Multi-radio medium access control protocol for wireless sensor networks

Junaid Ansari*, Xi Zhang and Petri Mähönen

Institute for Networked Systems,
RWTH Aachen University,
Kackertstrasse 9, D-52072,
Aachen, Germany
Email: jan@mobnets.rwth-aachen.de
Email: xzh@mobnets.rwth-aachen.de
Email: pma@mobnets.rwth-aachen.de
*Corresponding author

Abstract: This paper describes a dual-radio based medium access control protocol for wireless sensor networks. Our protocol combines the advantages of the two radios operating in different frequency bands to result in highly energy-efficient operation. The design effectively addresses the two dominant sources of energy consumption in sensor network communication, namely the idle listening and the ephemeral burst data traffic. This paper presents the design rationale and extensive empirical performance evaluation of the protocol in terms of power consumption and latency under various traffic loads and duty cycles. Experimental performance comparison with B-MAC show high gains of our approach. We derive analytical expression for the optimum transmit power level ratio of the two radios giving minimum energy consumption. We also model the mathematical relationship for the optimal duty cycle of the nodes to a given network traffic load and validate it through the prototype implementation on commercially available sensor nodes.

Keywords: wireless sensor MAC; low-power; multi-radio platform; performance evaluation.

Reference to this paper should be made as follows: Ansari, J., Zhang, X. and Mähönen, P. (2010) 'Multi-radio medium access control protocol for wireless sensor networks', *Int. J. Sensor Networks*, Vol. 8, No. 1, pp.47–61.

Biographical notes: Junaid Ansari is currently a PhD student and Research Assistant at the Institute for Networked Systems at RWTH Aachen University, Germany. He holds a BS degree in Electrical Engineering from National University of Sciences and Technology, Pakistan (2002) and MSc degree in Communications Engineering from RWTH Aachen University (2006). His current research interests include low-power design and energy-efficient networking solutions for embedded wireless networks. He has been actively working on various large-scale research projects related to wireless embedded networking funded by European Union and German government at the Institute for Networked Systems at RWTH Aachen University, Germany.

Xi Zhang is currently a PhD student and Research Assistant at the Institute for Networked Systems, RWTH Aachen University, Germany. She completed her MEng in Electrical and Electronic Engineering from Imperial College London, the UK, in 2007. Her current research interests include cross-layer optimisation of protocol stacks, mainly physical, medium access control and routing layer, in searching for low-power, real-time and intelligent solutions for wireless sensor networks. As a part of her research work she has participated actively in European research projects as well as research projects carried out in collaboration with industry.

Petri Mähönen is a Full Professor and Head of the Institute for Networked Systems at RWTH Aachen University. Before moving to Aachen in 2002, he has worked as a Professor and Research Director of networking at the Center for Wireless Communications, Oulu. He has been principal investigator in several international research projects, and consults several multinational companies on wireless and 4G research strategies. His current research with his students focuses on wireless internet, low-power communications, applied mathematical methods for telecommunications, and cognitive networks. He is also a research area coordinator and a principal investigator for UMIC research centre at RWTH.

1 Introduction

Wireless Sensor Networks (WSNs) have a broad range of applications. Since sensor nodes are battery powered and have therefore limited energy resource, keeping the network operational over long periods of time is challenging. Replacement of the batteries on sensor nodes is either highly cumbersome or impractical. In most of the WSN applications, the major expenditure on energy consumption comes from data communication. Therefore, energy efficient communication, especially the role of Medium Access Control (MAC) protocols becomes very important. Generally, traffic load in WSNs is very small and therefore the radio does not necessarily need to be active all the time. In an ideal case for conserving energy, sensor nodes switch on the radios only for data transmission/reception and turn them off when there is no traffic in the network. While it is easy to define when data transmission is required, reception is usually unpredictable to a sensor node. Sensor nodes therefore, need to listen to the channel periodically in anticipation of potential packets. If no packet is detected, listening to the channel goes wasted and is referred to as idle listening. Periodically turning on/off the radio for battery power conservation is called as the duty cycle operation. In order to exchange data, the transmitter and the receiver(s) must have their radios on at the same time. Different duty cycling MAC solutions have been proposed in the WSN research community for coordinating the active periods of the nodes and for establishing data communication (Demirkol et al., 2006). Each MAC protocol inflicts a different amount of control overhead for establishing data exchange, which has a direct impact on the energy consumption.

The radio activities in WSNs include idle listening to the medium and the actual data transmission/reception. The amount of data exchange varies from application to application, however idle listening remains dominant. In order to minimise the overall energy consumption of a WSN, both of the activities are required to be optimised. Since these two operations are very distinct, they impose different requirements on the radio capabilities. Idle listening or channel polling operation is carried out frequently in a periodic manner. Data exchange, on the other hand, is rare and sporadic. Many single radio platform based MAC solutions remain handicapped due to the lack of specialised ability to simultaneously cater these different needs. In this work, we describe a Multi-Radio MAC (MR-MAC) protocol, which is designed from this perspective. Our MAC uses a low power sniffer radio to effectively handle idle listening and a fast/bursty radio for actual data communication. The operating frequencies of the two radios are selected based on the characteristic features needed by the two operations. The low-power slow data rate sniffer radio, which is used for the exchange of control information to coordinate the actual data communication, operates in low frequency band. Owing to having a larger bandwidth in high frequency bands, data is transmitted over a fast rate radio operating in high frequency band. Our protocol design combines the advantages of the two radios very effectively and achieves a highly energy efficient performance as compared to the contemporary single radio based solutions as we show later in the paper. We opt for preamble sampling MAC approach in our protocol because it does not impose the need for an explicit synchronisation among the nodes in the network and is light weight (Demirkol et al., 2006). The preamble sampling technique also effectively handles the network dynamics such as mobility, old nodes dying and new nodes appearing in the network. We have presented the preliminary design approach of the protocol in the Workshop on Energy for Wireless Sensor Networks (Ansari et al., 2008). In the following, we present entirely new results and in depth evaluation of the protocol.

The rest of the paper is organised as follows: we give a state-of-the-art overview of the WSN MAC protocols in Section 2. Section 3 describes the protocol design in detail. Section 4 describes the prototypic details. In Section 5, we derive the analytical expression for the optimum duty cycle value and in Section 6, we express the optimal transmit power level ratios of the two radios. Section 7 presents the comprehensive performance evaluation of the protocol and its empirical performance comparison with B-MAC protocol in terms of power consumption and latency. In our experimental comparison, we consider the B-MAC implemented on both the radios used by our prototype platform. Finally, we conclude the paper in Section 8.

2 Related work

WSN MAC protocols exercise a duty cycling operation in order to minimise idle listening for conserving energy. In order for the nodes to communicate with each other, they must be active (radio turned on) at the same time. Different techniques are used in order to align the active periods of the nodes and hence to establish communication. One approach is to explicitly synchronise the nodes so that they follow a common sleep schedule. Many protocols such as S-MAC (Ye et al., 2002), T-MAC (van Dam and Langendoen, 2003), nanoMAC (Ansari et al., 2007a), etc., belong to this category. These protocols are IEEE 802.11 inspired contention based protocols and exercise RTS/CTS/DATA/ACK handshake, which is also used for synchronisation purposes. Establishing a common synchronised schedule by using explicit SYNC packets or exchanging RTS/CTS frames leads to a high control packet overhead. Control overhead is one of the major sources of energy wastage, especially when the traffic load is itself very low. In addition to the contention based protocols, conflict free TDMA based protocols such as BMA (Li and Lazarou, 2004), LMAC (van Hoesel and Havinga, 2004), etc., are also used in the WSN community. Time-slotting inherently has a duty cycling behaviour and conserves energy as the nodes are active only in the assigned slots. Slot assignment and maintenance has a high control overhead, which makes contention free protocols a less attractive choice. Furthermore, contention free protocols suffer from scalability and mobility problems since slots need to be updated when the network topology changes.

One very popular class of contention based MAC protocols is the preamble sampling (or channel polling) based MAC protocols. Preamble sampling protocols such as B-MAC (Polastre et al., 2004) transmit a long continuous preamble to ensure that all the potential receivers, sniffing the channel periodically, are able to detect the presence of the carrier. After detecting a preamble, receiving nodes keep on listening to the preamble and eventually receive the data that immediately follows the preamble. This way, asynchronously waking-up nodes are implicitly synchronised through the preamble sequence. In Micro-Frame Preamble MAC (Bachir et al., 2006), the continuous preamble is replaced by a series of micro-frames. Each micro-frame contains a full set of information about the upcoming data frame. Only the addressed node receives the data frame and the rest of the nodes go into the sleep mode. This scheme saves energy by avoiding the reception of the rest of the preamble sequence and irrelevant data frames. WiseMAC (El-Hoiydi and Decotignie, 2004) shortens the length of the preamble to be transmitted for unicast transmission by exploiting the sleep schedules of the nodes in the network. The sleep schedule of the nodes is transmitted in the acknowledgement packets. WiseMAC also incorporates the potential clock skews and drifts established in the neighbourhood sleep schedules of the nodes. Unlike WiseMAC, X-MAC (Buettner et al., 2006) uses preamble strobing technique for unicast transmission. It divides the monolithic preamble into small frames each containing the destination's address information. After transmitting a preamble frame, the transmitter expects an acknowledgment of the preamble frame. If the acknowledgement is not received within a certain timeout interval, subsequent preamble frame is sent. In the worst case, the preamble transmission duration becomes equal to the periodic channel check interval. On the contrary, upon receiving an acknowledgement of the preamble frame, data frame is sent immediately. This preamble shortening technique is useful only when maintaining the reliable sleep schedules of the neighbours is not difficult e.g. in case when there is only little node mobility or when the nodes do not appear or disappear very frequently. SCP-MAC (Ye et al., 2006) protocol uses synchronised channel polling and combines the features of schedule based protocols with preamble sampling. This approach is suitable for networks operating in low duty cycles with static characteristics. A few hybrid protocol designs have also been proposed such as Funnelling-MAC (Ahn et al., 2006) and Z-MAC (Rhee et al., 2005), which behave as CSMA type or TDMA type in the case of low and high traffic volumes, respectively. These protocols are shown to outperform B-MAC in high traffic load scenarios but have a fairly complex signalling and control overhead.

