

# $\mu$ -MAC: An Energy-Efficient Medium Access Control for Wireless Sensor Networks

Andre Barroso, Utz Roedig and Cormac Sreenan  
Mobile & Internet Systems Laboratory, University College Cork, Ireland  
Email: {a.barroso u.roedig c.sreenan}@cs.ucc.ie

**Abstract**—For the long-term deployment of wireless sensor networks, energy efficient MAC protocols are necessary. The transceiver of a sensor node should only consume energy while actively taking part in communication. Energy consumption in idle mode should be avoided as much as possible.

In this paper it is shown how application layer knowledge in the form of flow specifications can be used to improve the energy properties of a MAC protocol. A new protocol, named  $\mu$ -MAC, is proposed and evaluated through simulations.

## I. INTRODUCTION

Ideally, a MAC protocol for wireless sensor networks should consume energy only while transmitting or receiving useful data, not in idle (listening) mode. This property allows for an extended lifetime of nodes and therefore of the entire network. Long lifetimes is an essential requirement for the deployment of low-maintenance wireless sensor networks. This requirement led to the proposal of different MAC protocols targeting the special characteristics of wireless sensor networks. A review of some of these protocols is provided in Section II. Essentially, these MAC protocols reduce idle listening by periodically switching transceivers to sleep mode. By adjusting the radio duty-cycle, defined as the ratio between sending/ receiving/idle and sleeping time, MAC protocols can achieve better energy consumption characteristics.

Wireless sensor networks constitute a particular class of networks in which traffic presents much more predictable characteristics in several applications. Very often, the traffic is comprised of periodic packets flowing from different nodes towards a fixed set of nodes called base-stations. In a typical scenario, the base-station broadcasts a message (called an interest) containing a specification of the requested information to the network. Addressed nodes subsequently start reporting collected data back to the base-station. These reports are sent with a certain frequency and for a long time for monitoring

purposes. The knowledge of the periodicity and duration of these reports can be used to increase the efficiency of radio operation by the MAC layer in a sensor network. In this paper it is shown how flow specifications can be translated in reservation requests that are implemented in the MAC layer of a sensor network with the purpose of improving its energy consumption properties.

The rest of this paper is organized as follows: In Section II, related work in the area is outlined. Section III describes design goals and metrics used for the protocol proposed in this paper. In Section IV, a new media access control protocol for wireless sensor networks, named  $\mu$ -MAC, is presented. Section V gives evaluation results for  $\mu$ -MAC. The paper concludes with a summary and an outline of future work.

## II. RELATED WORK

S-MAC [1] is one of the best-known protocols designed specifically for wireless sensor networks. S-MAC is a contention-based protocol and achieves energy conservation through three basic techniques: nodes periodically sleep instead of constantly listening to an idle channel; radios are turned off while the shared media is used for transmission between other nodes (overhearing avoidance), and a message passing scheme is used to reduce contention latency for sensor network applications that require store-and-forward processing as data is moved through the network. Each node has a basic radio duty-cycle that can be tuned to trade-off bandwidth and latency for energy savings. Bandwidth reduces dramatically as the duty-cycle decreases since neighbours contend in smaller windows of time for transmission. Furthermore, very low duty cycles impose higher delays in the delivery of data and larger buffers to store data collected/received. S-MAC mitigates these problems when large messages need to be transmitted by fragmenting them and preventing a pair of nodes to sleep before all the fragments are transferred. The protocol

requires periodic synchronization among neighbours so they can agree on sleep/active schedules.

The use of radio fixed duty-cycles in sensor nodes can waste considerable amounts of energy since the communication subsystem is activated even though no communication will take place. During periods of high activity, the radio must be turned on long enough to handle the traffic. However, as the load decreases, the radio will remain essentially idle. T-MAC [2], another contention-based protocol, addresses this problem by employing an adaptive active/sleep duty cycle for the radio operation. The protocol reduces idle listening by periodically transmitting all messages in bursts of variable length, and sleeping between bursts. The active period is ended simply by timing out on hearing no activation event (e.g. messages). The time-out period defines the minimum amount of idle listening in a duty-cycle. T-MAC suffers the same high latency, large buffer size problem as S-MAC when the duty-cycle is very low.

TRAMA [3] is a schedule-based protocol that provides energy-efficient conflict-free channel access in wireless sensor networks. As a schedule-based protocol, TRAMA is able to deliver good energy savings at the expense of delay but it has a fairly complex mechanism for assigning transmission slots to nodes.  $\mu$ -MAC shares a common architecture with TRAMA, with the communication channel divided into a contention and a contention-free period.  $\mu$ -MAC uses, however, a different slot reservation mechanism and relies extensively on information provided by upper layers to improve radio utilization.

### III. DESIGN

In this section, the design rationale for  $\mu$ -MAC is presented and assumptions for the correct operation of the protocol are stated.

#### A. Design Metrics

A typical wireless sensor node consumes most of its energy in communication activities/hardware. Therefore, energy consumption can be greatly reduced by optimizing the communication components of a sensor. Transceiver hardware, MAC layer, routing layer or application communication patterns are all elements that can be optimized to achieve energy reduction. This paper focuses specifically on the MAC layer.

