



Telecommunications Group

Departamento de Engenharia Electrotécnica

Faculdade de Ciências e Tecnologia da Universidade Nova de Lisboa

Technical Report

Title A Wireless Sensor MAC Protocol for Bursty Data Traffic

Doc. Ref. **Telecommunications Group TR-2007-01**

Date **March 2007**

Authors	Luis Bernardo	Rodolfo Oliveira	Miguel Pereira	Mário Macedo	Paulo Pinto
Checked by					

Technical reports are available at <http://tele1.dee.fct.unl.pt/tech-reports>.
The files are stored in PDF, with the report number as filename. Alternatively, reports
are available by post from the above address.



ABSTRACT

This paper proposes MH-MAC, a new MAC protocol for wireless sensor networks capable of handling applications that generate infrequent huge peaks of traffic. Existing protocols are not adapted to this kind of applications. Asynchronous protocols are energy efficient for the long inactive periods, but fail to cope with the bandwidth and latency requirements of the peak when more than two nodes are sending data to a common sink. Synchronous protocols that support contention free slots provide good throughput for handling the load peaks, but consume unnecessary energy maintaining clocks synchronized for very long idle periods. MH-MAC is a multimode hybrid protocol that can be configured by the application to run in asynchronous mode or in synchronous mode, with or without contention, providing the best possible trade-off. MH-MAC is a single hop MAC, which supports multi hop applications through a cross-layering API. The paper includes simulation results with the energy consumption, latency and throughput for the operation modes of MH-MAC, showing the asynchronous-synchronous trade-offs and the state transition overhead.

Index Terms – **Wireless Sensor Network, MAC protocol, Multimode, Cross-layer support**



TABLE OF CONTENTS

Abstract	ii
1 Introduction.....	1
2 MH-MAC design	2
2.1 Asynchronous Mode	2
2.2 Synchronous Mode	3
2.3 Mode Swapping	5
3 Application Programming Interface	8
4 Performance Evaluation.....	9
5 Conclusions and Further Work	11
Acknowledgements	12
References	12



1 INTRODUCTION

Energy efficiency is a dominant concern on the design of the medium access control (MAC) layer protocols for wireless sensor networks (WSNs). Nevertheless MAC protocols must also satisfy the application delay and throughput requirements. Applications that generate infrequent huge peaks of traffic pose a challenging problem for the existing MAC protocols.

Standard WSN MAC protocols usually are designed for periodic traffic, or for seldom traffic, but not for applications where both characteristics are needed at different instants. Duty cycling is a common mechanism for achieving energy efficiency. Nodes periodically cycle between an awake state and a sleep state. Protocols designed for seldom traffic, such as B-MAC [1] and X-MAC [2], let nodes run their duty cycles independently. They rely on *low power listening* (LPL), also called preamble sampling, to link together the sender and the receiver asynchronously. Packet sending is preceded by a large preamble, or a sequence of small preambles, larger than the duration of a duty cycle period. Protocols designed for periodic sending, such as S-MAC [3], T-MAC [4], SCP-MAC [5], and Z-MAC [6], require the additional clock synchronization overhead. Nodes run synchronized duty cycle periods. S-MAC [3] is a periodic synchronous protocol, which runs a CSMA (Carrier Sense Multiple Access) MAC contention resolution protocol during the fixed duration of the awake state. T-MAC [4] improves S-MAC by adapting the awake state duration to the load. If the radio is inactive for more than a threshold time the node goes asleep before the end of the normal awake duration time. SCP-MAC [5] introduces the scheduled channel polling technique to achieve awake duty cycle values as low as 0.01%. During a very short awake time nodes scan the medium for energy. If energy is detected, nodes stay awake and wait for a packet reception. This mechanism requires very precise synchronization between the sender and the receiver, due to the small durations proposed for the awake state and for the sender's awake signals (packets). A common problem for CSMA contention based protocols are collisions with nodes two hop distant, called the *hidden terminal* problem. A common solution to the problem is the RTS/CTS exchange. However, this solution can incur in high bandwidth overhead [1]. Z-MAC [6] improves CSMA using a TDMA contention free mode when the traffic increases. It introduces the overhead of creating and maintaining a global slot schedule. Nonetheless, WSN applications often create a sink tree [7] where the trunk links demand much more throughput than the leaf links. Z-MAC fails to cope with such applications because it divides TDMA slots evenly amongst neighbor nodes. Bursty traffic presents a challenge to the synchronous WSN MAC protocols, due to the high