All the above mentioned protocols are designed to use single channel in a particular frequency band. Schurgers et al. (2002) have proposed in STEM the idea of using two radios operating in separate frequency bands to completely separate data transfer from wake-up. In the tone based approach of STEM, a long wake-up tone is sent to make sure that the destined receiver has awoken once. It is similar in some aspects to the preamble sampling approach used in MR-MAC. However, since MR-MAC uses meaningful fields in the preamble in contrast to the meaningless wake-up tone, non-addressed nodes can avoid receiving irrelevant data and preamble sequence. Furthermore, MR-MAC nodes are able to know the time duration of the channel activity and are able to extend their sleep intervals accordingly.

DCMA/AP (Dual Channel Multiple Access with Adaptive Preamble) (Ruzzelli et al., 2006) is a dual channel MAC protocol designed especially for WSNs and uses two separate channels in the same frequency band, which can be used simultaneously. The main idea is to conserve energy consumption by avoiding RTS/CTS control frames. The data channel is used for preamble and data packet transmissions while the control channel may indicate reception in progress to avoid hidden terminals and packet overhearing. MR-MAC uses the two channels on radios specialised for control and data. Furthermore, using preamble sampling drastically lessens the amount of control overhead needed for maintaining node synchronisation in order to align the active periods of the nodes for data communication.

RTWAC (Ansari et al., 2009) uses an external extremely low-power wake-up circuit attached to a sensor node to wake-up its main radio. Upon the wake-up, a node uses its main radio for data communication. Since the address information is included in the wake-up packet sent over an out-of-the-band modulated wake-up signal, non-addressed nodes can avoid turning their main on-board radios. This approach can help in avoiding idle listening unlike the case in duty cycled MAC protocols and is highly energy efficient when no communication is going on. However, RTWAC has a limited coverage range which limits its application areas.

3 Protocol design

MR-MAC is a *p*-persistent preamble sampling MAC protocol which uses dedicated high and low frequency bands for data and control, respectively. MR-MAC aims at combining the advantages of low and high frequency bands in order to achieve power efficiency. Generally, higher frequency bands have larger bandwidths which allow high communication data rates. The paper by Oppermann et al. (2004) shows that radios operating in higher frequency bands consume less energy per bit as compared to radios operating in lower frequency bands. This fact inspired us to use a radio operating in high frequency band for bursty data communication. Low frequency band radios, on the other hand, consume less energy in idle listening for a given receiver sensitivity threshold. Also, a radio operating in low frequency band requires a lower transmit power level to achieve a certain communication range than the transmit power required by a radio operating in high frequency band.

Instead of transmitting a monolithic preamble sequence with no useful information, MR-MAC transmits a number of small preamble frames containing control information. It also employs a number of preamble shortening techniques as described in Section 3.1. Preamble sequence transmission/reception in general consumes a large amount of energy for radios operating in higher frequency bands and supporting higher data rates than radios operating in lower frequency bands with low data rates. MR-MAC uses an extremely low power sniffer radio, operating in a low frequency band for control/preamble transmission and a high frequency radio for bursty data transmission. MR-MAC performs channel polling operations only in the low frequency band, while the high frequency transceiver is turned on only during data transmission/reception. Overall, this scheme leads to highly optimised utilisation of the radio resources. MR-MAC has the ability to transmit multiple data frames with a single preamble reservation. This allows MR-MAC to efficiently handle large amount of traffic loads in an energy efficient manner over the bursty radio. Figure 1 shows a simplified state machine of MR-MAC.

Figure 1 Simplified state-machine of MR-MAC. The preamble sampling operation is carried out over the Low Frequency Radio (LFR) while bursty data is transmitted over the High Frequency Radio (HFR) (see online version for colours)

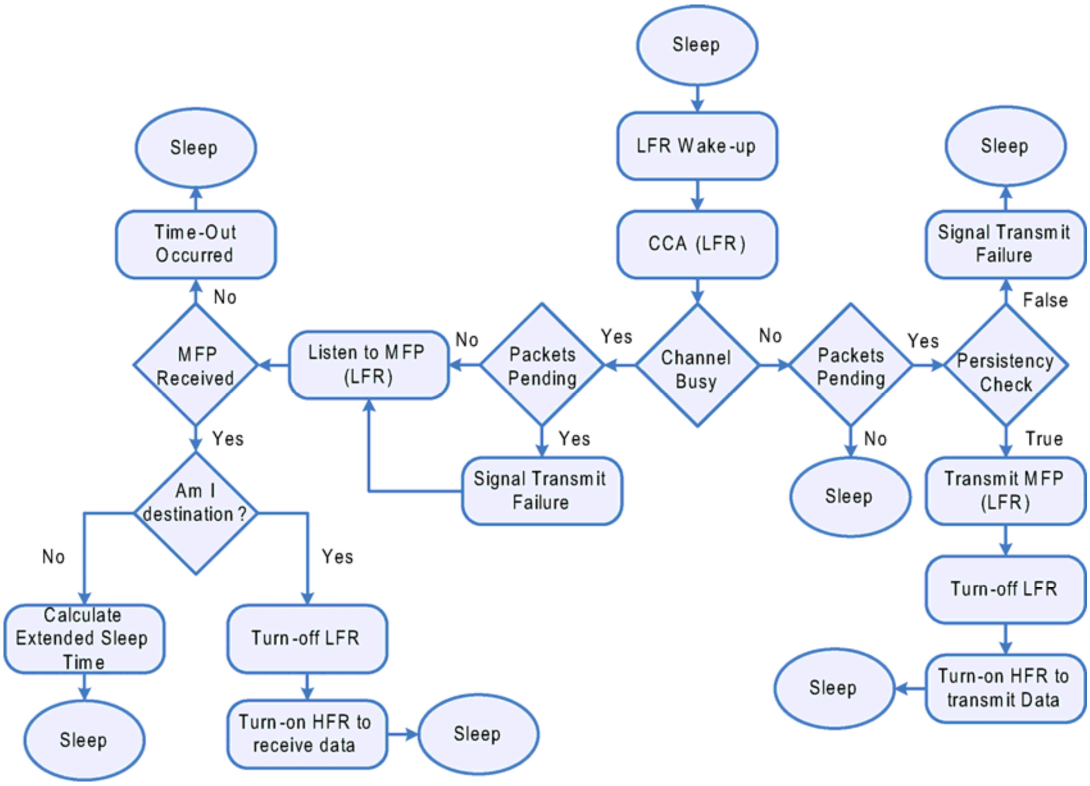

3.1 *Preamble optimisations*

In the following, we describe the preamble length optimisations applied in MR-MAC in order to achieve energy conservation.

3.1.1 *Preamble framelets*

MR-MAC divides the monolithic preamble sequence into small preamble frames which are transmitted back to back to form Micro-Frame Preamble (MFP). These small preamble frames are called as framelets in this paper. Each framelet contains control information such as the radio byte sequence for PLL locking followed by the synchronisation bytes for the receiver. The receiver makes bit offset adjustments based on the synchronisation bytes to correctly receive the rest of the preamble frame. The address information included in the framelet allows the non-addressed nodes to go to sleep without listening to the rest of the preamble sequence and data packets following the preamble. The addressed node goes to sleep since the rest of the framelets does not contain any useful information. A down counter value is transmitted in every framelet, which indicates the beginning of data transmission. An addressed node can therefore precisely estimate when to turn on its high frequency band radio. The size of the data packet is also included in framelets so that an addressed node is able to know how many data packets it needs to receive. A non-addressed node is able to estimate how long the medium is going to be busy based on these two fields and correspondingly extends its sleeping period to achieve additional energy savings.

Our experiments have shown that if the data size is smaller than a certain threshold, it is more energy efficient to piggyback data into the preamble framelets. Piggybacking data into the preamble frames forms Data Frame Preambles (DFPs) (Mahlknecht and Boeck, 2004). When DFPs are used in MR-MAC, the radio operating in high frequency band is not used at all. Figure 2 shows the operational cycle of MR-MAC for the case of MFP and DFP broadcast transmissions.

3.1.2 *Preamble optimisation based on the neighbourhood sleep schedules*

In the case of unicast transmission for preamble sampling based MAC protocols, the length of the preamble can be reduced drastically if the transmitter knows the sleep schedule of the receiver (El-Hoiydi and Decotignie, 2004). In MR-MAC, all the nodes maintain a neighbourhood sleep schedule information similar to WiseMAC. Each preamble framelet contains the information of a node's next wake-up offset. A node receiving the framelet updates the sleep schedule of its neighbour based on the next wake-up offset value. Perfect timing information of the destination node allows a transmitter to delay the transmission of the packet until the potential receiver wakes up. In this way, the transmitter needs to transmit only one preamble framelet. In contrast to the approach of WiseMAC protocol to announce the sleep schedules in the acknowledgment frames, MR-MAC announces the sleep schedules in the preamble frames, which also allows non-addressed nodes to update their neighbourhood timing information while overhearing a preamble framelet.

