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, School of Computer and Communication Sciences

CH-1015 Lausanne, Switzerland

phone: (+41) 21 693 6616, fax: (+41) 21 693 6610

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

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-Wideband (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, 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 alleviates the need for an exclusion scheme. 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 compared to alternative designs.

Keywords: Medium Access Control, Ultra-Wideband, 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.[1] At the same time, many applications require high data rates. Ultra-wideband (UWB) wireless networks have the potential to satisfy both requirements. The low emitted power of UWB limits very high data rates to sender-receiver distances of tens of meters. Longer distance networks can be achieved through appropriate coding and modulation schemes at the expense of the achievable data rate. UWB is characterized by an extremely broad use of the radio spectrum which makes it relatively robust against channel impairments such as multipath fading.

---

[1]Note that we do not address the issue of maximizing battery lifetime, which is typical for sensor networks.

It was shown in [27] that the optimal wide-band signaling consists of sending infrequent short pulses. We use a Time-Hopping UWB impulse radio (TH-UWB IR) physical layer as in [15, 28]. 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 (see Figure 1). 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 Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) [16], Time Division Multiple Access (TDMA), or a combination of both [10] 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 networks), 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 [4, 21]. 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 [7] 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. In [24], sources send to a central base station at full power, as soon as they have something to transmit, but adapt the channel code in order to allow the central destination to properly decode in the presence of interfering sources. A striking feature of the model in [24] is, that the optimal scheme does require any power control [2]. A specific mathematical analysis of an optimal design for ultra-wideband networks including exclusion, power control, and rate adaptation is given in [23]. Their findings are:

- The optimal design should not use power control. Sources should send at full power whenever they send.

- 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. However, the channel

---

[2]given their centralized scenario, and with specific assumptions on optimal coding

code (hence the rate) should be adapted to this interference.

- In contrast, interference from inside the exclusion region should be combated.

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 silencing sources within the exclusion region, we propose to use 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 4 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 severely impacts the rate achievable by other users, 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 is that the MAC should primarily manage access by adapting rate to interference, without attempting to exclude competing sources by a mutual exclusion protocol. To be able to handle the remaining interference, we replace the traditional repetition codes [28] with more sophisticated RCPC codes (Section 3.2), the rates of which we adapt dynamically to the current level of interference. In these aspects, our design radically differs from existing ones.

There still remains some exclusion to implement since we assume that a node can be engaged exclusively either in the reception or the transmission of one single packet. This is enforced by the "Private MAC" (Section 6).

Our main contribution is a system for pulsed TH-UWB with the following three components. (1) *Interference mitigation* is described in Section 4.2. (2) *Dynamic channel coding* continuously adapts the rate to variable channel conditions and interference (Section 5). 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. With a TH-UWB physical layer, it is difficult to do "carrier sensing" without actual

decoding. The challenge of absence of carrier sensing is solved by an elaborate signaling protocol that alternates between direct access to a receiver and an invitation-based scheme. We do not use any separate channel for control. Our protocol is fully implemented in ns-2. Simulation results show a significant increase in throughput compared to traditional protocol design.

2   Related Work

We already mentioned in the introduction some related work [23, 24] that suggests that rate control is preferable to power control.

Proposals for a distributed MAC layer for pulsed UWB ad-hoc networks can be found in [4, 18, 11, 3, 2, 13]. They are all based on the TH-UWB physical layer of [28] and use a fixed channel code. Essentially based on a combination of power control and mutual exclusion, [4] proposes a distributed control admission function. It is based on the evaluation of the potential interference a new link would cause on existing links. A source broadcasts 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 addressed. The approach is similar to [18].

In [11] an invitation based scheme to address the problem of contention for a destination is proposed. A node that is ready to receive broadcasts an invitation for other nodes to compete for access to it. However, no power or channel code control is performed. Note that in [4, 18, 11] a separate control channel is used to transmit control packets. Recently, [3, 2, 13] proposed traditional CSMA based MAC designs for pulsed UWB physical layers. Both [3] and [2] have a fixed channel coding and no power control. While they address the problem of carrier sensing through the use of a frequency-domain detection method [3], this greatly augments the complexity of the physical layer circuitry. A busy tone protocol is added in [2] to alleviate the hidden terminal problem inherent in CSMA.

