$ where g is the energy required for a sensor reading. Furthermore, integrating data aggregation is not difficult. Assume that the data can be aggregated with a factor f, indicating that $1/f$ packets can be merged into a single packet. Then a sphere S_i receives $(N - b_i)f$ packets and will transmit $(N - b_i + s_i)f$ packets in each iteration. Further, the earlier constant values r and t should be replaced by r_i and t_i due to increased size of the data packets when data is aggregated. Equation 2 must be rewritten as

$$m_i = \frac{(N - b_i)f}{s_i} r_i + \frac{(N - b_i + s_i)f}{s_i} t_i.$$

Despite its importance on energy consumption, we do not consider data aggregation here since it is highly application dependent.

IV. PRACTICAL EVALUATION OF THE MATHEMATICAL MODEL

In this section we present a practical evaluation and validation of the mathematical model presented in the previous section. For this part of our work we used the ESB nodes developed at FU Berlin [2] together with the GoldCaps as described in Section II.

A. Evaluation Setup

The mathematical model described in this paper focuses on the energy consumption of a readily set up wireless sensor network. It does not make any assumption about how a route is found but instead provides an upper bound of the lifetime that is based on the use of an optimal routing strategy. Therefore we implemented static routing based on the topologies described in the previous section. Each node is provided at startup with a routing table that contains its downstream neighbour.

As we wanted to measure the (topology dependent) impact of packet transmission and reception on the lifetime of the sensor network we kept the nodes sleeping all the time except for the very short packet transmission and receive cycles. This can be done using a time slotted system where nodes need not to be kept idle listening (note that listening nodes consume nearly as much energy as transmitting nodes). The fixed routing table provided at startup determines the send and receive slot for each node. This is done by directly relating the node ID to a transmission and reception slot. As our experiments were done using up to 10 nodes, we structured the time slots into iterations of 10 seconds. A node with ID 5 for example only listens for incoming packets at the beginning of the fifth second in each iteration of 10 seconds. As the node knows the ID of its downstream neighbour it can exactly meet the listening interval of this downstream node. Note that using a time-slotted architecture collisions are completely avoided and therefore no carrier sense or random backoff mechanisms have to be used. The resulting communication pattern minimizes the uptime of the nodes and proves to be implementable with COTS hardware.

Though in another context we successfully realized a temperature surveillance scenario using this communication pattern, for the evaluation of the mathematical model we do not transmit any real sensor values as the energy consumption of specific sensors is not of general interest. The size of the transmitted packets in our evaluation is 21 bytes (including source address, destination address and error check bytes), a quite reasonable size for the transmission of simple sensor values.

The measurements are performed in the following way: We charge all GoldCaps to their full capacity, switch all nodes to the GoldCaps as only power source and then send a broadcast synchronization signal using an additional monitoring node. After that, the time-slotted communication starts as described. The monitoring node logs all sent packets together with the packet source and destination.

B. Evaluated Topologies

We evaluate the two basic topologies given in Figure 3 and Figure 4. The routing scheme for these topologies can be seen in the Figures 6 and 7. The number of connecting lines between peer nodes denotes the number of packets that are sent in each iteration (as mentioned we do not use data aggregation schemes). The routing scheme in Figure 6 is an optimal routing scheme that we have chosen to implement for this topology. The routing scheme used in Figure 7 divides the traffic at node 4 in two paths. The two paths 4-5 and 4-9 are interleaved iteration by iteration, meaning that in the first iteration node 4 sends all 4 packets to node 5, in the next iteration it sends all 4 packets to node 9 and so on. Nodes 5 and 9 send one packet every iteration, and every second iteration four additional packets. Another option would be to let node 4 send 2 packets to node 5 and 2 packets to node 9 in each iteration. We implemented this option as well with results similar to those presented here.

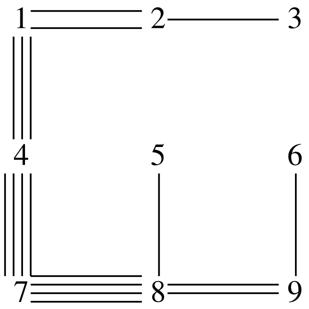

Fig. 6. Routing scheme for Figure 3

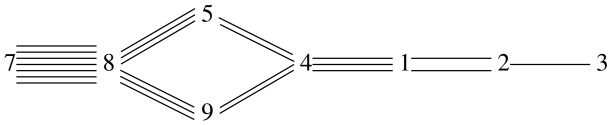

Fig. 7. Routing scheme for Figure 4

C. Lifetime Calculations

In order to compare the calculated and the measured lifetime of the sensor network we compare the number of packets that an individual node sends until it runs out of energy.

