

Tradeoff Analysis of PHY aware MAC in Low-Rate, Low-Power UWB networks

Alaeddine El Fawal and Jean-Yves Le Boudec and Ruben Merz and Božidar Radunović and Jörg Widmer and Gian Mario Maggio
EPFL Technical Report IC/2005/032

Abstract

We are interested in the design of physical layer (PHY) aware MAC for self-organized, low power, low data rate impulse-radio ultra-wideband (IR-UWB) networks. In such networks, energy consumption is much more of a concern than the achieved data rates. So far, a number of different architectures have been proposed for IR-UWB, but in the context of rate efficiency. The choices made for rate efficient designs are not necessarily optimal when considering energy efficiency. Hence there is a need to understand the design tradeoffs in low rate operation. Our aim is to present the different design alternatives, to evaluate their suitability and assess their costs and benefits for networks with very low energy consumption. We identify the four main functions a PHY-aware MAC design has to provide: (1) interference management, (2) access to a destination, (3) sleep cycle management, and (4) signal acquisition. Then we present a non-exhaustive list of the many possible *building blocks* and use it to analyze the design choices that existing proposals make. Finally, we review the performance implications of these design choices with respect to a very low power ultra-wideband networking architecture.

1. INTRODUCTION

Emerging pervasive networks assume the deployment of large numbers of wireless nodes, embedded in everyday life objects. In this type of networks, the focus is rather on minimizing energy consumption than maximizing rate. There exist numerous possibilities to implement a PHY-aware MAC design for low-rate, low-power UWB networks. Hence, there is a need to understand the design and implementation tradeoffs.

In traditional designs, there is a clear frontier between the medium access control (MAC) layer and the physical layer (PHY). The MAC layer primary goal is to coordinate access to the physical layer and to ensure that interference is limited to an acceptable minimum. Also, the MAC should permit nodes to sleep when no data communication is necessary. The physical layer is responsible for the actual transmission of information bits between the nodes that should communicate. It also controls rate and power of the transmission. In general, there is no interaction between the two layers and the MAC layer has no control over the power or rate used by the physical layer.

In PHY-aware MAC designs, the MAC has access to some or all of the physical layer parameters. Interference does not need to be completely prevented, but it needs to be managed (Section 2). For example, the rate or the power can be dynamically adapted to the level of interference. Examples of such schemes for UWB are [3] and [7]. In [7], rather than preventing interference, sources adapt their rate such that their destination can sustain the interference.

A very fundamental design decision is, whether to allow interference (i.e., permit concurrent transmission within the “same area”) or to enforce mutual exclusion. Also allowing random access or imposing some form of super-frame structure within which transmissions have to occur, whether to use power control, and how to coordinate nodes such that many of them can sleep are important design decisions. Many of these considerations have implications on the physical layer as well as the MAC layer. And as demonstrated in [3] and [7], a PHY-aware MAC protocol can significantly improve the performance.

We concentrate on large self-organized networks and do not address the case of Wireless Personal Area Networks (WPAN). We focus here on UWB impulse radio (IR) physical layer systems for low data-rate (LDR) applications. These systems make use of ultra-short duration (< 1 ns) pulses which yield ultra-wide bandwidth signals characterized by low duty cycle ($\simeq 1\%$) and extremely low power spectral densities [14]. The multiple-access interference (MAI) for such systems, unlike narrow-band systems, is caused by “collisions” occurring between pulses belonging to different simultaneous transmissions. The digital information to be transmitted may be encoded by using pulse position modulation (PPM) and/or binary phase shift keying (BPSK). UWB-IR systems are especially attractive for LDR wireless communications as they potentially combine low power consumption, immunity to multipath fading and location/ranging capability¹. A complete design targeting energy

A. El Fawal, J.-Y. Le Boudec, R. Merz, B. Radunović and J. Widmer are with EPFL, School of Computer and Communication Sciences, CH-1015 Lausanne, Switzerland. G.-M. Maggio is with STMicroelectronics, Advanced System Technology - R&I, Chemin du Champ-des-Filles 39, Case Postale 21, CH-1228, Geneva, Plan-les-Ouates, Switzerland. The contact author is R. Merz: ruben.merz@epfl.ch

¹In fact, a Task Group was formed in March 2004 to investigate a UWB alternative PHY to the IEEE 802.15.4 wireless standard, associated with Zigbee.

efficiency should also consider energy efficient routing [9]. However, for reasons of brevity we do not consider routing in this paper. We also do not consider ranging.

