

An adaptive TDMA based MAC Protocol for Mobile Wireless Sensor Networks

Stephan Mank
Distributed Systems
Operating Systems
BTU Cottbus
Cottbus, Germany

Reinhardt Karnapke
Distributed Systems
Operating Systems
BTU Cottbus
Cottbus, Germany

Joerg Nolte
Distributed Systems
Operating Systems
BTU Cottbus
Cottbus, Germany

Email: smank@informatik.tu-cottbus.de Email: karnapke@informatik.tu-cottbus.de Email: jon@informatik.tu-cottbus.de

Abstract—Radio transceivers for wireless sensor networks are typically cheap devices offering low bandwidth communication only. When physical events in the real world trigger spontaneous communication in many nodes the single communication channel is under heavy load and many messages are lost due to collisions. CSMA/CA schemes are well suited to spontaneous communication, but do not provide a high channel utilization under heavy load. TDMA protocols have some conceptual advantage here, but in the case of mobile sensor nodes they need to be adaptive and establish TDMA schedules on demand. MLMAC is a novel TDMA based MAC protocol that can react on changing radio neighborhoods in mobile networks. In this paper we introduce the MLMAC protocol and present the results of real experiments with a group of mobile sensor nodes based on RCX robots.

I. INTRODUCTION

Sensor networks are collections of small sensor nodes with wireless neighborhood broadcast facilities. Since sensor networks shall be deployed in large scales (possibly thousands of nodes [1], [2]), the overall cost dictates the use of cheap and simple radio transceivers for communication. These lack most of the common capabilities of WLAN or bluetooth networks. Even typical tasks like medium access control or the addressing of individual nodes in the direct radio neighborhood are left entirely to software layers [4]. In this paper we present Mobile LMAC (MLMAC), a novel TDMA MAC protocol for mobile wireless sensor networks. It is a modification of the lightweight medium access protocol (LMAC) [10] which has been adopted for our RCX robots and for a different scenario where all nodes are mobile. Therefore, finding a slot is more difficult and must be repeated under some conditions as nodes can enter or leave a radio neighborhood at any time. As far as we know, no previous implementation of LMAC has been used on real hardware.

This paper is structured as follows: Section II describes our MAC protocol while section IV describes the setup of our experiments and their results. We finish with related work in section V and a conclusion in section VI.

II. MOBILE LMAC (MLMAC)

A. LMAC

In [10] L.F.W. van Hoesel and P.J.M. Havinga introduced LMAC (lightweight medium access control), which is based

TABLE I
THE CONTROL MESSAGE USED IN LMAC

Description	Size (bytes)
Identification	2
Current Slot Number	1
Occupied Slots	4
Distance to Gateway	1
Collision in Slot	1
Destination ID	2
Data Size (bytes)	1
Total	12

on a TDMA scheme. Time is divided into frames and slots. Each node reserves a slot in which it can send, this slot reoccurs every frame. Every Slot is used to send a control message followed by data payload. Table I shows the contents of a LMAC control message. Its total size amounts to 12 Bytes. It contains the identity of the sender and its slot number followed by the most important field *Occupied Slots*, which represents a Bitmask of Slots. An unused slot is represented by a 0 while a 1 represents an occupied one. Thus it is possible for every node to determine unoccupied slots by combining the control messages of its neighbors. This is done by performing a simple OR operation on the fields *Occupied Slots* of all received control messages. The distance to the Gateway is also transmitted, along with information of overheard collisions. Finally, the ID of the destination and the size of the data unit are given. The initialization of nodes is started by the gateway, which defines its own slot and is used for synchronization. After one frame, all direct neighbors of the gateway know its slot and choose their own ones. This information is transmitted to their neighbors who synchronize on these messages. After each frame, a new set of nodes with a higher hop distance from the gateway are synchronized until every node knows its slot. These slots only need to be locally unique, as the nodes only compete with others up to 2 hops distant. To conserve node energy, a node's transceiver is turned off for the remainder of the current slot when it is not addressed in the control message. As slots are computed only once in LMAC, this protocol is not suitable for mobile sensor

TABLE II
THE CONTROL MESSAGE USED IN MOBILE LMAC (MLMAC)

Description	Size (bytes)
Identification	1
Slot number and Status	1
Occupied Slots	1-2
Identity of the Synchronization	1
Age of Synchronization	1
Total	5-6

networks, where nodes can enter and leave other nodes' radio neighborhood at any time.

