

Abstract Frames for Reducing Overhearing in Wireless Sensor Networks

Abdelmalik Bachir¹, Dominique Barthel¹,
Martin Heusse², and Andrzej Duda²

¹ France Telecom R&D, Meylan, France

² LSR-IMAG Laboratory, Grenoble, France

Abstract. We present a novel idea for energy saving by avoiding the reception of redundant broadcast frames. It is based on sending an abstract frame just before a data frame: the former contains a digest of the latter. We evaluate the energy savings of this scheme analytically and by means of simulation in ns-2. Although we have applied our approach to SMAC, the key idea is generic and can be used to reduce energy consumed by broadcast frames in a large variety of MAC protocols.

1 Introduction

Energy saving is crucial in designing long-lived sensor networks, mainly because nodes are powered by batteries that may be costly, difficult, or impossible to replace or to recharge. Measurements presented in the literature [1, 2] confirmed by our experiments with the MC 13192 SARD sensor node show that radio communication is a major source of energy consumption. Therefore, a node needs to switch the radio off whenever possible to save energy. At the same time, to be able communicate with other nodes, they need to activate their radios during some common active periods. Moreover, their competition for the radio channel should follow a set of rules defined in a MAC layer. Energy may be wasted due to the inherent behavior of the MAC layer because of the following problems:

- *Idle Listening*: since a node does not know when it will be the receiver of a frame, it must keep its radio in receive mode, which consumes energy.
- *Collisions*: the energy used to send and receive a frame is wasted when the frame collides with another frame.
- *Overhearing*: this happens when a sensor node receives and decodes an irrelevant frame (e.g. a broadcast frame that has already been received).
- *Protocol Overhead*: control frames do not carry useful information although their transmission consumes energy.

SMAC is an example of an energy efficient access method that tries to address all these issues [3].

We focus on the problem of reducing the impact of idle listening and overhearing. We observe that many applications require an one-to-all communication scheme such as a *network-wide broadcast*. A straightforward way to achieve this scheme is *flooding*: a node broadcast a message to its neighbors that forward it

further on so that all the nodes eventually receive it. Network-wide broadcast is usually used in management protocols such as route discovery of on-demand routing protocols [4] or during the interest propagation phase in Directed Diffusion [5], and in many application protocols. For a given node in the network, only the broadcast frame received first is actually relevant and all subsequent ones are redundant since they carry the same data contents. Redundant transmissions degrade the performance of the network: when combined with a contention based MAC protocol, redundant transmissions increase the collision rate as reported in the broadcast storm problem [6]. Moreover, redundant transmissions simply drain more energy than needed, so their reception should be avoided.

Since reducing energy consumption is the critical issue in sensor networks, many authors have proposed protocols for energy-efficient broadcast: CDS (Connected Dominating Sets) [7], MPR (Multi Point Relays) [8], or RNG (Relative Neighborhood Graphs) [9]. They try to reduce energy consumption of the network by limiting the number of required transmissions or by performing power control to reduce transmission power. However, to the best of our knowledge, there is no protocol that attempts to reduce the number of redundant receptions. Although limiting the number of required transmissions already implies less redundant receptions, we propose to reduce redundant receptions even further to save more energy.

To reduce energy consumption during the reception of broadcast messages, we propose *abstract frames*: an abstract frame is a small control frame sent before each data frame. It contains a digest of the subsequent data frame. A node listening to the channel uses the information in the abstract frame to identify and filter out redundant messages at the MAC layer. If the node has already received the data frame, it can switch its radio off and save energy.

The rest of the paper is organized as follows. In Section 2, we present the key idea of abstract frames. In Section 3, we analyze the performance of our abstract frames method in terms of lifetime extension compared to two MAC protocols: an ideal one that totally avoids idle listening, but does not filter out redundant messages at the MAC layer and a practical one, SMAC [3]. In Section 4, we report simulation results on the performance of abstract frames method when used with SMAC.