2. AN INTEGRATED VIEW OF THE PHYSICAL AND MEDIUM ACCESS CONTROL LAYERS

A PHY aware MAC layer globally *manages* the interference and medium access on a shared communication channel. The main goal is to maximize the overall lifetime of the network. Still there is the complementary goal which is to maximize the rate offered to each node while possibly remaining fair. In a PHY-aware MAC, the following set of functions must be provided:

- *Interference management*: a source can *control* the interference it creates by controlling the transmit power or the time when a packet is transmitted, or it can *adapt* to the existing interference (by reducing its rate to permit reliable reception at the destination).
- *Access to a destination*: we assume that a node can either send or receive from one source. Thus, an exclusion protocol is necessary to enforce that only one source communicates with the destination. This *private* exclusion protocol only involves the potential sources and the destination. This issue would not exist in a centralized network with a basestation using multiuser detection [13].
- *Sleeping*: although not strictly required, it is obviously necessary in a low power context. There is an important tradeoff between long sleep cycles that permit to efficiently save energy and short cycles that facilitate communication and improve responsiveness.

In addition, we also consider *signal acquisition*: this function belongs to the physical layer, but its performance is impacted by design choices of the PHY-aware MAC layer. It consists of acquiring the timing and code synchronization of the source, in order to detect whether the node under consideration is among the set of intended destinations. Typically, when a preamble is sent, channel estimation is performed as well. Signal acquisition can be critical for UWB-IR systems in dense multipath environments, since the signal energy is spread over several paths, making their detection challenging, especially in the presence of noise and other interfering users.

A. Available Building Blocks for a PHY-aware MAC

In the following, we list a set of building block that can be used to implement the functions described above.

a) *Rate control and adaptation*: Often, the transmit rate is adapted as a function of the channel condition (essentially the attenuation) between the source and the destination. However, the rate can also be adapted as a function of the interference created by other devices in the network.

Rate control can be done by controlling the modulation order, the time-hopping spreading gain, or the channel code rate used at the physical layer. The rate is normally adapted based on feedback from the destination. This feedback is based on statistics gathered at the receiver either in a predictive or in a reactive manner. For the former, a source inserts a pilot symbol in a packet and the channel is measured at the receiver based on the received pilot symbol. For the latter, the receiver typically looks at local statistics such as the likelihood ratios at the output of the receiver.

Note that rate control involves no nodes other than the source-destination pair.

b) *Power control*: the transmit power can be adjusted to keep the signal to noise ratio (SNR) at the destination constant, or to minimize the amount of interference created on the neighbors.

Contrary to rate control, power control requires interaction with other devices in the network. If a source increases its transmit power, it will create more interference on concurrent receivers. Hence, a source needs to know not only the minimum power required by its destination to ensure proper signal detection and decoding but also the maximum interference that ongoing transmissions in the vicinity of the transmitter can sustain.

c) *Mutual exclusion*: A mutual exclusion protocol prevents nodes from transmitting at the same time. It is often implemented by control packet signaling. Hence, the number of nodes affected depends on the transmit power of the control packets. For example, the RTS-CTS handshake procedure of 802.11 implements mutual exclusion.

Mutual exclusion is an effective way to reduce interference since it prevents concurrent transmissions. Most protocols use mutual exclusion to manage interference. However, as we show in Section 3.3.3, exclusion is not always necessary.

d) *Single channels versus multi-channel protocols*: In a multi-channel protocol, the transmission medium is separated into several orthogonal (or quasi-orthogonal) transmission channels. Since parallel transmissions can occur, there is a clear advantage in terms of rate increase. Still, a potential disadvantage is that it becomes impossible to overhear transmission from other active nodes on other channels.

By definition, two signals that send concurrently on two orthogonal channels will not or almost not interfere depending on the level of orthogonality. Orthogonal channels are inherent with a pulse based UWB physical layer thanks to time-hopping [14]. Note that for the channels created with time-hopping sequences to be perfectly orthogonal, a very accurate

synchronization is required among transmitters and the sequences need to be non-overlapping and aligned in time. Other possibilities are to separate the bandwidth into orthogonal sub-bands. Quasi-orthogonal and orthogonal channels inherently solve the traditional hidden-node terminal problem present in single-channel protocols. Still, in the case of quasi-orthogonal channels, appears the issue of the *near-far* effect. When an interferer is relatively distant from a receiver, the occasional interference due to the non orthogonality is often negligible. However, when the interferer is much closer to the receiver than its associated transmitter, the interference created becomes non-negligible for the receiver. Depending on the particular physical layer at use, the near-far effect can have more or less impact.

e) Multi-user reception techniques: With a single user decoder, all signals apart from the one coming from the user are considered to be noise. With a multiple-user receiver, signals coming from several users can be successfully received in a joint manner. For example, in case of a near-by, strong interferer, the interfering signal can be decoded, subtracted, and the total noise can be greatly reduced.