The mathematical model concentrates on the energy consumption of sending and receiving. Due to that reason we keep the nodes sleeping immediately before respectively after sending or receiving packets. Nevertheless as a sleeping node also consumes some amount of energy this has to be taken into account. Thus the overall energy budget (measured in mAs) of a node per iteration is given by

$$E = n_t t + n_r r + s, \tag{3}$$

where t is the energy consumption for transmitting one packet, n_t the number of sent packets, r as the energy consumption for receiving one packet, n_r the number of received packets, and s the energy consumption during the sleep phases. The energy consumption for sending a packet is the product of transmission time and current consumption during packet transmission. Similar the energy consumption for receiving a packet is the product of reception time and current consumption during packet reception.

The values that have to be filled into this equation result from oscilloscope measurements with our ESB nodes: The current consumption during sending and receiving averages to 7.2 mA, the current consumption during sleep mode to 270 µA. The transmission and reception times differ from the times that could be calculated using packet size and send rate. The reason for this difference is that the transceiver needs some start-up time and that for packet reception the receiving node's transceiver should take into account minimal de-synchronization and thus be switched on some ms before the local clock indicates the start of the receive slot. Thus, we fill in 20 ms for transmission time and 30 ms for reception (as opposed to the theoretical value of 8.75 ms for a packet of 21 bytes size at 19.2 kbit/s). Note that our oscilloscope measurements with the given hardware show that on average the energy consumption of a transceiver receiving packets does not differ from the energy consumption of a transceiver switched on without receiving a packet. Thus we get

$$\begin{aligned} E = n_t \times 0.02s \times 7.2\text{mA} + n_r \times 0.03s \times 7.2\text{mA} \\ +(10 - 0.02n_t - 0.03n_r)270s\mu\text{A}. \end{aligned} \tag{4}$$

The total amount of energy (in mAs) that the GoldCap provides is calculated according to Equation 1 to 1368 mAs ($U_{max} = 3.8V$, $U_{min} = 2$V (below 2V the voltage regulator does not work anymore), $C = 1$F). Dividing this value of 1368 mA by the result of Equation 4 gives the number of iterations and thus the maximum lifetime of a node.

D. Experimental Results

The presented averaged numbers of iterations (denoting the lifetime of a node) result from 10 test runs per topology.

The first three columns of the Tables show the node ID, the number of iterations and thus the number of packets that the corresponding node is able to transmit according to the mathematical model and to Equation 4. The fourth column gives the averaged measured number of packets transmitted before the node ran out of energy, resulting in the number of measured iterations given in column five. For a comparison between calculated and measured values, the percentage of measured packets to calculated packets (and thus for the iterations) is given, both averaged and the minimum and maximum value for each node over the 10 test runs.

The results are shown in Table I and II. Note that in Table I node 7 and in Table II node 8 are omitted since these are the base stations which we assume to have longer lasting energy supply. The results prove a reasonable match between the calculated results and the measured lifetimes of the nodes. While there are aberrations between nodes and between different experiment runs, the lifetime predictions and especially the order of the nodes dying due to energy loss perfectly agree with the mathematical model. In particular, for the network in Figure 7, Table II shows that node 4 in sphere S_2 runs out of energy first.

The aberrations are caused to a great extent by device tolerances. The de-facto capacity of a 1F-capacitor can vary around the nominal value, and this holds for all used electronic parts like resistors and the microprocessor of the ESB sensor nodes as well.

E. Lifetime Approximations of Wireless Sensor Networks

The experimental results using GoldCap capacitors validate the mathematical model for small networks. It is also straightforward to estimate the lifetime when long-term energy sources such as rechargeable batteries are used. As opposed to 1368 mAs of the GoldCap capacitor, AAA rechargeable batteries provide typically at least 1200 mAh. Thus the lifetime of the network increases with the factor of about 60.

In our experimental setup we let aside the question of long-term synchronization. As we are using a real-time clock, it is sufficient to synchronize the network once at startup using a broadcast sync packet. Deployed sensor networks with long-term energy sources require re-synchronization. Schemes that build hierarchical structures such as TPSN [10] fit very well with the presented mathematical model. To incorporate the energy consumption for the extra message exchanges required, one or more synchronization phases could be executed during the lifetime of the GoldCaps.