B. Differences and Similarities between LMAC and MLMAC

The main differences between LMAC and MLMAC is their intended scenario. While LMAC assumes a static sensor network where communication takes place between nodes and the gateway, in our scenario, all nodes are mobile and communicate among each other. There may be one or many gateways or maybe there is none. MLMAC does not depend on a gateway to start the synchronization, instead, it is fully dynamic.

The choice of scenario has a number of consequences for the developed MAC protocol. First, there may not be a gateway to start the synchronization (see next section II-B.1). Second, the chosen slots are not fixed in time. Due to mobility, it may become necessary for a node to choose a new slot when it enters a different radio neighborhood (see II-B.2). Third, MLMAC needs to differentiate between bidirectional and unidirectional links (see section II-C).

1) *Slot Synchronization*: The node that wants to send a packet first starts the synchronization. This removed the necessity to use the field *Distance to Gateway* for synchronization. Even when it is not used for synchronization, the field *Distance to Gateway* could be used to support routing decisions in stationary sensor networks. In mobile sensor networks however, this distance could not be determined only once and saved for further use, as the mobility of nodes will lead to a change in topology after a certain time. This time depends on the range of the transceivers and on the speed of the nodes of course, but eventually the change in topology will take place. The field *Distance to Gateway* is removed from the header of MLMAC. The fields *Destination ID* and *Collision in Slot* were not used either, because of the used hardware. There was no way to shut down the transceivers and the radio module used a built in checksum to discard faulty packets. Thus, collisions could not be detected directly and this part of LMAC was not implemented. As this decision is based solely on the used radio modules, it could be revoked in the future, when a different platform is used (see section VI).

MLMACs control message format is shown in table II. This control message is quite different from the one used in LMAC. Due to our small sensor network, we reduced the field containing the identity of the sender to one Byte. Slot

number and Status contains 5 Bit for the senders slot and 3 Bit for its status. The field *Occupied Slots* is used exactly as in LMAC, only its size is reduced.




Fig. 1. Occupied slots as seen by each node

The field *Occupied Slots* in the control message of a node contains the used slots of all its neighbors and itself. In the case of node 4 for example the slots 3, 4, 5, 6 and 8 would be marked as used, which results in a representation as 00111101. Note that in this example the third bit from the left represents the third slot. Figure 1 shows how slots are chosen with a simple example containing only eight nodes. In this example you can see that node 2 is not synchronized yet. It receives the control messages from its neighbors and combines them. 10000100 (from node 1) — 00111000 (from node 3) — 00111101 (from node 4) = 10111101 (seen on node 2).

This means that node 2 can choose between slots 2 and 7. If it chooses slot 2, its control message would contain the Bitmask 11110000 in the field *Occupied Slots*, as node 2 receives messages from nodes 1,3, and 4 and adds its own choice. If it chooses slot number 7 the field *Occupied Slots* would contain the Bitmask 10110010. Note that this method solves the hidden station problem. The number of slots can be chosen between 3 and 16 in our implementation, thus the size of *Occupied Slots* varies between 1 and 2 Byte. If more slots per frame are needed, the size of the field *Occupied Slots* grows. Thus far, the choice of slots is done in the same way as in LMAC, except for the fact that the synchronization started by a node rather than by the gateway.

2) *Slot changes*: The second difference is the fact that MLMAC stays adaptive even after slots are chosen. The last two fields of a control message are needed, because every node can start synchronization. Due to this fact, it is possible that two distant nodes start a synchronization separately, as both of them assume that they are the first to send. Their neighbors would synchronize with them and increase the *Age of Synchronization* by one before retransmitting. In this case, two different synchronizations would be flooding the net and meet somewhere in between the two starter nodes. At this point, nodes would realize that some of their neighbors use a different synchronization by comparing the field *Identity of the Synchronization*. Now the field

Age of Synchronization is compared. If the received value is equal to or higher than the local value saved in a node, this node becomes unsynchronized again and tries to find a new slot.

