This article was published in IEEE Communications Magazine, (42)-12, 2005. `http://infoscience.epfl.ch`

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

Alaeddine El Fawal, Jean-Yves Le Boudec, Ruben Merz, Božidar Radunović, Jörg Widmer, and Gian Mario Maggio

Abstract

We are interested in the design of physical-layer aware medium access control (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 achieved data rates. So far, a number of different solutions have been proposed in the context of data rate efficiency for IR-UWB. However, 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 very low power networks, which is the aim of this paper. To this end, we first identify *what* a PHY-aware MAC design has to achieve : (1) interference management, (2) access to a destination and (3) sleep cycle management. Second, we analyze *how* these functions can be implemented, and provide a list of the many possible building blocks that have been proposed in the literature. Third, we use this classification to analyze fundamental design choices. We propose a method for evaluating energy consumption already in the design phase of IR-UWB systems. Last, we apply this methodology and derive a set of guidelines; they can be used by system architects to orientate fundamental choices early in the design process.

I. INTRODUCTION

Emerging pervasive networks assume the deployment of large numbers of wireless nodes, embedded in everyday life objects. In these types of networks, the focus is more on minimizing energy consumption than maximizing

---

The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322, and by CTI contract No7109.2;1 ESPP-ES. A. El Fawal, J.-Y. Le Boudec, R. Merz, and B. Radunović are with EPFL, School of Computer and Communication Sciences, CH-1015 Lausanne, Switzerland. J. Widmer is with DoCoMo Euro-Labs, Landsberger Strasse 312, 80687 Munich, Germany. G. M. Maggio is with STMicroelectronics, Advanced System Technology, 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`

rate. There exist numerous possibilities to implement cross-layer design between the physical layer (PHY) and the medium access control (MAC) layer 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 MAC layer and the physical layer. In this case, the primary goal of the MAC layer is to coordinate access to the physical layer by enforcing mutual exclusion between concurrent transmitters. 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 the rate and power level 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. For example, interference does not need to be completely prevented, but it needs to be managed (see Section II-A). The rate or the power can be dynamically adapted to the level of interference. Examples of such schemes for UWB can be found in [1], [2]. In [2], rather than preventing interference, sources adapt their rate such that their destination can sustain the interference.

An important design decision for a PHY-aware MAC is whether to allow interference (i.e. , permit concurrent, interfering transmissions) or to enforce mutual exclusion. Other important design decisions are: allowing random access or imposing some form of temporal super-frame structure within which transmissions have to occur, deciding whether to use power control and how to coordinate nodes such that many of them can sleep. These choices have implications on the physical layer, as well as the MAC layer. As demonstrated in [1] and [2], a PHY-aware MAC protocol can significantly improve performances.

We concentrate on large self-organized networks; we do not address the case of Wireless Personal Area Networks (WPAN). We focus on IR-UWB physical layer systems for low data-rate (LDR) applications. These systems make use of ultra-short duration ($<$ 1ns) pulses that yield ultra-wide bandwidth signals. They are characterized by low duty cycle ($\simeq$ 1%) and extremely low power spectral densities [3]. Multi-user access is possible thanks to pseudo-random time hopping sequences (THS) that randomize the transmit time of each pulse. The multi-user interference (MUI) for such systems, unlike narrow-band systems, is caused by “collisions” of pulses of different simultaneous transmissions [3]. IR-UWB systems are especially attractive for LDR wireless communications as they potentially combine low power consumption, immunity to multipath fading and location/ranging capability [1]. A complete design targeting energy efficiency should also consider energy efficient routing. However, for brevity we do not consider routing in this paper. We also do not consider ranging, signal acquisition and channel estimation. These functions are out of scope of this paper. The reader can refer to [4] and [5] for detailed discussions.

In Section II, we explore the design space of PHY-aware MAC protocols; we discuss *what* functions a PHY-aware MAC design must provide (Section II-A), *how* to implement them (Section II-B) and how they are implemented in

---

[1] In fact, the Task Group 4a was formed in March 2004 to investigate a UWB alternative PHY to the IEEE 802.15.4 wireless standard, associated with Zigbee.

existing UWB designs (Section II-C). Note that even tough we focus on UWB designs in Section II-C, the rest of Section II is not specific to UWB. In Section III we analyze the performance implications of fundamental design choices. We propose a method for evaluating energy consumption in the design phase of IR-UWB systems (Sections III-A and III-B) and derive a set of guidelines that can be used by system architects to orientate fundamental choices early in the design process (Section III-C).

