

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/312941966>

# A Joint PHY/MAC Architecture for Low-Radiated Power TH-UWB Wireless Ad-Hoc Networks

Article · January 2004

CITATIONS

53

READS

19

4 authors, including:



[Joerg Widmer](#)

IMDEA Networks

247 PUBLICATIONS 10,331 CITATIONS

[SEE PROFILE](#)



[Jean-Yves Le Boudec](#)

École Polytechnique Fédérale de Lausanne

519 PUBLICATIONS 21,701 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



AmpliCot [View project](#)



Network Calculus [View project](#)

# A Joint PHY/MAC Architecture for Low-Radiated Power TH-UWB Wireless Ad-Hoc Networks

Ruben Merz      Jörg Widmer      Jean-Yves Le Boudec      Božidar Radunović

EPFL Technical Report IC/2004/61

School of Computer and Communication Sciences, EPFL

CH-1015 Lausanne, Switzerland

`{ruben.merz, joerg.widmer, jean-yves.leboudec, bozidar.radunovic}@epfl.ch`

July 8, 2004

## Abstract

Due to environmental concerns and strict constraints on interference imposed on other networks, the radiated power of emerging pervasive wireless networks needs to be strictly limited, yet without sacrificing acceptable data rates. Pulsed Time-Hopping Ultra-Wide Band (TH-UWB) is a radio technology that has the potential to satisfy this requirement. Although TH-UWB is a multi-user radio technology, non-zero cross-correlation between time-hopping sequences, time-asynchronicity between sources and a multipath channel environment make it sensitive to strong interferers and near-far scenarios. While most protocols manage interference and multiple-access through power control or mutual exclusion (CSMA/CA or TDMA), we base our design on rate control, a relatively unexplored dimension for multiple-access and interference management. We further take advantage of the nature of pulsed TH-UWB to propose an interference mitigation scheme that reduces the impact of strong interferers. A source is always allowed to send and continuously adapts its channel code (hence its rate) to the interference experienced at the destination. In contrast to power control or exclusion, our MAC layer is local to sender and receiver and does not need coordination among neighbors not involved in the transmission. We show by simulation that we achieve a significant increase in network throughput.

**Keywords:** Medium access control, TH-UWB, Multi-hop wireless networks, System design, Simulation.

# 1 Introduction

Emerging pervasive networks assume the deployment of large numbers of wireless nodes, embedded in everyday life objects. For environmental and health concerns as well as coexistence with other wireless technologies, it is important that the level of radiated energy per node be kept very small.<sup>1</sup> At the same time, many applications require high data rates. Ultra-wide band (UWB) wireless networks have the potential to satisfy both requirements. UWB is characterized by an extremely broad use of the radio spectrum which makes it relatively robust against channel impairments such as multipath fading.

It was shown in [23] that the optimal wide-band signaling consists of sending infrequent short pulses. Our physical layer model is based on Win-Scholtz's proposal [25] using pulse position modulation (PPM). Time is divided into chips of very short duration. Chips are aggregated into frames and a sender transmits one pulse in one chip per frame. Multi-user access is provided by pseudo-random *Time Hopping Sequences* (THS) that determine in which chip each user should transmit. Due to the non-zero cross-correlation between time-hopping sequences, time-asynchronicity between sources, and a multipath channel environment, TH-UWB is sensitive to strong interferers.

Existing wireless MAC protocols manage interference and multiple-access in two ways. (1) *Mutual exclusion* schemes such as CSMA/CA, TDMA, or a combination of both [9] avoid interference by allowing only one transmission at a time within the same collision domain. (2) *Power control* allows to manage interference in a more sophisticated way. It is used for example for Code Division Multiple Access (CDMA) networks. While in synchronous settings (cellular base station), CDMA networks manage multi-user interference primarily by means of power control, asynchronous settings (ad-hoc networks) require the use of both power control and mutual exclusion [2, 17].

All such schemes have a high practical overhead. The use of RTS/CTS handshakes and the possibility of collisions drastically affects the performance in ad-hoc environments [5] and adjusting the transmit powers of all nodes within a collision domain requires a significant amount of coordination among nodes.

A largely unexploited dimension is to let the rate vary with the level of interference. A mathematical analysis of an optimal design for wide-band networks including exclusion, power control, and rate adaptation is given in [19, 24]. It is proved in [19] that the optimal design should not use power control but that sources should send at full power whenever they send. Furthermore, it is optimal in terms of throughput to allow interfering sources to transmit simultaneously, as long as they are outside a well-defined *exclusion region* around destinations, and to adapt the channel code (hence the rate) to this interference; in contrast, interference from inside the exclusion region should be combated. Similar conclusions are drawn in [24]. These results indicate that in our case the optimal design should (1)

---

<sup>1</sup>Note that we do not address the issue of maximizing battery lifetime, which is typical for sensor networks.

allow sources to send at maximum power (2) allow interference outside the exclusion region but forbid it inside the exclusion region (3) let sources adapt their rate to the interference experienced at their destination. We use these findings as foundations for our design.

With a TH-UWB physical layer, interference at a receiver is most harmful when pulses from a close-by interferer collide with those of the sender. Instead of enforcing exclusion within the exclusion region, we propose a different form of interference management called *interference mitigation*. It is based on detecting and canceling the impact of interfering pulses that have a significantly higher energy than the signal received from the sender. In contrast to exclusion-based mechanisms or power control, this interference mitigation scheme does not require any coordination between senders.

Our analysis in Section 5 suggests that the exclusion region is negligible when interference mitigation is used. This might seem obvious because we use a multi-user (in some sense “multi-channel”) physical layer, but it is not. Indeed, even with interference mitigation, in near far scenarios the activity of one user (on one “channel”) severely impacts the rate achievable by other users (on other “channels”), thus multiple access must be controlled. This is also witnessed by the fact that all existing proposals for TH-UWB do incorporate a MAC (Section 2). Our main finding in Sections 4 and 5 is that the MAC should primarily manage access by adapting rate to interference, without attempting to exclude competing sources by a mutual exclusion protocol. In this, our protocol radically differs from existing ones. There still remains some exclusion to implement because we assume that a node can be engaged exclusively either in one reception or in one sending. This is enforced by the “Private MAC” (Section 7).

Our main contribution is a system for TH-UWB with the following three components. (1) Interference mitigation is described in Section 5. (2) Dynamic channel coding continuously adapts the rate to variable channel conditions and interference (Section 6). To avoid the problem of signal to interference and noise ratio (SINR) measurements, the optimum code is determined after packet reception and piggybacked in the acknowledgment to the sender. (3) Private MAC resolves contention for the same destination. The challenge of absence of carrier sensing is solved by a careful balance of invitation and signaling. We do not use any separate channel for control.

Our design is fully implemented in ns-2. Simulation results show a significant increase in throughput compared to traditional protocol design. Also, because the source constantly adapts to the varying channel conditions, mobility is well supported.

## 2 Related Work

We have already mentioned in the introduction the state of the art [19, 24] that suggests that channel code control is preferable to power control.