Due to the mobility of nodes, a node X may leave the radio range of node Y. Both nodes then realize that they do not receive any more control messages from each other and remove the other one from the neighbor list. When X moves into the radio range of another node Z which knows a different node W which uses the same slot as X, the control messages of X and W collide at Z. Therefore, Z does not receive any more control messages in that slot and marks it as unused. Nodes X and W receive the control message from Z and realize that there must have been a collision of control messages. After this, they give up their current slot and try to find a different one.

C. Handling of unidirectional links

To determine whether a link is unidirectional or bidirectional, a neighbor list is used. In this list a counter is stored for every neighbor. When a node X receives a control message from node Y which does not contain the slot of node X, X increments the counter for Y in its neighbor list. If the received control message contains X's slot, the counter is decreased. The range of the counter is 0 - N where N can be configured freely. Then, a threshold can be set, from which on the link will be counted as (partially) unidirectional. The need for this will become apparent when the state machine of MLMAC is discussed in section II-D

D. The finite state machine of MLMAC

MLMACs state machine is shown in figure 2. The rectangles represent states, the arrows transitions between them. The text on the arrows describes the necessary event for that transition and the action that is taken separated by a slash.

Initially, all nodes begin in the *Wait*-state. When they want to send a message without having received a control message yet, they change into the *Starter*-state. When only one node switches to starter, this is a stable state and the node remains there. If another node switched to the *Starter*-state earlier, this node gains knowledge of that fact after some time and switches to the *Sleep*-state from which it will return into the *Wait*-state after a certain time.

If a node received a control message from another node in *Starter*- or *Ready*-state while in the *Wait*-state it synchronizes its local time with that of the originator of the control message and switches into the *Unsync*-state. After waiting one frame to overhear all transmitted control messages and calculate used slots, it chooses its own one and transitions into the *Sync*-state. The next time, that node's slot is active, it starts to transmit its own control messages in every frame and changes to state *Slotverify*. This state is used to verify that no other node has chosen the same slot during the last frame. This would be indicated by a collision of control messages in this slot and lead to a change into the *Sleep*-state.

As said before, the used hardware does not enable us to detect collisions directly. Rather, a node X that transmitted a control message can determine if a collision occurred by listening to its neighbors' control messages. If no collision occurred, the neighboring nodes have added X's slot to the field *Occupied Slots* in their control messages. Otherwise they did not. When X receives control messages containing its slot, it knows that no collision occurred because no other node has chosen the same slot. Therefore, it switches into the *Ready*-state.

Like the *Starter*-state, this is a stable state as long as no collision occurs. If a collision occurs, there must have been a mistake in the process of choosing slots, and the node returns into the *Sleep*- and finally into the *Wait*-state to start over again.

Note that MLMAC also distinguishes between collisions on unidirectional and bidirectional links. If a collision on a unidirectional link occurred on a node in *Slotverify*- or *Ready*-state, this node stays in the same state.

III. IMPLEMENTATION

A. Hardware

For our implementation we used modified Lego RCX robots [6]. These robots feature a Renesas H8/300 micro-




Fig. 3. A modified RCX Robot

controller, which is a 16 Bit processor with a frequency of 16 MHz. The RCXs have three input channels and three output channels on top and an infrared module in the front. The Robots have been additionally equipped with a radio module of type ER400TRS [3] which we use instead of the included infrared module (IR) to enable broadcast protocols (see figure 3). The ER400TRS has a range of up to 250 meters which we reduced to about one meter, operates in the Pan-European 433 to 434MHz frequency band and uses a 3.6V power supply. The over the air data rate is 19200 Baud. To enable usage of the ER400TRS, an external connector to the serial port has been attached which allows us to connect either the IR or the radio module. It may seem that the choice of radio module was not the best, but when sensor nodes are designed, the chosen hardware is not always the best, either. Rather, it is normally the cheapest hardware that is chosen.