II. THE DESIGN SPACE OF PHY-AWARE MAC PROTOCOLS

A. *What functions should a PHY-aware MAC provide ?*

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 that is to maximize the rate offered to each node while possibly remaining fair. Hence, 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.
- *Sleeping Management*: It is of crucial importance in a low-power context. There exists an important tradeoff between long sleep cycles, that permits for efficient energy savings, and short cycles that facilitate communication and improve responsiveness.

B. *How can the functions of a PHY-aware MAC be implemented ?*

In this section we review how the functions above can be implemented, according to published designs. We give a list of 9 building blocks, each of them contributing to one or several functions. The mapping between building blocks and functions is given in Table I.

1) Rate Adaptation: Often, the transmission 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 often 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.

2) Power Control: The transmit power can be adjusted to keep the signal-to-interference-and-noise ratio (SINR) 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 tolerate.

3) Mutual Exclusion: A mutual exclusion protocol prevents nodes from transmitting at the same time. Most traditional protocols use mutual exclusion to manage interference, but, as we see in Section III-C.4, mutual exclusion is not always necessary in our setting. It is often implemented by control packet signaling (for example with an RTS-CTS handshake as in 802.11). The number of nodes affected depends on the transmit power of the control packets.

4) Multi-Channel: In a multi-channel protocol, the transmission medium is separated into several orthogonal or quasi-orthogonal transmission channels. Since simultaneous transmissions can occur, there is a clear advantage in terms of rate increase. Still, a potential disadvantage (e.g., for broadcast) is that it becomes impossible to overhear transmission from other active nodes on other channels. Quasi-orthogonal channels are inherent with an IR-UWB physical layer thanks to time hopping [3]. 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 non-overlapping 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, the issue of the *near-far* effect appears. 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.

5) Multi-user Reception: With a single-user receiver, 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 [6]. For example, a near-far interferer would be jointly received instead of being treated as interference. This annihilates the near-far effect and makes multi-user reception potentially attractive. However, it generally necessitates the receiver to be accurate synchronization with all the sources that it wishes to decode and furthermore knowledge of all their transmitted signal characteristics. In addition, the complexity of the decoding operation is excessively high. Nevertheless, thanks to the particular structure of IR-UWB signal, there exists several suboptimal techniques that are still worth considering, such as interference mitigation described in Section II-C.

6) Random versus Scheduled Access: Random access schemes are straightforward to implement in their simplest form (Aloha). They are often improved with some of the following components:

- Carrier-sensing avoids sending on the channel if the channel is already busy. Since there is no carrier, carrier sensing is not well defined with IR-UWB physical layers. One possibility to emulate carrier sensing with

IR-UWB is 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 is used to resolve collisions.
- A hand-shake procedure where nodes exchange RTS/CTS packets before each transmission is used to reserve medium access for data transmission. Since these packets are much shorter than data packets, the performance penalty in case of a collision is low. 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.

An alternative is *scheduled access*. A schedule decides when and which nodes are allowed to send. It can allow only a single node to transmit (TDMA), or it can allow for multiple transmissions, if they do not interfere significantly. 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 nodes do not all “hear” each other.

7) Time-slots: Slotted transmissions can reduce interference (as in slotted Aloha) or improve power saving since a node can sleep during unused slots. It also facilitates timing acquisition; with slots, the nodes are coarsely time synchronized. It can greatly help in achieving the nano-second level synchronization necessary with IR-UWB for two hosts to communicate.

8) Sleeping: Slotted versus Unslotted: Letting nodes sleep is the most effective way to conserve energy in a wireless network and thus maximize the lifetime. However, this requires a mechanism that allows nodes to be contacted even though they might sleep from time to time.

We consider two types of sleeping protocols. The former one is time slotted and uses a periodic beacon. As in the previous section, this beacon provides a coarse-level synchronization and denotes the start of a superframe. A superframe has two parts: a *reservation window*, during which potential senders announce transmission requests, and a data transmission window, during which the actual packet transmissions take place. Receivers can then sleep for most of the second part, except for the periods when announced transmissions occur.