Existing distributed MAC proposals for UWB ad-hoc networks can be found in [10, 2, 14]. They are all based on a TH-UWB physical layer of [25] and fixed channel code. [2, 14] are essentially based on a combination of power control and mutual exclusion. In [2] a distributed control admission function is based on the evaluation of the interference generated by each potential new link over active links. A source broadcast an RTS-like control packet before sending data. Every neighbor that receives this control packet responds to the source, adding information that allows the source to evaluate if the data transmission is admissible or not. However, the problem of contention for a destination is not mentioned. The approach is similar to [14]. An invitation based scheme to address the problem of contention for a destination is proposed in [10]. A node that is ready to receive broadcasts (on a common THS) an invitation for other nodes to compete for access to it. However, no power or channel code control is performed. Another power control protocol, although based on CDMA, can be found in [17]. A new transmission can proceed if it does not destroy any ongoing transmission in its vicinity. Information about neighbors is obtained by exchanging control packets on a separate (ideal) control channel.

Finally the IEEE 802.15 Task Group 3a has recently reviewed proposals for an alternate UWB physical layer for the IEEE 802.15.3 MAC [1]. The MAC is not distributed but based on the concept of piconets, where a piconet coordinator grants access to members of the piconet on a TDMA basis.

The concept of rate adaptation has been proposed for 802.11 networks [11, 20]. However they differ from our proposal in that rate adaptation is performed only to track the state of the channel (basically, the distance to the access point). Interference from other users is managed by an exclusion protocol or treated as collisions. In contrast, we use rate adaptation as a mechanism to support multiple-access.

### 3 System Assumptions

#### 3.1 Mathematical Model of the Physical Layer




Figure 1: UWB physical layer model. In this example,  $PRP = 8$ , the THS is  $\{\dots, 6, 2, \dots\}$ ,  $x_i = 0$  and  $x_{i+1} = 1$ . The multipath is *not* represented.

It was shown in [23] that the optimal wide-band signaling in the low-power regime consists of sending short infrequent pulses. Consequently, our physical layer is based on the widely used proposal of [25]. It is a multiple-access physical layer. Time is slotted in chips of duration  $T_c$ , and chips are grouped in frames of duration  $T_f$ . An active source sends *one* pulse per frame. Which chip to use in each frame is

specified by the so-called Time-Hopping Sequence (THS). A detailed explanation of the THS properties is given in Section 3.2. A source  $k$  uses binary pulse position modulation (2-PPM) to produce the signal

$$s^{(k)}(t) = \sqrt{E_p} \sum_{j=-\infty}^{\infty} p(t - c_j^{(k)} T_c - j T_f - x_j^{(k)} \delta) \quad (1)$$

where  $E_p$  is the pulse energy and  $p(t)$  is a unit energy pulse.  $T_c$  is the chip duration and  $T_f$  is the frame duration. Note that  $T_f \geq (PRP \cdot T_c) + T_g$  where  $PRP$  is the Pulse Repetition Period and  $T_g$  is the guard time duration. The sequence  $\{c_j^{(k)}\}_{j=-\infty}^{\infty}$  is the THS of the  $k$ th source. It specifies in each frame which chip should be used by  $k$ . If the bit  $x_j^{(k)} = 0$ ,  $p(t)$  is sent at the beginning of the chip. If  $x_j^{(k)} = 1$ ,  $p(t)$  is offset by  $\delta$ . The shape of  $p(t)$  is given by the second derivative of a Gaussian pulse [25]

$$p(t) = \left(1 - 4\pi \left(\frac{t}{\tau_p}\right)^2\right) \exp\left(-2\pi \left(\frac{t}{\tau_p}\right)^2\right) \quad (2)$$

where  $\tau_p$  represents a time normalization factor. The autocorrelation  $\Theta(x)$  of  $p(t)$  in (2) is

$$\Theta(x) = \left[1 - 4\pi \left(\frac{x}{\tau_p}\right)^2 + \frac{4\pi^2}{3} \left(\frac{x}{\tau_p}\right)^4\right] \exp\left[-\pi \left(\frac{x}{\tau_p}\right)^2\right] \quad (3)$$

Prior to the *modulator*, the data bits  $\{d_j^{(k)}\}_{j=-\infty}^{\infty}$  are fed to a *channel encoder* of rate  $R(t) \leq 1$  to produce the sequence  $\{x_j^{(k)}\}_{j=-\infty}^{\infty}$ . Note that the rate of the channel code is time-varying. The *variable* rate of a source is  $\frac{R(t)}{T_f}$ . Since in practice, the rate of a channel code is a discrete function, we denote by  $R_0 = 1 > R_1 > R_2 > \dots > R_N$  the set of rates offered by our channel code.

At a destination, the received signal  $r(t)$  is the sum of the signals received from all sources convolved with their respective impulse channel response

$$r(t) = \sum_{k=1}^{N_u} \left(h^{(k)} * s^{(k)}\right)(t - \tau_k) + n(t) \quad (4)$$

where  $n(t)$  is zero-mean white Gaussian noise of variance  $\frac{N_0}{2}$  and  $h^{(k)}(t)$  is the impulse response of the channel between a source  $k$  and the receiver. Let  $s^{(1)}(t)$  be the signal of interest.

$$\sum_{k=2}^{N_u} \left(h^{(k)} * s^{(k)}\right)(t - \tau_k) \quad (5)$$

is the Multi-User Interference (MUI) created by the  $N_u - 1$  other sources in the network.  $\tau_k$ ,  $k = 1, 2, \dots, N_u$  model the time-shifts between links in the network. The impulse response  $h^{(k)}(t)$  models the multipath channel:

$$h^{(k)}(t) = \sum_{i=0}^L \alpha_i^{(k)} \delta(t - \nu_i) \quad (6)$$

where  $\alpha_i^{(k)}$  is the attenuation coefficient taking into account path loss and random fading on the  $i$ th path and  $L$  is the number of multipath components. The signal  $r(t)$  is passed through a *rake receiver* and sampled to produce the received symbols  $\{y_j^{(0)}\}_{j=-\infty}^{\infty}$ . They are in turn passed to the *channel decoder* who will attempt to recover the transmitted data. The details of a rake receiver and of the channel decoding process are out of the scope of this paper. More details can be found in [18]. Finally, although only one signal can be received and decoded at a time, we can listen to several signals using different THS [25].

### 3.2 Properties of Time Hopping Sequences

Formally, a THS is a random sequence  $\{c_j\}_{j=-\infty}^{\infty}$  of integers uniformly distributed in  $[0, PRP - 1]$  [25]. The integer  $c_j$  specifies the chip to be used for transmission in the  $j$ th frame. In practice, THSs are pseudo-random and periodic.<sup>2</sup> Their use is twofold. First, by introducing a random shift in the transmitted pulse, they avoid peaks in the energy spectrum. Second, they permit multiple users to share the same channel.