Another power control protocol, although based on DS-CDMA, can be found in [21]. 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.

The IEEE 802.15 Task Group 3a 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.

In [13], a Direct-Sequence UWB (DS-UWB) physical layer with variable length spreading sequences is used. However, it differs 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. The concept of rate adaptation has also been proposed for 802.11 networks [12, 25] for the same purpose. In contrast, we use rate adaptation as a mechanism to support multiple-access.

It is also interesting to note that the interference mitigation scheme (Section 4.2) can be seen as an adaptive version of the multi-user blinking receiver in [5] in that it doesn't need the knowledge of all the THS in the network. A similar scheme is presented in [19].

3 System Assumptions

3.1 Modeling of the Physical Layer

Figure 1: UWB physical layer model. In this example, $PTP = 4$, the THS is $\{\ldots, 2, 1, \ldots\}$ (Multipath is not represented.) Interference Mitigation (IFM) reduces the impact of pulse collisions with strong interferers.

We use a time-hopping ultra-wideband impulse radio (TH-UWB IR) multiple-access physical layer as in [15, 28]. However, we exchanged the usual repetition code of [28] for a more sophisticated channel code (see Section 3.2). For our protocol we use Binary Phase Shift Keying (BPSK) modulation. However, our design is independent of a particular modulation scheme and can be applied to any TH-UWB IR

physical layer. A source k produces the signal

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

where E_p is the pulse energy, $p(t)$ is a unit energy pulse and $\left\{ x_j^{(k)} \right\}_{j=-\infty}^{\infty}$ is the information bearing sequence. T_c is the chip duration and T_f is the frame duration. Note that $T_f \geq PTP \cdot T_c$ where PTP is the average Pulse Transmission Period. We keep a fraction $\lfloor \epsilon \cdot PTP \rceil$ of chips unused at the end of a frame to account for the multipath delay spread. The sequence $\left\{ c_j^{(k)} \right\}_{j=-\infty}^{\infty}$ is the THS of the kth source. The received signal at the destination is

$$r(t) = \sum_{k=1}^{N_u} \left(h^{(k)} * s^{(k)} \right) (t - \tau_k) + n(t) \tag{2}$$

where $h^{(k)}(t)$ is the impulse response of the channel between a source k and the receiver, $n(t)$ is zero-mean white Gaussian noise of variance $\frac{N_0}{2}$ and N_u is the total number of users in the network. Since we consider a decentralized network with no global clock, τ_k, $k = 1, 2, \ldots, N_u$ model the time-shift between transmitters in the network. The impulse response $h^{(k)}(t) = \sum_{l=0}^{L} \alpha_l^{(k)} \delta(t - \nu_l^{(k)})$ models the multipath channel. The attenuation coefficient $\alpha_l^{(k)}$ takes into account path loss and random fading on the lth path, ν_l is the delay on the lth path, L is the number of multipath components and $\delta(x)$ is the Dirac delta function. The signal $r(t)$ is passed through a *maximum ratio combining* (MRC) rake receiver and sampled to produce the received symbols $\{y_j\}_{j=-\infty}^{\infty}$. We follow the steps and assumptions of [15] and generalize them for a Rake receiver. It is straightforward to show that the sampled output of the rake receiver is equal to $\{y_j\}_{j=-\infty}^{\infty}$ where $y_j = \sum_{l=1}^{L} y_{j,l} = \sum_{l=1}^{L} (S_{j,l} + I_{j,l}) + n_j$. Let $k = 1$ be the user of interest, $S_{j,l} = \left(\alpha_{j,l}^{(1)} \right)^2 x_j^{(1)}$, the multi-user interference (MUI) is

$$I_{j,l} = \alpha_{j,l}^{(1)} \sum_{k=2}^{N_u} x_j^{(k)} \sum_{m=1}^{L} \alpha_{j,m}^{(k)} \Theta \left(\varphi_k + c_j^{(k)} T_c + (\nu_m^{(k)} - \nu_l^{(1)}) \right) \tag{3}$$

where φ_k is uniformly distributed on $\left[-\frac{T_f}{2}, \frac{T_f}{2} \right)$, the filtered white noise is $n_j \sim \mathcal{N} \left(0, \frac{N_0}{2} \right)$ and $\Theta(\tau) = \left[1 - 4\pi \left(\frac{\tau}{\tau_p} \right)^2 + \frac{4\pi^2}{3} \left(\frac{\tau}{\tau_p} \right)^4 \right] \exp \left[-\pi \left(\frac{\tau}{\tau_p} \right)^2 \right]$ is the pulse[3] autocorrelation where τ_p is a time normalization factor. The output of the Rake receiver is in turn passed to the *channel decoder* who will attempt to recover the transmitted data.