The latter approach is unslotted: each receiver wakes up according to its own *listening schedule*. A transmitter that wants to communicate with a given receiver first needs to learn its listening schedule. Typically, if all nodes have the same sleeping scheme (but delayed in time), a transmitter simply has to send a *long preamble*, as long as the maximum sleeping time. The destination, sure to wake up at some time in between, will receive the preamble and answer to the transmitter.

9) Centralized Architecture: A design choice for all the above possibilities is to have either a fully-decentralized or a master-slave architecture where the network consists of one or several subnetworks, each controlled by a coordinator. A coordinator can be a base-station or an arbitrary node elected by a network[2].

---

[2] Coordinator election is out of the scope of this paper

C. *Which Building Blocks Are Used by Existing Designs ?*

We now use the building blocks to analyze several proposed designs for UWB. For each of the three functions described in Section II-A, we analyze which building blocks are used, and give more details. We summarize the results of this section in Table I. Since many of the concepts of UWB designs are borrowed from narrow-band designs, we add the IEEE 802.11 protocol [7], Bluetooth [8], IEEE 802.15.4 (Zigbee) [9] and a CDMA design [10] to Table I for comparison purposes.

		802.11	CA-CDMA	Bluetooth	802.15.4	MBOA	Power + Rate controlled [1]	UWB^2	DCC-MAC
Rate adaptation		I					I		I
Power control			I				I		
Mutual exclusion		I,A	I,A	I,A	I,A	I,A	I,A		
Multi-channel			I	I			I	I	I
Multi-user detection									I
Access	Random	A	A		A	A	A	A	A
	Scheduled			A	A	A			
Time slots		S		A,S	A,S	A,S			
Sleeping	Slotted	S				S			
	Unslotted								
Centralized architecture				S	S				

TABLE I

Each row is a building block of a PHY-aware MAC described in Section II-B. The table shows which function a building block contributes to providing in existing designs proposed in the literature. I: interference management; A: access to a destination; S: sleeping management.

In [1], the authors present a joint power and rate controlled design for a UWB physical layer. Interference is managed by a mixture of mutual exclusion, power and rate adaptation, and by taking advantage of the quasi-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. Indeed, for every receiver there exists an interference margin, which indicates by how much the interference can increase without destroying ongoing transmissions. 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 UWB^2 [11], 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 [2] is a design based on theoretical results from [12]. It uses rate adaptation but no power control. It proposes to take advantage of the infrequent nature of collisions at the pulse level by using *interference mitigation*; it consists in declaring as erasures the outputs of the receiver that are abnormally high (i.e. when a pulse collision occurs with a near-far interferer). Indeed, a receiver can estimate the average received energy from its current source; a pulse collision with a strong interferer can then be easily identified since the received energy when it occurs is much higher. The loss of information due to the erasures is recovered by the error-correcting code. At the cost of a small rate reduction, it greatly alleviates the effect of one or several near-far interferers. Note that interference mitigation does not necessitate any synchronization between the transmitters. It also does not increase the power consumption.

With interference mitigation, mutual exclusion becomes unnecessary. 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 sequences.

Note that we do not discuss sleeping management for the three previous protocols since they do not address it.

The IEEE 802.15.4 (Zigbee) protocol [9], although it is based on a narrowband physical layer, is also of interest since the IEEE 802.15.4a Task Group is currently working on the standardization of an alternative UWB physical layer for the 802.15.4 one. It is a single-channel protocol based on CSMA-CA (with an optional RTS-CTS mechanism). Its main operating mode[3] is the so-called beacon-enabled mode where the network is organized as a slotted *piconet*. A piconet coordinator periodically broadcasts beacons. Access inside a given slot uses exclusion and is arbitrated through CSMA-CA.

Hence, interference is managed entirely by mutual exclusion, thanks to the slotting procedure. Access to a destination is ensured by the optional RTS-CTS procedure or relies on collision detection and a backoff mechanism. For sleeping management, the centralized and slotted structure permits nodes to easily sleep and wake up when necessary.

The MBOA protocol[4] is very similar to 802.15.4. It recently emerged from the inconclusive effort of the IEEE 802.15.3a Task Group working on an alternative UWB physical layer for IEEE 802.15.3. There is only a beacon based mode, with slotted access in between beacon transmissions.