3.1.3 Preamble strobing

Preamble strobing (Buettner et al., 2006) is a technique used in unicast transmission where a transmitting node sends a preamble framelet and waits for its acknowledgement. If the acknowledgement is not received within a certain timeout duration, subsequent preamble frames are sent. If the acknowledgement is received, preamble frame transmission is stopped and data packet is transmitted immediately. The maximum number of preamble framelets required to be sent corresponds to the periodic channel check interval. This happens in the worst case, when either the receiving node's schedule has the maximum possible offset to the transmitting node's sleep schedule or the receiver is out of the transmission range. Preamble strobing has its advantages especially when the neighbourhood schedule is either unavailable or is unreliable, e.g., in networks with mobility. Furthermore, offsets might be introduced in the estimated sleep schedule information of the receivers due to clock jitter accumulation over extended period of time. In MR-MAC protocol, preamble strobing technique is combined with the neighbourhood sleep schedule based preamble optimisation. This way a node performs preamble strobing with additional information of the receiver's sleep schedule which result in highly improved performance even in the case of clock-drifts and mobility of nodes. Figure 3 shows the operational cycle of MR-MAC for unicast transmissions. Preamble optimisations based on the neighbourhood sleep schedules combined with the preamble strobing technique is only used in unicast transmissions. We have evaluated the performance of MR-MAC protocol with and without unicast optimisations.

Figure 2 (a) Operational cycle of broadcast transmission, where the long preamble is divided into smaller preamble frames. After receiving an MFP, both the receivers A and B are implicitly synchronised to receive the data packets over bursty radio. (b) Operation cycle for broadcast transmission, where the data is piggy backed with inside the preamble framelets. The two receivers A and B receive one DFP packet and sleep during the rest of preamble transmission

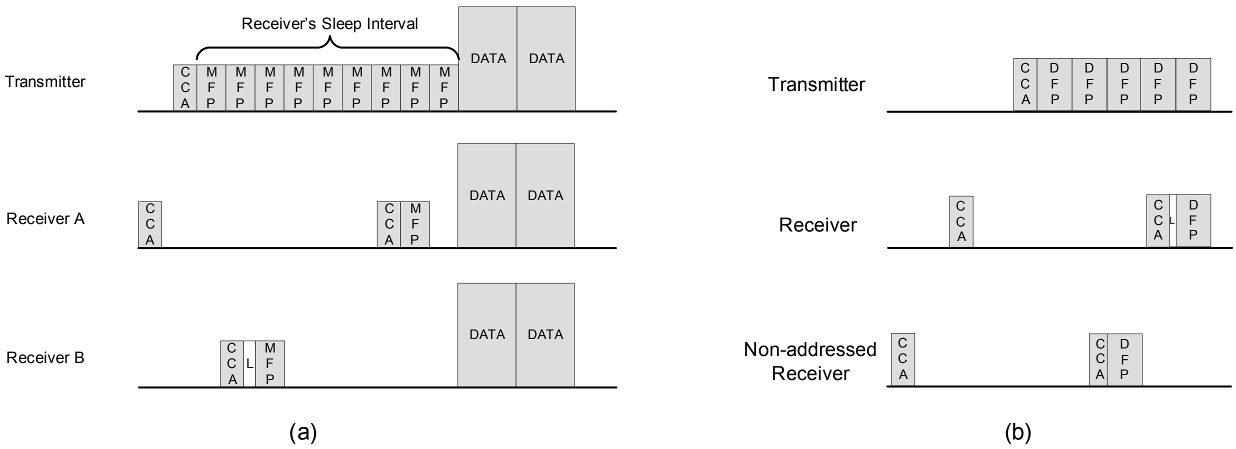

Figure 3 Operational duty cycle of unicast transmission where preamble strobing is combined with neighbourhood-based sleep schedule. The imprecise sleep schedule estimate of the receiver is compensated through the preamble sampling operation. (a) Unicast transmission for the case of MFP and upon receiving the acknowledgement of an MFP, data frames are immediately transmitted over the bursty radio. A non-addressed node on the other hand, quickly goes to sleep upon receiving an MFP. (b) Unicast transmission case of a DFP. A receiving node acknowledges a DFP and transmitter stops sending further DFPs. A non-addressed node on the contrary goes to sleep

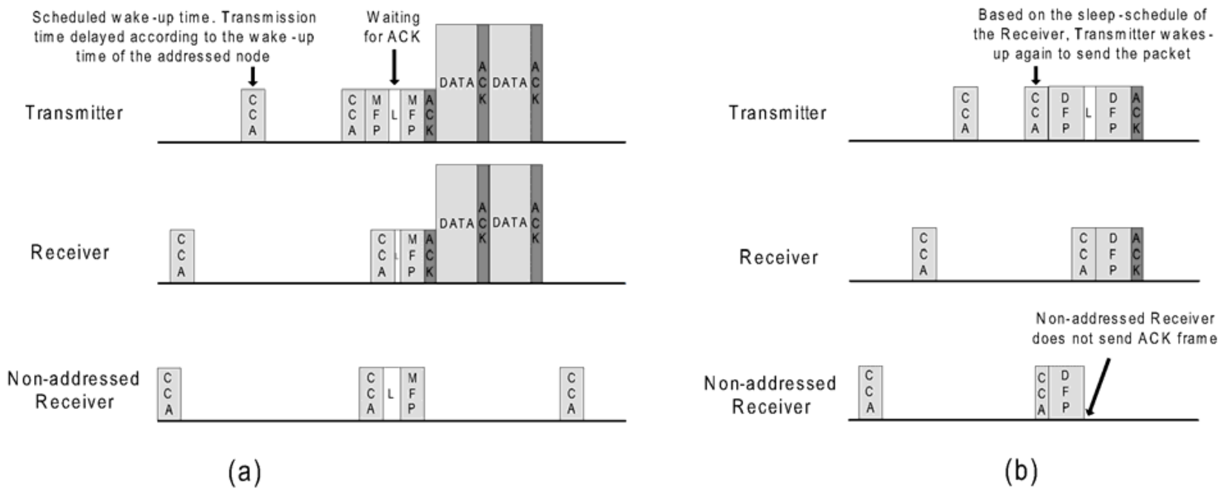

3.2 Support for spectrum agility

Higher frequency bands have generally larger bandwidth and have more number of available communication channels. Due to this characteristic, a large number of devices/networks are attracted to operate in higher frequency bands thereby making the spectrum crowded and prone to interference. Since WSNs are energy constrained, they remain handicapped in competing for the same channel against more powerful networks. Any foreign potential transmission can disrupt the sensor network communication severely and therefore, it becomes desirable that sensor networks wisely select the communication channel in environments prone to interferences. Spectrum agile solutions become important in this context for WSNs. Agreement on the communication channel to be used can either be achieved through a decentralised channel selection algorithm or through a dedicated control channel (Cormio and Chowdhury, 2009). MR-MAC provides support for spectrum agility where, a transmitter first finds a less interfering high frequency channel and announces it over the control channel in the MFP frames for data communication. Since in our case, the low frequency channel is less congested and has little interference as observed by different spectral occupation measurement campaigns (Henry, 2005; Islam et al., 2008; Lopez-Benitez and Casadevall, 2008), our scheme suits well to the current spectrum occupation characteristics.

3.3 Flexible parameters

MR-MAC protocol exposes a number of tunable parameters, which can be set on-the-fly for the (re)-configuration of the protocol. These parameters include local duty cycle, preamble length, CCA thresholds and duration, maximum allowed size of a data packet, initial backoff window size and the persistency values. Some of these parameters are inter-related and their configuration automatically leads to the re-adjustments of the dependent parameters.

4 Prototypic implementation

Our hardware prototype platform (as shown in Figure 4) consists of one Moteiv Inc.'s TelosB sensor node and one Texas Instruments' CC1000 radio chip. TelosB node has an on-board MSP430 series microcontroller and a CC2420 radio transceiver from Texas Instruments. We interfaced an external CC1000 radio module to the microcontroller. Both the configuration and signalling interfaces of the CC1000 radio chip are connected through the extension connectors on TelosB. The IEEE 802.15.4 compliant radio chip, CC2420, operates in 2.4 GHz band. This packetised radio acts as the burst radio while CC1000 radio, operating in 433 MHz band, acts as the low frequency sniffer radio. CC1000 radio provides byte level interface and consumes much less power in idle mode and is therefore suitable for preamble sampling operation and control purposes. CC1000 radio supports multiple data rates ranging from 0.6kbps to 76.8 kbps. CC2420 radio chip offers a data rate at 250 kbps and is therefore good for burst data transmission. We have implemented MR-MAC protocol using nesC programming language in TinyOS 2.x operating system and followed the strict hardware abstraction design philosophy presented in (TinyOS, 2009). We have demonstrated the working prototype in ACM Conference on Embedded Networked Sensor Systems 2007 (Ansari et al., 2007b).

Figure 4 Snapshot of the prototype sensor node platform (see online version for colours)

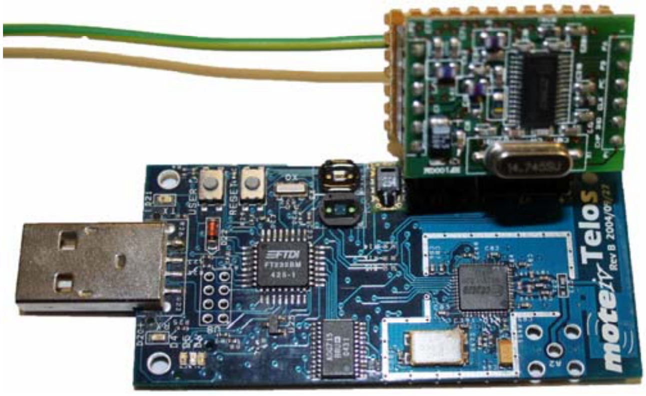

5 Analytical expression for optimum energy consumption

In this section, we present the analytical model of MR-MAC protocol and derive an expression for the optimum duty cycle giving minimum energy consumption for a given traffic load in a network neighbourhood. Since the control overhead and hence energy consumption directly depends on the type of the traffic, we analyse the cases for unicast and broadcast traffic types separately in Section 5.2 and Section 5.3, respectively.