As such multi-user detection is potentially very attractive. However, this procedure generally necessitates to be accurately synchronized with all the sources we wish to decode and furthermore to know all their transmitted signal characteristics. In addition, the complexity of the decoding operation is excessively high. Nevertheless, thanks to the particular structure of pulse based time-hopping UWB signal, there exists several suboptimal techniques that are still worth considering. In particular it is possible to take advantage of the infrequent nature of collisions at the pulse level [7]; a receiver can estimate the average received energy from its current source. A pulse collision with a strong interferer can be easily identified since the received energy is then much higher. The resulting received symbol is then replaced by an erasure symbol and the loss of information due to erasures are recovered by the error correcting code.

f) Random access or scheduled access: Random access is straightforward to implement in its simplest form which is Aloha. However, at high utilization, the throughput becomes extremely low. As such, random based access protocols are often improved with some of the following components:

- Carrier-sensing is a mechanism that permits to avoid sending on the channel if it is already busy. Note that carrier sensing is impossible with UWB physical layers, since there is no carrier. One possibility to emulate carrier sensing with UWB is then to actively decode. This is especially complex in a network with multiple time-hopping sequences, since a node has to sense for all possible time-hopping sequences.
- A back-off procedure with timer management used to resolve collisions.
- Hand-shake procedure where nodes exchange RTS/CTS packets before each transmission to reserve medium access for data transmission. RTS/CTS are still transmitted using random access. However, these packets are much shorter than data packets, hence the performance penalty in case of a collision is lower. Such a hand-shake procedure can be private between a source and its destination or can involve more nodes.

Random access is typically used in ad hoc networks since it requires none or very few coordination among nodes.

The other possibility is to use centralized scheduling. One node, a coordinator, decides in each slot which nodes are allowed to send during that slot. It can allow only a single node to transmit (TDMA), or it can allow multiple transmissions if they do not interfere significantly. A coordinator can be a base station, or an arbitrary node elected by a network. Although this approach is more efficient from a medium access point of view, it is very difficult to implement in large, self-organized networks where not all nodes can hear each other.

g) Time-slotted transmission: Slotted transmission can reduce interference (as in Aloha) or improve power saving since a node can sleep during unused slots. It also facilitates signal acquisition. Synchronizing slots on the nano-second-level (fine-grain synchronization) required for signal acquisition is possible only in a centralized scenario with a single base-station. In our decentralized model, such a level of synchronization cannot be achieved due to the different propagation times between nodes. Still, the time slot structure provides a coarse synchronization. From this coarse synchronization, signal acquisition can be performed much more efficiently than without any time common reference.

h) Sleeping protocols: long preambles versus periodic beaconing: Letting nodes sleep is the most effective way to conserve energy in a wireless network. However, this requires a mechanism that allows nodes to be contacted even though they might be sleeping from time to time.

We consider two types of sleeping protocols. One is slotted and the other is unslotted. Slotted schemes assume there is a dynamically elected coordinator that periodically sends beacons (we do not discuss how to elect a coordinator). This beacon provides a coarse-level synchronization and denotes a start of a superframe. A superframe comprises of two parts. The first part is a reservation part. During that part senders announce transmission requests and receiver listens for these requests. The actual packet transmissions take place during the second part of the superframe. A receiver can then sleep for most of the second part, except for the periods when the announced transmissions will occur.

In the unslotted approach, each receiver wakes up according to its own schedule. A transmitter that wants to transmit to a certain receiver first needs to learn its schedule. Typically, in order to learn a schedule, a transmitter has to transmit a

preamble for the maximum sleep time. The destination is sure to wake up some time in between, will receive the preamble and then stay awake to receive the actual data packet afterward.

i) *Centralized architecture*: A design choice for all the above possibilities is to have a fully decentralized versus master-slave architecture where the network consists of one or several subnetworks, each controlled by a coordinator.

B. How and what are the building blocks used in existing designs

	Aloha	802.11	M-MAC	CA-CDMA	Bluetooth	802.15.4	MBOA	[3]	UWB ²	DCC-MAC
Rate adaptation		I	I					I		I
Power control				I				I		
Mutual exclusion		I,A	I,A	I,A	I,A	I,A	I,A	I,A		
Multi-channel			I	I	I			I	I	I
Multi-user detection										I
Access										
Random	A	A	A	A		A	A	A	A	A
Scheduled					A	A	A			
Slotted		S	A,S		A,S	A,S	A,S			
Sleeping										
Beaconing		S	S				S			
Long preamble										
Centralized					S	S				

TABLE I

BUILDING BLOCKS USED BY EACH PROTOCOLS FOR INTERFERENCE MANAGEMENT (I), ACCESS TO A DESTINATION (A) AND SLEEPING MODE (S)

We now present a classification of several proposed designs and existing standards. For each of the four functions previously described in Section 2 we analyze which building blocks were used and how they were used. We summarized the results in Table I. We do not limit ourselves to UWB designs since many of their concepts are borrowed from narrow-band designs.