An important property of THSs is their average Hamming cross-correlation  $H(x_1, x_2)$  defined in [13]. The smaller the Hamming cross-correlation, the smaller the bit error rate (BER) in a multi-user environment.  $H(x_1, x_2)$  is a measure of the number of hits between two THS  $x_1$  and  $x_2$  averaged over all possible time shifts between  $x_1$  and  $x_2$ . A detailed explanation of Hamming cross-correlation can be found in [13].

Interestingly, if the time shift between two transmissions is large enough, it does not even matter whether the THSs are the same or not. Assume  $S_1$  transmits data to  $D$  using THS  $C_D = \{c_1, c_2, \dots, c_p\}$  of period  $p$ . If another source  $S_2$  uses the same THS  $C_D$  but with a time shift  $\tau$  larger than the multipath delay, the interference created at  $D$  by  $S_2$  will be no different from the interference that would be created with a different THS  $C'_D$  [25].

It is possible to generate THSs that have a better cross-correlation than uniform random THSs. However, in practice only little can be gained in terms of BER reduction due to the asynchronicity between different sources and the multipath environment [13].

Although THSs share conceptual similarities with CDMA spreading codes, they are *not* equivalent. Whereas in TH-UWB the transmission of pulses is infrequent (one per frame), it is continuous in CDMA. Two concurrent transmissions in TH-UWB only interfere when pulses overlap, whereas they always interfere in CDMA. Therefore, finding good time hopping sequences is almost trivial, while finding low cross-correlation spreading sequences for asynchronous CDMA systems is a difficult prob-

---

<sup>2</sup>The period  $p$  is orders of magnitude larger than the transmission time of a packet. For a given  $p$ , there are  $PRP^p$  possible THSs.

lem [18]. They cannot be computed on the fly and assignment protocols is necessary [22]. In our case, no assignment protocol is required. The details of how THSs are used in our protocol are discussed in Section 7.

### 3.3 Practical Aspects of the Physical Layer

Recall that we are interested in low-radiated power, of the order of  $1\mu\text{W}$ . We assume that the pulse shape  $p(t)$  has a width  $T_p = T_c - \delta$ , and *peak power*  $P_{\text{peak}} = E_p/T_p = 0.28\text{ mW}$  [9]. The *radiated power*  $P_{\text{rad}} = \frac{P_{\text{peak}}}{PRP \cdot T_c + T_g}$  is defined as the average power during transmission. Since the chip time  $T_c = 0.2\text{ ns}$  (roughly a 5 GHz bandwidth corresponding to values proposed for future UWB devices [9]) and  $T_g = 20T_c$ , we need  $PRP = 260$ . Therefore, the maximum rate is equal to  $\frac{1}{PRP \cdot T_c + T_g} = 18\text{ MB/s}$ .

The distribution of the channel coefficients  $\alpha_i^{(k)}$ ,  $i = 1, \dots, L$  is given by [6]. We consider  $L = 5$ . The simple repetition coding scheme of [25] is replaced by a more sophisticated variable rate channel coding scheme. We use so-called rate compatible punctured convolutional (RCPC) codes [8, 4], in particular the family in [4]. There are  $N = 31$  possible rate:

$$\{1, 8/9, 8/10, 8/11, \dots, 8/32, 1/5, 1/6, \dots, 1/10\}$$

More details, in particular on their rate compatibility feature in conjunction with incremental redundancy, are given later in Section 6.1. Nevertheless, only *one* decoder is necessary for all the possible rates. In addition, an interleaver is used [18].

### 3.4 Synchronization

Synchronization of the physical layer is required only between a source and a destination (for unicast), and is performed at the destination only. It relies on the presence of a synchronization preamble at the beginning of each packet. We assume that synchronization can be maintained over the whole duration of a packet, and can be re-established for each data packet [15]. There is no global synchronization.

## 4 Exploring the Design Space

Power control and exclusion are the most common schemes used to manage interference and multiple access. However, as mentioned before, a largely unexploited dimension is rate adaptation. In [19], a mathematical analysis of an optimal design including power control, exclusion, and rate adaptation is performed. Their findings are the following:



- The authors prove that instead of using power control, sources should send at full power whenever they send.
- Furthermore, it is optimal to allow interfering sources to transmit simultaneously, as long as they are outside a well-defined *exclusion region* [19] around destinations.
- In contrast, every source inside the exclusion region of a destination should be silenced.
- Finally, sources should continuously adapt their rate to the level of interference experienced at the destination.

We build our design around these findings. Clearly, the most important parameter is the size  $\gamma$  of the exclusion region. Hence, the rest of this section characterizes  $\gamma$  for the particular TH-UWB physical layer we use.

In networks with arbitrary topology, all destinations  $D_1, \dots, D_n$  have a different  $\gamma_k$ . In such scenarios, the computation of each  $\gamma_k$  per  $D_k$  is a hard problem [19]. Therefore, we have to resort to the symmetric topology of Figure 2 where  $\gamma = \gamma_1 = \dots \gamma_n$ . Links  $\{S_1, D_1\}, \dots, \{S_n, D_n\}$  are placed in an alternate way on a cylinder of length  $L$ . The distance between two links is  $d$ .

If  $\gamma > d$ , then the optimal is to have all nodes sending at the same time. The achievable rate on a link in this case is  $R_{all}(d)$ . If  $\gamma < d$ , then it is optimal to have only the nodes on one side of the cylinder sending at the same time (half of the sources). In this case, the achievable rate on a link is  $R_{excl}(d)$ . In the limit case, when  $\gamma = d$ , we have  $R_{all}(d) = R_{excl}(d)$  [19]. We use this property to determine the size of the exclusion region. We compute  $R_{all}(d)$  and  $R_{excl}(d)$  for various values of  $d$  and determine where  $R_{all}(d) = R_{excl}(d)$ .

To determine the best  $\mathcal{R}_{all}(d)$  (and  $\mathcal{R}_{excl}(d)$ ) that  $S_1$  can achieve in the presence of interferers  $S_j$ ,  $j = 2, \dots, n$ , we first fix a BER threshold  $\mu^*$ . A typical value for  $\mu^*$  in a wireless environment is  $10^{-5}$ . We then determine the maximum possible rate  $R_i$ ,  $i = 1, 2, \dots, N$  that drives the BER below  $\mu^*$ . While modeling the MUI in equation (5) as Gaussian allows to use an analytical expression [25] for  $\gamma$  in an uncoded TH-UWB physical layer, it is well known that this Gaussian assumption does not hold for TH-UWB [3]. We further use convolutional channel codes and a multipath channel and therefore have to turn to Monte Carlo simulations to derive  $\gamma$ .

The physical layer model is the one described in Section 3.1 with the parameters of Section 3.3. In particular, *no* Gaussian assumption on the MUI is done. We simulate the discrete-time equivalent channel of the continuous-time model of Section 3.1. Without loss of generality, we consider the received signal on the  $i$ th path of the rake receiver.

$$r_i(t) = \sum_{k=1}^{N_u} \alpha_i^{(k)} s^{(k)}(t - \tau_k - \nu_i) + n(t) \quad (7)$$




Figure 2: Multiple interferers scenario: nodes are symmetrically distributed on the edges of a cylinder. Corresponding peers are located on adjacent disks. There are  $n = 6$  links in total and every second link is inverted such that each destination is close to an interfering source. The distance between a source and a destination is the length of the cylinder  $L$ , and the distance between a destination and the adjacent interfering source is  $d$ .