III. PERFORMANCE ANALYSIS OF THE DIFFERENT DESIGN CHOICES

In this section we use the classification developed in the previous sections to evaluate several important design choices for low-power, low-rate IR-UWB networks. Our results are obtained either by review of the literature, or by

---

[3] There is also a distributed mode where communication occurs on a point-to-point basis using a CSMA-CA MAC

[4] The UWB physical layer adopted by MBOA is not based on impulse radio but on a multiband OFDM radio.

ad-hoc analysis and simulations. We derive six facts that can be used as guidelines. But first we define the energy consumption model and performance metrics used in the analysis.

A. *Energy Consumption Model*

Our goal is to define an energy model that can be applied early in the design process, *before* an actual hardware is developed and can be instrumented. This is a serious challenge, but we can take advantage of the nature of IR-UWB to derive a generic model, which is flexible enough to account for a large set of options.

With IR-UWB, time is divided into frames of N_c short duration chips[5]. We use this 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 between two 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, we model the energy consumption by considering the energy *per chip* for each state. An energy consumption model is defined by the vector

$$\vec{q} = [q_{tx}\ q_{rx}\ q_{ao}]$$

where q_{tx} is the cost for transmitting a pulse, q_{rx} receiving a pulse and q_{ao} for being in the active-off state. Since the same transceiver elements are used for signal acquisition and reception, the acquisition energy consumption is also equal to q_{rx}. The cost while sleeping is negligible. It is currently impossible to give precise figures for $\vec{q}$, but only relative values are relevant to our performance evaluation. It is thus possible to limit our analysis to a small set of scenarios, as shown on the top of Table II.

We now show on an example how our energy model is used. The energy consumption E_{packet} to receive a packet of 127 bytes (including a synchronization preamble of 20 bytes) using binary modulation (one pulse carries one bit) is

$$E_{packet} = 8 \cdot \left(\underbrace{20 \cdot N_c \cdot q_{rx}}_{\text{Energy for the preamble acquisition}} + \underbrace{107 \cdot q_{rx}}_{\text{Energy when a pulse is present}} + \underbrace{107 \cdot (N_c - 1) \cdot q_{ao}}_{\text{Energy in the active-off state}} \right)$$

where the factor eight appears since we consider bytes. With this model, the energy consumed for each received or transmitted packet can be easily computed. The lifetime of a node is then the time necessary to consume all the energy contained in the battery of the node.

B. *Performance Metrics and Simulation Parameters of the Performance Analysis*

We use two metrics to be consistent with the goal of maximizing network lifetime while keeping rates as high as reasonably possible. In addition, fairness issues are taken into account. The metrics are the sum of logs of node lifetimes and the sum of logs of average link rates. Indeed, log utility metrics are known to achieve a good tradeoff between efficiency and fairness [12].

---

[5] Only one pulse is transmitted per frame

The remaining assumptions about the physical layer parameters and the topology for the simulations are given in Table II. Note that the physical layer supports several transmission rates (from 100kbit/s to 1Mbit/s). For the simulations, all nodes have an identical physical layer and the same initial battery power.

Energy consumption models $\vec{q} = [q_{tx}\ q_{rx}\ q_{ao}]$	1	$\vec{q} = [1\ 1\ 1]$	Baseline model
	2	$\vec{q} = [1\ 5\ 1]$	Higher cost for reception
	3	$\vec{q} = [1\ 1\ 0.5]$	Lower cost for active-off
	4	$\vec{q} = [1\ 5\ 0.5]$	Higher cost for reception, lower cost for active-off
Physical layer parameters	Frame length $N_c = 1000$ chips		
	Chip duration $T_c = 1$ ns		
	Rates (using puncturing) $R = \left\{1, \frac{8}{9}, \frac{8}{10}, \dots, \frac{8}{32}, \frac{1}{5}, \frac{1}{6}, \dots, \frac{1}{10}\right\}$Mbit/s		
	Energy per pulse $E_p = 0.2818$ mW		
	802.15.4a channel model		
Sleeping protocols parameters[a]	$T_b = 50\mu$s, $T_{fa} = 10\mu$s		
	$T_{RTS} = T_{CTS} = T_{ACK} = 800\mu$s		Packet size is 20 bytes
	$T_{DATA} = 10200\mu$s		Packet size is 127 bytes
Topology for the simulations	Randomly distributed on a 20m x 20m square		
	Links chosen randomly		

TABLE II

ENERGY CONSUMPTION MODEL, TRAFFIC LOAD MODEL, PHYSICAL LAYER PARAMETERS AND ASSUMPTIONS FOR THE PERFORMANCE ANALYSIS

[a] Packet lengths are computed assuming the smallest data rate

C. Conclusion from the Performance Analysis: Guidelines About the Optimal Design

We conducted our performance analysis by analyzing existing literature and by performing extensive simulations when needed. More details and our simulation code can be found in [13]. This leads us to the following six facts about the optimal design for low-rate, low-power UWB networks.