Equation (3) allows to accurately model the MUI for the simulations in Section 4. To simulate the channel, we use the distribution of channel coefficients $\alpha_i^{(k)}$, $i = 1, \ldots, L$ given in [8] with $L = 5$.

---

[3]We consider a second derivative Gaussian shaped pulse

3.2   Practical Aspects of the Physical Layer

We assume that $T_c = 0.2$ ns (which results in roughly 5 GHz of bandwidth), $E_p/T_c = 0.28$ mW [10] and $\epsilon = 0.1$. To obtain an average radiated power of 1μW we need $PTP = 280$. The rate compatible punctured convolutional (RCPC) codes of [6] replace the repetition code of [28]. Let $R_0 = 1 > R_1 > \ldots > R_N$ be the set of code rates. Data encoded at rate R_i is obtained by *Puncturing* (i.e., by removing certain bits) from data encoded at rate R_N. Moreover, with *rate compatibility* puncturing is done such that data with rate R_i is a subset of the data at rate $R_j, j > i$. Hence, RCPC code can provide *incremental redundancy*: whenever data at rate R_i cannot be decoded, it is only necessary to send the punctured bits that are part of R_j and not of R_i. The available code rates of the RCPC code are $R_n = \{1, 8/9,\ 8/10,\ 8/11, \ldots, 8/32,\ 1/5,\ 1/6, \ldots, 1/10\}$, $n = 0, \ldots, 31$, hence the maximum rate of our physical layer is $\frac{R_0}{PTP \cdot T_c} = 18$ Mbps.

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 (possibly with the help of additional synchronization sequences within the packet), and can be re-established for each data packet. There is no global synchronization.

3.3   Impact of Time Hopping Sequences on the MAC design

A THS determines in which chip each user should transmit. Hence a THS is simply a random sequence of integers uniformly distributed in the interval $\big[0, \lfloor(1-\epsilon)PTP\rfloor - 1\big]$. Although THSs share conceptual similarities with Direct Sequence (DS) spreading codes used in CDMA systems, they are *not* equivalent. Whereas in TH-UWB the transmission of pulses is infrequent (one per frame), it is continuous in DS systems. Two concurrent transmissions in TH-UWB only interfere when pulses overlap, whereas they always interfere in DS systems. Therefore, finding good time hopping sequences is almost trivial, while finding low cross-correlation spreading sequences for asynchronous DS systems is a difficult problem [22]. They cannot be computed on the fly and an assignment protocol is necessary [26]. In our case, no assignment protocol is required. The THS of a node is generated using its MAC address as seed for a pseudo-random number generator and is identical for every packet.

4 Exploring the Design Space

As indicated in [23], the optimal design in our case should allow interference outside the exclusion region but forbid it inside the exclusion region. Clearly, the most important parameter is the size γ of the exclusion region. Hence, the next section characterizes γ for the particular TH-UWB physical layer we use. In particular, we want to show that γ is negligible in our case.

4.1 Computing the size of the exclusion region

In networks with arbitrary topology, all destinations $D_1, \ldots, D_n$ have a different exclusion region γ_k. In such scenarios, computing γ_k is known to be a hard problem [23]. We therefore resort to the symmetric topology of Figure 2 where $\gamma = \gamma_1 = \ldots = \gamma_n$. The links $\{S_1, D_1\}, \ldots, \{S_n, D_n\}$ are placed in an alternate way on a cylinder of length $\mathcal{L}_c$. The distance between two links is d. If $\gamma > d$, then it is optimal 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 $R_{all}(d) = R_{excl}(d)$ [23]. We use this property to determine the size of the exclusion region. For a given n and $\mathcal{L}_c$ 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, $j = 2, \ldots, n$, we first fix a BER threshold μ^*. A typical value for μ^* in a wireless environment is 10^{-5} [22]. We then determine the maximum possible rate R_i, $i = 1, 2, \ldots, N$ that drives the BER below μ^*. Modeling the MUI in equation (3) as Gaussian would permit to use an analytical expression [28] for γ in an uncoded TH-UWB physical layer. However, in our settings, this assumption does not hold for TH-UWB [15]. Since we further use convolutional channel codes and a multipath channel, we turn to simulations to derive γ. The physical layer model is the one described in Section 3.1 with the parameters of Section 3.2. We furthermore consider two types of channel decoding policy, hard-decision and soft-decision. With a hard-decision policy, only the sign of the demodulator output is passed to the channel decoder, whereas the sign *and* the amplitude is passed in the case of soft-decision decoding. Usually, the soft-decision policy performs better than a hard-decision policy [22]. However,