The signal of interest is  $s^{(1)}(t)$  and  $x_j^{(1)} = 0$ . We will further assume  $c_j^{(1)} = 0$  for all  $j$  [3]. Assuming perfect synchronization with the reference signal, the decision statistic for the first path is

$$y_{j,i} = \int_{jT_f}^{(j+1)T_f} r_i(t) v(t - \tau_1 - \nu_i - jT_f) dt \quad (8)$$

where  $v(t)$  is the correlation template. With 2-PPM,  $v(t) = p(t) - p(t - \delta)$  [25]. The correlation of the template  $v(t)$  with a time-shifted pulse is defined as

$$\tilde{\Theta}(x) = \int_{-\infty}^{\infty} p(t - x) v(t) dt = \Theta(x) - \Theta(x - \delta) \quad (9)$$

Substituting (7) into (8) yields

$$y_{j,i} = S_i + I_i + n \quad (10)$$

where  $n \sim \mathcal{N}(0, N_0 \tilde{\Theta}(0))$ .  $S_i = \pm \alpha_i^{(k)} \tilde{\Theta}(0)$  depends on user 0's signal bit  $x_j^{(0)}$ .  $I_i$  is the MUI in the discrete-time equivalent channel. We model the difference of time shifts asynchronism as [25, equa. 55]

$$\tau_k - \tau_1 = j_k T_f + \alpha_k, \quad -\frac{T_f}{2} \leq \alpha_k \leq \frac{T_f}{2} \quad (11)$$

where  $j_k$  is the value of the time difference  $\tau_k - \tau_1$  rounded to the nearest frame time and  $\alpha_k$  is uniformly distributed on  $\left[-\frac{T_f}{2}, \frac{T_f}{2}\right)$ . Based on the assumption [25, equa. (57)], we then note that only one pulse from each interfering source in each frame contributes to the interference term. We can write  $I_i$  in the form [25, equa. (76)]

$$I_i = \sum_{k=2}^{N_u} \alpha_i^{(k)} \tilde{\Theta} \left( \alpha_k + c_j^{(k)} T_c + \nu_i + x_j^{(k)} \delta \right) \quad (12)$$

Finally, the  $L$  paths are gathered to yield the decision statistic  $y_j = \sum_{i=1}^L y_{j,i}$ . To recover the original data and compute the BER, the  $y_j$ 's are passed the decoder. We consider two types of channel decoding policy, hard-decision and soft-decision. With a hard-decision policy, only the sign of the demodulator




Figure 3: We show  $R_{all}(d)$  (all sources send at the same time) as well as  $R_{excl}(d)$  (one side of the cylinder sends at the same time) vs. interferer distance  $d$ . We determine  $R_{all}(d)$  for interference mitigation decoding, soft-decision decoding, and hard-decision decoding. The intersection between  $R_{all}(d)$  and  $R_{excl}(d)$  gives  $\gamma$ , the size of the exclusion region.

output is passed to the channel decoder, whereas the sign *and* the amplitude is passed in the case of soft-decision. Usually, the soft-decision policy performs better than a hard-decision policy [18]. However, the underlying assumption in this case is that the total interference (MUI and noise) has a Gaussian density. Since the MUI is not Gaussian for a TH-UWB physical layer, this justifies our interest in studying the performance of a hard-decision policy.

We use the discrete-time equivalent channel model to find  $\gamma$  for the topology given in Figure 2. We consider  $n = \{4, 8, 16\}$  links and use link lengths of  $L = \{2, 6, 12, 18\}$  to obtain varying signal and interference intensities. The results are shown in Figure 3. In the hard-decision case, there is no exclusion region (i.e.,  $R_{all}(d) > R_{excl}(d) \forall d$ ). However, the performance is very poor for large values of  $d$ , when interferers are distant and the dominant interference is Gaussian noise. In the soft-decision decision case, an exclusion region of 1 to 4 meters is present depending on  $L$  and  $n$ . Although the probability of collision  $P_{col}(n) = 1 - \left(1 - \frac{2T_c}{T_f}\right)^n$  ( $< 1\%$  for  $n = 1$ ,  $7.5\%$  for  $n = 10$ ) with interfering pulses is low, they have a large impact in the case of nearby interfering sources. When they occur, the amplitude of  $|y_j|$  is much larger than the regular Gaussian. With soft-decision decoding, a large amplitude sample  $|y_j|$  propagates over the decoding of several subsequent samples  $|y_{j+1}|, |y_{j+2}|, \dots$ . It significantly deteriorates the decoding process and causes several decoding errors. This does not happen with a hard-decision policy since only the sign of the output sample is used. However, soft-decision decoding clearly outperforms hard-decision decoding large values of  $d$ , as is to be expected [18].<sup>3</sup>

Intuitively, the optimal decoding policy, should be an adaptive combination of hard-decision when strong interferers are present and soft-decision otherwise. Motivated by [21], we build on these observations to propose a simple, yet efficient scheme to reduce the effect of strong interferers with a soft-decision policy. This reduces  $\gamma$  and still avoids the complexity of an exclusion scheme.

<sup>3</sup>The result is similar when using more powerful codes like turbo codes.

## 5 Interference Mitigation

**Mitigation of Interference by Erasures:** We propose an interference mitigation (IFM) scheme at the physical layer to reduce the effect of strong interferers. We cancel the samples  $y_j$  resulting from a collision with pulses of a strong interferer and replace them with erasures (i.e. we skip them in the decoding process). That is

$$y_j = \begin{cases} \varepsilon, & \text{if } |y_j| > \mathcal{B}(t) \\ y_j, & \text{otherwise} \end{cases} \quad (13)$$

where  $\varepsilon$  is an erasure and  $\mathcal{B}(t)$  is the *erasure threshold*. Since  $P_{col}(n)$  is low, only a small percentage of erasures is produced and the channel decoder can recover from them.

**Setting the Threshold:** We should ideally set  $\mathcal{B}(t)$  such that an erasure is only declared when there is a collisions, and not due to Gaussian noise. The optimal value of  $\mathcal{B}(t)$  depends both on the average received power from the source and on the Gaussian noise. Whereas a too large  $\mathcal{B}(t)$  is equivalent to the case without erasures, a small  $\mathcal{B}(t)$  will declare too many erasures. We set

$$\mathcal{B}(t) = E[|Y_s(t)|] + kE[|N(t)|] \quad (14)$$

$E[|Y_s|]$  and  $E[|N|]$  are the estimates of the mean absolute amplitude of the signal of interest and of the mean absolute noise amplitude respectively. We use  $k = 2.8$ . The optimal choice of  $\mathcal{B}(t)$  remains to be further analyzed.

It has been shown through indoor channel measurements that variations in the received signal power are typically caused by shadowing rather than fast fading [21]. Hence, a receiver can track the strength of the received signal during several chips, and estimate its average over time.