A MAC protocol directly drives the transceiver operation of a sensor node and therefore it ultimately defines how efficiently the transceiver is used. A modern

transceiver chip (e.g. [4]) can be operated in the following four states: **send**, **receive**, **idle** (listening) and **sleep**. In terms of energy consumption, the send, receive and idle states consume approximately the same amount of energy (in the order of mA). Time spent in any one of these three states is referred as awake time. In sleep mode, the transceiver consumes significantly less energy (in the order of  $\mu$ A). Thus, energy can be significantly saved by the MAC layer if the transceiver is kept as much as possible in sleep mode. An important goal of a MAC protocol is to increase the percentage of sleep time of the radio. Clearly, this goal should not be achieved by completely disregarding delivery rate and communication latency. These metrics are basic to evaluate any MAC protocol designed for wireless sensor networks.

#### B. Design Approach

The goal of this paper is to define a MAC protocol with high sleep ratios while keeping message latency and reliability at acceptable levels.

To achieve this goal, the following two structural design choices were made:

- 1) Adoption of a schedule-based approach to multiplex access to the shared medium;
- 2) Use of application layer knowledge in the MAC layer in the form of flow specifications.

Schedule-based protocols avoid energy costly overhearing and retransmissions due to collisions. On the other hand, such protocols have stronger time synchronization requirements. Application layer knowledge in sensor networks can be easily integrated in the operation of the MAC layer because of the traffic nature of such systems. Reports in sensor networks are normally generated periodically and the rate and duration of a flow are commonly specified in the interests issued by the base-station. Flow specifications are used to perform bandwidth reservation in the MAC layer. The reservation mechanism restricts idle listening thus improving energy efficiency.

#### C. Preconditions

The correct and efficient operation of  $\mu$ -MAC depends on a few preconditions. As a schedule-based protocol, clock synchronization between nearby nodes is required. An external clock synchronization scheme, such as a beacon source, is assumed, although in-band clock synchronization is possible. Switching intervals between scheduled slots compensate for errors in the clock synchronization mechanism. The sensor field should not

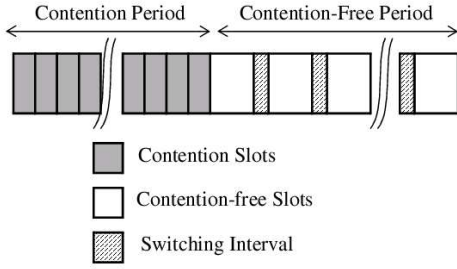


Fig. 1. Time slot organization

be subject to frequent topological changes. Contention periods, responsible for network organization, incur a large overhead and must take place infrequently. Adaptation to highly frequent structural changes in the field is therefore not possible. Last, knowledge of the traffic pattern must be available. In many operation scenarios, the base-station generates a request for data containing the desired traffic information. For these scenarios, the precondition can be fulfilled.

#### IV. THE $\mu$ -MAC PROTOCOL

$\mu$ -MAC assumes a single time-slotted channel, organized as depicted in Fig. 1. Protocol operation alternates between a contention and a contention-free period. The contention period purpose is to build a network topology and to initialize transmission sub-channels, *i.e.*, collection of related slots in the contention-free period. The contention-free period is used to transfer data between nodes. Transmission in slots of this period are scheduled to avoid collisions.

During the contention-free period, data is transferred in collections of related slots called *sub-channels*.  $\mu$ -MAC differentiates between two classes of sub-channels: *general traffic* and *sensor reports*. This distinction reflects the asymmetry existent in many sensor networks, where the data flow from sensors to the base station is predominant over all other traffic. General traffic is comprised mainly of interests from the base-station and routing setup to collect sensor data. In general, this traffic has no preferential flow direction and is broadcasted in the network. Sensor reports, on the other hand, flow towards the base-station and are generally comprised of periodic data packets transmitted during long periods of time.

In  $\mu$ -MAC, the general traffic sub-channel is formed by a number of approximately equal spaced contention-free slots. Given the nature of the traffic carried, every node must listen to packets transmitted during these

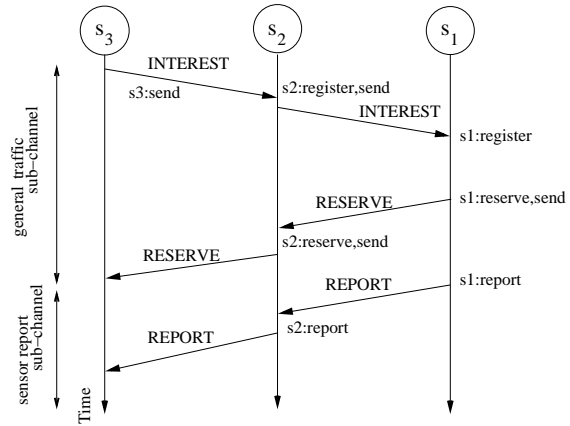


Fig. 2. Protocol Operation

slots. The bandwidth assigned to this sub-channel must be set to meet a trade-off between energy consumption and latency. Sensor reports sub-channels, on the other hand, have a variable bandwidth that is automatically adjusted in each hop according to the traffic expected.

##### A. Flow Reservation Mechanism

In order to allow a reservation for upcoming flows the MAC layer provides the functions **send**, **report**, **register** and **reserve**). These functions might be used in a transparent way by the application layer by encapsulating them within operating system functions used to open and maintain sockets.

$\mu$ -MAC provides two commands for data communication in the network: **send** and **report**. The **send** command transfers data using the general traffic sub-channel, while the **report** command uses sensor report sub-channels.

Besides data communication,  $\mu$ -MAC provides two commands involved in the construction of sensor report sub-channels: **register** and **reserve**. In providing the capability for upper layers to build sub-channels according to their needs,  $\mu$ -MAC is able to generate radio duty-cycles that are not over or under dimensioned.