2 Abstract Frames

An abstract frame is a small frame sent immediately before each broadcast frame. It contains a digest of the data in the subsequent broadcast frame. A node uses the information in the abstract frame to learn about the subsequent data contents. If a node learns from the abstract frame that the data frame has already been received, it can switch its radio off, because the subsequent data is redundant as shown in Fig. 1. In this way, a node only overhears redundant abstract frames instead of overhearing redundant data frames, which contributes to save more energy since abstract frames are expected to be far shorter than data frames.

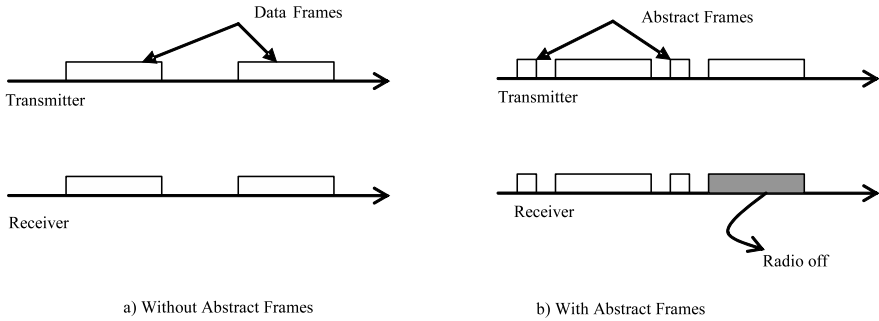


Fig. 1. Avoiding redundant frame reception by means of Abstract Frames

An abstract frame has a field that contains either a unique identifier, or a hash of the data contained in the subsequent broadcast frame. When the MAC layer needs to transmit a frame, it constructs and transmits the corresponding abstract frame before transmitting the broadcast frame. It inserts the hash field of the abstract frame in a table to avoid receiving it again from its neighbors. This table logs frames that have been recently seen so the MAC layer may switch the radio off when it expects a redundant reception.

According to this procedure, the MAC layer always receives an abstract frame before a data frame for broadcast communications. It first checks in its table whether there is an entry with the same hash value. If the entry exists, then the node switches its radio off to avoid receiving the same data again. If there is no such entry in the table, then the node continues to listen to the channel in order to receive the subsequent data frame. Once it has received the data frame, the node updates its table to avoid receiving redundant transmissions of the received data frame.

One can argue that there will be collisions in computing the hash value leading to false positives so that a node may ignore a data frame that has not been previously received. However, we think that this situation is unlikely for the following reasons. First, hash-field entries in the table of the MAC layer are not permanent, but cleaned after a timeout value. Second, we can choose a hash function and the size of the digest so that collisions are very rare. A frame will only be missed if it involves two simultaneously active colliding broadcasts during the timeout. Note that because broadcasts are not acknowledged, they are usually unreliable anyway.

There are several ways for reducing energy consumption caused by broadcasts. The most immediate one is to reduce the number of transmitted frames by avoiding redundant transmissions. Many proposed protocols select only a subset of nodes to flood a message while ensuring that all nodes eventually receive the message: CDS (Connected Dominating Sets) [7], MPR (Multi Point Relays) [8], or RNG (Relative Neighborhood Graphs) [9]. Another approach tries to optimize the transmission range by seeking a good trade-off between consuming more energy required to reach farther nodes or having more retransmissions [10, 11, 12]. We can apply the abstract frame approach to all such protocols, because it reduces energy

consumption at the MAC layer. As all of them try to reduce broadcast traffic, the more efficient they are, the less abstract frames are necessary.

3 Theoretical Performance

Although the use of abstract frames results in less energy consumption during the reception of redundant frames, it also increases the energy drained per transmitted data frame. Therefore, we propose to analyze the performance of abstract frames taking into account these two parameters together. We propose to compare the lifetime of a node running a MAC protocol without abstract frames, which we call protocol P , to the lifetime of the same node when running protocol P with abstract frames, which we call P' .