Using the same simulation model and parameters as in the previous section, the rates achieved with interference mitigation are depicted in Figure 3. With interference mitigation, we take full advantage of the soft-decision policy for large values of  $d$ . For low values of  $d$ , interference mitigation considerably reduces the effect of collisions with pulses from strong interferers. With up to 8 links, there is no exclusion region. For 15 links, a small exclusion region is present for link distances of 12 and 18 meters. However, the rate difference between the exclusion case and the case when all sources send together is small. All in all, we find that the size of the exclusion region size is negligible.

## 6 Rate Adaptation Protocol

### 6.1 Rate Compatibility and Incremental Redundancy

We use the notation  $\overline{\mathcal{D}}$  to represent a block of data and  $\overline{\mathcal{C}}_k$  to represent the block of coded bits produced by encoding  $\overline{\mathcal{D}}$  with rate  $R_k$ . The RCPC codes that we introduced in Section 3.3 provide a variable en-

coding rate by means of *puncturing*; whenever a block of data  $\overline{\mathcal{D}}$  needs to be encoded at rate  $R_i > R_N$ ,  $\overline{\mathcal{D}}$  is first encoded with the lowest code rate  $R_N$  to produce  $\overline{\mathcal{C}}_N$ . Then  $\overline{\mathcal{C}}_i$  is generated by removing (i.e. puncturing) appropriate elements from  $\overline{\mathcal{C}}_N$ . Another property of the RCPC codes is their *rate compatibility* feature. It means that  $\overline{\mathcal{C}}_i$  is a subset of  $\overline{\mathcal{C}}_{i+1}$ . The rate compatibility feature permits to use the RCPC code in an Automatic Retransmission reQuest (ARQ) scheme with Incremental Redundancy (IR); whenever the decoding of some  $\overline{\mathcal{C}}_i$  fails, a lower rate  $R_j < R_i$  is required (i.e., more information needs to be sent by the source). With IR, the source does not have to send the complete  $\overline{\mathcal{C}}_j$ , but only the additional symbols that permit to obtain  $\overline{\mathcal{C}}_j$  from  $\overline{\mathcal{C}}_i$ . All of the transmitted symbols are used for decoding.

## 6.2 Dynamic Channel Coding

The goal of dynamic channel coding is to constantly adapt to the highest rate code that still allows decoding of the data packet at the receiver. For this, we exploit the feature of our codes that a destination that can decode can also determine the highest possible code rate.<sup>4</sup> Channel code adaptation works as follows:

- A source  $S$  keeps in a variable `codeIndex( $D$ )` the code index to use for communication with  $D$ . Initially or after an idle period,  $S$  uses the lowest rate code with `codeIndex( $D$ )` =  $N$ .
- When  $D$  sees that a packet is sent but cannot decode it, it sends a NACK back to  $S$ .
- As long as  $S$  receives NACKs, further packets with punctured bits (each time up to the size of the original packet) are sent, until the transmission succeeds or no more punctured bits are available. In the latter case,  $S$  may attempt a retransmission at a later time.
- As soon as  $D$  can decode, it computes the smallest index  $j$  that could have been used as described below.  $D$  returns index  $j + 2$  in the ACK to  $S$ .
- When a source with `codeIndex( $D$ )` =  $i$  in the cache receives an ACK with index  $j + 2$ , if  $j + 2 < i$  then `codeIndex( $D$ )` =  $i - 1$ , else `codeIndex( $D$ )` =  $j + 2$ .
- If  $S$  receives neither an ACK nor a NACK, it is likely that  $D$  is not listening (see Section 7). In this case,  $S$  will abort the transmission (without sending incremental redundancy) but may retry at a later time.

Since it is hard to measure the SINR in UWB, we determine the optimum code after packet reception. It is piggybacked in the acknowledgment to the sender. For good performance and a short transmission delay, sending redundant information should rarely be necessary. It is more important that the transmission succeeds directly without having to send additional punctured bits than using the highest possible code rate.

---

<sup>4</sup>Note that in contrast to the data part of a packet, the MAC header is always encoded with at rate  $R_N$  so that a receiver can determine that it received a packet even if it is not able to decode the data.

Decoding of a data packet encoded with rate  $R_i$  is performed by step-wise traversal of the trellis of the Viterbi decoder [18]. At each step a trellis branch is chosen, where a branch corresponds to a specific decoded bit. The packet is then reproduced from the bits corresponding to the sequence of selected branches. Hence, as soon as the outcome of a decoding step for a higher rate code  $R_j > R_i$  differs from that of the actual channel code, code  $R_j$  can be eliminated. Because of the rate compatibility feature of RCPC codes, this allows to also eliminate all codes with  $R_k > R_j$ . The highest rate code that remains is still powerful enough to decode the packet.

Ideally, the more stable the channel conditions, the closer the code used for the next transmission should be to this highest rate code. In practice, we use a safety margin to reduce the probability of retransmission when channel condition deteriorate. We find that the heuristic of using a channel code rate  $R_{i+2}$  if the highest possible code rate is  $R_i$  performs sufficiently well. The code  $R_{i+2}$  is indicated to the sender in the ACK. The same calculations are performed for all subsequent data transmissions to maintain the same safety margin. If conditions improve and the safety margin is larger than 2, the code index is reduced and if the safety margin was violated the code index is increased accordingly.

## 7 Private MAC

With the proposed physical layer, many senders may communicate simultaneously within the same collision domain and a sender cannot know if the intended receiver is idle or busy other than by actively listening for packets to or from it. To design an efficient, low overhead MAC layer, a careful orchestration of the transmissions of the nodes is required. Our MAC layer is based on a small amount of signaling between communicating nodes and careful selection of timeout values and THSs to listen on.

We use receiver-based THSs which means that data packets are transmitted using the receiver's THS. A node listens to up to three THSs at the same time.<sup>5</sup> It always listens on the broadcast THS, which is the same for all nodes, and on its own THS. When sending data to another node, it further listens on the THS of the destination. We denote by  $THS(S)$  the THS of node  $S$  and by  $THS(B)$  the broadcast THS.

**Successful Transmission:** A successful data transmission consists the actual data packet, an ACK, and an idle signal. The code used for the data packet depends on previous channel conditions whereas ACK and idle are always coded with the lowest rate code.

Assume a node  $S1$  has data to transmit to a node  $D$ , and  $D$  is idle, as in the first transmission shown in Figure 4.  $S1$  will send the data packet using  $THS(D)$  and will also start listening on  $THS(D)$ . As

---

<sup>5</sup>Remember that a node can *listen* on more than one THS but can only *receive* from one node at a time. Furthermore, a node can either send or receive but not both.

soon as  $D$  can decode (see Section 6.1), it sends back an ACK on its own  $THS(D)$ . The ACK carries an idle flag. It is set if  $D$ 's interface queue is empty (i.e.,  $D$  was the final destination of the current packet and it neither has an own packet to transmit nor another packet to forward). While  $S1$  is waiting for the ACK from  $D$ , it disables listening on its own THS to avoid receiving a data packet and therefore missing the ACK. Upon reception of the ACK,  $S1$  transmits an idle signal on its own  $THS(S1)$ , ceases to listen on  $THS(D)$ , and starts listening on its own THS again.