B. Software

The authors of [11], [12] introduced Reflex, a novel event driven operating system for deeply embedded systems. We

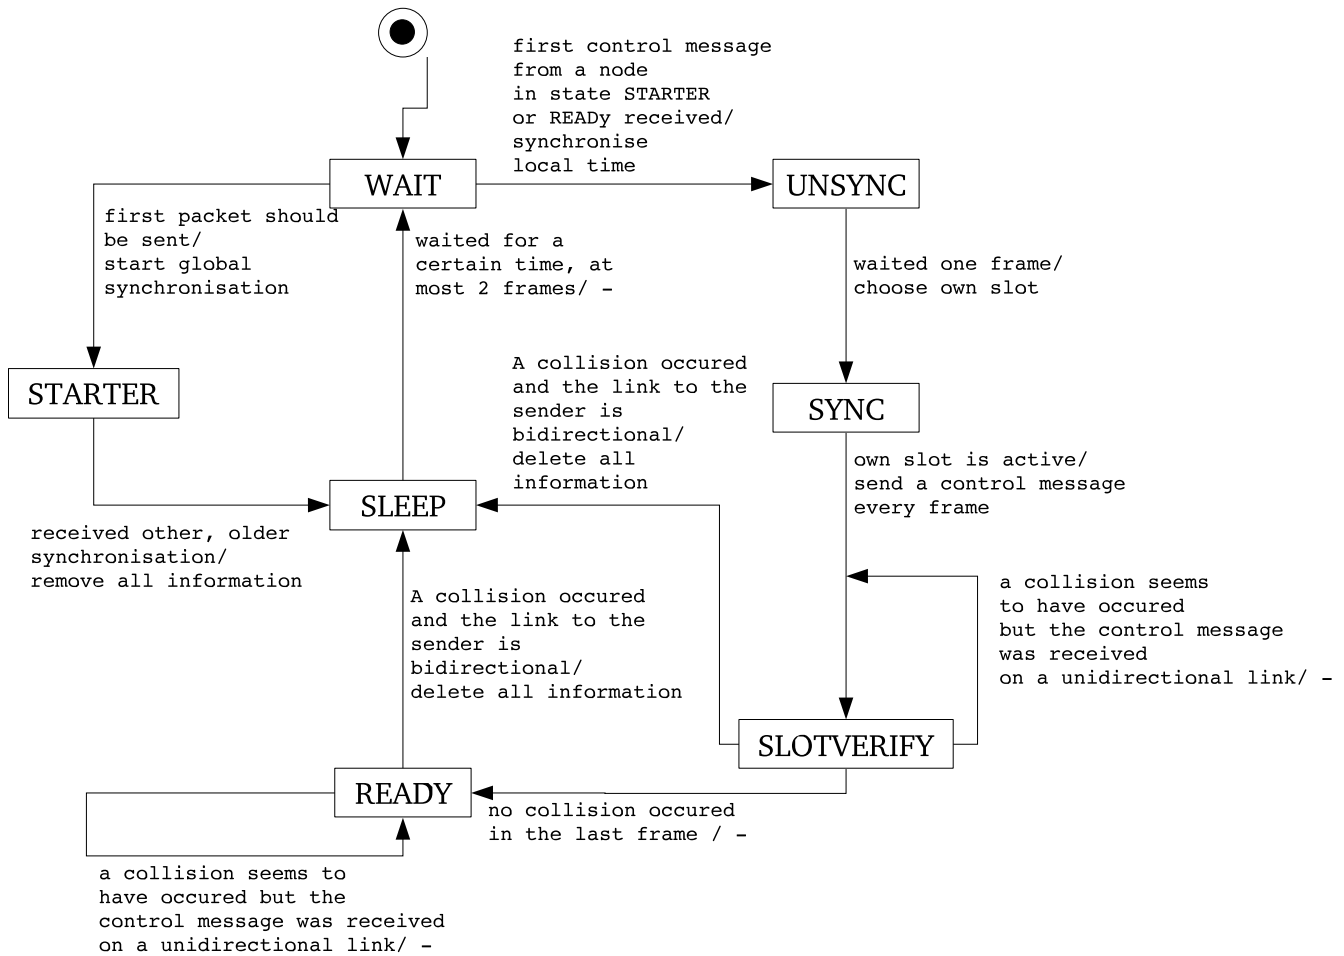


Fig. 2. The finite state machine used in MLMAC

decided to use this operating system for 2 reasons. First, the resource constraints of wireless sensor networks are not much different from those of deeply embedded systems. Second, reflex was already ported for the H8/300 processor family when we started our work on the MLMAC. Therefore, we were able to concentrate on the features of MLMAC, without having to worry about the underlying platform.