synchronization overhead paid during the idle periods, when no packets are flowing. On the other hand, asynchronous protocols delay packet sending and limit the maximum throughput.

This paper proposes MH-MAC, a multimode hybrid MAC protocol that is capable of running in asynchronous mode, contention synchronous mode, and contention-free synchronous mode. MH-MAC allows applications to operate in the asynchronous mode for most of the time, and change to contention-free synchronous mode during the data traffic peaks, optimizing the overall performance.

In the following, Section 2 presents the MH-MAC protocol. Section 3 describes MH-MAC application programming interface. Section 4 evaluates the protocol performance using TOSSIM [8] simulations, and Section 5 provides a discussion of future work and our conclusions.

2 MH-MAC DESIGN

Multimode hybrid MAC protocol (MH-MAC) is designed to support cross-layering applications for packetizing radios, like the Chipcon CC2420. MH-MAC is optimized for convergecast applications, where mobile sink nodes collect data asynchronously from sensors of isolated fixed WSNs. MH-MAC can be in one of three states: asynchronous state; synchronous state; or the full on state, where the node does not sleep. By default MH-MAC state is asynchronous, but applications can change it to full on, or synchronous. In the full on state, data packets are preceded by an RTS/CTS exchange, and are acknowledged. In the synchronous state, temporarily contention-free slots can be reserved for the communication with neighbor nodes, trading off energy for throughput and delay. The following subsections present the MH-MAC operation modes associated with the two duty cycling states, and the state transition algorithms.

2.1 Asynchronous Mode

The MH-MAC asynchronous mode runs a LPL algorithm similar to the X-MAC protocol [2]. Senders send a sequence of short preambles during at most twice the time of the duty cycle period, before sending the data packet. This assures that the receiver is awake when the data packet is sent. The preambles contain the destination address and are separated by pauses. Like in X-MAC, MH-MAC allows unicast receivers to send early preamble acknowledgments as soon as they are awake and receive a preamble. This shorts the senders' total preamble duration (Fig. 1). However, for broadcast, the full preamble is required. MH-MAC improves X-MAC overhearing protection, by also including in the preamble the data packet sending time. Broadcast receivers (Fig. 2) use this



information to schedule a radio sleep until the beginning of the data packet transmission. Unicast receivers use it to schedule a precise sleep period. X-MAC uses random sleep intervals when a node is waiting to send a packet and it receives preambles not destined to itself.

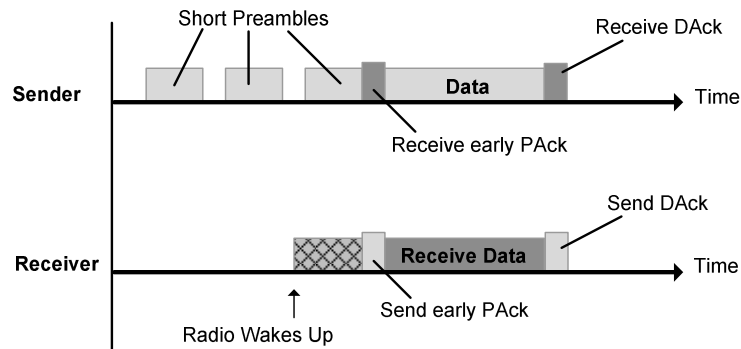


Figure 1: Unicast asynchronous transmission.