9

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 underlying assumption in this case is that the total interference (MUI and noise) has a Gaussian density. This justifies our interest in analyzing the performance of a hard-decision policy.

The results are shown in Figure 3. We consider $n = \{4, 8, 16\}$ links and use link lengths of $\mathcal{L}_c = \{2, 6, 12, 18\}$ to obtain varying signal and interference intensities. When hard-decision is used, there is no exclusion region (i.e., $\mathcal{R}_{all}(d) > \mathcal{R}_{excl}(d) \, \forall \, d$. 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 case, an exclusion region of 1 to 4 meters is present depending on $\mathcal{L}_c$ and n. Although the probability of collision $P_{col}(n) = 1 - \left(1 - \frac{2T_c}{T_f}\right)^n$ with interfering pulses is low ($< 1\%$ for $n = 1$, 7.5% for $n = 10$), they have a large impact in the case of nearby interfering sources. When such collisions occur, the amplitude of $|y_j|$ is much larger than when only regular Gaussian noise is present. With soft-decision decoding, a large amplitude sample $|y_j|$ propagates over the decoding of several subsequent samples $|y_{j+1}|, |y_{j+2}|, \ldots$ 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 at large values of d[4].

---

[4]The result is similar when using more powerful codes like turbo codes.

10

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 γ, the size of the exclusion region.

11

Intuitively, the optimal decoding policy, should be an adaptive combination of hard-decision when strong interferers are present and soft-decision otherwise. 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 γ and still avoids the complexity of an exclusion scheme.

4.2 Replacing Exclusion by Interference Mitigation

We demonstrate how an interference mitigation (IFM) scheme at the physical layer reduces the effect of strong interferers and how it can replace an exclusion scheme. An IFM scheme cancels the samples y_j resulting from a collision with pulses of a strong interferer and replaces them by erasures (i.e. skips them in the decoding process). That is

$$
y_j = \begin{cases} \varepsilon, & \text{if} \quad |y_j| > \mathcal{B}(t) \\ y_j, & \text{otherwise} \end{cases}
\tag{4}
$$

where ε 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.

Ideally, $\mathcal{B}(t)$ should be set 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 on the average received power from the source, the number of fingers used by the Rake receiver and on the Gaussian white 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 $\mathcal{B}(t) = E[|Y_s(t)|] + kE[|N(t)|]$ where $E[|Y_s|]$ 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 $k = 2.8$. This value was found by simulations. 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 [17]. 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 IFM, 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.

5 Rate Adaptation by Dynamic Channel Coding with Incremental Redundancy

We propose a dynamic channel coding (DCC) procedure where the rate is constantly adapted to the level of interference experienced at a receiver. Ideally, DCC should always use the highest code rate that still allows decoding of the data packet. For this, we exploit the feature of our codes that a destination that can decode can also determine the highest possible code rate.[5] Dynamic channel coding works as follows:

- A source S keeps `codeIndex`(D) in a variable, where `codeIndex`(D) is the code index to be used for communication with D. Initially or after a sufficiently long 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 such that rate R_j would have allowed to successfully decode. 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 6). In this case, S will abort the transmission (without sending incremental redundancy) but may retry at a later time.

---

[5]Note that in contrast to the data part of a packet, the MAC header is encoded with rate R_N so that a receiver can determine that it received a packet even if it is not able to decode the data. Also control packets always use rate R_N.

Figure 4(a) gives an example of the rate adaptation. Since it is hard to measure the SINR in UWB, we determine the optimum code after packet reception. Decoding of a data packet encoded with rate R_i is performed by step-wise traversal of the trellis of the Viterbi decoder [22]. 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.