5.1 Models and parameters

Consider a network neighbourhood size of $n+1$ nodes and assume that each node transmits data packets of length l_{pkt} periodically at the rate r_{data} per second. Each node consumes power in the operations: radio setup, carrier sensing, transmitting control information, transmitting data packets, receiving control information, receiving data packets, channel polling and sleep state denoted by P_{radio_setup}, P_{cs}, P_{tx_1}, P_{tx_2}, P_{rx_1}, P_{rx_2}, P_{poll} and P_{sleep}, respectively. Table 1 lists the terms used in our model.

Table 1 Parameters and their notations

Notation	*Meanings*
$t_{\text{poll\_once}}$	Single channel polling interval
$t_{\text{radio\_setup\_once}}$	Single radio setup interval
$T_{\text{poll\_period}}$	Channel sampling period
L_{mfp}	Length of a microframe preamble
N_{mfp}	Number of microframes
$t_{\text{cs\_once}}$	Channel carrier sensing interval
l_{pkt}	Length of the data packet
t_{b1}	Bit duration corresponding to low frequency radio data rate
t_{b2}	Bit duration corresponding to high frequency radio data rate
$l_{\text{ack\_1}}$	Length of the acknowledgement packets over the low frequency radio
$l_{\text{ack\_2}}$	Length of the acknowledgement packets over the high frequency radio

5.2 Broadcast

The overall energy consumption at a node is the sum of the energy spent in each operation and is given by,

$$E = E_{\text{radio\_setup}} + E_{\text{poll}} + E_{\text{rx}} + E_{\text{cs}} + E_{\text{tx}} + E_{\text{sleep}} \tag{1}$$

Expressing energy as the product of power and time using the variables defined gives,

$$\begin{aligned} E = {} & P_{\text{radio\_setup}} t_{\text{radio\_setup}} + P_{\text{poll}} t_{\text{poll}} + P_{\text{rx\_1}} t_{\text{rx\_1}} + \\ & P_{\text{rx\_2}} t_{\text{rx\_2}} + P_{\text{cs}} t_{\text{cs}} + P_{\text{tx\_1}} t_{\text{tx\_1}} + P_{\text{tx\_2}} t_{\text{tx\_2}} + P_{\text{sleep}} t_{\text{sleep}} \end{aligned} \tag{2}$$

Normalising the individual time intervals in different operations by the channel sampling period, $t_{\text{poll\_period}} = L_{\text{mfp}} N_{\text{mfp}} t_{\text{b1}}$.

$$\begin{aligned} t_{\text{radio\_setup}} &= \frac{t_{\text{radio\_setup\_once}}}{T_{\text{poll\_period}}} \\ t_{\text{poll}} &= \frac{t_{\text{poll\_once}}}{T_{\text{poll\_period}}} \\ t_{\text{rx\_1}} &= n r_{\text{data}} \left(1.5 L_{\text{mfp}} t_{\text{b1}}\right) \\ t_{\text{rx\_2}} &= n r_{\text{data}} l_{\text{pkt}} t_{\text{b2}} \\ t_{\text{cs}} &= r_{\text{data}} t_{\text{cs\_once}} \\ t_{\text{tx\_1}} &= r_{\text{data}} N_{\text{mfp}} L_{\text{mfp}} t_{\text{b1}} \\ t_{\text{tx\_2}} &= r_{\text{data}} l_{\text{pkt}} t_{\text{b2}} \\ t_{\text{sleep}} &= 1 - t_{\text{radio\_setup}} - t_{\text{poll}} - t_{\text{rx\_1}} - \\ & \quad t_{\text{rx\_2}} - t_{\text{cs}} - t_{\text{tx\_1}} - t_{\text{tx\_2}} \end{aligned}$$

It may be noted that MR-MAC requires 1 to 2 micro preamble frames in order to decide about the destination and timings of the data packet, in the best and worst case. This gives an average reception of 1.5 micro preamble frames. In this derivation, we consider the radio setup time of only the sniffer radio (radio 1) since it is significant. It accounts for the time needed to send the configuration commands over the SPI or UART interface to start the radio in the desired configuration. It also accounts for the time required by the clock crystal to stabilise. The high frequency radio (radio 2) configuration setup is carried out in parallel to the transmission/reception of the preamble frames. Since it does not cause any extra delays, the setup time for high frequency radio is not modelled.

Our target is to find the relationship of the sampling period (directly related to duty cycle) which leads to the minimum energy consumption. Since the sampling period is directly related to the length of the preamble, we find the number of microframes corresponding to the minimum energy consumption. Plugging the terms in equation (2) and taking the derivative w.r.t. N_{mfp} gives,

$$\begin{aligned} \frac{dE}{dN_{\text{mfp}}} = & -\frac{P_{\text{radio\_setup}} t_{\text{radio\_setup\_once}}}{L_{\text{mfp}} N_{\text{mfp}}^2 t_{\text{b1}}} - \frac{P_{\text{poll}} t_{\text{poll\_once}}}{L_{\text{mfp}} N_{\text{mfp}}^2 t_{\text{b1}}} + \\ & P_{\text{tx\_1}} r_{\text{data}} L_{\text{mfp}} t_{\text{b1}} + \frac{P_{\text{sleep}} t_{\text{radio\_setup\_once}}}{L_{\text{mfp}} N_{\text{mfp}}^2 t_{\text{b1}}} + \\ & \frac{P_{\text{sleep}} t_{\text{poll\_once}}}{L_{\text{mfp}} N_{\text{mfp}}^2 t_{\text{b1}}} - r_{\text{data}} L_{\text{mfp}} t_{\text{b1}} P_{\text{sleep}}. \end{aligned}$$

Putting $\frac{dE}{dN_{\text{mfp}}} = 0$ and simplifying the terms gives the optimum number of microframes $\left(\hat{N}_{\text{mfp}}\right)$ for the minimum energy consumption in equation (3).

$$\hat{N}_{\text{mfp}} = \left(\frac{t_{\text{poll\_once}} \left(P_{\text{poll}} - P_{\text{sleep}}\right)}{L_{\text{mfp}}^2 t_{\text{b1}}^2 r_{\text{data}} \left(P_{\text{tx\_1}} - P_{\text{sleep}}\right)} + \frac{t_{\text{radio\_setup\_once}} \left(P_{\text{radio\_setup}} - P_{\text{sleep}}\right)}{L_{\text{mfp}}^2 t_{\text{b1}}^2 r_{\text{data}} \left(P_{\text{tx\_1}} - P_{\text{sleep}}\right)} \right)^{\frac{1}{2}} \tag{3}$$

Since $\hat{N}_{\text{mfp}} \in$ Natural numbers, we take the ceiling value. It may be noted that the expression for the optimum number of microframes is independent of the number of nodes in the network because we consider a congestion free case where all the nodes are able to transmit their queued packets. The network size (n + 1) governs a lower bound, $n \le \frac{1}{r_{\text{data}} \left(L_{\text{mfp}} t_{\text{b1}} N_{\text{mfp}} + l_{\text{pkt}} t_{\text{b2}}\right)}$.

The optimal sampling time expression is,

$$S.T._{opt} = \hat{N}_{\text{mfp}} L_{\text{mfp}} t_{\text{b1}}. \tag{4}$$

5.3 Unicast

The total energy consumption at a node in the case of unicast transmission takes the same form as equation (2). Unlike the broadcast transmission, a node happens to be the destination for k packets out of the total n packets it hears from its neighbours. MR-MAC tries to optimise the number of microframes to be sent for unicast transmission by using the neighbourhood sleep schedules. In the best case, only one microframe right at the instant when the destination is scheduled to wake-up is required to be sent. In the absence of any timing information, the number of microframes needed to be sent depends upon the offset between the sleep schedules of the transmitting and receiving nodes. In the worst case, the number of microframes needed to be sent becomes the same as in the case of broadcast transmission, given by equation (3). The expressions for time durations are given by:

$$t_{\text{radio\_setup}} = \frac{t_{\text{radio\_setup\_once}}}{T_{\text{poll\_period}}}$$

$$t_{\text{poll}} = \frac{t_{\text{poll\_once}}}{T_{\text{poll\_period}}}$$

$$T_{\text{poll\_period}} = L_{\text{mfp}} N_{\text{mfp}} t_{\text{b1}}$$

$$t_{\text{rx\_1}} = n r_{\text{data}} \left(1.5 L_{\text{mfp}} t_{\text{b1}}\right) + r_{\text{data}} l_{\text{ack\_1}} t_{\text{b1}}$$

$$t_{\text{rx\_2}} = k r_{\text{data}} l_{\text{pkt}} t_{\text{b2}} + r_{\text{data}} l_{\text{ack\_2}} t_{\text{b2}}$$

$$t_{\text{cs}} = r_{\text{data}} t_{\text{cs\_once}}$$

$$t_{\text{tx\_1}} = r_{\text{data}} N_{\text{mfp}} L_{\text{mfp}} t_{\text{b1}} + k r_{\text{data}} l_{\text{ack\_1}} t_{\text{b1}}$$

$$t_{\text{tx\_2}} = r_{\text{data}} l_{\text{pkt}} t_{\text{b2}} + k r_{\text{data}} l_{\text{ack\_2}} t_{\text{b2}}$$

$$t_{\text{sleep}} = 1 - t_{\text{radio\_setup}} - t_{\text{poll}} - t_{\text{rx\_1}} - t_{\text{rx\_2}} - t_{\text{cs}} - t_{\text{tx\_1}} - t_{\text{tx\_2}}$$

5.4 Comparison of analytical results to the implementation results

Power consumption measurements are carried out for the MR-MAC implementation on our prototypic platform in order to verify the optimum duty cycle deduced from our mathematical model. From equation (4), we can see that data transmission rate r_{data} is a variable while other terms are fixed either due to radio properties or protocol design. The lowest power consumption per node should be achieved at the optimum duty cycle. We plotted the power consumption curves at different data transmission rates in a network using the analytical expression. The parameters that we used for the analytical expression are based on our measurements and are listed in Table 2. The plots for power consumption per node at different duty cycles are concave in shape and with a unique minimum point.