##### B. Protocol Operation

A typical operational scenario for  $\mu$ -MAC can be divided into three stages. A graphical representation of these stages is shown in Fig. 2. The figure depicts three nodes:  $s_3$  is the base station issuing an interest;  $s_2$  is an intermediary node; and  $s_1$  reacts to the interest by generating reports.

In the first stage, the base-station requests data from the network by issuing interests. An interest contains the type of data requested, the frequency the data should be reported and the duration of the interest. The routing

Type	SourceAddr	DestAddr	DeleteNum	AddNum	Deleted NodeIDs
Added NodeIDs		GT_SlotID_#1	...		GT_SlotID_#n
SR1_SlotID_#1	...	SR1_SlotID_#n	...	SRN_SlotID_#1	...

(a) Control Header

Type	SourceAddr	DestAddr	InterestID	SlotID_#1	...	SlotID_#n
------	------------	----------	------------	-----------	-----	-----------

(b) Reserve/Beacon/Data Header

Fig. 3. Packet Header Formats

layer broadcasts interests through calls to the send command initiating the first stage. As the path towards the base-station is established across the entire network, the routing layer registers with  $\mu$ -MAC the MAC address of the next hop towards the base-station and the associated interest. The register command causes the protocol to bind the interest with an available outgoing sensor report sub-channel in the node. This binding is used in the next stage when bandwidth for the interest needs to be reserved.

If a sensor is able to produce the data requested in the interest, it issues a reserve command to  $\mu$ -MAC triggering the second stage. The protocol activates the sub-channel adjusting its bandwidth according to the frequency of reports. The protocol also initiates a cascading process of sub-channel activation until a continuous sub-channel exists between the sensor producing the data and the base-station collecting them.

After requesting the reservation of a sensor report sub-channel, the sensor generating data may start reporting them towards the base-station. Reports are transferred between nodes through calls to the report command (third stage).

Within the interest a duration is specified. This duration is also used for the reservation. After this time period, reservations expire. A reservation might also time out if a channel is not used, even if a reservation exists. In the case a reservation request cannot be fulfilled, a reject message is reported back to the sink, indicating that the interest can not be handled by the network.

In the remainder of this section, the mechanisms involved in the operation of  $\mu$ -MAC are presented in greater detail.

### C. Contention Period

The purpose of the contention period is twofold: building a network topology and initializing transmission sub-channels. These objectives are met by the exchange

of packets in time slots subject to collisions. In order to transmit a packet, a node chooses a slot and sends it. If two or more packets are transmitted in the same slot, collision occurs and the packets are lost. Packet delivery reliability can be increased by multiple retransmissions.

*Topology Discovery.* The transmission scheduling mechanism of  $\mu$ -MAC, which prevents collisions during the contention-free period, is based on the Neighbourhood-aware Contention Resolution (NCR) algorithm [5]. In order to schedule collision-free transmissions, NCR requires each node to be aware of the local topology of the network, up to a distance of two hops away.  $\mu$ -MAC applies the Neighbour Protocol of TRAMA [3] for local topology discovery.

Fig. 3a depicts the control packet used by  $\mu$ -MAC during the contention period for achieving the objectives of topology discovery and sub-channel initialization.

*Sub-Channel Initialization.* A second purpose of the contention period is to initialize sub-channels used for data transmission.  $\mu$ -MAC supports one general traffic sub-channel and up to  $N$  sensor report sub-channels.  $N$  is a protocol parameter defined at design time and defines the maximum number of simultaneous distinct interests the sensor network can service.

Sub-channels are collections of scheduled slots used in collision-free transmission. The bandwidth of a sub-channel  $S$  is defined by the minimum separation  $P_s$  between any two of its slots during a contention-free period. Sub-channel initialization is comprised of two steps: *slot assignment*, i.e., defining the first  $n$  slots of the sub-channel observing their minimum separation, where  $n$  is a parameter of the protocol. This parameter is associated with the sub-channel fault-tolerance as explained later; and *sub-channel tuning* or communicating the selected slots to neighbour nodes so that they can turn on the radio for receiving packets at the appropriate times.

*Slot Assignment.* A sub-channel initialization follows the general principle of slot reservation adopted by  $\mu$ -MAC and depicted in Fig. 4. The Neighbourhood-aware Contention Resolution (NCR) algorithm defines transmission priorities among nodes for each scheduled slot. When a node learns that it has the highest priority amongst possible interfering nodes, it is said the node owns the slot. In order to identify the first slot for a sub-channel  $S$ , a node looks for the first unassigned slot it owns in the contention-free period. The second slot can




Fig. 4. Slot reservation mechanism for a node  $u$  and sub-channel  $S$

be found by the same procedure, however, the search must begin  $t + P_S$  time units ahead of the first slot, assuming the slot starts at time  $t$ . Similarly, all  $n$  slots are identified.

Sensor report sub-channels differ from the general traffic sub-channel regarding the way their bandwidth are set. The general traffic sub-channel  $GT$  has a bandwidth defined at design time and although this value may be changed,  $P_{GT}$  is expected to be fixed during the network operation. On the other hand, a sensor report sub-channel  $S$  has the value  $P_S$  defined by the traffic generated according to the interest it services. Therefore, if a sub-channel is not servicing any interest,  $P_S$  is zero.