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. 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 conditions deteriorate. We find that the heuristic of using a channel code rate R_{j+2} if the highest possible code rate is R_j performs sufficiently well. The code R_{j+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 is violated the code index is increased accordingly.

6 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.[6] It always listens on the broadcast THS, which is

[6]We assume that node can *listen* on more than one THS but can only *receive* from one node at a time. Furthermore, a node

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 of 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(b). $S1$ will send the data packet using $THS(D)$ and will also start listening on $THS(D)$. As soon as D can decode, it sends back an ACK on it's 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 it's 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(b). 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 [9]. We found that directly alternating between sending and receiving instead of doing it 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

---

can either send or receive but not both.

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 3.3, 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 transmissions, 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 (i.e., after the 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 it's own packet transmission and when this idle signal is received, $S2$ can resume with it's 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 it's 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 it's 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.

7   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 in this section is to investigate if our protocol works as expected under more realistic 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 payed to accurately model radio interference of concurrent transmissions. For signal propagation we use a UWB-specific propagation model proposed in [8], which is derived from indoor UWB measurements. Further details of the ns-2 implementation are described in [20].

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 [21]. 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: 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. This corresponds to an ideal TDMA protocol where the schedule is obtained at no cost.

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 [4, 18].

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.

7.1 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 sender sends UDP packets at the highest possible rate to the receiver. The number of senders varies between 1 and 32.

For up to 8 senders, power control performs almost as well as the DCC-MAC since the adaptation of transmit power allows the nodes to 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. 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.

7.2 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. Again we consider scenarios with 1 to 32 senders. The distance L between sender and receiver varies from 1m to 20m but for reasons of brevity we only show the worst case graph with $L = 20$m. The distance d from a receiver to the closest interferer is 1m.

The arrangement of the nodes in a near-far topology results in much stronger interference than in the random scenario. As can be seen from Figure 5(b), the DCC-MAC clearly outperforms the other MAC solutions. There is a moderate drop in rate from 2900 Kb/s to 1800 Kb/s when we increase the number of senders from 1 to 32. For the other MAC protocols, the drop in rate with an increasing number of senders is much more pronounced. Again, power control comes closest to DCC-MAC performance since it allows for a limited amount of concurrent transmissions. However, the difference in performance is more obvious than in the random scenario. It achieves between 75% and 30% of DCC-MAC's rate in a near-far setting. Both exclusion-based protocols, TDMA and random access, have very similar performance which is significantly worse than that of power control and the DCC-MAC. The improvement in SINR and the resulting higher channel code rates cannot compensate for the loss in transmission time due to exclusion.

7.3 Multi-hop Scenario

Multi-hop forwarding in wireless networks has been extensively studied and was shown to be difficult for exclusion-based MAC protocols (see for example [7, 14]). As is usually done, we investigate multi-hop performance of the different MAC protocols using a topology where nodes are equally spaced along a line. 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 simulation results are shown in Figure 6. In general, TCP throughput is lower than UDP throughput since TCP data packets compete with acknowledgments traveling on the return path.[7] 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 with 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

---

[7]For TCP we use the TCP Sack implementation of ns-2 version 2.27 with default parameters.

in between that of the DCC-MAC and the exclusion-based protocols.

We also performed a number of simulations with node mobility. The results do not differ significantly from the ones presented above and are therefore omitted. In all cases, the variable channel codes were sufficient to cope with the channel variations caused by mobility.

8   Conclusion

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

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

For our design we use BPSK modulation. Other, non-coherent modulation schemes are discussed for UWB for example in [17]. It seems that our MAC protocol would apply with little change to such modulation schemes, but this also remains 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] N. August and D. Ha. An efficient UWB radio architecture for busy signal MAC protocols. In *IEEE SECON*, pages 325–334, October 2004.

[3] N. August, H.-J. Lee, and D. Ha. Pulse sense: A method to detect a busy medium in pulse-based ultra wideband (UWB) networks. In *IEEE UWBST*, pages 366–370, May 2004.

[4] 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.

[5] E. Fishler and H. Poor. Low-complexity multiuser detectors for time-hopping impulse-radio systems. *IEEE Transactions on Signal Processing*, 52(9):2561–2571, September 2004.

[6] 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.