A node may do a backoff between 0 and the maximum backoff time  $t_{max}$  before sending. To ensure that any node that wants to send to  $S1$  can do so after the idle signal,  $S1$  waits for a time interval of  $t_{max}$ . Only if no node sends a packet to  $S1$  during this time interval,  $S1$  is allowed to send the next packet. Otherwise, it first has to receive a data packet from another node as shown in the example in Figure 4. It can then send an ACK with the busy flag set to indicate that it will now send the next data packet. In the example,  $S1$  has to further forward the packet it received from  $S2$  and will do so immediately after the transmission of the ACK.

This scheme ensures that nodes alternately send and receive (unless there is nothing to send or to receive). It is vital for a fair sharing of resources (i.e., access to nodes). A probability of 50% for sending and receiving is near the optimal operating point for the relay simulations considered in [7]. We found that directly alternating between sending and receiving instead of randomly improves forwarding performance with our MAC layer.

**Failed Transmission:** A node  $S1$  is only allowed to immediately send data to a destination  $D$  if none of the previous transmission attempts to  $D$  failed (or if  $S1$  and  $D$  did not communicate at all for a certain amount of time and  $D$  is idle). If  $D$  is busy, such a transmission attempt will fail, but will usually only cause a small amount of interference and will not disrupt  $D$ 's communication (as indicated by the bubbles in the graph). If  $D$  is already receiving data on  $THS(D)$ , a data packet from  $S1$  will be sent on the same THS. Due to the low auto-correlation of THSs discussed in Section 3.2, a transmission on the same THS will create the same interference as a transmission on a random THS, unless the phase difference between the overlapping transmissions is less than a few chip times. Only in the case of two almost simultaneous transmission the interference is likely to result in a packet loss. In any case, any further transmissions from  $S1$  to  $D$  are only possible after  $S1$  receives the corresponding idle signal. If instead  $D$  is sending to another node, communication will take place on a different THS and  $S1$  will only cause some interference. A node may repeat a failed transmission a certain number of times (in our simulations we use 4).

**Deferred Transmission:** In the example, node  $S2$  transmits to  $S1$  while  $S1$  itself is transmitting to  $D$ . Therefore,  $S2$ 's transmission will fail. After the transmission of the data packet,  $S2$  will start listening on  $THS(S1)$  for the ACK. If it does not immediately receive an ACK (or NACK) after the time it takes to send the ACK and twice the maximum propagation delay (expiration of the *send* timer), it knows that

the transmission failed. It will then set a *wait\_for\_idle* timer to the duration of a packet transmission with the lowest rate code, the transmission time of an ACK and twice the maximum propagation delay. When this timer expires, the data packet is resent.

If during this time *S2* receives an idle signal or an ACK with the idle flag set from *S1*, as shown in the example, it will cancel the *wait\_for\_idle* timer and start the backoff timer. If the backoff timer was still paused from a previous transmission attempt, it will resume the backoff with the current value of the backoff timer. When the backoff timer expires, *S2* sends a data packet. If it sees a data packet for *S1* before the timer expires, it would pause the backoff timer and restart the *wait\_for\_idle* timer. If instead *S2* were to receive an ACK from *S1* with the busy flag set, it would know that *S1* will transmit a packet and would therefore start the idle timer anew and continue to listen for the next idle signal. *S1* has to issue an idle after its own packet transmission and when this idle signal is received, *S2* can resume with its backoff.

This is shown in the example for the transmission of data packet 3 and 4. Both *D* and *S2* have a packet to transmit to *S1* and their backoff timers are running. *D*'s timer expires first. Assume that *S2* can decode the MAC header encoded at with the lowest rate code (but will not necessarily be able to decode the data part of the packet). *S2* will pause its backoff timer and set the *wait\_for\_idle* timer. In the example, the timer is started anew after *S2* receives an ACK from *S1* with the busy flag set. In case *S2* cannot even decode the MAC header, it will send a packet after the expiration of the backoff timer but as mentioned before, this transmission will usually only create some interference.

In the special case where a node wants to send to the node it just received a packet from, as is the case with data packet 3 from *D*, the node *always* has to wait for the idle signal even if it would otherwise be allowed to send immediately. This is necessary to prevent that the data packet is sent at the same time as the idle signal and is therefore lost. It is further possible to piggyback data onto the ACK packets. For simplicity, this is not done in the example figure but it significantly improves performance when two-way communication is common (e.g., when TCP is used as transport protocol). After such an exchange, both nodes have to issue an idle signal or an ACK with the idle flag set to allow other nodes to contact them.

While a node is waiting for an idle signal, it will listen on the destinations THS to receive the idle, as well as its own. In case a data packet is received, it will reply with an ACK and then resume waiting for the idle signal.

A node may resume sending without waiting for an idle signal (or an ACK with the idle flag set), when the idle timer expires and the following transmission succeeds (i.e., no idle signal is received for the maximum transmission time but the destination was in fact idle). Dynamically switching between immediate transmission and an invitation-based scheme allows to keep access delays low in a lightly loaded network and at the same time provides fair access to nodes as soon as there is contention. Futile



packet transmissions to destinations that are busy are almost completely avoided.

## 8 Simulations

Thus far, we analyzed the basic properties of our protocol in very simple scenarios by means of Matlab simulations. The main goal of the simulations is to investigate if our protocol works as expected under more realistic network conditions. To this end, the well-known network simulator ns-2 has been significantly extended by incorporating a model for a UWB physical layer as well as new MAC layer protocols. Since interference plays an important role, much attention has been paid to accurately model radio interference of concurrent transmissions. For signal propagation we use a UWB-specific propagation model proposed in [6], which is derived from indoor UWB measurements. Further details of the ns-2 implementation are described in [16].

The following protocols are compared to our dynamic channel coding-based MAC protocol (DCC-MAC):

**Power Control:** The power control MAC is based on the CA/CDMA protocol proposed in [17]. We adjusted the protocol to work together with a UWB physical layer instead of CDMA for which it was originally designed. While our implementation abstracts from some protocol details, it captures the main aspect of adjusting the power instead of the channel code. We define a minimum signal-to-interference ratio that is necessary to achieve a given probability of error. The transmission power of the packet is then set so as to achieve the desired SINR plus a safety margin, which allows for a limited amount of future transmissions to overlap with the current transmission. If the required power level exceeds the maximum power limit at the sender or the interference margin of ongoing transmissions, the sender defers from transmitting and retries after a random backoff.

**Mutual Exclusion with Random Access (RA):** All nodes use the same time hopping sequence. Therefore, if a node is transmitting, all other nodes within communication range will receive the packet and cannot send (since a node cannot send and receive at the same time). All nodes but the destination discard the packet. If a node has a packet to transmit while another node is sending, it retries after a backoff.