- The Aloha protocol is the simplest design where interference (and access to the same destination) are not managed at all.
- The 802.11 protocol [5] is based on CSMA-CA with an optional RTS-CTS mechanism. Several different transmit rates coupled with an automatic repeat request (ARQ) mechanism allow the MAC to adapt to the channel condition between the source and the destination. Furthermore, the 802.11h standard [6], an evolution of the 802.11a standard, performs transmit power control in order to reduce interference. Nevertheless, interference management and access to a destination are based on mutual exclusion, with a narrow-band physical layer a collision is destructive and must be avoided. The impact of mutual exclusion depends on the RTS-CTS mechanism and the carrier sensing threshold. To enable sleeping, the access point broadcasts beacons that permit devices to easily and regularly switch to doze mode. In the case of an ad hoc mode configuration, each device broadcasts its own beacon. It is then up to a source to listen to the beacon of a wanted destination.
- Multi-channel designs based on 802.11 such as [12] take advantage of the several channels available for 802.11 networks to operate in. Hence, interference is managed by orthogonal channels and exclusion (in case two nodes use the same channel for their transmission). Access to a destination is enforced using a modified RTS-CTS procedure happening on a dedicated control channel.
- Bluetooth (or IEEE 802.15.1) [6] is based on a piconet paradigm where any communication occurs in a master-slave fashion. Bluetooth is a slotted protocol where only one node can transmit in any hopping slot. It uses neither rate nor power adaptation. For sleep management, the centralized and slotted structure permits nodes to easily sleep and wake up when necessary.
- 802.15.4 (Zigbee) [1] is a single-channel protocol based on CSMA-CA (with an optional RTS-CTS mechanism) for a narrow-band physical layer. The protocol supports two operating modes: (1) a so-called beacon-enabled mode where the network is organized as a slotted piconet. A piconet coordinator periodically broadcasts beacons (note that any node can operate as a piconet coordinator). Access inside a given slot uses exclusion and is arbitrated through CSMA-CA. (2) A distributed mode where communication occurs on a point-to-point basis using CSMA-CA for medium access control.

With (1), interference is managed entirely by exclusion with the help of a slotting procedure. Access to a destination is ensured by the optional RTS-CTS procedure or relies on collision detection and a backoff mechanism. Sleeping management is similar to that of Bluetooth.

- Power controlled CDMA based design (CA-CDMA) [8]: inherent in every direct sequence CDMA based designs is the need for power control. Even when orthogonal sequences are used, transmissions may interfere with each other and are subject to near-far effects due to the asynchronicity of the devices. Hence, interference is managed by a combination of mutual exclusion, power control and pseudo-orthogonal channels. Any attempt to communicate starts with a handshake between the source and its potential destination on a control channel. If no admissible power is found, the data transmission does not occur. As in [3], mutual exclusion has a variable impact depending on how much a destination can sustain interference. Devices use CTS-like packets on the control channel to communicate their power margin. For a given receiver, the smaller the power margin, the larger the number of nodes prevented from transmitting concurrently around the receiver.

We now present several existing designs for UWB physical layers.

- The MBOA MAC is very similar to 802.15.4 and emerged from the inconclusive effort of the IEEE 802.15.3a group working on a UWB physical layer for the IEEE 802.15.3 protocol². The main difference is that there is only one operating mode, a beacon based one, with slotted access in between beacon transmission.
- A joint power and rate controlled design for an UWB physical layer: in [3], the authors solve a joint power and rate link assignment problem. Based on the solution to the optimal optimization problem, the authors propose a suboptimal distributed algorithm.

Interference is managed by a mixture of mutual exclusion, by adapting the transmit power as well as the rate and by taking advantage of the pseudo-orthogonal channel due to time-hopping sequences. If after the distributed handshake procedure, there exists no satisfying power and rate assignment, no data communication occurs (exclusion). The number of nodes affected by mutual exclusion is variable. For every receiver there exists an interference margin, which indicates by how much the interference can increase without destroying the ongoing transmission. The smaller the interference margin, the larger the number of nodes prevented from sending. Access to a destination is enforced by the same RTS-CTS type of handshake that is used for finding the power and rate assignment.