In our analysis, we consider that all the communications are broadcast and nodes forward frames according to the flooding algorithm. As a candidate for protocol P , we take two examples. The first protocol is an ideal MAC protocol that does not have idle listening. The second protocol is SMAC that reduces idle listening via active/sleep schedules. For the sake of simplicity, we do not consider collisions in the following analysis.

We call E_P (resp. $E_{P'}$) the energy drained during a complete local flooding operation when nodes use protocol P (resp. P'). To quantify the ratio of lifetime extension by protocol P' compared to protocol P , we calculate gain G_P defined in (1) for the two candidate MAC protocols.

$$G_P = \frac{E_P}{E_{P'}} \tag{1}$$

3.1 Ideal MAC

To calculate the lifetime of a node, we consider the complete local flooding operation consisting of the reception of all frames from its neighbors and forwarding the broadcast frame exactly once. Therefore, if the node has n neighbors, then the energy drained during the flooding operation is:

$$E_{ideal} = nTP_r + TP_t, \tag{2}$$

where T is the transmission time of the data frame and P_t (resp. P_r) is the power drained by a transmission (resp. a reception).

When abstract frames are used, the node receives all abstract frames, but only one data frame. This is because the node discards the other data frames since they are redundant. In addition to that, the node transmits one abstract frame that precedes the data frame. The energy drained in this case is thus:

$$E_{ideal'} = (nA + T)P_r + (T + A)P_t, \tag{3}$$

where A is the transmission time of an abstract frame. Finally, the lifetime extension is the following:

$$G_{ideal} = \frac{E_{ideal}}{E_{ideal'}} = \frac{(n + \rho)T}{(1 + \rho)T + (n + 1)A}, \tag{4}$$

where $\rho = P_t/P_r$.

From Eq. (4), we conclude that the lifetime extension increases when data size increases and it converges to $\frac{n+\rho}{1+\rho}$ when $t \rightarrow \infty$. Also, when the number of neighbors increases, the gain increases up to $\frac{T}{A}$ when $n \rightarrow \infty$.

Note that the performance obtained with the use of abstract frames depends on the ratio $\frac{A}{T}$ of the abstract frame transmission time to the transmission time of data frames: the smaller A compared with T , the larger lifetime extension we get. We can calculate A_{max} , the maximum value of A beyond which there is no gain in using abstract frames: we need that $G_{ideal} > 1$, which leads to the following:

$$A_{max} = \left(\frac{n - 1}{n + 1}\right) T \tag{5}$$

3.2 SMAC

We follow the same methodology to evaluate the lifetime extension ratio for SMAC. Fig. 2 shows an example of a node with three neighbors. In SMAC, a node alternates active periods during which nodes can communicate and sleep periods during which nodes switch their radios off to save energy. The ratio of the period durations is controlled by the MAC duty-cycle parameter. The duration of the active period depends on a couple of parameters such that the data transmission time and the slot time used in the backoff procedure when nodes contend for the channel. Note that SMAC protocol chooses carefully the duration of the active period so that it is large enough to hold a data transmission including contention. However, there is no guarantee that the local flooding operation fits a single active period¹.

Let us call D the duration of the active period. The local flooding operation may fit one active period or more, depending on many parameters like the number of nodes n . Let us assume that a node needs k active periods, ($k > 0$) to carry out the local flooding operation. Therefore, we have the following relations (see Fig. 2):

$$E_{smac} = nTP_r + TP_t + [kD - (n + 1)T]P_i \tag{6}$$

and,

$$E_{smac'} = (T + nA)P_r + (T + A)P_t + (n - 1)TP_s + [kD - n(T + A) - (T + A)]P_i, \tag{7}$$

where P_i (resp. P_s) is the power drained during the idle (resp. sleep) mode. In general, the power drained in the idle mode, in which the radio is ready to receive,

¹ For the sake of simplicity, Fig. 2 shows that the local flooding operation fits one active period.