*Sub-Channel Tuning.* The slots assigned to a specific sub-channel by a node are all used for transmission of its own packets. Since  $\mu$ -MAC seeks to reduce energy consumption by turning off the radio of nodes as much as possible, it is necessary to inform neighbours of the slots used for transmission so they will listen to them. The first  $n$  slots of each sub-channel of a node are transmitted to neighbours in the final part of the control packet shown in Fig. 3a. The first sequence of  $n$  fields are reserved to the general traffic sub-channel, followed by  $N$  sequence of  $n$  fields, each one for a different sensor report sub-channel.

When a node receives a control packet from a neighbour, it schedules its radio to wake up and listen to transmissions at the slots identified in the packet.

1) *Contention-Free Period* : The purpose of the contention free-free period is to transmit application level data between network links. The transmission scheduling mechanism of  $\mu$ -MAC guarantees that packets sent in time slots of this period are collision-free. Before sending a packet, however, a node has to activate a communication sub-channel initialized during the contention period. Every data packet is communicated through slots belonging to some sub-channel. Transmission scheduling and sub-channel

activation in  $\mu$ -MAC are outlined in the following paragraphs.

*Transmission Scheduling.* The process of assigning transmission priorities among nodes in each time slot of the contention-free period is referred as transmission scheduling.  $\mu$ -MAC assigns transmission priorities according to the Neighbourhood-aware Contention Resolution (NCR) algorithm [5]. NCR guarantees that transmissions in any given time slot are collision free provided that local topology information is accurate. In any given wireless network, interference with transmissions of a node  $u$  can only be caused by one-hop or two-hop neighbours. Nodes farther away cannot reach receivers of  $u$ 's transmission. NCR is able to distribute transmission priorities among nodes in the network such that each node in the one or two-hop neighbourhood of  $u$  refrains from transmitting data every time  $u$  considers itself the holder of the highest priority. Therefore, transmissions are always collision-free.

Transmission priorities among nodes for a given slot are assigned through a hash function computed locally. The priority of a node  $u$  over a slot  $t$  is given by:

$$p(u, t) = \text{Rand}(u \otimes t) \otimes u \quad (1)$$

where function  $\text{Rand}(x)$  is a pseudo-random number generator that produces a uniformly distributed random number using seed  $x$ . The expression  $a \otimes b$  denotes concatenation between strings  $a$  and  $b$ . Node  $u$  is able to compute transmission priorities for all possible interfering nodes because the neighbour protocol provides the required topology information. If a node  $u$  has the highest priority among the possible interfering nodes in a given slot, it can transmit without collisions.

When the contention period fails to inform the nodes on the correct local topology due to packet losses, collisions are possible in the subsequent contention-free period. Therefore, it is important for the contention-free period to be long enough to ensure high delivery rates.

*Sub-channels Activation.* When a sub-channel is initialized it only provides up to  $n$  slots to be used by a transmitting node. These slots are concentrated close to the beginning of the contention-free period and have no packet scheduled for transmission. Sub-channel activation is the process of extending the sub-channel along the whole contention-free period. If a packet needs to be transmitted, it can be scheduled to one of many slots of the sub-channel.

The activation process follows the same procedure of slot assignment and tuning employed during the sub-channel initialization. The main difference between initialization and activation resides in the fact that every sub-channel needs to be initialized while activation takes place on demand. Another difference is that tuning for activation is carried out by data packets (Fig. 3b) instead of control packets.

Activation is triggered in distinct ways for the general traffic and sensor report sub-channels. The former, is activated immediately after initialization while the latter depends on the issue of interests by the base station. Therefore, sensor report sub-channels remain initialized but inactive until more traffic is generated in the network. This characteristic is important because once a sub-channel is activated it uses ALL slots reserved for it independent off the traffic generated by the upper layers. Note however that since slot reservation matches the traffic to be generated, no waste of resources takes place once the sub-channel is activated and reports start to flow towards the base station. The following paragraphs introduce the elements involved in the activation of each type of sub-channel.

*General Traffic.* Immediately after the general traffic is initialized it is activated by scheduling beacon or data packets in their slots. Beacon packets are used in the absence of any transmission data from the upper layers. Both beacon and data packets share the same header (Fig. 3b).  $\mu$ -MAC performs slot assignment dynamically: every time a beacon or data packet is scheduled for transmission, the next  $n$  slots of the sub-channel are computed and placed in the header for tuning of the receivers. In fact, only one new slot needs to be computed per scheduled packet since  $n-1$  were computed previously. The slot assignment follows the procedure depicted in Fig. 4.

A node must listen to the general traffic sub-channel of every neighbour. The sub-channel is kept activated by the constant stream of beacon or data packets flowing between neighbour nodes. If this stream is interrupted, the sub-channel is lost until it is re-initialized in the next contention period. In order to reduce the probability of such events, each data or beacon packet carries the information of the next  $n$  slots of the channel. Therefore,  $\mu$ -MAC tolerates up to  $n-1$  packet losses before disconnecting temporarily a sub-channel.

*Sensor Report.* The activation process of sensor report sub-channels takes place on demand and depends




Fig. 5. Periods of slots from sensor report sub-channels associated with a common interest

on interests issued by the base-station. Interests are transmitted hop by hop using the general traffic sub-channel. In each node receiving the interest, an upper layer issues a register command to the  $\mu$ -MAC protocol. This command binds the interest with an available sensor report sub-channel and informs  $\mu$ -MAC on the MAC address of the next hop towards the base station.