- With UWB² [4], interference is managed by pseudo-orthogonal channels and access to a destination is managed by a handshake procedure. Whenever a device wants to talk to a particular destination, it starts an RTS-CTS exchange on a common channel. If the destination is not busy, it answers on the common channel and includes a particular dedicated time-hopping sequence in the CTS packet. The subsequent data transmission uses the particular time-hopping sequence proposed in the CTS packet.
- DCC-MAC [7] is a design based on theoretical results from [10]. It uses rate adaptation but no power control. It proposes to use interference mitigation (erasure of pulses that are well above the expected received power). It further builds on the fact that mutual exclusion is not needed in the case of a time-hopping UWB physical layer. Hence interference is managed by rate adaptation, pseudo-orthogonal channels through time-hopping sequences and a suboptimal multiuser type of receiver. DCC-MAC was designed to avoid the need for a control channel. As such, the problem of access to a destination is managed by a subtle control of timers and careful use of time-hopping sequence.

3. PERFORMANCE ANALYSIS OF DIFFERENT DESIGN CHOICES

In this section we formulate several important design choices and we analyze their performances. We first define performance metrics and energy models that we will use in the analysis. The remaining assumptions are summarized in Table II.

A. Performance Metrics

An important issue in choosing a performance metric is fairness. If one bases a metric solely on efficiency, it may favor very unfair network designs where a few nodes may receive zero rates, or may use up their energy much earlier than the rest of a network. Hence, we use log utility performance metrics, which are known to achieve a good tradeoff between efficiency and fairness [10]. We choose metrics consistent with the goal of maximizing network lifetime while keeping rates as high as reasonably possible. Thus our metrics are (1) the sum of logs of node lifetimes and (2) the sum of logs of average flow rates.

²The UWB physical layer used by MBOA is not based on impulse radio but on a multiband radio.

Topology	Randomly distributed on a 20m x 20m square. Links are chosen randomly
Load Models	(1) regular sensor readings (very predictable communication pattern where nodes individually regularly send out sensor readings)
	(2) emergency case, bursty (traffic due to an emergency situation, resulting for example in network flooding)
	(3) emergency case, delay sensitive (isolated, non predicted, but very urgent data, for example for shutting down a life threatening device)
Physical layer parameters	Pulse repetition period $PRP = 1000$ chips, chip duration $T_c = 1$ ns, Energy per pulse $E_p = 0.2818$ mW
Channel model	fading and multipath according to the 802.15.4a model

TABLE II
MODEL ASSUMPTIONS FOR THE PERFORMANCE ANALYSIS

B. Energy Consumption

We take advantage of the structure of the physical layer to define a chip-level model of energy consumption. During a chip, the physical layer can either transmit a pulse, receive a pulse, perform signal acquisition, be in an active-off state, or sleep. The active-off state occurs due to time-hopping. When a node is in between subsequent pulse transmissions or receptions, energy is consumed only to keep the circuit powered up but no energy is used for transmitting or receiving pulses.

Hence, to model the energy consumption, we break it down to consider the energy *per chip* for each state. Since it is difficult to give exact numbers for each of those states, we use relative values normalized to the energy for pulse transmission. These values are given in Table 3.2; q_{tx} is the cost for transmitting a pulse and q_{rx} is the cost when receiving a pulse. Since the same transceiver elements are used for signal acquisition and reception, the acquisition energy consumption is also equal to q_{rx} . The cost in the active-off state is q_{ao} , and the cost while sleeping is negligible compared to transmission or reception. In Table 3.2, we consider four scenarios to investigate the impact of reception being more costly than transmission as well as active-off being less costly than transmission and reception. We assume it takes $50\mu s$ for signal acquisition when coarse

	q_{tx}	q_{rx}	q_{ao}
Model 1	1	1	1
Model 2	1	5	1
Model 3	1	1	0.5
Model 4	1	5	0.5

TABLE III
VALUES OF POWER QUANTA FOR THE DIFFERENT MODELS USED IN THE EVALUATION. VALUES ARE BASED ON BEST GUESSES OBTAINED BY INTERVIEW OF IMPLEMENTERS.

and fine synchronization are required. When only fine synchronization is necessary, we assume $10\mu s$.

For example, the energy consumption to receive a packet of 107 bytes (plus a preamble of 20 bytes) is (we use the parameters in Table II): $20 \cdot 8 \cdot PRP q_{rx} + 107 \cdot 8 q_{rx} + 107 \cdot 8 \cdot (PRP - 1) q_{ao} = 20 \cdot 8 \cdot 1000 q_{rx} + 107 \cdot 8 q_{rx} + 107 \cdot 8 \cdot 999 q_{ao}$.