Table 2 Parameter values measured on our prototypic platform for analytical expression

Parameter	*Value*
Power in radio setup ($P_{\text{radio\_setup}}$)	13 mW
Power in channel polling (P_{poll})	25 mW
Power in transmit mode ($P_{\text{tx\_1}}$)	31 mW
Power in sleep mode (P_{sleep})	1.78 mW
Time for radio setup ($t_{\text{radio\_setup\_once}}$)	5 ms
Time for one time channel polling ($t_{\text{poll\_once}}$)	3.5 ms

For the hardware measurements, we considered a network size of 3 nodes to present a non-congested network environment. The length of the preamble-frame, l_{frame} is 64 bits. Figure 5 shows the average power consumption of a node at different duty cycles for different data packet rates in the case of broadcast transmission. It may be observed from Figure 5 that the minimum average power consumptions obtained for 1 packet per second, 0.5 packet per second and 0.3 packet per second are at around 5%, 3.5% and 2%, respectively. These values correspond closely to the results listed in Table 3 which we have obtained from the analytical expression using equation (4). We can see that from both the implementation results and analytical results that the optimum duty cycle values decrease as the data transmission rate decreases.

Figure 5 Power consumption at a sensor node with different duty cycles for different data traffic rates (see online version for colours)

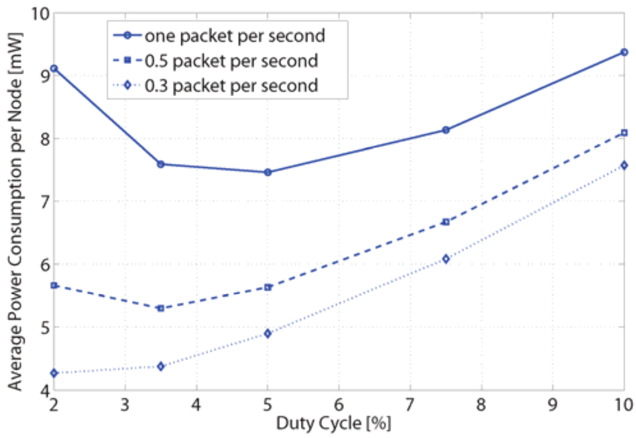

Table 3 Analytical values of the optimal sampling periods at different data packet rates

Data packet rate [s^{-1}]	*Optimal duty cycle [%]*
1.0	5.10
0.5	3.61
0.3	2.80

6 Optimum transmit power levels on the dual-radio platform

In dual-radio platforms, the transmission power levels of both the radios are set high enough so that both the radios on the receiving node are in their receiving range. If either of the two radios is unable to reach the receiving node, the communication remains unsuccessful. If the control channel radio does not reach a destination node, data communication cannot be established at all while in the converse case, data packet remains undelivered. As a crude design principle, the transmit range of the bursty data radio should at least be the same as that of the control channel in order to allow the bursty radio to deliver the data packets which have already been announced on the control channel. However, setting the transmit power level of the data channel and hence its range to be longer leads to energy wastage. Ideally, the transmit powers of the two radios should be set so as to achieve the same coverage range. Owing to the different nature of the radios on a dual-radio platform, blindly setting the transmit power is not enough and it is also necessary that the two radios give the same packet delivery rates. Since the packet error rate depends upon the packet size and the two radios have different packet lengths, optimally setting the transmit power levels of the two radios becomes challenging.

In the following, we derive an analytical expression for the transmit power level ratios of the two radios. The received power P_r at a certain distance d is given by the Friis transmit equation, $P_r = \frac{P_t G_t G_r \lambda^2}{(4\pi d)^2}$, where P_t is the transmit power, λ is the wavelength of the radio wave while G_r and G_t represent the gains of the receive and transmit antennas, respectively. The ratio of the transmit powers of the two radios (1 and 2, respectively) for an equal range is given by, $\frac{P_{t1} G_{t1} G_{r1} \lambda_1^2}{P_{r1}} = \frac{P_{t2} G_{t2} G_{r2} \lambda_2^2}{P_{r2}}$. Let's assume that the receive and transmit gains of each of the antennas are the same, i.e. $G_{t1} = G_{r1} = G_1$ and $G_{t2} = G_{r2} = G_2$. Thus, the ratio of the transmit powers is related to the ratio of the received power of two radios with equation $\frac{P_{t1}}{P_{t2}} = \frac{P_{r1} G_2^2 \lambda_2^2}{P_{r2} G_1^2 \lambda_1^2}$.

The two radios are targeted for different goals (low-power sniffing and bursty communication) on the dual-radio platform and therefore, potentially use different modulation schemes and have different receiver sensitivities. This leads to different Bit Error Rates (BER) (and hence different Packet Error Rates (PER), $PER = 1 - (1 - BER)^L$, where L is number of bits in one packet) on the two radios at a particular received power level. The optimal transmit power level ratios of the two radios can be determined by obtaining the ratio of the received power levels of the two radios at the receiving node giving the desired PER. The received power can be obtained from the Signal to Noise Ratio (SNR) versus BER curves through the relation,

$$SNR = \frac{P_r}{N_o} = \frac{E_b}{N_o} R .$$

Here N_o represents the noise power spectral density, E_b represents the energy per bit and R is the bit rate. The noise spectral density depends upon the temperature of the antenna and is given by $N_o = k_B T_0$, where k_B is the Boltzmann constant and T_0 is the temperature at the antenna. Therefore, P_r can be expressed as $P_r = \frac{E_b}{N_o} R(k_B T_0)$, and the ratio of the optimal transmit powers of the two radios is given by,

$$\frac{P_{t1}}{P_{t2}} = \frac{\left(\frac{E_b}{N_o}\right)_1 R_1 T_{o1} G_2^2 \lambda_2^2}{\left(\frac{E_b}{N_o}\right)_2 R_2 T_{o2} G_1^2 \lambda_1^2} \tag{5}$$

We can choose $\frac{E_b}{N_o} = SNR$ using the SNR-BER curves, which leads to the same PER value on both the radios.

On our prototype, the control channel radio operating in the lower frequency band uses FSK modulation while the bursty data radio operating in higher frequency band uses O-QPSK modulation. The probability of bit errors for the two modulation schemes in AWGN channels is given by $BER_{FSK} = Q\left(\sqrt{\frac{E_b}{N_o}}\right)$ and $BER_{OQPSK} = Q\left(\sqrt{\frac{2E_b}{N_o}}\right)$. Since the packet error rates of the two radios are required to be the same, $PER_1 = PER_2$. Substituting the equations for the two radios on our prototype leads to

$$\left(\frac{E_b}{N_o}\right)_1 = \left(Q^{-1}\left(1 - \left(1 - Q\left(\sqrt{2\left(\frac{E_b}{N_o}\right)_2}\right)\right)^{\frac{L_2}{L_1}}\right)\right)^2 . \tag{6}$$

After selecting the SNR values of the two radios satisfying equation (6), we determine the optimal transmit power level ratios the two radios on our prototype board using the equation (5).

7 Performance evaluation

We carried out extensive performance evaluation of our prototypic MR-MAC protocol implementation in terms of power consumption and latency. We also compared and analysed our experimental results to the widely used B-MAC protocol both on TelosB and MICA2 sensor node platforms. TelosB has an onboard CC2420 radio whereas MICA2 has a CC1000 radio. These radios are respectively the high frequency band and low frequency band radios on our prototype platform. Using same radio models for performance comparison gives an insight on the advantages of dual-radio MAC schemes on dual-radio platform against those using single radio platforms. Comparative experiments for the three platforms were carried out under the same traffic loads, duty cycles, radio transmit power levels and the network size.

7.1 Power consumption analysis of MR-MAC

Power consumption is one of the most important performance metrics for WSNs. Power consumption heavily depends on the operational duty cycle of a MAC protocol. If there is no traffic in the network, energy is drained by periodically polling the channel. Low duty cycles obviously result in less energy consumption in this case. When there is traffic in the network, power consumption directly depends on the amount of control overhead required to explicitly or implicitly synchronise the nodes and the amount of traffic to be delivered. The amount of control overhead associated with node synchronisation in preamble sampling protocols is directly related to the operating duty cycle values. We measured the power consumption of MR-MAC on our prototyped sensor node platform and that of B-MAC protocol on both TelosB and MICA2 platforms. We consider different duty cycle values and different amount of traffic loads for a thorough evaluation.