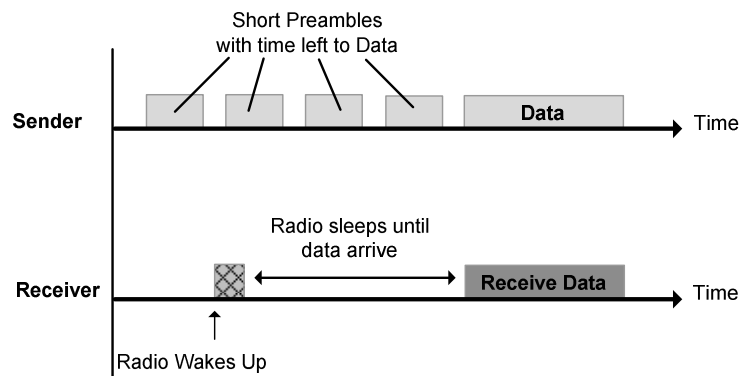


Figure 2: Broadcast asynchronous transmission.

MH-MAC defines an additional "SHUT-UP" packet, for handling preamble collisions. If a receiver hears more than one preamble, it can send a SHUT-UP packet to one of the senders, with a configurable probability p . THE SHUT-UP packet includes the preamble time remaining until transmission, allowing the receiving sender to sleep until the end of the other packet transmission. Inter-preamble time includes a jitter that improves the probability of not having collisions between to concurrent senders.

Only unicast data packets are acknowledged. After a data packet reception, the receiver stays awake for a short period, waiting for possible new packets before returning to sleep.

2.2 Synchronous Mode

The MH-MAC synchronous mode was designed to optimize data collection from a distributed set of sensor nodes into a single sink, over a sink tree. Therefore, it provides basic flooding,



neighbor detection, and slot conflict resolution functionalities. MH-MAC organizes the duty cycle period into a sequence of fixed length slots (100 ms). Each slot is capable of carrying more than eleven data packets with 112 bytes on 802.15.4 radios. Nodes run synchronized duty cycle periods. Each node has one or more public slots, and zero or more dedicated slots to communicate with specific neighbors.

During public slots, nodes run a contention-based protocol similar to the one proposed in T-MAC [4]. Senders run a back off timer and scan the network before sending an RTS or Data packet, respectively for unicast and broadcast transmissions. Unicast packets are preceded by a RTS/CTS exchange and are acknowledged. If no energy is detected in the channel in a public slot for 20 ms, the receiver node goes into sleep. Dedicated Slots are reserved for unicast communications between two nodes. Senders run a short random back off timer before scanning the networks' energy and sending the data packets. Data sent through dedicated slots is also acknowledged, but no RTS/CTS are sent.

In order to maintain synchronization, nodes send sparse SYNC packets on their public slots, defining the beginning of the duty cycle period. The SYNC packet contains the local address, the local slot assignment plan, and a hop counter. Initial slot assignment is done radially, from the sink to the farthest located nodes, taking into account the slots occupied by neighbor nodes. Each node keeps track of its public slot, its dedicated slots, the slots occupied by other neighbor nodes, and other neighbor public slots. A SYNC packet has a maximum validity of 180 seconds, and its information is discarded after that time. Nodes are asleep during empty slots or slots occupied by neighbors.

Nodes located at most three hops away can use the same public slot. This distance allows a message to go through two nodes, and reach the third node before the end of the public slot inactivity time. Every three hops, nodes must use different public slots, organized in a staggered wakeup schedule that minimizes the source to sink delay [9]. Nodes in the border have two public slots: one for communicating with its ascendant node connecting it to the sink; and another one connecting it to their descendent nodes. The contention-based access algorithm is run for both slots, with two distinct groups of nodes.

Nodes in asynchronous or full on modes can exist in the neighborhood of nodes in synchronous mode. On these cases, asynchronous and full on nodes must maintain a table with the public or dedicated slots of each synchronous neighbor node, and send packets to them starting the transmission on their public slots. When the destination node state is unknown, synchronous nodes

must send a preamble preceding the data packets. If the awake periods are aligned, the preamble overhead is minimized by the immediate reception of an early preamble acknowledge. Notice that collisions might occur on dedicated slots when non-synchronous nodes exist in the neighborhood of synchronous nodes. It always occurs during the transition to synchronized state, but can be avoided afterwards on a static WSN network if all nodes are set to the same state.