C. Design choices

In the following, we analyze a number of claims on design tradeoffs, either by review of the literature, or by ad-hoc analysis.

1) *Power adaptation is not needed*: Different power adaptation strategies for low-power UWB network are discussed in [11]. One of the analyzed power adaptation strategies is 0/PMAX: whenever a node transmits data, it transmits with the maximum allowed transmission power. It is shown that any feasible rate allocation and energy consumption can be achieved with this simple power control strategy, hence power adaptation is not needed. An intuitive explanation lies in the fact that, since the signal to interference plus noise ratio (SINR) is non-linear in interference, when increasing the transmit power of a node, this has more effect on increasing the received signal than on increasing interference on other nodes. Therefore, it is more beneficial for a node to transmit with maximum power and a high rate to finish a transmission quickly and then let other nodes transmit, rather than use a low power, which prolongs the duration of the transmission. Transmission power is only a fraction of the total power consumed to send and receive a packet. In the above model, only the transmitted power is modeled. If we consider total consumed power, it becomes even more beneficiary to transmit with maximum power. When the power is high, the rate is high as well, and the transmission time is short. This way, we use the circuits for a shorter period of time, and we further decrease consumed power.

An important assumption in the above results is that the rate adaptation and mutual exclusion protocols are optimal. This is not always the case. When it is not true, the conclusions do not necessarily hold. However, the results suggest that power adaptation beyond 0/P_{MAX} is not needed.

2) *A suboptimal but simple form of multi-user detection can be beneficial:* Optimal multiuser detection is an efficient way to manage multiple access, however it remains currently impractical, especially for low-complexity devices. Nonetheless, there are clearly benefits in a sub-optimal solution such as interference mitigation as proposed in [7]. It consists in declaring as erasures the outputs of the RAKE receiver that are abnormally high. Its effect remains important even in a non line-of-sight, multipath environment. It greatly alleviates the effect of one or several near-far interferers on a receiver.

3) *Mutual exclusion is not needed when interference mitigation is applied:* Even with very low power, in case of near-far scenarios (interferers close to a non intended destination), it may seem desirable to enforce some form of mutual exclusion. However, if interference mitigation is applied, we may hope that a large part of the interference is eliminated. In this section we investigate whether mutual exclusion is needed when interference mitigation is present.

We first consider the rate metric. We assume each active receiver has a mutual exclusion region of radius r around it. During reception, no node inside the exclusion region is allowed to transmit. We vary the value of r and for each value we find the mutual exclusion strategy that satisfies the exclusion region constraints, and that maximizes the metric. We plot the average rates achieved in the optimal cases for different r in Figure 1, on the left. As one can see from the figure, it is optimal to let all nodes transmit concurrently at all times. The main reason why we do not need mutual exclusion is interference mitigation. Without interference mitigation, the optimal exclusion region size is approximately 2 meters. However, with interference mitigation we successfully eliminate interference from nearby nodes, and there is no more need for exclusion. At the same time, we benefit by allowing concurrent transmissions and increasing spatial reuse.




Fig. 1. 20 nodes (10 links) are randomly placed on 20m x 20m square. Each link has to achieve a minimum rate of 100 kbps. x -axis: size r of the mutual exclusion region. Rate constraints are low enough such that they can be satisfied even with large exclusion regions. Left: average network rates, as a fraction of rates achieved when there are no exclusions. Right: average network lifetime for different r , again relative to the average network lifetime when no exclusions are enforced ($r = 0$). The results are shown for load model 1; similar results are for load models 2 and 3.

The story is different for the lifetime metric. When rate constraints are low, each node transmits only during a small fraction of time. It is then obvious that the energy is minimized if there is no interference. We evaluate the optimal r using numerical simulations. The results are depicted in Figure 1, on the right. We see that with large r we can only slightly increase the lifetime of the node. The interference mitigation again handles most of the interference, and there is no need to implement an exclusion protocol.

4) *Rate control is needed:* If the rate (thus modulation and coding) is fixed to some predefined value, this value has to be small enough to be feasible on long links. This in turn imposes the same small rates on short links, which is highly inefficient from a rate or lifetime viewpoint. If transmission rates are low, packet transmissions last longer, and more energy is consumed to keep circuits running. We can thus conclude that rate control is needed.

5) *What is the optimal rate adaptation method ?:* We discuss two different rate adaptation techniques. One is proactive: a transmitter sends a pilot signal before every transmission to estimate a channel. The other one is reactive: a transmitter estimates a channel using feedbacks from previous transmissions.