[7] 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 *IEEE INFOCOM*, volume 3, pages 1744–1753, March-April 2003.

[8] S. Ghassemzadeh, R. Jana, C. Rice, W. Turin, and V. Tarokh. Measurement and modeling of an ultra-wide bandwidth indoor channel. *IEEE Transactions on Communications*, 52(10):1786–1796, October 2004.

[9] M. Grossglauser and D. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, August 2002.

[10] D. Hélal and P. Rouzet. ST Microelectronics Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 139r5, July 2003.

[11] 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.

[12] 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.

[13] W. Horie and Y. Sanada. Novel CSMA scheme for DS-UWB ad-hoc network with variable spreading factor. In *IEEE UWBST*, pages 361–365, May 2004.

[14] H.-Y. Hsieh and R. Sivakumar. IEEE 802.11 over multi-hop wireless networks: Problems and new perspectives. In *IEEE VTC-Fall*, volume 2, pages 748–752, September 2002.

[15] B. Hu and N. Beaulieu. Accurate evaluation of multiple-access performance in th-ppm and th-bpsk uwb systems. *IEEE Transactions on Communications*, 52(10):1758–1766, October 2004.

[16] L. Kleinrock and F. A. Tobagi. Packet switching in radio channels: Part 1–carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Transactions on Communications*, 23(12):1400–1416, December 1975.

[17] R. Knopp and Y. Souilmi. Achievable rates for uwb peer-to-peer networks. In *International Zurich Seminar on Communications*, pages 82–85, 2004.

[18] S. Kolenchery, J. Townsend, and J. Freebersyser. A novel impulse radio network for tactical military wireless communications. In *IEEE MILCOM*, volume 1, pages 59–65, October 1998.

[19] W. Lovelace and J. Townsend. Chip discrimination for large near far power ratios in uwb networks. In *IEEE MILCOM*, volume 2, pages 13–16, October 2003.

[20] R. Merz, J. Widmer, J.-Y. L. Boudec, and B. Radunoviⱬ Ultra-wide band MAC and PHY layer implementation for ns-2, http://icapeople.epfl.ch/widmer/uwb/index.html, 2004.

[21] A. Muqattash and M. Krunz. CDMA–based MAC protocol for wireless ad hoc networks. In *MOBIHOC*, pages 153–164, June 2003.

[22] J. G. Proakis. *Digital Communications*. McGraw–Hill, New York,NY, 4th edition, 2001.

[23] 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.

[24] S. Raj, E. Telatar, and D. Tse. Job scheduling and multiple access. DIMACS Series in Discrete Mathematics and Theoretical Computer Sciences, 2003.

[25] 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.

[26] 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.

[27] E. Telatar and D. Tse. Capacity and mutual information of wideband multipath fading channels. *IEEE Transactions on Information Theory*, 46(4):1384–1400, 2000.