**Mutual Exclusion with TDMA:** is the ideal mutual exclusion without overhead. We do not actually implement this protocol in ns-2. Instead, we simulate transmission of every link independently of others, and obtain the rate for each one. We assume each link has the channel access for the equal fraction of the time, and from that we calculate the average data rate per link.

While MAC protocol details differ, the principles on which the implemented power control MAC is based are the same as the ones of other power control protocols proposed for UWB, such as [2, 14].

Similarly, the MAC layer proposed for 802.15.3 [1] can be seen as a combination of TDMA and the exclusion-based random access MAC.

For all of the simulated MAC protocols, the *same* UWB physical layer model is used. The parameters of the physical layer (such as peak power and capacity) are the ones described in Section 3. Since we are interested in very low-power MAC protocols, we allocate the same maximum power limit for the exclusion-based MAC protocols as for the DCC-MAC. We analyze the average data throughput achieved by all nodes, taking into account the loss in bit rate due to channel coding and the overhead due to the transmission of control packets.

## 8.1 Simulation Results

**Generalized Near-Far Scenario.** The near-far scenario we used for the simulations is an “unfolded” two-dimensional version of the one shown in Figure 2, since ns-2 does not allow for three-dimensional simulation topologies. We consider networks with 2 to 16 senders. The distance between sender and receiver varies from 1m to 20m but for reasons of brevity we only show the worst case graph with a distance of 20m.

Simulations with a varying number of interfering nodes are depicted in Figure 5. The sender-receiver distance is 20m for all of the communicating pairs of nodes. The DCC-MAC clearly outperforms the other MAC solutions. There is only a moderate drop in rate from 2300 Kb/s to 1800 Kb/s when we increase the number of nodes from 2 to 16 (i.e., 1 to 15 interferers). For the other MAC protocols, the drop in rate with an increasing number of senders is more pronounced. Power control comes closest to DCC-MAC performance since it allows for a limited amount of concurrent transmissions. It achieves between 75% and 30% of DCC-MAC’s rate. Both exclusion-based protocols, TDMA and random access, have very similar performance which is significantly worse than that of power control or DCC-MAC. The improvement in SINR and the resulting higher channel code rates cannot compensate for the loss in transmission time due to exclusion.

**Random Scenario.** In this scenario, nodes are randomly placed on a square surface of 20m×20m. Source-destination pairs are randomly chosen such that each node is either a source or a destination of exactly one link. The number of senders varies from 1 to 32.

With random node placement, the probability that there are many strong interferers is much lower than in the constructed near-far scenario. For up to 8 senders, power control performs almost as well as the DCC-MAC since the adaptation of transmit power allows that the nodes send concurrently for most of the simulated topologies. However, for 16 or more senders, the performance of power control quickly drops to that of the exclusion based protocols, since the increased interference exceeds the allocated interference margins. For the exclusion-based protocols we see that the achieved throughput is inversely

proportional to the number of senders. As before in the near-far scenarios, the DCC-MAC only has a slight decrease in rate for larger numbers of senders due to the dynamic code adaptation that becomes important when the number of nodes (and therefore interference) is high.

**Multi-hop Scenario.** Multi-hop forwarding in wireless networks has been extensively studied and was shown to be difficult (see for example [5, 12]). As is usually done, we investigate multi-hop performance of the different MAC protocols using a simple line topology. Source and destination are at either ends of the line of nodes; intermediate nodes forward packets between them. The distance between nodes is 20m.

The results of simulations are shown in Figure 7. In general, TCP throughput is lower than UDP throughput since TCP data packets compete with acknowledgments traveling on the return path. The most apparent drop in throughput occurs when the number of hops increases from one to two. The intermediate node in a 2-hop topology can either send or receive which necessarily halves the throughput for all of the protocols. What is striking is, that the DCC-MAC is able to maintain this rate when the number of hops increases beyond 2. For UDP, there is a small drop in throughput from 2 hops to 3 hops and from there on the rate remains constant. Also TCP, the decrease is on the order of a few percent. This excellent performance of the DCC-MAC is mainly due to the good interplay of timers and signals which results in close to optimal schedules. (No piggybacking of data is used.) For power control, there are a number of schedules that allow concurrent transmissions over at least a few of the hops; throughput is therefore in between that of the DCC-MAC and the exclusion-based protocols.

**Impact of Mobility.** Finally, we analyze the impact of mobility on the performance of the network. We consider the random scenario from Section 8.1 and let nodes move according to the random way-point model. Node speed varies between 2m/s and 10m/s with 0 pause time. To isolate the effect of mobility on the MAC protocol from the performance of a particular routing protocol, we do not use multi-hop routing.

Comparing the achieved network throughput in the mobile scenario given in Figure 8 with the throughput of the static network, we observe that our MAC protocol is very resilient to mobility. A change of channel conditions due to mobility is compensated by our channel code adaptation mechanism. The adaptation is sufficiently fast compared to the node speed to prevent a degradation of the rate. The same holds true when TCP is used instead of UDP. Even for TCP, the variations in channel code caused by mobility are not sufficiently high to result in a perceptible decrease in throughput.

## 9 Conclusion

We have presented a joint PHY/MAC architecture for very low power UWB. We assume that all nodes have simple receivers and transmitters (single user decoding, only one receiver per node, send and receive cannot be simultaneous) and all have the same value of PRP. Future work should focus on removing these restrictions.

Our scheme works very well for very low power UWB, i.e., when PRP is large. Our initial results indicate that even for medium values of PRP (around 100) the performance remains similar. For very low PRP, interference mitigation is not possible. Exclusion mechanisms such as TDMA or CSMA/CA are required. Given the high spatial reuse of our protocol when PRP is large, it is not clear that there is a large benefit of allowing PRP to be small, in other words, to allow more radiated power. Further research is needed to clarify this issue.

We use PPM modulation. Other, non coherent modulation schemes are also discussed for UWB [21]. It seems that our MAC protocol would apply with little change to such modulations, but this is also for further study.

Finally, we have developed a protocol guided by the idea of arranging the physical layer and the MAC protocol such that collisions may be replaced by rate reduction. This idea is optimal for our setting, but it could prove interesting in other settings as well. The optimal MAC protocol in narrowband systems is likely to be a combination of dynamic channel coding and mutual exclusion. Mutual exclusion has severe performance problems, as witnessed by the intense research on improving the 802.11 MAC protocol for use in ad-hoc and mobile networks. In contrast, dynamic channel coding does not appear to have these problems, since it is a private affair between a source and a destination, Therefore it would be interesting to add this component to existing MACs.

## References

- [1] IEEE 802.15.3 MAC standard, available at <http://www.ieee.org>.
- [2] 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.
- [3] G. Durisi and S. Benedetto. Performance evaluation of TH-PPM UWB systems in the presence of multiuser interference. *IEEE Communications Letters*, 7(5):224–226, May 2003.
- [4] P. Frenger, P. Orten, T. Ottosson, and A. Svensson. Rate-compatible convolutional codes for multirate DS-CDMA systems. *IEEE Transactions on Communications*, 47(6):828–836, June 1999.
- [5] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla. The impact of multihop wireless channel on TCP throughput and loss. In *INFOCOM*, 2003.