A major benefit from a predictive approach is a more accurate channel estimate. Its channel estimate is more recent, as it comes from the pilot. In the reactive approach, it comes from the previous packet, which was not necessarily very recent. However, the pilot estimate is not always helpful. If an interfering transmission starts after the pilot, but during the packet

transmission, it will not be detected. The rate will be adapted to the state of the channel without interference, and a packet loss is inevitable. Similar effects may happen due to fast channel fading.

The predictive approach is needed only if there is a long pause between consecutive packets. In this case, the channel might have changed completely (e.g. due to mobility), and it is beneficial to estimate the channel before transmitting. Also, the predictive approach is useful when it comes for free, as in the case of protocols with RTS/CTS-like handshake. If the time between consecutive transmissions is short, the reactive approach can provide an accurate channel estimate. Using incremental redundancy further decreases the performance penalty of a wrong channel estimate [7].

6) *Slotted sleeping is better than unslotted if occasional bursts need to be supported.*: As described in Section 2.1, we consider two sleeping protocols: slotted and unslotted. They are depicted in Figure 2. In the slotted case, a reservation window is at the beginning of a superframe and each node knows its starting time. In the unslotted case, if a transmitter wants to transmit to a destination, it first needs to learn its listening schedule.




Fig. 2. Sleeping protocols: the *Slotted sleeping protocol* is depicted on the top. We consider a transmitter and a receiver (and assume this example, that a transmitter is also a coordinator). The protocol defines a superframe. A superframe begins with a beacon sent by a coordinator. It is followed by a reservation window. If a transmitter wants to transmit, it sends a RTS in a reservation slot on the TH code of the receiver (hence concurrent reservations for different receivers are possible). The receiver replies with a CTS if it accepts the reservation. If a reservation is successful, the actual data transmission occurs in the corresponding data slot, and is followed by an ACK. In the beginning of a superframe, nodes need a long beacon to achieve a coarse level synchronization. Afterward, there is only a short preamble before every packet. In the *unslotted sleeping protocol* (bottom), every node has a regular interval of duration T , called listening schedule. Reservations are done during the listen window at the beginning of the interval, and if successful, are followed by packet transmission. Since a pause between two reservation periods can be long, we need a long preamble at the beginning. One packet at most can be received during T .

The main objective of sleeping protocols is energy minimization, hence lifetime-maximization. In this section we analyze which protocol is more efficient with respect to the lifetime metric, subject to being able to sustain an occasional burst of data. More precisely, we assume the network is designed to occasionally sustain a traffic load λ_{max} per receiver during burst intervals (load model 2 in Table II), and compute its lifetime assuming that most of the time it is subject to load model 1 (with average traffic intensity $\lambda_0 < \lambda_{max}$). In the slotted case, a receiver can serve $\gamma S_A / T_{sf}$, and in the unslotted case γ / T packets per second, where γ is the network utilization and S_A , T_{sf} and T are defined on Figure 2. A network with utilization close to 100% is unstable. In order to guarantee network stability, we take $\gamma = 0.7$. Note that if two requests overlap, one of them is very likely to be accepted due to the nature of time hopping sequences and the signal acquisition procedure. Therefore, we can assume that the total submitted traffic is close to λ_{max} per receiver.

We compare the lifetimes achieved with slotted and unslotted protocols in Figure 3 (left). Networks with slotted sleeping protocols have 15%-50% longer lifetimes. If a network lifetime is around one year, it can be increased for 2-6 months, which is significant. We can conclude that if a coarse-level synchronization comes at a low cost, or for free (as in a master-slave system like bluetooth), it is optimal to use it. If this is not the case, in order to compare the two protocols, we need to compare their implementation overheads. The main overhead of a slotted protocol is in electing a coordinator [2], and managing the cases when communicating nodes hear several different superframes. The main overhead of an unslotted protocol is the learning time when a node needs to learn schedules of neighbors, either due to a topology change or due to a clock drift.




Fig. 3. Lifetime comparison for slotted and unslotted sleeping protocols. On the left, we consider the scenario when a network has to sustain given maximal rate. On the x axis we give different maximal rate λ_{max} per user. On the y axis we see the average lifetime for the slotted protocol divided by the average time for the unslotted one. In both cases we assume each receiver has a long term average ingress traffic $\lambda_0 = 10\text{kbps}$. We compare the two approaches for the different energy models, described in Table 3.2. We also take two extreme values of S_A . In all cases, the slotted protocol outperforms the unslotted one by 15%-30%. On the right we consider networks design for maximum access delay. On the x axis we plot maximum access delay. Other parameters are the same as in the previous case. Here we see that the unslotted protocol outperforms the slotted one.