When a sensor able to generate the data requested receives the interest, it calls  $\mu$ -MAC s reserve command. Reserve has two arguments: an interest identification and the frequency of reports to be generated.  $\mu$ -MAC updates the bandwidth of the sub-channel associated with the interest and sends a reserve packet to the next hop towards the base station. The sub-channel is activated by scheduling beacon or data packets transmissions in the sub-channel for the next contention-free period after the reserve packet was sent. Each node receiving a reserve packet also updates the bandwidth of the sub-channel associated with the interest and sends a reserve packet to the next hop towards the base station. Activation always takes place in the next contention-free period after the reserve packet is sent. Eventually, the base-station is reached and a path of active sub-channels is formed between the sensor generating data and the base-station. Sensor reports are kept active by sending beacons if no data is scheduled for a given slot. The procedure of slot assignment and tuning is the same as the one used for the general traffic sub-channel.

A node may receive a reserve packet associated with an interest for which it has already an active sub-channel. In this case, the node is still required to generate a reserve packet towards the base station in order to update the bandwidth of sub-channels along the way (Fig. 5).

Because the sub-channel activation process from source to base-station is not instantaneous, reports are likely to be dropped along the way during a period of time. This period lasts until the full path is established and the bandwidth necessary for the reports is reserved. The sensor application therefore may choose not to send reports instantaneously after an interest is received but to add an initial delay (referred as delay before report).

*Impact of Nodes Leaving/Joining the Network.*  $\mu$ -MAC has presently no mechanism to dynamically rearrange reserved channels based on topological changes. Only upper layers can trigger such operations. The base-station must therefore monitor the incoming traffic and re-broadcast the interest if losses are observed. Addition of new nodes also requires the same action.

## V. ANALYSIS

The performance of the  $\mu$ -MAC protocol was assessed through simulation experiments. For comparison purposes, S-MAC experiments were also conducted. S-MAC is a well known protocol for wireless sensor networks briefly reviewed in Section II. The following paragraphs describe the simulation environment, the experimental setup and the results.

### A. Simulation Environment and Experimental Setup Simulation

Simulation studies were performed using the ns-2 [6] simulator, given its broad acceptance in the research community and the benefit of reusing a wide code base available. The experiments are characterized by parameters in four different classes: *sensor field structure, traffic load, MAC protocols* and *simulation*.

- *Sensor Field Structure.* The sensor field is comprised of 6 sensor nodes arranged in a straight line according to Fig. 6. The base-station is the left-most node. The radio range is adjusted so that only immediate adjacent nodes can communicate. Each node has therefore a left and/or a right neighbour. This topology is made deliberately simple to rule out from the study the impact of choosing different routing protocols.
- *Traffic Load.* In each experiment, the base-station issues an interest requesting nodes  $s_1, s_2, s_3, s_4$  and  $s_5$  to send periodic reports on the data collected. Each sensor sends 50 bytes to the base station per report generated. Therefore, while node  $s_1$  only handles the communication flow of reports generated locally, node  $s_5$  handles the communication of five report flows (the one generated locally, and four other coming from remote nodes). Radios have a transmission rate of 20Kbps.
- *MAC Protocol Parameters.* Packets received from upper layer protocols are initially placed in a buffer by the MAC protocols simulated. This buffer has a fixed maximum capacity and trespassing this

limit cause packets to be dropped. Unless otherwise stated, this buffer is able to hold 20 packets. Parameters specific to  $\mu$ -MAC include the duration of contention-free period, duration of contention period, size of slots, duration of switching interval, number of retransmissions during contention periods, time between general traffic slots, maximum number of sub-channels available per node, delay before reporting and the minimum number of MAC packets dropped in a sub-channel before it is deactivated.

In every experiment involving  $\mu$ -MAC, the slot size during the contention period is 20ms. The duration of the contention period and the number of control packet retransmissions are chosen to guarantee a delivery ratio of 99%. According to [7], this ratio can be attained in a network with an average of  $N$  one/two-hop neighbours by retransmitting each packet 7 times over  $7 \cdot 1.44 \cdot N$  slots chosen randomly. For the sensor field used in our experiments,  $N$  is overdimensioned as 25 and the duration of the contention period is set to  $7 \cdot 1.44 \cdot 25$  or 252 contention slots (approximately 5s).

The slot size during the contention-free period is 40ms. At a transmission rate of 20Kbps, packets of up to 100 bytes can be fit into contention-free slots. The duration of slots includes time for switching intervals. In all experiments, each node has only one sensor report sub-channel available besides the general traffic sub-channel. This means nodes can service only one interest at a time. The minimum number of MAC packets dropped in a sub-channel before it is deactivated is set to one as channels are assumed to be reliable. The pseudo-random number generator used to implement NCR is MD5 [8].

Besides  $\mu$ -MAC, experiments were also conducted with S-MAC for comparison purposes. Parameters specific to S-MAC include the listen/sleep duty cycle, duration of the listen period, number of retransmissions before dropping packets. In all experiments, the duration of the listen period is 143 ms, with 55 ms dedicated to the exchange of schedules between neighbour nodes. S-MAC retransmits a packet up to five times if no ACK is received from a neighbour before dropping it. No fragmentation of messages is assumed (one message = packet).

- *Simulation.* Each experiment simulates the operation of a sensor field until each reporting sensor transmits 500 report packets. A warm-up period for local topology discovery is defined in experiments




Fig. 6. Sensor field Structure used in the evaluation

involving both MAC protocols. Similarly, a flush period is defined to empty the network before the end of each simulation. The results were averaged over multiple runs and 95% confidence intervals for the averages computed.