Fig. 4 illustrates a possible slot allocation schedule for the sink tree shown in Fig. 3. An eleven slot duty-cycle period is used and dedicated slots were allocated for all connections. All nodes are within two hops, and share the same public slot (0). Fig. 4 presents the slots occupied by neighbors in grey, and the local dedicated slots in black. Node A uses the slots 3, 6, and 10 dedicated respectively to communications with nodes B, D and C. Node A has slots 5, 8, and 9 occupied by neighbor nodes C and D. Applications running on branch nodes can allocate slots that minimize packet propagation delay to the sink node, like in the Fig. 4. The algorithm looks for the nearest free slot before a reference slot. Usually, applications specify the address of a node connected to the sink. In Fig. 4, C would provide A's slot when it assigns slots to F and G. Otherwise, MH-MAC distributes slots randomly. Notice that MH-MAC only manages synchronous connection at a local level. Using MH-MAC API at each node, an application running on all nodes can create synchronous connections covering the entire WSN.

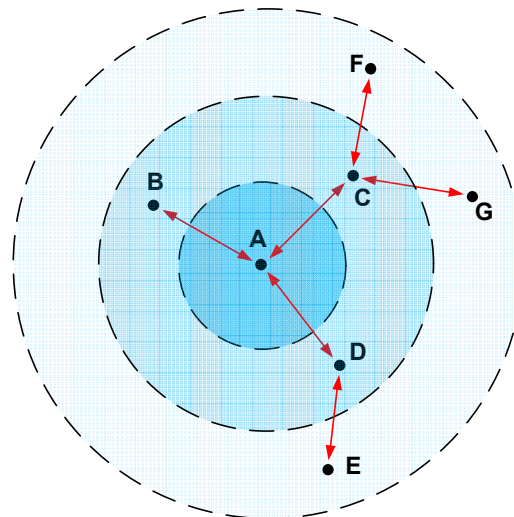


Figure 3: A synchronous WSN topology example.

2.3 Mode Swapping

Applications can modify the MH-MAC state, using the application programming interface presented below. State changes are notified to the neighbor nodes using *Hello* packets. These



packets have an MH-MAC state field, and a *slot reserve bit* (SRBit). *Hello* packets are always broadcasted using the asynchronous mode algorithm, to allow all neighbor nodes to detect the state changes.




Figure 4: Example slot assignment for Fig.3's WSN.

The most demanding state change occurs from the asynchronous state to the synchronous state. Fig. 5 illustrates the packets exchanged during this transition. Node A broadcasts the *Hello* packet preceded by a sequence of preambles, signaling the synchronous state. After receiving this packet, nodes B and C send an event to the application, together with the SRBit and the sender address. The application can decide to accept and store the A's public slot schedule, or to ignore it. When a node accepts the synchronization request, it starts a random back off timer, senses the network for other transmissions, and finally transmits a *Request* packet. If the node does not receive an *OK* packet, it restarts the back off timer and repeats the procedure, until a maximum time of 1100ms after the reception of the *Hello* packet.

The *Request* packet also carries a SRBit. The *Request/OK* exchange is used to assign a dedicated slot when the SRBit in the *Request* packet is set. When the synchronous node sends a *Hello* packet with the SRBit set, it requests every neighbor to ask for dedicated slots. The synchronization node assigns the dedicated slots for each *Request* packet received. After waiting for 1100ms, the synchronization process ends, and an event is generated to the sender's application with the list of neighbors and the dedicated slots assigned. A SYNC packet is also broadcasted, defining the instant when the neighbors can also start their synchronization process. The application on each neighbor starts its synchronization independently. The total duration of the state transition from the asynchronous state to the synchronous state on each node is equal to the preamble duration (twice the duty cycle period) plus 1100 ms.




Figure 5: Asynchronous to synchronous state change.