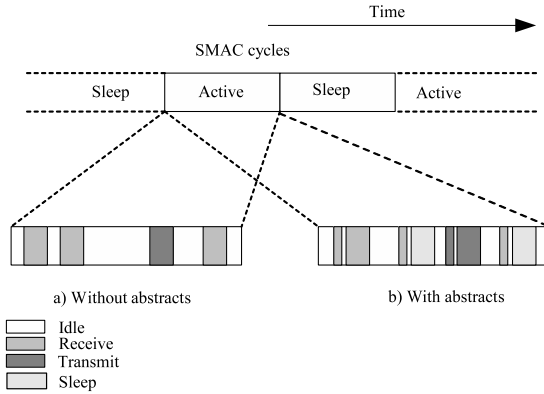


Fig. 2. Operation of SMAC with and without Abstract Frames

is slightly less than the power drained during the receive mode. However, to simplify the comparisons, we will assume that $P_i = P_r$. Also, the power drained during the sleep mode is negligible compared to other modes, so we assume that $P_s = 0$. By denoting $\rho = P_t/P_r$, we obtain the following lifetime extension for SMAC:

$$G_{smac} = \frac{E_{smac}}{E_{smac'}} = \frac{T\rho + kD - T}{(T + A)\rho + kD - nT - A} \tag{8}$$

Eq. (8) shows that the gain depends on kD , the duration of active periods needed for the complete local flooding. These active periods include idle listening. To show the effect of idle listening on lifetime extension, we propose to rewrite Eq. (8) in function of I , the amount of idle listening during the complete local flooding:

$$I = kD - nT - T \tag{9}$$

Thus, Eq. (8) can be rewritten as:

$$G_{smac} = \frac{(n + \rho)T + I}{(1 + \rho)T - A(1 - \rho) + I} \tag{10}$$

Eq. (10) shows that the gain decreases when idle listening increases. This is quite expected, because when idle periods dominate, then we will not get significant lifetime extension since the power drained during idle listening will be the same whether abstract frames are used or not.

Note that Eq. (10) implicitly includes the effect of traffic load on lifetime extension because idle listening depends on traffic load. Indeed, when traffic load is high, nodes spend more time in transmit and receive modes, which decrease the amount of idle listening. Therefore, we conclude that lifetime extension increases with traffic load. However, an excessive traffic load causes collisions that actually decreases the lifetime extension. We study this factor through simulation in Section 4.

Using the same approach as in Section 3.1, we calculate A_{max} for SMAC. We have,

$$A_{max} = \left(\frac{n-1}{\rho-1} \right) T \quad (11)$$

Interestingly, when $\rho \leq 1$, we have no constraint on the abstract frame transmission time to obtain a lifetime extension. Eq. (10) shows that the gain will always be larger than 1, provided $\rho \leq 1$. In this case, the power drained in the transmit mode is less than the power drained in the idle mode and transmitting abstract frames saves more energy than by staying idle.

4 Simulation

We have used ns-2 [13] to quantify the lifetime extension achieved with the use of abstract frames by simulation. We have chosen SMAC to represent a low power MAC protocol mainly because its code is public and seems to be stable. However, the idea of avoiding redundant frames reception by means of abstract frames may apply to a large variety of MAC protocols.

We compare the lifetimes achieved by two MAC protocols: SMAC without abstract frames and SMAC', which is SMAC with abstract frames. We carry out simulations to get more insight into the energy saving ratio since we have not taken all the parameters into account in the mathematical analysis.