7) *Unslotted sleeping is better than slotted if occasional delay sensitive loads need to be supported* : We now consider a variant of the previous paragraph. We continue to assume that, most of the time, the network is subject to load model 1 (with average traffic intensity λ_0), but now we assume, that, occasionally, it has to support a small number of unpredicted, but very urgent messages (load model 3). When a node generates a packet, it cannot send it immediately, but it first needs to wait for a reservation period to obtain access to a destination. In the case of slotted protocol, in the worst case, a node has to wait T_{sf} to send a packet. In the unslotted case, the worst case delay is T . We assume that the worst case is limited by application constraints to T_{ad} . We then want to compare energy savings for the two approaches assuming the same delay constraint.

In Figure 3 (right) we can see network lifetimes for different delay constraints and different energy models. We see that the conclusions are now reversed. The unslotted protocol always performs better or equal, since the unslotted protocol has only one reservation slot per time T_{ad} , while the slotted one has S_A reservation slots and every node has to listen for an RTS during these S_A slots.

8) *Is it optimal to have long or short frames?*: Due to hardware constraints, it is not possible to increase a pulse energy when a frame size is larger, even though it would be allowed by regulations. On the contrary, if a frame is short, a node is

able to transmit more pulses per time unit. Potential issues with the short frame size are higher interference and a possible inter-symbol interference.

We consider a short frame size $T_f = 100ns$ and a long one $T_f = 1000ns$, and we verify by simulations that for the first one we obtain roughly 8 times higher rates than for second. This in turn decreases the energy spending since circuits have to be activated for shorter periods of times.

REFERENCES

- [1] E. Callaway, P. Gorday, L. Hester, J. Gutierrez, M. Naeve, B. Heile, and V. Bahl. Home networking with IEEE 802.15.4: a developing standard for low-rate wireless personal area networks. *IEEE Communications Magazine*, 40(8):70–77, August 2002.
- [2] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: an energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. *ACM Wireless Networks Journal*, 8(5), September 2002.
- [3] F. Cuomo, C. Martello, A. Baiocchi, and C. Fabrizio. Radio resource sharing for ad hoc networking with UWB. *IEEE Journal on Selected Areas in Communications*, 20(9):1722–1732, December 2002.
- [4] M.-G. Di Benedetto, L. Nardis, M. Junk, and G. Giancola. (UWB)²: Uncoordinated, wireless, baseborn, medium access control for UWB communication networks. *To appear, Mobile Networks and Applications special issue on WLAN Optimization at the MAC and Network Levels*, 2005.
- [5] A. Doufexi, S. Armour, M. Butler, A. Nix, D. Bull, J. McGeehan, and P. Karlsson. A comparison of the HIPERLAN/2 and IEEE 802.11a wireless LAN standards. *IEEE Communications Magazine*, 40(5):172–180, May 2002.
- [6] E. Ferro and F. Potorti. Bluetooth and wi-fi wireless protocols: a survey and a comparison. *IEEE Wireless Communications*, 12(1):12–26, February 2005.
- [7] R. Merz, J. Widmer, J.-Y. Le Boudec, and B. Radunovic. A joint PHY/MAC architecture for low-radiated power TH-UWB wireless ad-hoc networks. *To appear, Wireless Communications and Mobile Computing Journal, Special Issue on Ultrawideband (UWB) Communications*, 2005.
- [8] A. Muqattash and K. Marwan. CDMA-based MAC protocol for wireless ad hoc networks. In *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC'03)*, pages 153–164, June 2003.
- [9] M. Neely, E. Modiano, and C. Rohrs. Dynamic power allocation and routing for time varying wireless networks. In *Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies, Proceedings*, volume 1, pages 745–755. INFOCOM 2003, 2003.
- [10] B. Radunovic and J. Y. Le Boudec. Optimal power control, scheduling and routing in UWB networks. *IEEE Journal on Selected Areas in Communications*, 22(7):1252–1270, September 2004.
- [11] B. Radunovic and J.-Y. Le Boudec. Power control is not required for wireless networks in the linear regime. In *WoWMoM*, June 2005.
- [12] J. So and N. H. Vaidya. Multi-channel mac for ad hoc networks: handling multi-channel hidden terminals using a single transceiver. In *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, pages 222–233, 2004.
- [13] L. Tong, V. Naware, and P. Venkatasubramaniam. Signal processing in random access. *IEEE Signal Processing Magazine*, 21(5):29–39, September 2004.
- [14] M. Z. Win and R. A. Scholtz. Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications. *IEEE Transactions on Communications*, 48(4):679–691, April 2000.