Table 4 lists the energy and power consumption breakdown for the basic operations of both the CC1000 and the CC2420 radio chips on our platform working at 3V. The power consumption of CC2420 chip in active mode is approximately twice that of CC1000 radio. This fact strongly supports our idea of using the CC1000 radio chip for control packets.

Table 4 Energy and power consumption breakdown for our prototype sensor node platform

Operation breakdown	*Energy [μJ]*
Turning on CC1000 to Tx mode	109
Turning on CC1000 to Rx mode	65.2
Turning off CC1000	6.22
Turning on CC2420 to Tx mode	15.6
Turning on CC2420 to Rx mode	15.7
Turning off CC2420	0.313
Operation breakdown	*Power [mW]*
CC1000 in Transmission Mode	31
CC1000 in Reception Mode	25
CC2420 in Transmission Mode	52.4
CC2420 in Reception Mode	55.4

We measured the power consumption of MR-MAC in unicast transmission when the transmitter can precisely estimate the sleep schedule of the receiver. Figure 6(a) shows the power consumption of MR-MAC protocol for unicast transmission with and without unicast preamble optimisations. It can be observed that combining preamble strobing with neighbourhood sleep schedule-based preamble optimisations results in a significant amount of energy saving, especially at low operating duty cycles. With the perfect knowledge of the receiver's sleep schedule, the transmitter wakes up to send preamble framelet when the receiver starts channel polling. A sequence of synchronisation bytes (0x55 or 0xAA) enables a receiver to detect the presence of the carrier. Consequently, the receiver starts searching for the radio locking sequence (0x33CC) following the synchronisation bytes. The actual preamble frame containing the control information follows the radio locking sequence. In case of perfect knowledge about a receiver's sleep schedule, a transmitter needs to transmit only one framelet. The length of channel polling, also known as clear channel assessment duration, is long enough to cover the waiting duration for framelet acknowledgement between adjacent preamble framelets. Furthermore, both the transmitter and the receiver have to pay the price of receiving and transmitting the acknowledgement of the preamble frame, respectively. The slight power consumption difference between the transmitter and the receiver is caused by the differences in power consumption of radios in different radio states, i.e. transmitting and receiving. The power consumption curve of the transmitter follows the trend of the receiver and increases with increasing duty cycle due to the need for more frequent channel polling. This argument is supported by the fact that the increasing slopes of the addressed receiver, non-addressed receiver and low power listening node (without any transmitter in the vicinity) are approximately the same. A number of factors contribute to the difference between the receiver with and without unicast optimisation. Although the receiver without optimisation saves energy in acknowledgement transmission and less radio states switching (from Receive-to-Transmit and Transmit-to-Receive), it needs to listen to an average of 1.5 preamble framelets instead of one due to the lack of implicit synchronisation between nodes. Overall, the power consumption of receiver with unicast optimisation is slightly lower than without the unicast optimisations. The offset between the addressed receiver and non-addressed receiver indicates the energy consumed in data reception for the case of addressed receiver. The power consumed by the non-addressed receiver is lower than that of the low power listening node because the non-addressed receiver calculates the ongoing transmission time and prolongs its sleeping time accordingly till the end of data transmission.

Figure 6 (a) Power consumption performance of multi-radio MAC running on the prototype platform at different duty cycles. The transmission rate was chosen to be 1 Hz, the traffic pattern is unicast and the data packet size was taken to be 1000 bytes. (b) Power consumption performance of multi-radio MAC running on the prototype platform at different data sizes. The frequency of transmission is 1 Hz, the traffic pattern is unicast and the duty cycle is 1% (see online version for colours)

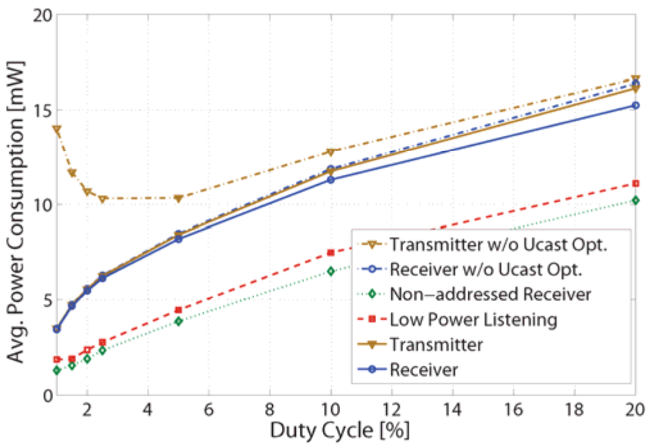

(a)

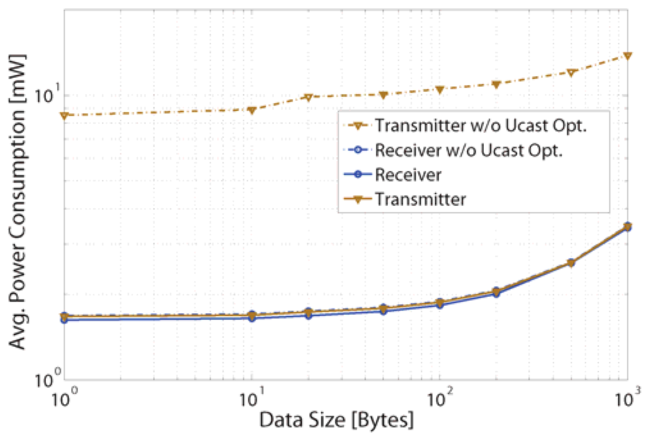

(b)

MR-MAC is evaluated by varying data sizes under the same duty cycle operation. In Figure 6(b), we can see a smooth increase of energy consumption with increasing data size. When the data size is small, data frame preamble is used to avoid the use of high frequency radio. When the data size is large, bursty radio is used to deliver large data packets very quickly and energy efficiently. The threshold of switching data frame preamble to micro frame preamble is found to be 20 bytes. This figure is found out analytically (and is also confirmed by empirical studies) in order to achieve lowest receiver power consumption possible.

7.2 Power consumption comparison of MR-MAC to B-MAC

In order to quantify the advantages of the dual-radio MAC protocol more effectively, we carried out the energy performance comparison of MR-MAC with B-MAC protocol using both the radios as used by MR-MAC. Comparison results against the widely used B-MAC protocol as the reference will help the sensor network community to understand the energy conservation gains of the MR-MAC design in a better way. The experimental performance comparison is carried out from two aspects: by analysing the power consumption at different data sizes while keeping the duty cycle constant, and vice versa. MR-MAC by default uses the maximum supported baud rate of 76.8 kbps on the CC1000 radio. All the experiments for standalone evaluation measurements of MR-MAC are carried out at this rate. For comparisons with B-MAC, we lower down the baud rate of CC1000 chip on our prototype sensor node to 19.2 kbps to be consistent with that of B-MAC implementation on MICA2 for a fair comparison. Figure 7(a) and Figure 7(b) show the power consumption of transmitters with varying duty cycles and data sizes respectively, when broadcast transmission is carried out. Since the preamble length of a broadcast transmission cannot be shortened, no improvement is achieved by preamble optimisation methods in MR-MAC. The receiver power consumption behaviour of a broadcast transmission is the same as the receiver without unicast optimisations. Unlike MR-MAC, which is capable of transmitting a train of data frames with only a single preamble reservation, B-MAC can transmit only one packet each time it seizes the channel. This accounts for the bigger difference between B-MAC and MR-MAC when the data size increases. The maximum packet size of the B-MAC is limited to 255 bytes for MICA2 platform and 122 bytes for TelosB platform. In Figure 7(a), we can see that at low data sizes, the offset between B-MAC and MR-MAC is almost constant. The power consumption of B-MAC on MICA2 shoots up when data size goes above 250 bytes while for B-MAC on TelosB a significant step up is observed at data sizes greater than 100 bytes. Although the maximum data size used in this experiment is 1000 bytes, MR-MAC can support up to 4095 bytes with a single preamble reservation. When the data size is below 100 bytes, MR-MAC still outperforms B-MAC mainly due to the energy saved by using low frequency sniffer radio for preamble transmission (as compared to B-MAC on TelosB) and using data frame preamble when data size is small and high frequency burst radio when data size is large (as compared to B-MAC on MICA2). While analysing the performance at various duty cycle, the data size is kept constant to be 100 bytes as shown in Figure 7(b). At this data size, B-MAC does not need multiple packet transmissions. The general behaviour of MR-MAC is similar to B-MAC due to the nature of preamble sampling MAC protocols. The power consumption decreases initially with an increase of duty cycle since the required preamble becomes shorter. When the duty cycle goes above 5%, the energy spent in channel polling starts dominating over the energy spent in preamble transmission and the overall power consumption starts increasing.