On top of Reflex, the authors of [5] implemented Copra, a configurable communication framework. Its layered structure supplies an easy way of switching between different implementations for a single layer, which we used in our tests to switch between MLMAC and the simple probabilistic collision avoidance MAC we used for comparison.

IV. EXPERIMENTS

We decided to experiment with a real sensor network right from the start for two reasons. First, as far as we know, LMAC has only been simulated. Second, we have observed drastic changes in link quality even in stationary sensor networks, as the authors of [9] have. As knowledge of these changes is vital for a MAC protocol, but we have not seen a realistic simulation model thereof, we think it is necessary to experiment with

real sensor nodes. For the same reason we did not simulate MLMAC but went to the real nodes right away.

In our experiments, we compared MLMAC to a competition based protocol, a simple collision avoidance protocol. As MLMAC is a TDMA based protocol, it should avoid collisions completely after a start period. A collision avoidance protocol on the other hand allows collisions but should produce a better throughput, as nodes do not have to wait for their turn to send. Thus, latency should be smaller.

In CSMA/CA the channel is sensed before transmission, and if it is occupied the transmission is delayed for a random time. This reduces the probability of collisions. Due to the transceivers we used, the implemented MAC could not be a real CSMA/CA. The radio modules use an inbuilt checksum and only packets that were received flawlessly are handed to the software. The radio modules only deliver whole packets, thus it is not possible to really listen to the medium. Therefore the channel is always assumed to be busy, and a random time is waited before transmission of a packet. In this paper, the resulting protocol is called CA-MAC.

A. Static single-hop Experiments

In the static single-hop experiments all robots were within direct radio range. The two MAC protocols were evaluated using 3, 5, and 9 robots. Every node wanted to transmit 50 Packets with a size of 49 Byte to all other nodes. Node 1 started the transmission and every other node followed suit as soon as it received an initial packet. The number of slots was configured in MLMAC to match the number of nodes in the experiments. As every slot had a length of 0.17 seconds, this resulted in frame lengths of 1.53 seconds (for 9 nodes), 0.85 seconds (5 nodes) and 0.51 seconds (3 nodes). The CA-MAC sent at random intervals, but over the whole length of the experiment it transmitted the same number of packets as were sent in the MLMAC experiments.




Fig. 4. Average of received packets in the static experiment

Figure 4 shows the results of the static experiments. The experiments with 3 and 5 nodes were repeated three times, the ones with 9 nodes two times. The numbers shown on the figure are the averaged results. As you can see the MLMAC always received all packets after the initialization phase was over. When using only three nodes, MLMAC even received all packets during initialization. Only during initialization with 5 or 9 nodes a few packets were dropped. The CA-MAC does not need to initialize, therefore no distinction is made. Due to the nature of the collision avoidance, the collision rate was high, though. Between a third and half of the packets were lost. This confirms that MLMAC fulfills all requirements for this scenario.

B. Static multi-hop line Experiments

In preparation for the mobile experiments we continued with two different experiments with a line topology. For these experiments all 9 robots were arranged in a row with nodes 1 and 9 at the ends. To find proper distances between nodes, one node transmitted packets periodically and the next node was moved until it was unable to receive those packets. Then, it was placed a step nearer to the transmitting node. The experiments were performed outdoors, on a concrete road. In the first experiments, node 1 transmitted 50 packets and node 9 counted how many arrived. Each node in between counted how many were forwarded and how many were dropped. The packets were sent in time intervals equaling the length of one,

three and five slots of MLMAC in three different experiments. In the second line experiment node 1 started transmitting 50 packets and once the first packet reached node 9, it started to transmit 50 packets, too. For the MLMAC, this was only done with sending intervals of 3 or 5 slots, as it is impossible to transmit 2 packets every slot without using some further means like e.g. aggregation. The packets were forwarded using a simple duplicate suppression algorithm. When a node receives a message it remembers the originator (node 1 or 9) and the sequence number of that packet. When a packet with the same or a lower sequence number was received, it was discarded, if its number was higher it was retransmitted. All experiments were conducted with MLMAC and with the CA-MAC and repeated 9 times for each of them. Again, the CA-MAC sent at random intervals, but over the whole length of the experiment these intervals evened out to those used in the MLMAC experiments.