Nodes must change their state to synchronous, starting from the sink node, to synchronize the duty-cycle periods within the sink tree. In order to avoid packet collisions, only one node can run the state change protocol within one hop radio range at each instant. Interferences can occur at two hop distant nodes or more [10]. However, nodes use a RSSI interference detection method to avoid collisions. Even if the initial synchronization fails, nodes can still change their state when a SYNC packet is latter received. The maximum time to change all nodes on a WSN to the synchronous state is given by (1) if the radio interference effects are ignored. v_{max} represents the maximum number of neighbors, r represents the depth of the sink tree, and T the duty cycle period. The effective time can be lower because some nodes have less than v_{max} neighbors. In order to have a faster transition to the synchronous state, the cycle time period must be shorter, resulting in less dedicated slots available and a more awake time. For the eleven slot duty cycle period presented in Fig. 4, T is equal to 1100ms. The time to set the seven nodes presented in Fig. 3 to synchronous mode would be 19.8 seconds.

$$TotalSyncTime \leq (2 + r(v_{max} - 1)) \times (2T + 1100) [ms] \quad (1)$$

Dedicated slots can also be canceled, created or reassigned after the initial synchronization setup phase, running the "Request" / "OK" packet exchange during the public slot time. Due to limited number of slots available (10 by default) some neighbors may not have dedicated slots assign to



them. These neighbors can use the public slots for data transfer. However, MH-MAC allows applications to reassign the slots. When the synchronization protocol ends, MH-MAC sends an event with the information about the neighbor lists and the slots allocated to the application. The application can then use the API to reassign the slots, to adapt it to the sink tree defined by the application.

3 APPLICATION PROGRAMMING INTERFACE

In order to support the cross-layering interaction between MH-MAC and the application and transport layers, the application programming interface (API) defines the set of commands (com.) and events presented in Table 1. The commands interface is implemented by MH-MAC and the event interface is implemented by the applications.

Table 1: Application Programming Interface

Type	Command
com.	<i>Asynchronize</i>
com.	<i>FullOn</i>
com.	<i>Synchronize (addr, dedic, [ref])</i>
com.	<i>StopSynchronize(addr)</i>
com.	<i>IsSynchronous (addr)</i>
event	<i>AsynchronizeDone</i>
event	<i>FullOnDone</i>
event	<i>SynchronizeDone(neig[], dedic[])</i>
event	<i>SynchronizeReq(addr, dedic)</i>

The *Asynchronize*, *FullOn*, and *Synchronize* commands are used to modify the MH-MAC state. The associated events *AsynchronizeDone*, *FullOnDone*, and *SynchronizeDone* are generated when the state change protocol ends. The *Synchronize* command starts the synchronization protocol presented in the previous section. It also requests *dedic* dedicated slots assignment to a unicast address, or one to all addresses, when *addr* is the broadcast address and *dedic* is not zero. The optional *ref* parameter defines where the search for free slots begins. The *SynchronizeDone* event returns the neighbor list and the slots assigned to each neighbor. Applications can use the *IsSynchronous* command to know if *addr* is synchronized with the local node. MH-MAC sends the application an event *SynchronizeReq* when a neighbor sends a synchronization request. The application can use the *StopSynchronize* to reject the request, or the *Synchronize* command to accept it. Additional commands are available to control the duty cycle period duration.



4 PERFORMANCE EVALUATION

The MH-MAC prototype was implemented in TinyOS 2.0 [11] and was tested on Xbow Telos B motes. However, due to the small number of motes available for this project, the performance evaluation was done using the TOSSIM simulator [8]. The current TOSSIM version does not support the CC2420 radio stack used by the LPL (low power listening) library. Therefore, we emulate the CC2420 radio stack and modified TOSSIM interface implementations to resolve the synchronization problems that occurred when the radio interface is turned off. Additionally, meters were placed on the MAC code to measure the number of milliseconds used for data transmissions, for data receptions, and the time spent in active and radio sleep states. Using the current consumption specifications shown in Table 2, we were able to estimate the total current consumption for the three states of MH-MAC. Following [2][5], we considered that in idle or receiving state the mote has the consumption of operation MCU+Radio RX, in radio sleep has the consumption of operation MCU Active, and during packet transmissions has the consumption of operation MCU+Radio TX.