Figure 7 (a) Power consumption performance of transmitters running MR-MAC, B-MAC on MICA2 and TelosB at different data sizes. The transmission rate is 1 Hz, the traffic pattern is broadcast and the duty cycle is 1%. (b) Power consumption performance of transmitters running MR-MAC, B-MAC on MICA2 and TelosB at different duty cycles. The transmission rate is 1 Hz, the traffic pattern is broadcast and data size is 100 bytes (see online version for colours)

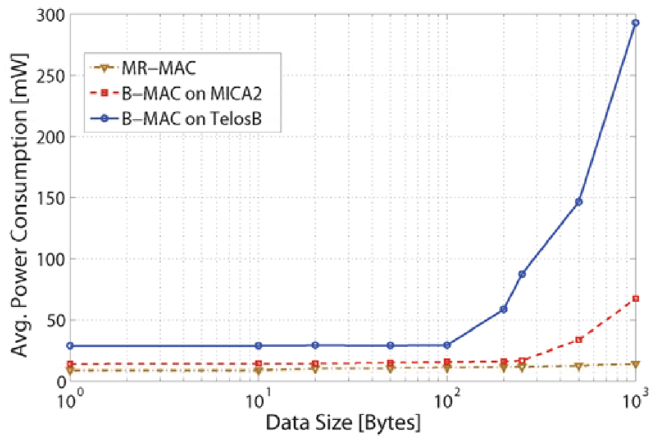

(a)

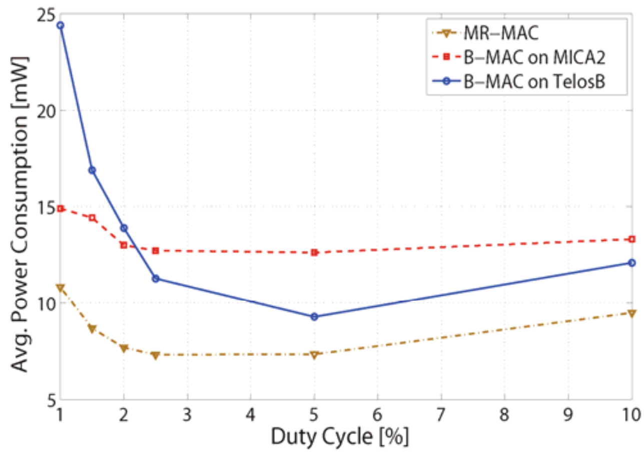

(b)

Figure 8(a) shows the performance of MR-MAC and B-MAC in the case of unicast transmission with different duty cycles. Comparing to Figure 7(b), there is no observable difference in terms of power consumption for B-MAC. However, there is a significant amount of power savings for MR-MAC especially at low duty cycle due to the preamble optimisation techniques. We can observe that as the duty cycle increases, the power consumption difference between MR-MAC and B-MAC on TelosB gets smaller. Since the preamble length decreases as duty cycle increases, the difference between non-optimised preamble length and optimised preamble length decreases. As the preamble length approaches to one preamble-frame size, the advantages brought by preamble optimisation techniques become less significant at the transmitter as well as at the receiver. In Figure 8(b), we can see that the power consumption in the low power listening mode of B-MAC on MICA2 is greater than the other two at low duty cycle values. Since CC1000 radio chip is used by both MICA2 and our prototype platform for channel polling, the power consumption by the radio chip should be the same for these two setups. The difference observed here is due to the different microprocessors used by the platforms. Atmel's 128L used by MICA2 is not as energy efficient as Texas Instrument Inc.'s MSP430 used by TelosB and our dual-radio platform. Although CC2420 radio offers more energy efficient switching between radio states, it consumes higher energy while sniffing the channel. Although at low duty cycle, the power consumption difference between B-MAC on TelosB and MR-MAC is very insignificant, the difference accumulates as duty cycle increases, i.e. the frequency of clear channel assessment activity increases. The non-addressed nodes using B-MAC on both the platforms suffer from receiving meaningless preambles and data packets. For longer preamble, more energy is wasted at the non-addressed receiving node.

Figure 8 (a) Power consumption performance of transmitters and receivers running MR-MAC, B-MAC on MICA2 and TelosB at different duty cycles. The transmission rate is 1 Hz, the traffic pattern is unicast and data size is 100 bytes. (b) Power consumption performance of non-addressed receivers and low power listening nodes running MR-MAC, B-MAC on MICA2 and TelosB at different duty cycles. The data size is 100 bytes (see online version for colours)

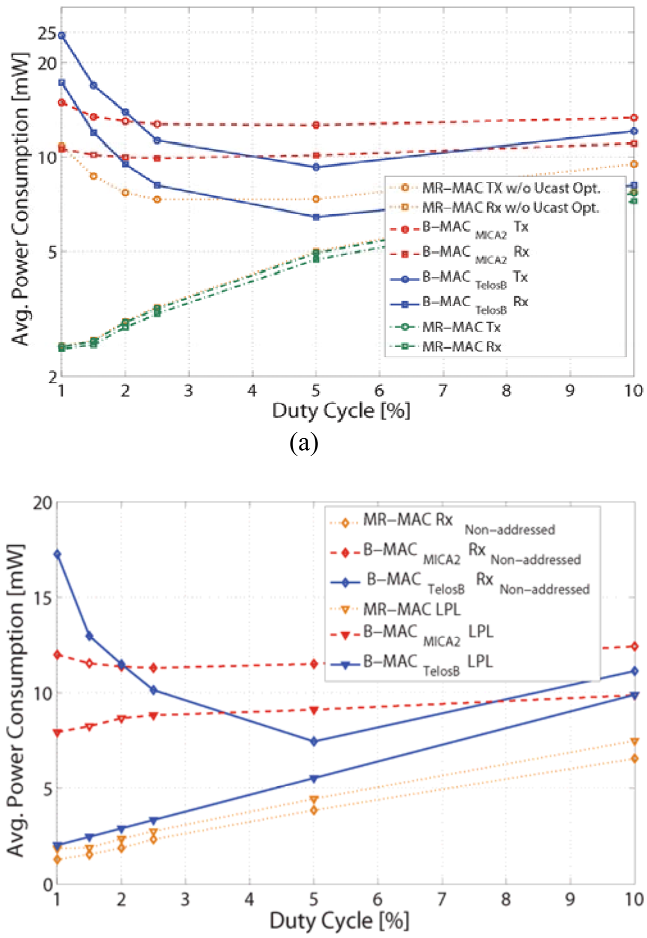

(b)

7.3 *Latency*

Latency is defined as the duration between when the packet is put into the waiting queue of transmitter and when it is successfully received by the receiver. It is an important metric in evaluating the performance of MAC protocols because in most of the applications, data has its shelf-life. If not delivered on time, data loses its significance. Latency requirements are highly dependent on the applications. Event detection or object tracking applications poses strict and short deadline on data delivery while many of the environment monitoring applications can afford loose deadlines. WSN MAC protocols have an inherent latency because of the duty cycle operation. Latency is induced in packet forwarding since the packet can only be delivered when the receiver wakes up. This latency exists in every hop and thus in multihop networks, the end-to-end latency increases with the number of hops.

Several optimisation techniques can be applied to duty cycle based MAC protocols in this regard. Data aggregation at nodes reduces the waiting time of packets in the transmission queue when all the packets in the queue can be transmitted back-to-back after the transmitter seizes the channel. With prior knowledge about the network topology, wake-up schedules of the nodes can be adjusted to optimise latency performance (Basagni et al., 2004; Miladi et al., 2006). Adaptive duty cycles can help in achieving a balance between energy consumption and latency requirement by adopting lowest duty cycle possible which satisfies the current latency requirement. In this paper, we measured the latency of MR-MAC over multiple hops at different duty cycles and data sizes. An n-hop circular network consisting of *n* nodes is used. Data are sent from one particular node and routed around a circle back to this node. When a node receives a packet, it immediately forwards the packet to the next node. The latency is measured by running a timer at the node which initiates the traffic. The multi-hop latency is taken as the time interval between the starting point of transmitting the packet (preamble inclusive) and the end of the reception of the packet. The rate of packet generation is low so that no packets are queued at the transmitters. Furthermore, the medium is made sure to be free of congestion and any interference.

From the measurements shown in Figure 9(a), we can observe that latency increases with decreasing duty cycle values since the average waiting time of transmitters for their receivers to wake-up increases with decreasing duty cycle. Latency is also affected by packet sizes since bigger packets require longer transmission time. Furthermore, assuming the same offset in the wake-up schedule between nodes in all the experiments, longer transmission results in higher possibility that a forwarding node, having to receive the packet during the wake-up period of its destination node, delays the transmission till the next wake-up period of the destination node. For example, node C wakes up in 30 ms after node B wakes up. For one transmission which requires 10 ms, node B can start transmission to node C in 20 ms after packet reception. However, if the packet transmission requires more than 30 ms, node B needs to wait for the next wake-up time of node C to start transmission. This delay increases the latency by a period length. Therefore, the offset between lines in Figure 9(a) with same duty cycle and different data sizes is not constant. It includes the difference in data transmission time, as well as the extra latency introduced by the difference in data transmission time.

Experiments were carried out in order to compare the latency of MR-MAC over multi-hop links to B-MAC on TelosB as well as MICA2. Figure 9(b) shows the latency of MR-MAC, B-MAC on TelosB and MICA2 at an operating duty cycle of 1%. Since receivers' duty cycles are the same for all three cases, channel polling time determines the wake-up period and thus the latency. The packet size in this experiment is set to be 100 bytes so that only one data packet transmission is required. From the figure, it is shown that B-MAC on MICA2 has a lower latency than B-MAC on TelosB because it has a shorter clear-channel assessment time. The channel polling time of MR-MAC is longer than B-MAC implementation on MICA2 since it has to cover the acknowledgement waiting gap between the preamble-frames. Although MR-MAC uses a shorter period of time to transmit 100 bytes than B-MAC on MICA2, the latency performance at low duty cycle MACs is still dominated by channel polling time. The curve for MR-MAC without unicast preamble optimisations is shown here to support the previous statements. MR-MAC with preamble optimisations outperforms B-MAC due to its preamble strobing feature. In B-MAC, data transmission takes place after the entire preamble has been transmitted. In MR-MAC, upon receiving the acknowledgement of a preamble frame, data transmission can take place immediately afterwards on the data channel. Preamble strobing technique approximately reduces the latencies to half.