- [6] S. Ghassemzadeh and V. Tarokh. UWB path loss characterization in residential environments. In *RFIC Symposium*, pages 501–504, June 2003.
- [7] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4), Aug. 2002.
- [8] J. Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Transactions on Communications*, 36(4):389–400, April 1988.
- [9] D. H  lal and P. Rouzet. ST Microelectronics Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 139r5, July 2003.
- [10] A. Hicham, Y. Souilmi, and C. Bonnet. Self-balanced receiver-oriented MAC for ultra-wide band mobile ad hoc networks. In *International Workshop on Ultra Wideband Systems*, June 2003.
- [11] G. Holland, N. H. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *Proc. ACM/IEEE MOBICOM’01*, pages 236–251, 2001.
- [12] H.-Y. Hsieh and R. Sivakumar. IEEE 802.11 over multi-hop wireless networks: Problems and new perspectives. In *IEEE VTC-Fall*, Vancouver, Canada, Sept. 2002.
- [13] M. Iacobucci and M.-G. Di Benedetto. Multiple access design for impulse radio communication systems. In *IEEE ICC*, volume 2, pages 817–820, 2002.
- [14] S. Kolenchery, J. Townsend, and J. Freebersyser. A novel impulse radio network for tactical military wireless communications. In *IEEE MILCOM*, pages 59–65, 1998.
- [15] Y. Ma, F. Chin, B. Kannan, and S. Pasupathy. Acquisition performance of an ultra wide-band communications system over a multiple-access fading channel. In *UWBST*, pages 99–103, 2002.
- [16] R. Merz, J. Widmer, J.-Y. L. Boudec, and B. Radunovic. Ultra-wide band MAC and PHY layer implementation for ns-2, 2004.
- [17] A. Muqattash and K. Marwan. CDMA-based MAC protocol for wireless ad hoc networks. In *Proceedings of MOBIHOC’03*, pages 153–164, June 2003.
- [18] J. G. Proakis. *Digital Communications*. McGraw–Hill, New York, NY, 4th edition, 2001.
- [19] B. Radunovic and J. Y. Le Boudec. Optimal power control, scheduling and routing in UWB networks. *To appear in the IEEE Journal on Selected Areas in Communications*, December 2004.
- [20] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic media access for multirate ad hoc networks. In *8th Annual Int. Conf. on Mobile Computing and Networking*, 2002.
- [21] Y. Souilmi, R. Knopp, and G. Caire. Coding strategies for UWB interference-limited peer-to-peer networks. In *WiOpt’03*, March 2003.
- [22] E. Sousa and J. Silvester. Spreading code protocols for distributed spread-spectrum packet radio networks. *IEEE Transactions on Communications*, 36(3):272–281, March 1988.
- [23] E. Telatar and D. Tse. Capacity and mutual information of wideband multipath fading channels. *IEEE Transactions on Information Theory*, 46(4):1384–1400, 2000.
- [24] S. Toumpis and A. Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE/ACM Transactions on Wireless Communications*, 24(5):736–748, May 2003.
- [25] 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.

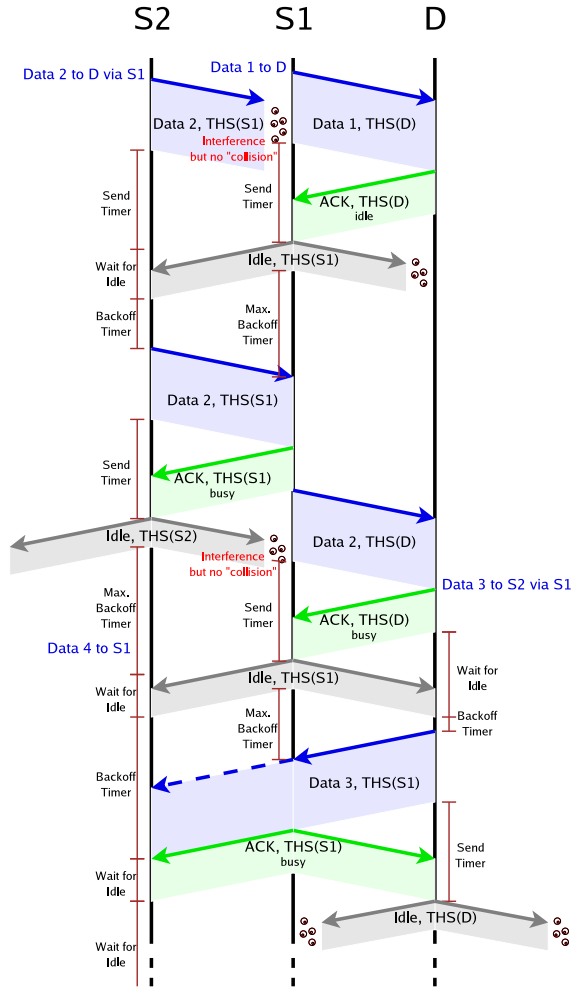


Figure 4: Multi-hop scenario: Transmission from  $S2$  to  $D$  via  $S1$  fails since  $S1$  is already sending to  $D$ .  $S2$  retries after receiving an idle signal from  $S1$ . After reception at  $S1$ , the packet is immediately forwarded to  $D$ .  $D$  then sends data back to  $S2$  (again via  $S1$ ). The interplay of *wait\_for\_idle* timer and backoff timer results in short idle times and forwarding delays and very few unnecessary transmissions to a destination that is busy.




Figure 5: Near-far scenario with a link length of 20m. We show average rate per user vs. number of (mutually interfering) senders.




Figure 6: Random scenario with nodes placed on  $20m \times 20m$  square. The number of nodes is given on x-axis, and the average rate is given on the y-axis.

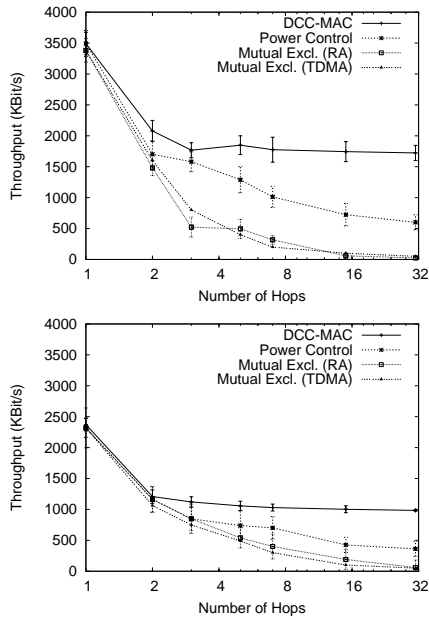


Figure 7: Throughput on the multi-hop network for UDP (left graph) and TCP (right graph). We show throughput vs. number of hops. There is almost no drop in throughput for the DCC-MAC as the number of hops increases.




Figure 8: Mobile network with random waypoint mobility model. The x-axis is gives the number of nodes in the network and on the y-axis is the average rate per user.