Fig. 5. forwarded Packets for each node. Packets originate at node 1 and stop at node 9

Figure 5 shows some results from the first experiments. If each node would only be able to receive messages from its direct neighbors, all graphs should have been steadily decreasing or stayed at the same level as a packet that was lost at e.g. node 3 could never be received at the following nodes 4 to 9. Instead the figure shows that e.g. when using the CA-MAC with a transmission interval of roughly 5 slots node 3 received only 8 messages while its neighbors received 19 (node 2) and 28 (node 4). This reveals a strong variation in radio range of the sending nodes. In a few cases, node 9 even received packets from node 1 directly. The overall low throughput can be explained by the same fact because the number of slots per frame was set to low for the MLMAC, as we assumed that 3 or at most 4 slots would be enough. For the CA-MAC the variation in radio range produced more collisions as more nodes were able to disturb each other. Still, the results of MLMAC are much better than those of the CA-MAC when messages are transmitted every 3 or 5 slots.

Figure 6 shows the number of packets originating at node 1 forwarded by each node in the second experiments, while figure 7 shows the same for those originating at node 9. The results are similar to each other and to those of the first line experiments. A big difference can be seen in the results using




Fig. 6. forwarded Packets for each node. Packets originate at node 1 and are forwarded to node 9. Node 9 sends to node 1 at the same time.

MLMAC when sending every 3 slots. Obviously, node 1 had problems finding its slot as 42 messages were lost from node 1 to node 2. Otherwise, they confirm the results obtained earlier. The radio range is not stable and therefore the number of slots was too small.




Fig. 7. forwarded Packets for each node. Packets originate at node 9 and are forwarded to node 1. Node 1 sends to node 9 at the same time.

C. Static multi-hop square Experiments

After we saw the variations in link quality and radio range in the line experiment, we performed some experiments to determine the stability of links in time. Therefore, we arranged nine robots in a square of three times three with a distance of ten meters between neighbors. This distance was determined as above. After placement, we tried to build a routing tree and send messages from various starter nodes. As we assumed that the node in the middle would be heard by all, we only started building the trees from the nodes on the corners. The nodes one, three, four and seven which were on the corners of the square repeatedly tried to build a routing tree. Once this tree was established 10 messages were sent. After this, the number of received messages on each node was determined.

On figure 8 you see two routing trees, which were obtained in our experiments. On the left side you can see a routing tree that nearly follows the theory. Direct neighbors of node 1 can communicate directly, the nodes 8 and 4 that are farthest away from 1 need a second hop via node 5. Node 6 experienced some problems and was not able to transmit to node 1 even




Fig. 8. Two RCX routing trees started from node 1 at different points in time

though it can hear node 1 directly. This is shown as a dashed lines which represents an unidirectional link which would normally have been discarded. On the right side you see a different routing tree obtained from the same starting node a little time later. Now, all nodes can communicate directly with node 1 except for nodes 6 and 9 which have no connectivity at all. While these two routing trees are not optimal, they somehow still represent the expected layout.




Fig. 9. Two routing trees started from node 4 at different points in time and one started from node 3

On figure 9 you see the same experiment again, this time started from node 4. The left side of figure 9 is very much the same as the right side of figure 8. All nodes communicate directly with node 4, except for node 7 which does not communicate at all and node 1 which is connected through node 9. The middle however, shows a completely different picture. There are only three nodes which communicate directly with node 4, the nodes 3,5 and 7. Node 2 is connected via node 3 which is acceptable, even that node 1 uses node 3 to reach node 4 is no problem. The problem is, that node 8, which should be a direct neighbor of node 4 needs three hops to connect. This was no single effect either - when we started the protocol from node 3 which can be seen on the right side, node 1 needed three hops to connect, too. The trees obtained from node 7 were not much better, either. Some routing trees we obtained did not even reach half the nodes.

As said before, after building the trees we sent 10 messages from each source and counted the number of arriving messages. For these experiments, only the best two routing trees from each node were used. Note that the ones discussed before are those best trees, some others were quite disastrous. The number of received messages varied between 1 and 10 out of 10. In the case of the routing tree shown in Figure 8 on the left side node 4 received only one message from node 1. All results for both routing trees can be seen in figure 10, the

italic numbers represent the starter nodes. In the case of the routing tree shown in figure 9 on the left side, node 8 received three messages from node 4. This seems strange as the range of the longest hop is 22,4 meters instead of 28.3 meters but the number of hops increased from one to three. When node 3 was master and node 1 had a distance of three hops with a longest hop length of 20 meters eight packets were received out of ten. The total number of received packets for all three routing trees is shown in figure 11.