Figure 9 (a) Latency performance of MR-MAC with various data sizes and duty cycles. (b) Comparison of latency between MR-MAC and B-MAC at 1% duty cycle (see online version for colours)

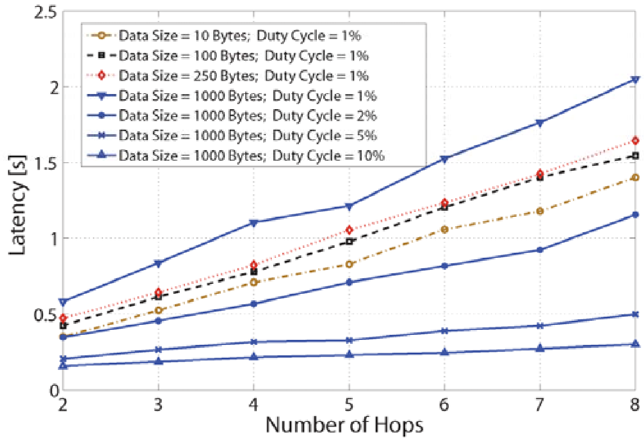

(a)

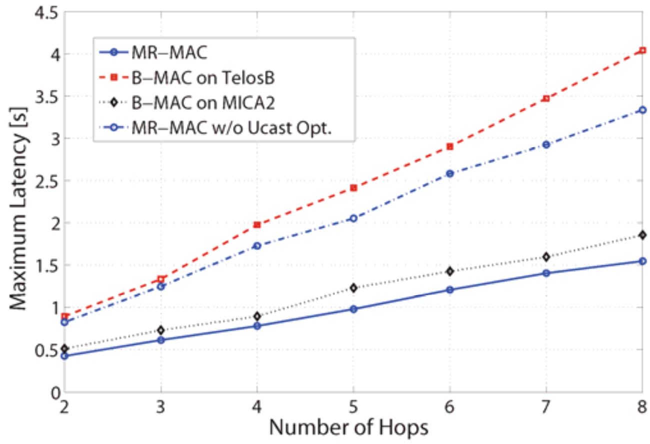

(b)

8 Conclusions

In this paper, we have presented a new MAC protocol for dual-radio platforms. Our protocol is able to handle the two main radio activities, namely the idle listening and the short-timed burst data communication, using a combination of two specialised radios to result in highly energy efficient manner. The protocol is based on the preamble sampling technique, which implicitly synchronises the asynchronously waking up nodes only when data communication is required. It does not inflict any extra overhead in order to maintain synchronisation among the nodes and efficiently handles the dynamism in the network caused by the mobility and appearance/disappearance of nodes. Our MAC protocol exercises various types of preamble shortening techniques in order to conserve energy.

The use of a dedicated control channel in a less congested low frequency band allows having spectrum agile medium access in more crowded and interference prone high frequency band. We have carried out the implementation of the protocol on a prototypic platform consisting of COTS sensor node and a radio module. We have analytically modelled the optimal duty cycle expression for a given amount of traffic load in a particular sized network and have shown empirically that our real node test-bed implementation fully adheres to the optimal duty cycle expression. We have also derived an expression for the optimal ratio of transmit power levels of the two radios on our prototype platform. The paper presents an extensive experimental performance evaluation of the protocol in terms of power consumption and latency over different duty cycle values and under various amounts of traffic loads. We have also performed similar comparative studies against the most widely used B-MAC protocol implemented on TelosB and MICA2 platforms. These two single radio platforms have CC2420 and CC1000 radios on-board, which are the two radios used by the MR-MAC prototype hardware platform. The comparison gives us an insight in quantifying the advantages of the dual-radio MAC approach to a MAC using a single radio. Our MAC protocol achieves remarkably better performance as compared to the popular B-MAC protocol.

Acknowledgements

We would like to thank financial support from European Union (project IST-034963-WASP), UMIC Research Centre and RWTH Aachen University.

References

Ahn, G-S., Hong, S.G., Miluzzo, E., Campbell, A.T. and Cuomo, F. (2006) 'Funneling-MAC: a localized, sink-oriented MAC for boosting fidelity in sensor networks', *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys'06)*, Boulder, Colorado, USA.

Ansari, J., Riihijärvi, J., Mähönen, P. and Haapola, J. (2007a) 'Implementation and performance evaluation of nanoMAC: a low-power MAC solution for high density wireless sensor networks', *International Journal of Sensor Networks*, Vol. 2, No. 5, pp.341–349.

Ansari, J., Zhang, X. and Mähönen, P. (2007b) 'Demo abstract: multi-radio medium access control protocol for wireless sensor networks', *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems (SenSys'07)*, Sydney, Australia.

Ansari, J., Zhang, X. and Mähönen, P. (2008) 'Multi-radio medium access control protocol for wireless sensor networks', *Presented at the Workshop on Energy in Wireless Sensor Networks Adjunct to IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'08)*, Santorini Islands, Greece.

Ansari, J., Pankin, D. and Mähönen, P. (2009) 'Radio triggered wake-ups with addressing capabilities for extremely low power sensor network applications', *International Journal of Wireless Information Networks*, Vol. 16, No. 3, pp.118–130.

Bachir, A., Barthel, D., Heusse, M. and Duda, A. (2006) 'Micro-frame preamble MAC for multihop wireless sensor networks', *Proceedings of the IEEE International Conference on Communications (ICC'06)*, Vol. 7, pp.3365–3370.

Basagni, S., Carosi, A. and Petrioli, C. (2004) 'Sensor-DMAC: dynamic topology control for wireless sensor networks', *Proceedings of the 60th IEEE Vehicular Technology Conference*, September, Los Angeles, USA.

Buettner, M., Yee, G.V., Anderson, E. and Han, R. (2006) 'X-MAC: a short preamble MAC protocol for duty-cycled wireless sensor networks', *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys'06)*, pp.307–320.

Cormio, C. and Chowdhury, K.R. (2009) 'A survey on MAC protocols for cognitive radio networks', *Elsevier Journal of Ad Hoc Networks*, Vol. 7, No. 7, pp.1315–1329.

Demirkol, I. Ersoy, C. and Alagoz, F. (2006) 'MAC protocols for wireless sensor networks: a survey', *Communications Magazine, IEEE*, Vol. 44, No. 4, pp.115–121.

El-Hoiydi, A. and Decotignie, J-D. (2004) 'WiseMAC: an ultra low power MAC protocol for the downlink of infrastructure wireless sensor networks', *Proceedings of the 9th International Symposium on Computers and Communications*, pp.244–251.

Henry, M.M. (2005) *NSF Spectrum Occupancy Measurements Project Summary*, Shared Spectrum Company, Technical Report.

Islam, M., Koh, C., Oh, S., Qing, X., Lai, Y., Wang, C., Liang, Y-C., Toh, B., Chin, F., Tan, G. and Toh, W. (2008) 'Spectrum survey in Singapore: occupancy measurements and analyses', *Proceedings of the 3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom).*

Li, J. and Lazarou, G.Y. (2004) 'A bit-map-assisted energy-efficient MAC scheme for wireless sensor networks', *Proceedings of the International Conference on Information Processing in Sensor Networks (IPSN'04)*, pp.55–60.

Lopez-Bentez, A.U.M. and Casadevall, F. (2008) 'Evaluation of spectrum occupancy in Spain for cognitive radio applications', *Proceedings of the 67th IEEE Vehicular Technology Conference.*

Mahlknecht, S. and Boeck, M. (2004) 'CSMA-MPS: a minimum preamble sampling MAC protocol for low power wireless sensor networks', *IEEE International Workshop on Factory Communication Systems.*

Miladi, M., Ezzedine, T. and Bouallegue, R. (2006) 'Latency of energy efficient MAC protocols for wireless sensor networks', *Proceedings of the International Conference on Digital Telecommunications (ICDT'06)*, Cote d'Azur, France.

Oppermann, I., Stoica, L., Rabbachin, A., Shelby, Z. and Haapola, J. (2004) 'UWB wireless sensor networks: UWEN-a practical example', *IEEE Communications Magazine*, Vol. 42, pp.527–532.

Polastre, J., Hill, J. and Culler, D. (2004) 'Versatile low power media access for wireless sensor networks', *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys'04)*, Baltimore, Maryland, USA.

Rhee, I., Warrier, A., Aia, M., Min, J. and Sichitiu, M. (2005) 'Z-MAC: a hybrid MAC for wireless sensor networks', *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems (SenSys'05)*, San Diego, California, USA.

Ruzzelli, A., O'Hare, G., Jurdak, R. and Tynan, R. (2006) 'Advantages of dual channel MAC for wireless sensor networks', *International Conference on Communication System Software and Middleware*, New Delhi, India.

Schurgers, C., Tsiatsis, V., Ganeriwal, S. and Srivastava, M. (2002) 'Optimizing sensor networks in the energy-latency-density design space', *IEEE Transactions on Mobile Computing*, pp.70–80.

TinyOS (2009) *TinyOS 2.x Enhancement Proposals*. Available online at: http://www.tinyos.net (accessed on 15 January 2010).

van Dam, T. and Langendoen, K. (2003) 'An adaptive energy-efficient MAC protocol for wireless sensor networks', *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys'03)*, pp.171–180.

van Hoesel, T. and Havinga, P. (2004) 'A lightweight medium access protocol (LMAC) for wireless sensor networks', *International Conference on Networked Sensing Systems (INSS'04)*.

Ye, W., Heidemann, J. and Estrin, D. (2002) 'An energy-efficient MAC protocol for wireless sensor networks', *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, pp.1567–1576.

Ye, W., Silvia, F. and Heidemann, J. (2006) 'Ultra-low duty cycle MAC with scheduled channel polling', *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys'06)*, pp.321–334.