1) 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 for all channel conditions. This in turn imposes the same small rates on good channel conditions. If transmission rates are low, packet transmissions last longer, and more energy is consumed to keep circuits running. This is highly inefficient from a rate or lifetime viewpoint [2]. Therefore, we can conclude that rate control is needed.

2) Power adaptation is not needed: Different power adaptation strategies for low-power UWB networks have been discussed in [12] and [14]. One of them is $0/P_{max}$; whenever a node transmits data, it is with the maximum allowed transmission power P_{max}. It is shown that any feasible rate allocation and energy consumption (hence lifetime) can be achieved with this simple power adaptation strategy. Hence power adaptation is not needed.

Intuitively, since the signal-to-interference-and-noise ratio with impulse radio is convex in interference, increasing the transmit power of a source has more effect on the received signal at the destination than on interference on other nodes. As such, it is beneficial for a node to transmit with maximum power. This ensures a high rate and data transmissions terminate quickly to let other nodes transmit. In contrast, using a lower transmit power prolongs the transmission duration which is detrimental to reducing power consumption. With a short transmission time, we use the circuits for a shorter period of time and thus increase the lifetime of a node.

3) A suboptimal and simple form of multi-user detection is beneficial: Optimal multiuser detection is an efficient way to manage multiple access. However it remains impractical in our settings; we consider low-complexity devices and it would require synchronization among transmitters. Nonetheless, there are clear benefits for IR-UWB when using *sub-optimal* solutions such as interference mitigation as demonstrated in [2] and [15]. At the cost of a small rate reduction, it greatly alleviates the effect of one or several near-far interferers.

4) Mutual exclusion is not needed when interference mitigation is applied: In case of near-far scenarios (even with a very low rate), it might seem desirable to enforce some form of mutual exclusion. However, if interference mitigation is applied, a large part of the interference is eliminated. We simulate the impact of mutual exclusion on rate and lifetime when interference mitigation is present.

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. For each value of r between 0 and 30 meters, we find all the subsets of nodes and rate of these nodes that maximize the rate metric and satisfy the exclusion region constraints. We use the baseline energy model (model 1). The results are similar with the other energy models.

The average rate achieved for different r is depicted on Figure 1. It can be observed that it is optimal to let all nodes transmit concurrently at all times (the maximum is reached for $r = 0$). Without interference mitigation, the optimal exclusion region size is approximately 2 meters. Thanks to interference mitigation, no mutual exclusion is required. The rate reduction due to interference mitigation is traded for an increased spatial reuse due to the absence of mutual exclusion

In the case of the lifetime metric, we evaluate the optimal r using numerical simulations. The results are depicted in Figure 1. With large r, the lifetime of the node is only slightly increased. When rate constraints are low, each node transmits only during a small fraction of time. This in turn reduces the energy consumed to keep the circuits running. Hence the total interference created is small and the energy consumed is minimized. Furthermore, interference mitigation handles most of the interference, and there is no need to implement an exclusion protocol.

5) Slotted sleeping is better than unslotted if occasional bursts must be supported: We consider a slotted and an unslotted sleeping protocol (Section II-B.8) as depicted in Figure 2. We analyze which protocol is more efficient with respect to average node lifetime.