Our application consists of a simple flooding agent that forwards a new broadcast message it receives exactly once (we assume that the size of data frames is fixed). The simulation scenario consists of one source node broadcasting messages according to a given traffic load. The other nodes flood the received broadcast message. The source node assigns a different message identifier for each new broadcast message. This identifier is used as our digest in the abstract frame. When the SMAC' agent receives an abstract frame, it checks whether it has already seen (sent or received) the data frame with the same message identifier. If the identifier is new, the SMAC' agent adds this identifier to an internal table and keeps the radio on to receive the subsequent data frame. However, if the identifier has been already seen, then the following data is redundant and the SMAC' agent switches the radio off to save energy.

The application agent counts the number of non-redundant received messages. We use this number to quantify the lifetime of a node. The ratio of the number of received messages with SMAC' out of the number of received messages with SMAC determines the lifetime extension. We have considered three situations to evaluate lifetime extension. These situations are the lifetime extension of the most vulnerable node, the lifetime extension of the most robust node, and the lifetime averaged over all the nodes. Results show that abstract frames extend the lifetime of SMAC with very close ratios in all these three situations. Therefore, we have chosen to only analyze the results corresponding to the average number of messages in the rest of the discussion.

We measure the impact of data payload, traffic load, and network density on the lifetime extension. We use a simple energy model to simulate low power

radio. In our energy model, the transmit mode uses 96mW, the receive mode uses 111mW and the ready-to-receive mode, that we also called idle mode, uses the same power as the receive mode.

4.1 Simple Star Network

For the topology, we take a simple star network with the source node placed at the center of the simulation area. The number of neighbors of the source node determines the density of the network.

For each simulation run, we calculate the lifetime extension as defined above. After some simulation runs, we calculate the average lifetime extension and the confidence interval corresponding to 95% of values. Note that the confidence interval is fairly large even with a very large set of simulation runs, which reflects the fluctuations in lifetime extension ratios due to the SMAC characteristics. In SMAC, the active period is composed of two periods: synchronization period and data period. Nodes use the synchronization period to exchange SYNC frames from time to time in order to synchronize on a common sleep/wakeup schedule. Note that SMAC does not guarantee that all nodes synchronize on a single common schedule. Nodes that share a common schedule form a virtual cluster and the network may contain one or several virtual clusters. Some nodes, called border nodes, may belong to more than one virtual cluster. There are many options to manage the active/wakeup schedules of the border nodes. In the implementation of SMAC we are using, a border node keeps active during all the active periods of the virtual clusters it belongs to. As a consequence, the actual MAC duty cycle of a border node increases, which decreases the efficiency of abstract frames since idle listening increases with the MAC duty cycle. Note that each time we run a simulation with a different random seed, we get different numbers of virtual clusters formed in the network. We have taken care to compare the lifetimes of SMAC and SMAC' in similar conditions, *i.e.*, for the same number of virtual clusters.

Fig. 3(a) shows the histogram of the lifetime extension for various data packet sizes. The variability is due to the number of virtual clusters formed during each

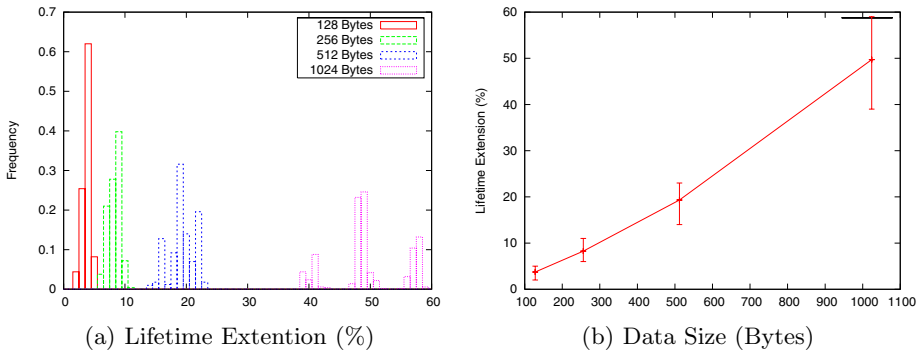


Fig. 3. Lifetime extension with respect to the data size




Fig. 4. Lifetime extension with respect to the traffic load and the number of neighbors

run. Fig. 3(b) presents a more legible view of the same data. Likewise, Fig. 4(a), 4(b), 5 show the lifetime extension as a function of other variables.