10	10	10
10	10	0
10	10	10

10	3	1
0	10	0
10	10	10

Fig. 10. Number of received messages out of 10 using the routing trees in figure 8

0	10	10
3	10	10
2	10	10

8	3	10
0	8	0
8	8	8

0	8	10
8	9	10
8	10	10

Fig. 11. Number of received messages out of 10 using the routing trees in figure 9

Similar effects have been observed in [9]. The authors describe an outdoor experiment in which they used 24 ScatterWeb ESBs [8] in an area of 80 times 140 meters. Every hour, a routing tree was built, and information sent to a sink. The experiment ran for three weeks. Often neighboring nodes did not hear each other while distant nodes did. In some cases these did not even have line of sight but were blocked by a building.

We have not yet seen any simulator that could produce such behavior, and we have not found any model to describe these effects either. All these results confirmed that we needed to evaluate MLMAC on a real sensor net in the mobile scenario.

D. Mobile Experiments

In the mobile experiment the RCX robots moved around a room of 6 times 8 meters containing normal bureau furnishing with a speed of 6 meters per minute in a nearly random manner. Initially, they only moved forward but every time they hit an obstacle, they moved back a pace, turned for a random time in a random direction and started moving again. As all nodes started with a different facing, the mobility soon led to a multi-hop environment. The application was the same as in the static experiment. 50 Packets of 49 Bytes were transmitted by each node and node 1 started the experiment. Due to mechanical failure, we could use only 8 robots, not 9 as in the static experiment. This resulted in a frame length of 1.36 seconds. Also, the Experiment with 3 nodes was ignored, as there would be no gain from it.

Figure 12 shows the results of the mobile experiments. Again, after the initialization, MLMAC delivered all packets flawlessly. During the initialization phase a few packets have been lost as was expected. The CA-MAC delivered between 18 and 27 percent of the packets. When the results are compared

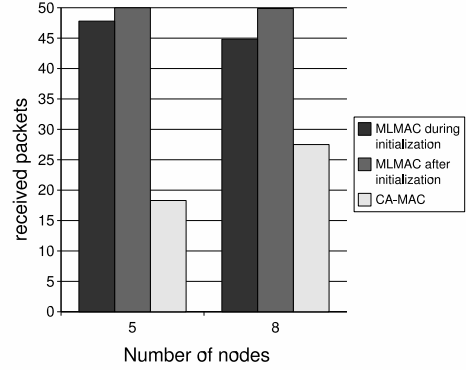


Fig. 12. Average of received packets in the mobile experiment

TABLE III
PROGRAM CODE SIZES OF A SIMPLE APPLICATION USING MLMAC OR CA-MAC, INCLUDING OPERATING SYSTEM AND DRIVERS

MAC Protocol	text	data	bss	total
MLMAC	14072	456	3522	18050
CA-MAC	10116	402	3376	13888

to the static experiment, it can be seen that mobility can be a challenge for medium access, the CA-MAC suffers heavily. MLMAC on the other hand performed nearly the same when mobility was used.

E. Program Code Size

Table III shows the size of the whole program code of a simple application that uses MLMAC or CA-MAC on the RCX robots. Note that this size depends on the architecture in use as primitive data types and pointers have different sizes on different platforms. The important part that you can see in the table is that MLMAC needs only about 4 kilobyte more than CA-MAC. This is less than a third of the size of CA-MAC. The overhead in program code size is acceptable for the gain in performance as shown earlier. Note also that the size measured is for our feasibility study implementation and we are positive that it could be reduced by careful re-implementation.

V. RELATED WORK

The authors of [10] have introduced LMAC, the lightweight medium access protocol for wireless sensor networks, which is the basis of this work. While some of their assumptions do not hold for our scenario, the main idea of representing the slots used by the neighboring nodes as a Bitmask can be found in MLMAC, too. One of their features, turning radio modules off to save energy, is not implemented in the current version of MLMAC. It could be added without much afford, though, if different hardware was used.