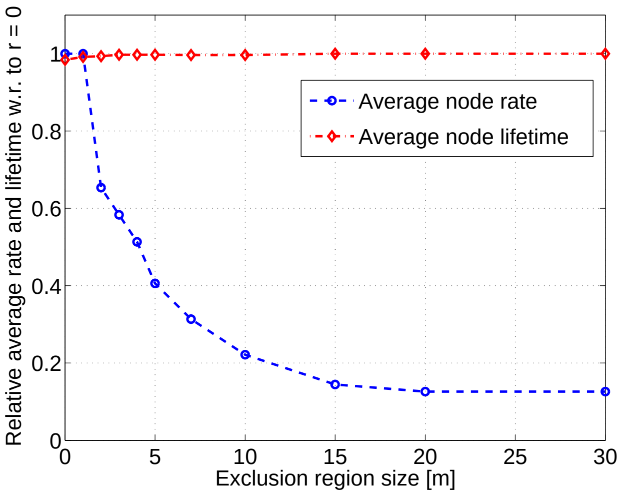

Fig. 1. Average node rate (dashed blue curved) and average node lifetime (red dot dashed curve) relative to the values at $r = 0$ versus the size of the exclusion region r. We use the baseline energy model (model 1), the results are similar with the other energy models. No exclusion region is required from a rate point of view. The presence of an exclusion region has negligible impact on the lifetime.

For the slotted protocol, transmission requests are carried out by sending an 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. For the unslotted protocol, reservations are done during the listening window at the beginning of the interval, and if successful, are followed by the packet transmission. Since a pause between two reservation periods can be long, we need a long preamble at the beginning.

We compute the lifetime assuming that most of the time the node is subject to a load λ_0. However, the network is designed to occasionally sustain a traffic load $\lambda_{max} > \lambda_0$ per receiver during burst intervals.

Let us define by γ the network utilization. In the slotted case, a receiver can receive $\gamma \frac{S_A}{T_{SF}}$ packets per second where S_A is the number of reservation slots in the reservation window and T_{SF} is the superframe length. In the unslotted case it can receive $\gamma \frac{1}{T_L}$ where T_L is the time interval between two listening windows. One packet at most can be received during T_L. Since a network with utilization close to 100% is unstable, we take $\gamma = 0.7$ to guarantee stability. Note that if two requests two the same destination overlap, one is very likely to be accepted due to time hopping and the signal acquisition procedure. Therefore, we assume that the total submitted traffic is close to λ_0 per receiver.

For two extreme values of S_A and the four energy models, we compare the lifetimes achieved with slotted and unslotted protocols. The parameters T_{SF} and T_L are chosen to sustain the bursty maximum load λ_{max}. The lifetime is then computed assuming a load $\lambda_0 = 10$kbit/s. The ratios of the lifetime in the slotted over the unslotted case are plotted on Figure 3(a). With slotted sleeping protocols, the lifetime is 15%-50% longer. If the lifetime is around one year, it can be significantly increased by 2 to 6 months. If the slotted structure comes at a low cost, or for free (as in a master-slave system like bluetooth), its use is optimal. If this is not the case, we need to compare the