### B. Simulation Results

Simulation experiments were conducted to assess the performance of  $\mu$ -MAC according to four metrics:

- *Percentage of Sleep Time*: the ratio between the total time radios are in sleep mode and the total time measured. This metric provides a measure on the energy efficiency of a MAC protocol. The more a radio is kept in sleep mode, the more energy efficient the protocol. Since the energy consumption during sleep mode can be orders of magnitude lower than in active states, even small differences in the amount of sleep can impact significantly the lifetime of nodes.
- *Delivery Rate*: the percentage of application packets that were delivered at the neighbour destination. Low duty cycles can be obtained at the expense of communication reliability. This metric indicates whether the MAC protocol is reliably transferring packets from node to node.
- *Queue Delay*: Low duty-cycles impact the time to transfer data inside the network. Queue delay is the time a packet awaits transmission computed from the moment it is received by the MAC layer.  $\mu$ -MAC has two different classes of sub-channels each one with a distinct delay behaviour.
- *Correlation between upper layer traffic and radio duty cycle*: this metric indicates the dependency between upper layer traffic and radio duty cycle. It is defined as the slope of the straight line fitted in the set of points generated by measuring the amount of awake time per traffic generated or received in a node. The awake time measured is normalized for the awake time of the lowest transmission/reception rate. A low slope indicates that the awake time is largely independent from communication. The larger the slope, the more the duty cycle depends on the upper layer traffic. A MAC protocol with low

■ No delay ■ 1 contention-free period ■ 2 contention-free periods

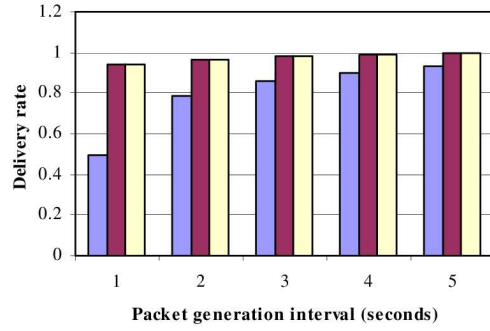


Fig. 7. Average packet delivery rate for  $\mu$ -MAC for different values of delay before report

overhead should present a high slope.

The remainder of this section presents and discusses the experimental results.

*Experiment 1: Impact of Delay Before Report on Delivery Rate.* The sub-channel activation process from source to base-station is not instantaneous. The reserve packets flow through the general traffic channel and command each node in the path towards the base station to activate/update a report sub-channel. In  $\mu$ -MAC, activation or bandwidth modification of sub-channels only happens in a contention period. Therefore, if reports are generated instantaneously after the reception of an interest, packets are likely to be dropped for lack of bandwidth.

Fig. 7 shows the effect of different values of delay before report on the delivery rate of  $\mu$ -MAC. When no initial delay is allowed the delivery rate is reduced. This reduction is proportional to the traffic generation rate as expected: the smaller the packet generation interval, the lower the delivery rate. When the delay is equal to the duration of one contention-free interval,  $\mu$ -MAC has enough time to configure the sub-channel for reporting and the delivery rate is greatly improved. Increasing the delay brings no additional benefits.

Preventing nodes from applying a proper delay before reporting has another potential undesirable consequence besides the initial loss of packets.  $\mu$ -MAC is designed to adapt precisely to the flow of reports. For each packet to be transmitted,  $\mu$ -MAC has a slot reserved before another packet is generated. This means that an initial burst of packets without slots will fill the buffer and keep it nonempty for the duration of the interest. This fact increases the queue delay of the packets



transmitted in the sub-channel. The proper delay before report avoid this state of affairs. Choosing a bandwidth slightly higher than needed also prevents long queues at the expense of energy consumption.

*Experiment II: Impact of Relative Duration Between Contention and Contention-Free periods on Sleep Percentage.* A precise local topology is crucial for the correct behaviour of the NCR algorithm used by  $\mu$ -MAC for transmission scheduling. Consequently, packets during the contention period need to be retransmitted several times to guarantee delivery with high probabilities. Multiple retransmissions require time and contention periods may for this reason be quite large. Since nodes keep their radios in active state during the entire contention period, the effect of its extension on energy consumption can be significant.

In order to compensate for the high energy consumption during contention periods, the interval between them should be as large as possible. Fig. 8 shows the effect of the relative duration between contention and contention-free periods on the percentage of sleep time in the simulated system. In the figure,  $\mu$ -MAC  $x\%$  means that the contention period duration is  $x\%$  the duration of the contention-free period. The figure also portrays the percentage of sleep obtained for S-MAC with duty cycles of 10%. The graph only represents points for which 100% delivery rate was achieved.

As expected, the lower the ratio between contention and contention-free intervals, the more nodes in the system sleep. This effect is a consequence of the reduced impact of the contention-period on the overall sleep. At high ratios, however, the sleeping time is dominated by the amount of application level traffic being transported and overhead from the contention-free period. Reducing the ratio beyond such point does not contribute to the increase of sleeping. The curved shapes of the  $\mu$ -MAC graphs also demonstrate that the percentage of sleep is affected by the traffic load in the system. This behaviour differs from S-MAC which is largely indifferent to the packet generation interval.