We notice that the lifetime extension is small for small data sizes, because the amount of time during which we switch the radio off to avoid redundant data reception becomes negligible (around 4%) compared to the time the radio is on. We have used the application duty-cycle that corresponds to the traffic load of 0.1 messages per second. However, when the data payload size increases to 1024 bytes, the lifetime extension increases by 40%-60%. This large interval is due to variations of the number of virtual clusters in each simulation run. As expected, the 40% ratio corresponds to the situation with many virtual clusters and the 60% ratio corresponds to fewer virtual clusters. Note that the formation of several virtual clusters decreases the lifetime extension because idle listening becomes significant. In [14], Li et al. pointed out this problem and proposed GSA (Global Schedule Algorithm) to make the network converge to one virtual cluster. We expect abstract frames to have better performance with GSA.

In Fig. 4(a), we have varied the traffic load from 0.0125 to 0.1 messages per second *i.e.*, from one message every 10 seconds to one message every 80 seconds. The payload of the messages is 512 bytes. The figure shows that the lifetime extension increases when the traffic load increases. In this case, the time during which data are exchanged becomes non negligible compared to the duration of idle listening. Hence, the duration during which SMAC' exploits abstract frames to switch off the radio becomes more significant. Note that we do not present results beyond 0.1 messages per second, because we have observed a considerable increase of collision rates for the traffic load larger than 0.1 and negligible energy saving ratios for the traffic load less than 0.0125. We argue that this is rather SMAC-dependent and not a result showing intrinsic low performance of abstract frames. We expect to get better performance with other low power MAC protocols that manage idle listening in a better way like TMAC [15], WiseMAC [16] and BMAC [17].

In Fig. 4(b), we have varied the number of neighbors of the source node to get different network degrees. This gives us a precise idea on what SMAC' is able




Fig. 5. Lifetime extension with respect to the number of nodes in the network

to achieve in situations in which the channel is not saturated and collisions are rare. For these simple star topology networks, we have observed a collision ratio less than 1%, which allows us to see the effect of network density on lifetime extension independently from collisions. As we expected, the lifetime extension increases with network density.

4.2 Realistic Network

In Fig. 5, we have measured the lifetime extension ratios for more realistic topologies. We have generated five networks with node positions distributed uniformly in a square area except for the source node always placed in the center. The networks are connected and the degree of a network with less nodes is less than the degree of a network with more nodes. The average densities of the networks are: 1.8, 2.7, 3.3, 3.6 and 4.12 for networks with 10, 20, 30, 40, and 50 nodes respectively. With the traffic load of 0.1 messages per second, we have observed collision rates increasing considerably.

Collisions mitigate the efficiency of abstract frames. Therefore, as opposed to what one may expect, the lifetime extension ratio may not systematically increase with the network density as shown in Fig. 5. This is because in most of the cases, collisions happen between abstract frames transmitted in the same time slot and receivers cannot correctly decode these abstract frames. Hence, it is not possible for receivers to know about the subsequent data contents so that they cannot switch their radios off. Note that this issue is rather related to the performance of the broadcast in 802.11-inspired MAC protocols. The way commonly used to decrease collision rates in these protocols is to increase the contention window size. However, this may be inefficient in SMAC, because increasing the contention window also increases idle listening, which is not desirable. Another reason of smaller gain ratios of abstract frames with SMAC is related to the formation of many virtual clusters. Indeed, when the density of nodes in network increases, border nodes will belong to more virtual clusters. Therefore, border nodes will be awakened during all the wakeup schedules of the virtual clusters they belong to, which increases idle listening and then mitigates

the gain obtained with the use of abstract frames. For example, in networks with 10 nodes (resp. 20, ..., 50) we have used in the simulation, each node belongs on the average to 1.3 (resp. 1.6, 1.8, 2.1 and 2.3) virtual clusters. We think that these mitigated gain ratios are rather due to SMAC and we expect larger gain ratios with protocols that manage idle listening better.