Using GoldCaps it is not possible to consider the relaxation effect of batteries. This effect enables batteries to recover a portion of their lost capacity when the discharge current is cut off or reduced [17].

In our experiments, packet loss due to transmission errors is nearly non-existent. As we use time slots, packet collisions do not occur. Packet loss might increase if the distance between neighbour nodes is at the edge of their transmission range. In this case it depends on the sensor network application if packets have to be retransmitted or if some losses are tolerable.

The lifetime of a sensor network can be further increased. As mentioned, GoldCaps can be used in connection with energy scavenging devices like solar cells. The solar cell can charge the GoldCap during the daylight and in the night the GoldCaps provide the energy to the sensor nodes. The presented mathematical model and its experimental validation allow to determine the parameters of such an autonomous sensor network, like the achievable lifetime between two

node ID	modeled iterations	calculated transm.	measured avg. transm.	measured iterations	percentage	min	max
4	353	1412	1159	290	82%	66%	95%
8	353	1412	1289	322	91%	75%	104%
1	387	1161	1046	349	90%	82%	100%
2	429	858	727	364	85%	72%	101%
9	429	858	730	365	85%	76%	98%
5	482	482	368	368	76%	68%	90%
3	482	482	378	378	78%	70%	90%
6	482	482	399	399	83%	74%	98%

TABLE I
RESULTS FOR FIGURE 6

node ID	modeled iterations	calculated transm.	measured avg. transm.	measured iterations	percentage	min	max
3	482	482	422	422	88%	84%	97%
2	429	858	762	381	89%	84%	98%
1	387	1161	1104	368	95%	90%	109%
4	353	1412	1282	320	91%	85%	102%
5	386	1158	1128	376	97%	85%	93%
9	386	1158	1063	354	92%	86%	105%

TABLE II
RESULTS FOR FIGURE 7

daylight periods depending on the amount of traffic in the network.

V. RELATED WORK

In this section, we present related work that deals with lifetime bounds of sensor networks or lifetime maximizations of networks or routings. To the best of our knowledge, none of these approaches has been validated using real sensor hardware.

Bhardwaj et al. provide bounds on the lifetime of a sensor network for basic data gathering scenarios [4]. In later work, the authors extend their analysis by including data aggregation and network topology [3]. While their paper does not provide any practical way to achieve these bounds [16], we demonstrate by experimentation that the bounds derived using our mathematical model are actually achievable. In contrast to [3], Duarte-Melo et al. have developed a modeling framework that is not based on the precise location of the sensor nodes but on probability distributions of the node densities and data rates over the sensing field [9]. Another approach to upper bounds on network lifetime is called cell based energy conservation techniques by Blough and Santi [5]. Also these techniques are hard to realize in practice since they assume an underlying perfect load balancer.

Kalpakis et al. provide bounds on the lifetime of a sensor network [13]. In contrast to our model, their model inherently comprises data aggregation. One major limitation of their model is that they assume that all nodes can reach any other node and the base station in a single hop.

There is also related work that deals with lifetime bounds of heterogeneous clustered networks. Duarte et al. study energy consumption and lifetime of heterogeneous wireless networks [8]. They quantify the optimal number of cluster heads and determine the energy allocation between the different types of sensor nodes. Mhatre et al. consider a heterogeneous sensor network whose task is the surveillance of an area over a certain time and minimize the overall cost using optimization [15], [16]. Based on a similar scenario as Mhatre et al., Heinzelman et al. have studied the performance of a homogeneous clustered network, minimizing the total energy spent in the network [12].

Chang and Tassiulas pursue a network lifetime maximization problem in unicast routing providing algorithms that converge to the optimal solution for single power levels and close to optimal when there are multiple power levels [6]. Kang and Poovendran [14] study a similar issue but related to broadcast routing based on a graph theoretic approach. Coleri et al. both perform a power analysis of a node running TinyOS and determine the lifetime of a tree sensor network of TinyOS motes [7]. Their power analysis of the sensor network highlights the fact that nodes closer to the base station have the lowest lifetime.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we have presented a hardware methodology for experimental evaluation of lifetime bounds for wireless sensor networks. The methodology is also useful for real-life comparison of the energy-efficiency of different medium access schemes, routing and other communication protocols. Using the presented methodology we have evaluated a simple mathematical model and shown that the lifetime bounds by this model actually are achievable.

The mathematical model can thus be used to predict the lifetime of sensor networks and to identify the bottlenecks in terms of energy provisioning in a network with given topology and routing scheme. Once these bottlenecks are identified, either additional nodes can be used or energy scavenging methods can be deployed, like attaching solar cells to these especially energy-critical nodes in a sensor network.