AI-LMAC is introduced in [7]. It is an enhancement of LMAC which allows dynamic reallocation of slots, depending on the network load. The authors assume a routing tree which

leads to a sink and optimize the slot usage along the branches of this tree. This is realized by the usage of so called Data Distribution Tables, which are used to determine the network load which results after a query from the sink. With this information, slots can be reserved according to the presumed needs. This approach does not fit into our scenario however, as we do not want to use a single sink. Rather, all nodes communicate among each other. Another difference is that MLMAC allows free choice of routing protocol.

VI. CONCLUSION AND FUTURE WORK

In this paper we have introduced Mobile LMAC, a novel TDMA protocol for mobile wireless sensor nodes. Mobile LMAC is based on the LMAC protocol but in our scheme each node can spontaneously establish a TDMA schedule on demand or join/leave existing schedules while nodes are moving. Thus a high channel throughput can be achieved for mobile sensor nodes even in heavy load situations. We have shown the feasibility of our protocol by means of experiments with a real mobile sensor network based on modified RCX robots. Furthermore, we compared the results with similar experiments using a simple collision avoidance MAC. The latter is reasonably well suited for sporadic communication between mobile nodes and does not require the exchange of additional protocol data. However, in the case of high load the Mobile LMAC achieved a far better result. In the future we will derive simulation models from our experiments and investigate the effectiveness of Mobile LMAC in large multi hop settings. We also plan to port MLMAC to the ScatterWeb Embedded Sensor Board (ESB) [8], which we can mount on top of the RCX robots (see figure 13). This way, the RCX robots will only be the platform for mobility and the radio module of the ESB can be used to experiment with power savings as in LMAC.




Fig. 13. An RCX with an mounted ESB

REFERENCES

- [1] I. Chatzigiannakis, S. Nikolettseas, and P. Spirakis. Smart dust protocols for local detection and propagation. In *POMC '02: Proceedings of the second ACM international workshop on Principles of mobile computing*, pages 9–16, New York, NY, USA, 2002. ACM Press.
- [2] I. Chatzigiannakis, S. Nikolettseas, and P. G. Spirakis. Efficient and robust protocols for local detection and propagation in smart dust networks. *Mob. Netw. Appl.*, 10(1-2):133–149, 2005.
- [3] circuit design inc. low power radio solutions. www.lprs.co.uk/main/product.info.php?productid=154.
- [4] J. M. Kahn, R. H. Katz, and K. S. J. Pister. Next century challenges: mobile networking for "smart dust". In *MobiCom '99: Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 271–278, New York, NY, USA, 1999. ACM Press.
- [5] R. Karnapke and J. Nolte. Copra - a communication processing architecture for wireless sensor networks. In *Euro-Par 2006 Parallel Processing*, pages 951–960. Springer, 2006.
- [6] H. Patterson-McNeill and C. L. Binkerd. Resources for using lego mindstorms. In *Proceedings of the seventh annual consortium for computing in small colleges central plains conference on The journal of computing in small colleges*, pages 48–55, , USA, 2001. Consortium for Computing Sciences in Colleges.
- [7] P. H. S. Chatterjea, L.F.W. van Hoesel. Ai-lmac: An adaptive, information-centric and lightweight mac protocol for wireless sensor networks.
- [8] J. Schiller, A. Liers, H. Ritter, R. Winter, and T. Voigt. Scatterweb - low power sensor nodes and energy aware routing. In *Proceedings of the 38th Hawaii International Conference on System Sciences*, 2005.
- [9] Turau, Renner, and Venzke. The heathland experiment: Results and experiences. In *Proceedings of the REALWSN'05 Workshop on Real-World Wireless Sensor Networks.*, Jun 2005.
- [10] L. van Hoesel and P. Havinga. A lightweight medium access protocol (lmac) for wireless sensor networks: Reducing preamble transmissions and transceiver state switches. In *INSS, Japan*, Jun 2004.
- [11] K. Walther, R. Hemmerling, and J. Nolte. Generic trigger variables and event flow wrappers in reflex. In *ECOOP - Workshop on Programming Languages and Operating Systems*, Jun 2004.
- [12] K. Walther and J. Nolte. Event-flow and synchronization in single threaded systems. In *Proceedings of First GI/ITG Workshop on Non-Functional Properties of Embedded Systems (NFPES)*, Mar 2006.