*Experiment III: Correlation between upper layer traffic and radio duty cycle.* In order to quantify the impact of application traffic on the radio duty cycle, a separate experiment was conducted. In this experiment only two nodes are simulated, one transmitting application packets and the other receiving them. For each node, the amount

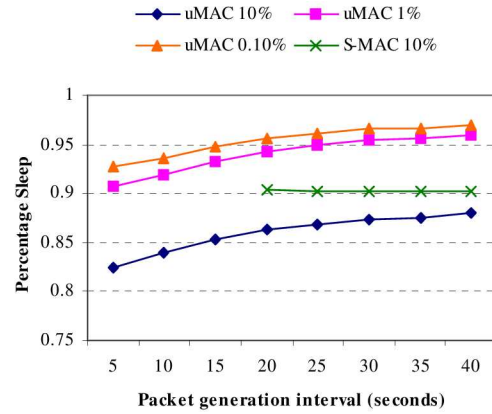


Fig. 8. Average percentage sleep versus packet generation interval

of awake time is plotted against the packet transmission/reception rate. The amount of awake time is normalized against the percentage of awake time for a packet transmission/reception rate of 0.05. A straight line is obtained through linear regression on the points generated. The results for  $\mu$ -MAC 10%,  $\mu$ -MAC 1% and S-MAC 10% are depicted in Fig. 9 for the receiving node. A similar behaviour was observed for the transmitting node.

The slope of the curves in the graph gives an indication of the correlation between the application traffic generated and the radio duty cycle in the node. The figure indicates that the duty cycles of S-MAC 10% and  $\mu$ -MAC 10% are largely insensitive to application traffic when compared to  $\mu$ -MAC 1%. For  $\mu$ -MAC 10%, the awake time for the receiving node at a packet receiving rate of 0.35 is 1.2 times the percentage awake time at a rate of 0.05. For  $\mu$ -MAC 1% this value increases to 2 times. A similar result is obtained for the transmitting node.  $\mu$ -MAC 1% is very efficient in turning on the radio only when more communication traffic is present.

*Experiment IV: Impact of Traffic on Queue Delay.*  $\mu$ -MAC has two different classes of sub-channels each one with a distinct delay behaviour. The general sub-channel has a fixed and predefined bandwidth. This bandwidth is expected to be defined at design time and although it is possible to change it later, it is traffic independent. Report sub-channels on the other hand have their bandwidth adjusted with the traffic rate they are expected to carry. In order to be highly energy efficient, the sub-channel should contain exactly one slot for each packet to be transmitted.  $\mu$ -MAC tries to achieve this behaviour by exploring information from

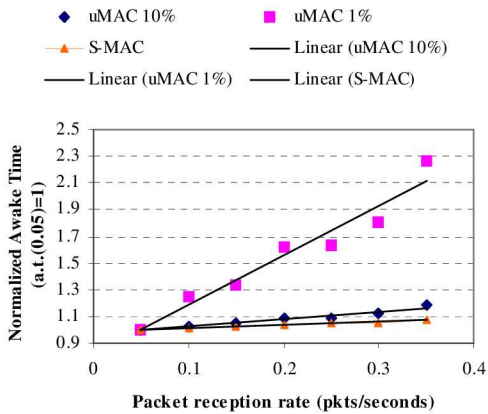


Fig. 9. Normalized awake time for the receiving node

predictable traffic when reserving slots as described in Section IV-C.1. The downside of this approach is that delay becomes also dependent on the traffic rate. In fact,  $\mu$ -MAC makes no attempt to place reserved slots close to the instant they are needed. It only guarantees that slots are present frequently enough to handle the incoming traffic. Therefore, sub-channels serving low rate traffic are expected to impose greater queue delay than traffic intensive sub-channels with more reserved slots.

Fig. 10 and Fig. 11 show queue delays associated respectively with the report and general traffic sub-channels of  $\mu$ -MAC. The queue delay for the report sub-channel is depicted for each node in the sensor field. Since each node handles the traffic of a different number of report flows, each outgoing sub-channel offers very distinct queue delay figures. The more report flows a node handles, the lower the queue delay. Similarly, the larger the packet generation interval, the higher the delay. This behaviour contrasts with the delays for interest packets (transmitted in the general traffic sub-channels) and packets transmitted using S-MAC. The queue delay for these packets are very independent on packet generation interval as shown in Fig. 11 and Fig. 12.

*Experiment V: Impact of Buffer Size on Delivery Rate.*  $\mu$ -MAC uses application level information on predictable traffic to guarantee high delivery ratios with minimum buffering. The previous knowledge on the shape of traffic as conveyed in interests makes it possible for slots to be reserved at a similar rate that packets needs to be transmitted. In MAC protocols where the traffic information is not taken into consideration, communication may not be possible before the arrival of more packets

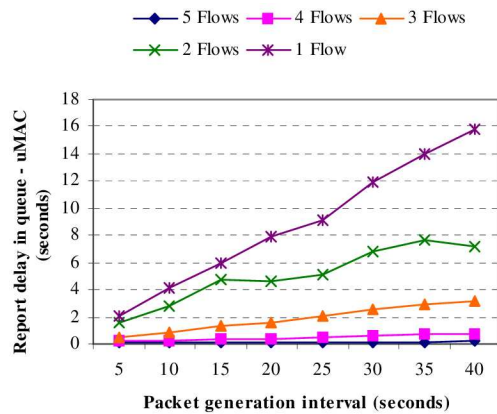


Fig. 10. Average queue delay for packets in report sub-channels

and buffering requirements increase. Fig. 13 shows that  $\mu$ -MAC can achieve 100% delivery rates for even very small buffer sizes.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper a schedule-based media access protocol for wireless sensor networks called  $\mu$ -MAC was presented and its performance was analyzed through simulation studies.  $\mu$ -MAC addresses the problem of achieving very low radio duty-cycles in sensor nodes while providing good delivery rates, low buffering requirements and delay characteristics that can be tolerated in a range of real world applications.