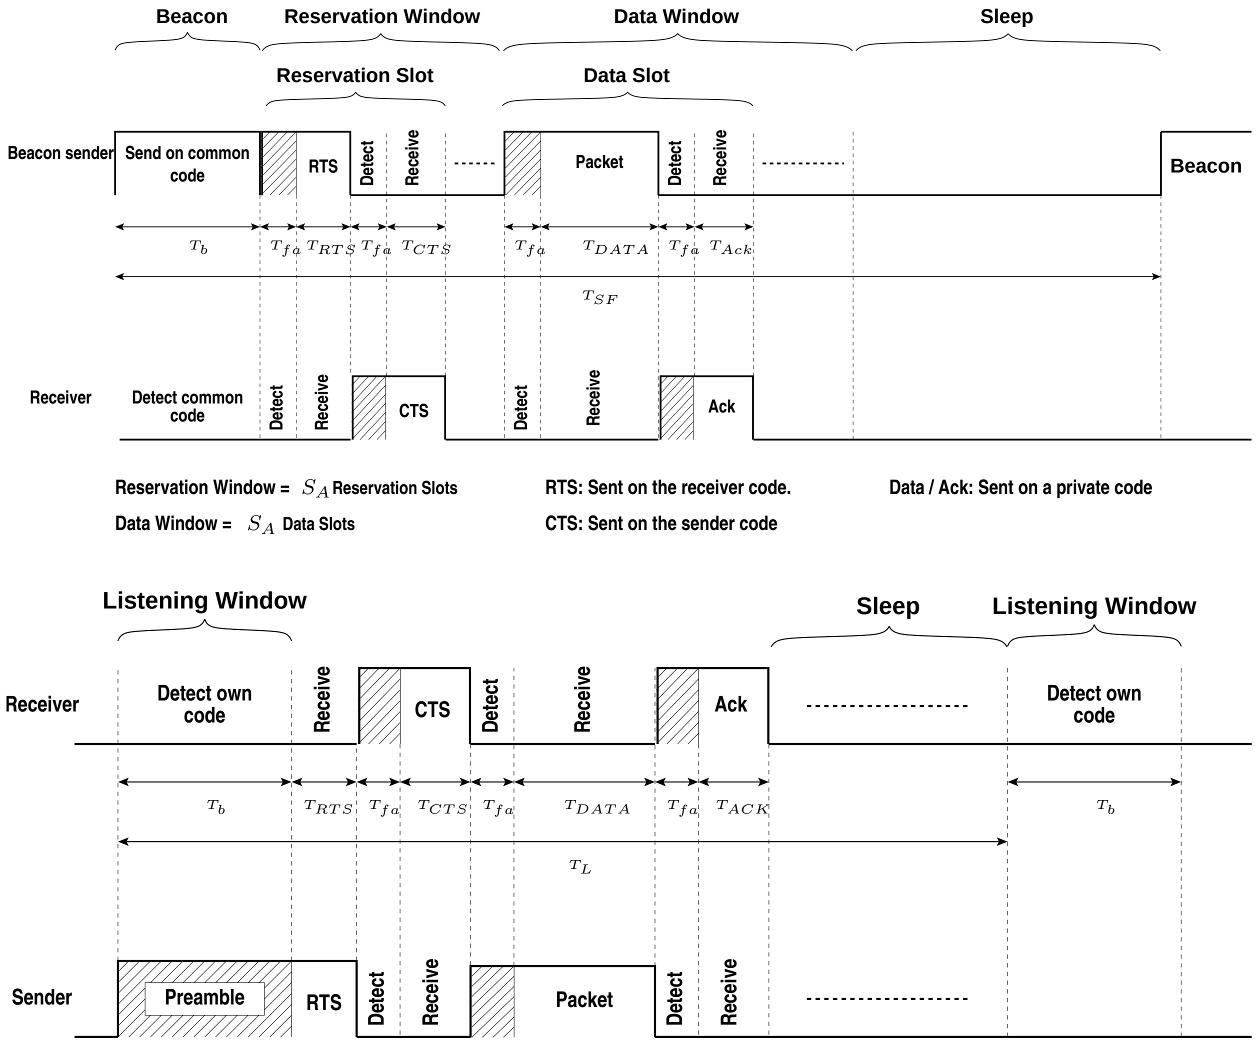

Fig. 2. The *slotted sleeping protocol* is depicted on the top; T_b is the length of the superframe beacon necessary to achieve coarse acquisition. Afterward, there is only a short preamble of length T_{fa} before every packet. The *unslotted sleeping protocol* is depicted on the bottom; T_L is the time interval between two listening windows.

implementation overheads to compare the two protocols. The main overhead of a slotted protocol is distributing the beacon 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.

6) Unslotted sleeping is better than slotted if occasional maximum latency must be supported: We consider a variant of the previous section. We still assume that most of the time, the network is subject to an average traffic load λ_0. However, it has to occasionally support a small number of unpredicted, but very urgent messages instead of a bursty high load.

When a node generates a packet, it cannot send it immediately. For the slotted protocol a node has to wait at