We believe that experimental validation of lifetime models for sensor networks is a very important topic. Therefore, we are currently discussing which of the more complex models we will next validate with our hardware approach.

ACKNOWLEDGMENTS

This work was partly financed by VINNOVA, The Swedish Agency for Innovation Systems.

Thanks to Joakim Eriksson and Niclas Finne for valuable help with the experiments that verified the equality of the lifetime of GoldCap-equipped sensors.

REFERENCES

[1] J. Alonso, A. Dunkels, and T. Voigt. Bounds on the energy consumption of routings in wireless sensor networks. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Cambridge, UK, March 2004.

[2] CST Group at FU Berlin. Scatterweb Embedded Sensor Board. Web page. Visited 2004-06-19. `http://www.scatterweb.com/`

[3] M. Bhardwaj and A. Chandrakasan. Bounding the lifetime of sensor networks via optimal role assignments. In *IEEE Infocom*, New York, NY, USA, June 2002.

[4] M. Bhardwaj, A. Chandrakasan, and T. Garnett. Upper bounds on the lifetime of sensor networks. In *IEEE International Conference on Communications*, Helsinki, Finland, June 2001.

[5] D. Blough and P. Santi. Investigating upper bounds on network lifetime extensions for cell-based energy conservation techniques in stationary ad hoc networks. In *MOBICOM '02*, Atlanta, USA, September 2002.

[6] J.H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *IEEE Infocom*, 2000.

[7] S. Coleri, M. Ergen, and T.J. Koo. Lifetime analysis of a sensor network with hybrid automata modeling. In *International Workshop on Wireless Sensor Networks and Applications*, Atlanta, Georgia, USA, September 2002.

[8] E. J. Duarte-Melo and M. Liu. Analysis of energy consumption and lifetime of heterogeneous wireless sensor networks. In *IEEE Globecom*, Taipei, Taiwan, November 2002.

[9] E. J. Duarte-Melo, M. Liu, and A. Misra. A modeling framework for computing lifetime and information capacity in wireless sensor networks. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Cambridge, UK, March 2004.

[10] S. Ganeriwal, R. Kumar, and M. Srivastava. Timing-sync protocol for sensor networks. In *The First ACM Conference on Embedded Networked Sensor Systems (SenSys 2003)*, Los Angeles, California, November 2003.

[11] EnOcean GmbH. EnOcean:radio switches and sensors without batteries. Web page. Visited 2004-06-19. `http://www.enocean.com/`

[12] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. An application-specific protocol architecture for wireless microsensor networks. *IEEE Transactions on Wireless Communications*, 1(4), October 2002.

[13] K. Kalpakis, K. Dasgupta, and P. Namjoshi. Maximum lifetime data gathering and aggregation in wireless sensor networks. In *IEEE Networks'02 Conference*, Atlanta, GA, USA, August 2002.

[14] I. Kang and R. Poovendran. Maximizing static network lifetime of wireless broadcast ad-hoc networks. In *International Conference on Communications*, Anchorage, Alaska, USA, May 2003.

[15] V. Mhatre, C. Rosenberg, D. Kofman, R. Mazumdar, and N. Shroff. Design of surveillance sensor grids with a lifetime constraint. In *First European Workshop on Wireless Sensor Networks (EWSN 2004)*, Berlin, Germany, January 2004.

[16] V. Mhatre, C. P. Rosenberg, D. Kofman, R. R. Mazumdar, and N. B. Shroff. A minimum cost surveillance sensor network with a lifetime constraint. *to appear in IEEE Transactions on Mobile Computing*, 2004.

[17] V. Raghunathan, C. Schurgers, S. Park, and M. Srivastava. Energy aware wireless microsensor networks. *IEEE Signal Processing Magazine*, 19(2):40–50, March 2002.

[18] R. Szewczyk, J. Polastre, A. Mainwaring, and D. Culler. Lessons from a sensor network expedition. In *First European Workshop on Wireless Sensor Networks (EWSN 2004)*, Berlin, Germany, January 2004.

[19] S. Tilak, N. Abu-Ghazaleh, and W. Heinzelman. A taxnonomy of wireless micro-sensor networks. *ACM Mobile Computing and Communications Review*, 6(2), 2002.

[20] T. Voigt, H. Ritter, and J. Schiller. Solar-aware routing in wireless sensor networks. In *International Workshop on Personal Wireless Communications (PWC 2003)*, Venice, Italy, September 2003.