[28] 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.
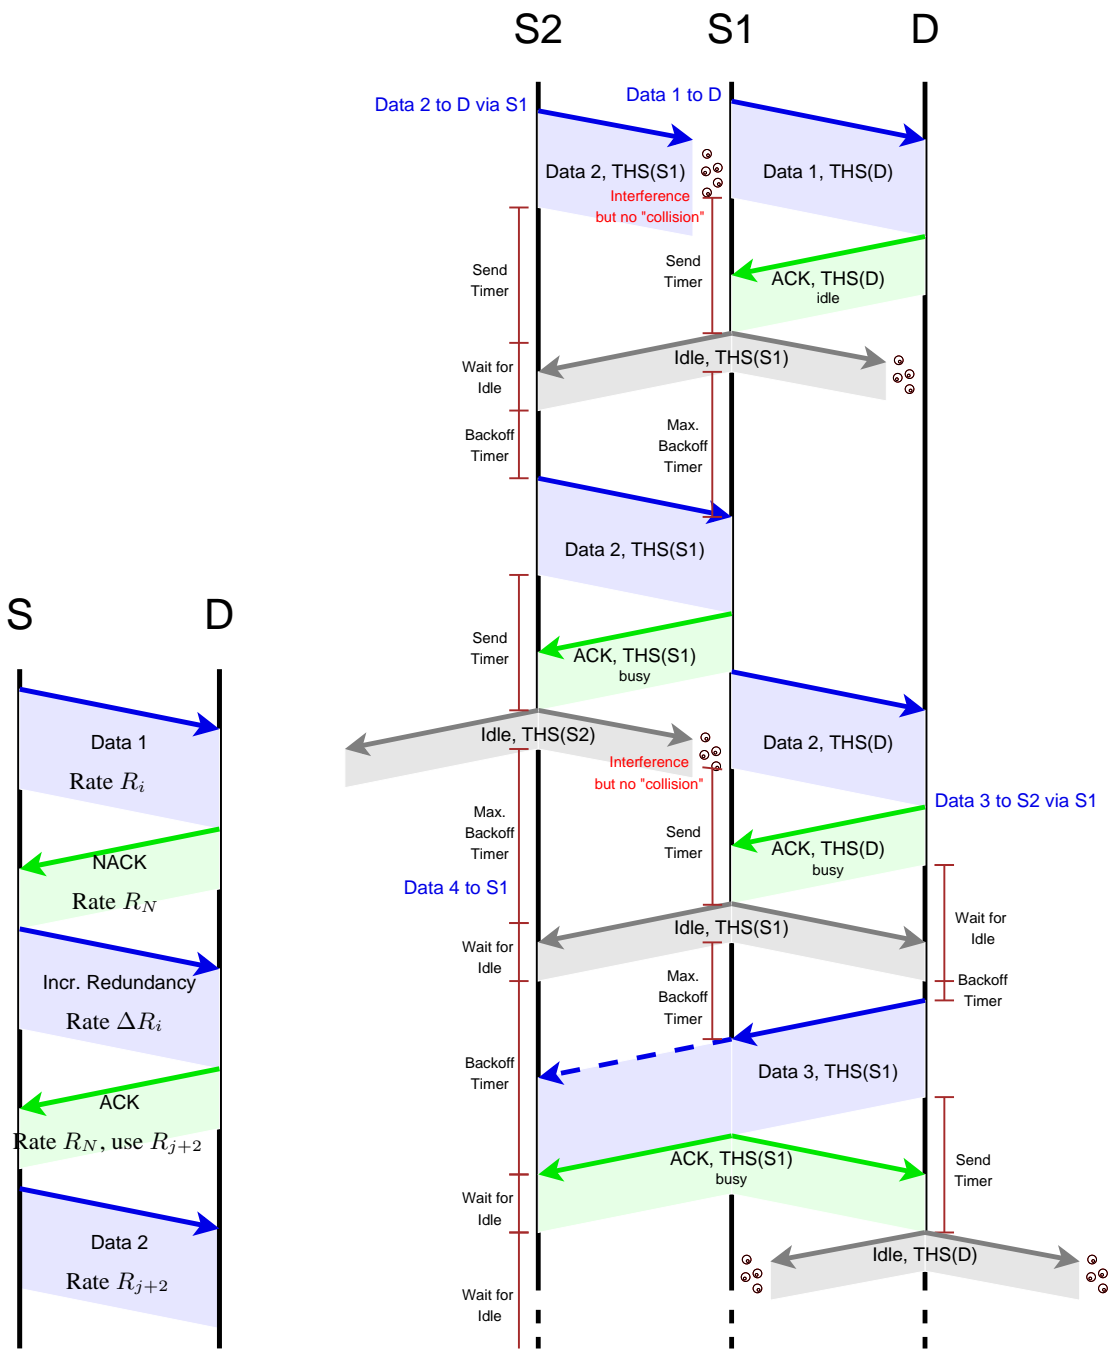
(a) Dynamic Channel Coding example: The first transmission attempt fails and additional data has to be transmitted so that the packet can be decoded. The rate to be used for subsequent packets is adjusted accordingly.

(b) 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 4: Organization of the MAC layer with Dynamic Channel Coding (left) and Multi-Access Control (right)

(a) Random scenario with nodes placed on 20m×20m square.

(b) Near-far scenario with sender-receiver distance of 20m and receiver-interferer distance of 1m.

Figure 5: Average throughput per user vs. number of senders for random and near-far scenarios with single-hop communication. The performance benefits of the DCC-MAC are particularly apparent in scenarios with strong interference.


Figure 6: UDP (left) and TCP (right) throughput achieved in the multi-hop scenario. We show average throughput vs. number of hops. There is almost no drop in throughput for the DCC-MAC as the number of hops increases.