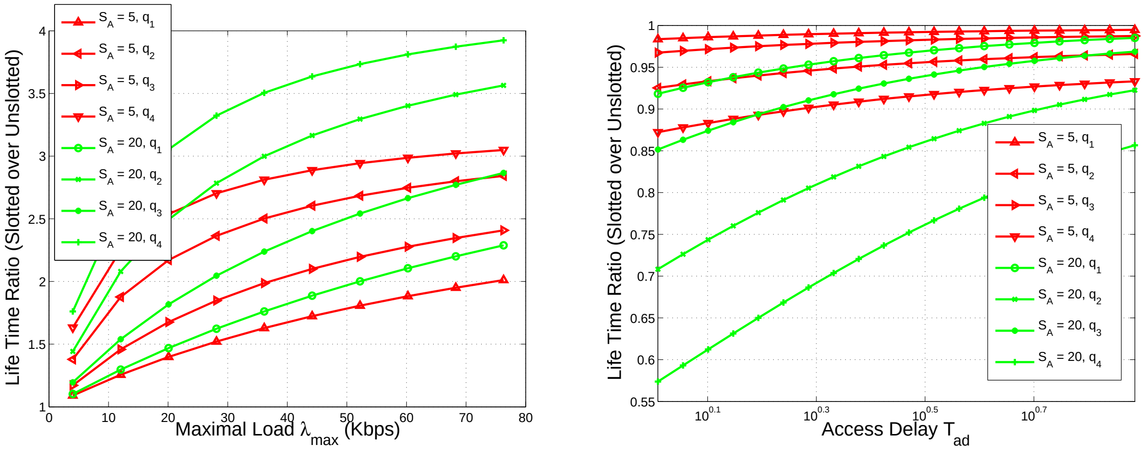

(a) Ratio of the average node lifetime in the slotted case over the unslotted case with respect to the maximal load λ_{max} (q_i stands for energy model i). In all cases, the slotted protocol outperforms the unslotted one by 15%-30%.

(b) Ratio of the average node lifetime in the slotted case over the unslotted case with respect to the access delay T_{ad} (q_i stands for energy model i). In this case, the unslotted protocol outperforms the slotted one.

Fig. 3. Lifetime comparison for slotted and unslotted sleeping protocols under various traffic constraints. We compare the performance for for S_A equal to 5 and 20 and all energy models (Table II), q_i stands for energy model i. In all cases $\lambda_0 = 10$Kbp/s.

most T_{SF} to send a packet. For the unslotted one, the worst case delay is T_L. In both cases, we assume that the worst case is limited by application constraints to T_{ad}. We then compare the energy savings for the two approaches as a function of T_{ad} for the different energy models.

The ratios of the lifetime in the slotted case over the unslotted case are plotted on Figure 3(b). The conclusions are the opposite of the previous section: the unslotted protocol always performs better or equal to the slotted protocol. Indeed, the unslotted protocol has only one listening window per time T_{ad}, whereas the slotted one has S_A reservation slots and every node has to listen for an RTS during these S_A slots.

IV. CONCLUSIONS

In this paper, we first explored the design space of PHY-aware MAC protocols. We described their functions and the various ways they can be implemented. This is directly useful for protocol designers to understand and exploit the large range of possibilities they have to design PHY-aware MAC protocols for UWB or other physical layers.

In the second part of this paper, our performance analysis lead us to formulate six guidelines for the design of low-rate, low-power IR-UWB networks. We also developed a new energy consumption model for impulse radio systems. The guidelines clearly call for an uncoordinated and decentralized protocol using rate adaptaiton and no power control.

REFERENCES

[1] Francesca Cuomo, Christina Martello, Andrea Baiocchi, and Capriotti Fabrizio. Radio resource sharing for ad hoc networking with UWB. *IEEE Journal on Selected Areas in Communications*, 20(9):1722–1732, December 2002.

[2] 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. *Wireless Communications and Mobile Computing Journal, Special Issue on Ultrawideband (UWB) Communications*, 5(5):567–580, August 2005.

[3] Moe Z. Win and Robert 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.

[4] S. Gezici, Zhi Tian, G.B. Giannakis, H. Kobayashi, A.F. Molisch, H.V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *IEEE Signal Processing Magazine*, 22(4):70–84, July 2005.

[5] Liuqing Yang and G.B. Giannakis. Ultra-wideband communications: an idea whose time has come. *IEEE Signal Processing Magazine*, 21(6):26–54, November 2004.

[6] Lang Tong, V. Naware, and P. Venkitasubramaniam. Signal processing in random access. *IEEE Signal Processing Magazine*, 21(5):29–39, September 2004.

[7] 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.

[8] 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.

[9] E. Callaway, P. Gorday, L. Hester, J.A. 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.

[10] Alaa Muqattash and Krunz 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.

[11] M.-G. Di Benedetto, L. Nardis, M. Junk, and G. Giancola. $(\mathrm{UWB})^2$: 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.

[12] 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.

[13] UWB research at EPFL-IC, simulation code. `http://icawww1.epfl.ch/uwb/code`, 2005.

[14] B. Radunovic. *A Cross-Layer Design of Wireless Ad-Hoc Networks*. PhD thesis, EPFL, Ph.D thesis 3301, June 2005.

[15] A. El Fawal and J.-Y. Le Boudec. A power independent detection method for UWB impulse radio networks. In *IEEE International Conference on Ultra-Wideband (ICU)*, 2005.