Table 2: Xbow Telos B current consumption [12]

Operation	Current
Mote Standby (RTC on)	5,1 μ A
MCU Idle (DCO on)	54,5 μ A
MCU Active	1,8 mA
MCU + Radio RX	21.8 mA
MCU + Radio TX (0dBm)	19,5 mA

We analyzed a single hop scenario where several nodes send packets to a single sink, for the three MH-MAC modes. A duty cycle period with eleven slots is used on the synchronous mode, supporting one public slot plus ten dedicated slots. In the synchronous state nodes generate SYNC packets every 60 seconds. Nodes generate 100 bytes data packets spaced with an average inter packet time (IPT). The interval between packet transmissions is a random variable with a uniform distribution in the interval $[0.5 \times \text{IPT}, 1.5 \times \text{IPT}]$. The load is uniformly distributed over the nodes.

Fig. 6 shows how the current consumption depends on the IPT value, with four active senders. Results show that when the interval between data packets is large, asynchronous mode (X-MAC) optimizes energy savings because it maximizes the time the nodes are asleep. However, when IPT is small, its current consumption increases significantly due to the preamble overhead. Synchronous mode presents the best energy efficiency for high data rate conditions, where IPT is very short.




Figure 6: Current consumption with 4 nodes and one sink.

Fig. 7 shows the maximum throughput measured when a varying number of nodes send a burst of 20 packets to a sink node at the MH-MAC mode's maximum speed. Results show that the asynchronous mode (X-MAC) is very efficient when one or two neighbor nodes compete, as reported in [1][2]. However, they also show that the asynchronous mode's throughput, and therefore B-MAC and X-MAC's throughput, collapse when three or more senders compete for the access to a single sink. For more than four senders, the synchronous mode outperforms all other modes, maintaining a controlled current consumption.

Fig. 8 shows how the latency evolves with the number of nodes, for an IPT of 1 second. It clearly shows that asynchronous mode is only effective for up to two active senders in the simulated conditions. For more than two active senders, the network latency increases fast with the number of senders. Full on mode is clearly the mode that minimizes the network latency, at the cost of also maximizing the current consumption. Therefore, synchronous mode presents the best trade-off for a tree shaped network, with more than two active neighbors sending data to a single sink node.




Figure 7: Throughput for a burst of 20 packets per node.