5 Conclusion

We have presented a novel idea for energy saving by reducing overhearing redundant copies of broadcast frames. It is based on abstract frames containing a digest of a subsequent data frame. We have evaluated of this scheme analytically and by means of simulation in ns-2. Although we have applied our approach to SMAC, the key idea is generic and can be used in a large variety of MAC protocols to further enhance energy savings of existing optimized broadcast protocols. We continue this work by evaluating the efficiency of abstract frames with other MAC protocols.

References

1. Ember Corporation. EM250 Single-Chip ZigBee/802.15.4 Solution, Data Sheet . 2005.
2. J. Polastre, R. Szewczyk, D. Culler. Telos: Enabling Ultra-Low Power Wireless Research. *Proceedings of IPSN/SPOTS*, pages 302–11, Los Angeles, CA, April 2005.
3. W. Ye, J. Heidemann, and D. Estrin. An energy-efficient MAC protocol for wireless sensor networks. *Proceedings of the IEEE Infocom*, pages 1567–76, New York, NY, July 2002.
4. J. N. Al-Karaki, A. E. Kamal. Routing techniques in wireless sensor networks: a survey. *IEEE Wireless Communications* , 11(6):6–28, Dec 2004.
5. C. Intagoniwat et al. Directed diffusion for wireless sensor networking. *IEEE/ACM Trans. on Networking*, 11(1):2–16, February 2003.
6. S.Y. Ni et al. The Broadcast Storm Problem in Mobile Ad Hoc Network. *Proceedings of the IEEE/ACM Mobicom*, 1999.
7. J. Wu and H. Li. On Calculating Connected Dominating Set for Efficient Routing in Ad Hoc Wireless Networks. *Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 7–14, Aug 1999.
8. A. Qayyum, L. Viennot, A. Laouiti. Multipoint Relaying for Flooding Broadcast Messages in MobileWireless Networks. *Proceedings of IEEE HICSS*, Big Island, HI, Jan 2002.
9. G. Toussaint. The Relative Neighborhood Graph of a Finite Planar Set. *Pattern Recognition*, 12(4):261–8, 1980.
10. J. Cartigny, D. Simplot and I. Stojmenovic. Localized Minimum-Energy Broadcasting in Ad-hoc Networks. *Proceedings of the IEEE Infocom*, San Francisco, CA, April 2003.
11. J. Cartigny, et al. Localized LMST and RNG Based Minimum-energy Broadcast Protocols in Ad hoc Networks. *Ad Hoc Networks*, 3(1):1–16, 2004.

12. F. Ingelrest, D. Simplot-Ryl and I. Stojmenovic. Target Transmission Radius over LMST for Energy-Efficient Broadcast Protocol in Ad Hoc Networks. *Proceedings of IEEE ICC*, Paris, France, June 2004.
13. <http://www.isi.edu/nsnam/ns/>.
14. Y. Li, W. Ye, and J. Heidemann. Energy and Latency Control in Low Duty Cycle MAC Protocols. *Proceedings of the IEEE WCNC*, New Orleans, LA, March 2005.
15. T. van Dam and K. Langendoen. An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks. *Proceedings of the ACM Sensys*, pages 171–80, Los Angeles, CA, November 2003.
16. Enz, C.C.; El-Hoiydi, A.; Decotignie, J.; Peiris, V. WiseNET: An Ultralow-Power Wireless Sensor Network Solution. *IEEE Computer*, 37(8):62–70, August 2004.
17. J. Polastre, J. Hill and D. Culler. Versatile Low Power Media Access for Wireless Sensor Networks. *Proceedings of the ACM SenSys*, 2004.