Low radio duty-cycles is a fundamental element in achieving energy-efficient MAC protocols for wireless sensor networks since communication operations are power intensive. Ideally, the radio should be activated if and only if packets are to be transmitted and/or received. The lack of knowledge on the traffic pattern of neighbour nodes, however, prevent an ideal 100% efficient operation of the radio to take place.

Wireless sensor networks constitute a particular class of networks in which traffic presents much more predictable characteristics in several applications. Very often, the traffic is comprised of periodic packets flowing from different nodes towards a fixed set of nodes called base-stations.

$\mu$ -MAC leverages the predictability of traffic behaviour in wireless sensor networks to increase the efficiency of radio utilization thus reducing duty-cycles. Specifically, the protocol uses traffic flow specifications present in the interests broadcasted by the base-station to adjust a node's communication bandwidth.

In the future, an implementation of  $\mu$ -MAC for TinyOS and an assessment of its performance in a prac-

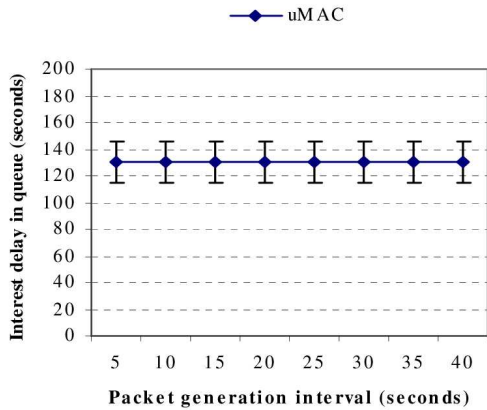


Fig. 11. Average queue delay for packets in general traffic

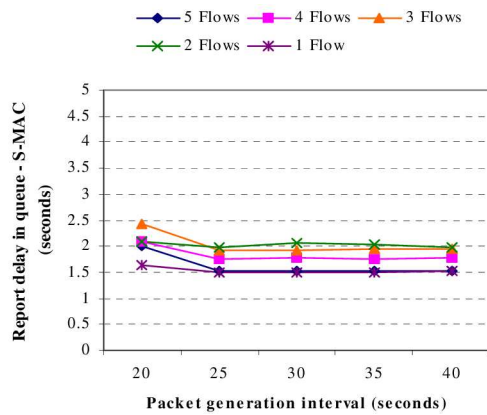


Fig. 12. Average queue delay for packets in S-MAC

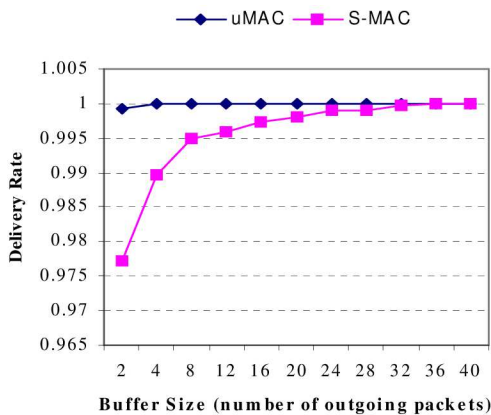


Fig. 13. Delivery rate versus buffer size

tical deployment is to be made available. Similar ideas in exploring traffic-level information to improve the energy characteristics of contention-based MAC protocols for wireless sensor networks are also being considered.

#### ACKNOWLEDGMENT

The support of the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged.

#### REFERENCES

- [1] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *IEEE Infocom 2002, Jun 23-27 2002*, vol. 3. Information Science Institute (ISI), Univ. of Southern California (USC), Los Angeles, CA, United States: Institute of Electrical and Electronics Engineers Inc., 2002, pp. 1567–1576.
- [2] T. van Dam and K. Langendoen, "An adaptive energy-efficient MAC protocol for wireless sensor networks," in *The First ACM Conference on Embedded Networked Sensor Systems (SenSys 2003)*, Los Angeles CA, November 2003.
- [3] V. Rajendran, K. Obraczka, and J. Garcia-Luna-Aceves, "Energy-efficient, collision-free medium access control for wireless sensor networks," in *The First ACM Conference on Embedded Networked Sensor Systems (SenSys 2003)*, Los Angeles CA, November 2003.
- [4] "nRF2401 Single Chip 2.4GHz Radio Transceiver Data Sheet," Nordic, <http://www.nvlsi.no/index.cfm>, 2002. [Online]. Available: <http://www.nvlsi.no/index.cfm>
- [5] L. Bao and J. Garcia-Luna-Aceves, "A new approach to channel access scheduling for ad hoc networks," in *7th Annual International Conference on Mobile Computing and Networking, Jul 16-21 2001*, Computer Science Department, University of California, Santa Cruz, CA 95064, United States, 2001, pp. 210–220.
- [6] "The NS manual," <http://www.isi.edu/nsnam/ns/ns-documentation.html>. [Online]. Available: <http://www.isi.edu/nsnam/ns/ns-documentation.html>
- [7] L. Bao and J. Garcia-Luna-Aceves, "Hybrid channel access scheduling in ad hoc networks," in *Proceedings 10th IEEE International Conference on Network Protocols, 12-15 Nov. 2002*. Comput. Sci. Dept., California Univ., Santa Cruz, CA, USA: IEEE Comput. Soc, 2002, pp. 46–57.
- [8] R. L. Rivest, "The MD5 message-digest algorithm," Internet Request for Comments, Apr. 1992, rFC 1321.