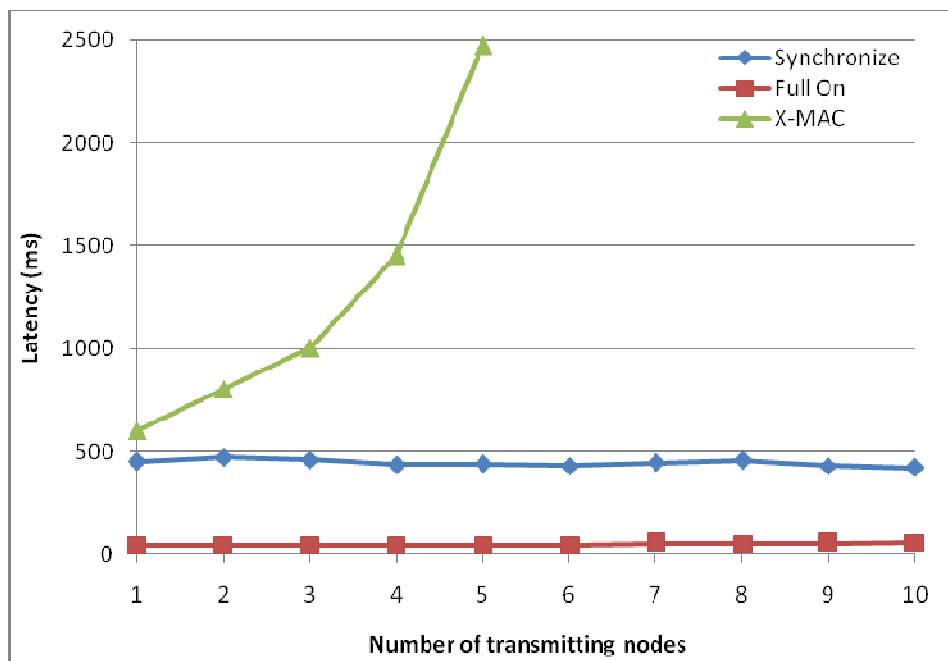


Figure 8: Latency for a burst of 20 packets per node.

5 CONCLUSIONS AND FURTHER WORK

In this paper we propose MH-MAC, a new multimode MAC protocol designed specifically for cross-layering applications. Simulation results show that each mode has a different scenario where it is advantageous: asynchronous for low power sporadic communication; full on for very low



delay; synchronous for handling data bursts. The synchronous mode setup delay can be amortized if large bulks of data are transferred. MH-MAC is being used for the development of asynchronous sensor monitoring and alarm applications in scattered WSNs. Mobile vehicles drive around the WSNs from time to time, collecting all the data stored since the last visit. MH-MAC improves the energy efficiency.

Future work in MH-MAC includes the improved handling of radio interference (inter-slot and on synchronization deployment), and the reduction of the asynchronous to synchronous transition overhead. We also envision improved WSN mobility support.

Acknowledgements

We thank David Moss for the support on the development of MH-MAC. This work was partially supported by the *Fundação para a Ciência e Tecnologia* under the project SIGAPANO POSC/EIA/62199/2004.

REFERENCES

- [1] J. Polastre, J. Hill, D. Culler, "Versatile low power media access for wireless sensor networks," in: Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys), 2004, pp. 95–107.
- [2] M. Buettner, G. Yee, E. Anderson, R. Han, "X-MAC: A Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks", The 4th ACM Conference on Embedded Sensor Systems (SenSys), 2006, pp. 307-320.
- [3] W. Ye, J. Heidemann, D. Estrin, An energy-efficient MAC protocol for wireless sensor networks, in: Proceedings of the Joint Conference of the IEEE Computer and Communications Societies (InfoCom), Vol. 3, 2002, pp. 214–226.
- [4] JT. van Dam, K. Langendoen, "An adaptive energy-efficient MAC protocol for wireless sensor networks," in: Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys), 2003, pp. 171–180.
- [5] Wei Ye, Fabio Silva, and John Heidemann, "Ultra-Low Duty Cycle MAC with Scheduled Channel Polling," in: Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys), 2006, pp. 321-334.
- [6] II. Rhee, A. Warrior, M. Aia, J. Min, "Z-MAC: A hybrid MAC for wireless sensor networks," in: Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys), 2005, pp. 90–101.
- [7] K. Akkaya, and M. Younis, "A survey on routing protocols for wireless sensor networks," Ad Hoc Networks, vol. 3, no. 3, pp. 325-349, May 2005.
- [8] P. Levis, N. Lee, M. Welsh, and D. Culler, "TOSSIM: Accurate and Scalable Simulation of Entire TinyOS Applications," in: Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys), 2003, pp. 126-137.
- [9] G. Lu, B. Krishnamachari, and C. Raghavendra, "An Adaptive Energy-Efficient and Low-Latency MAC for Data Gathering in Wireless Sensor Networks," in: Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS), 2005, pp. 224- 231.
- [10] J. Grönkvist, A. Hansson, "Comparison between Graph-Based and Interference-Based STDMA Scheduling", in: Proceeding of ACM Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc), 2001, pp. 255-258.
- [11] TinyOS 2.0 Documentation. <http://www.tinyos.net/tinyos-2.x/doc/>.
- [12] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling Ultra-Low Power Wireless Research," in Proceeding of International Symposium on Information Processing in Sensor Networks (IPSN), 2005, pp. 364- 369.