

A MAC protocol for UWB Very Low Power Mobile Ad-hoc Networks based on Dynamic Channel Coding with Interference Mitigation

EPFL Technical Report ID: IC/2004/02

Jean-Yves Le Boudec, Ruben Merz, Bozidar Radunovic, Joerg Widmer

LCA-I&C-EPFL, CH-1015 Lausanne, Switzerland

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

January 26, 2004

Abstract

We propose a MAC protocol for very low radiated power (1 microwatt) ultra-wide band (UWB) mobile ad-hoc networks. Unlike traditional approaches, our protocol fully utilizes the specific nature of the physical layer of UWB. This makes it possible to realize a fully distributed protocol that reaches a reasonably high bit rate per source in spite of the low power. We use dynamic channel coding as an efficient way to adapt to varying channel conditions due to mobility and interference. Our design is based on existing theoretical findings that state that it is optimal to allow interfering sources to transmit simultaneously, as long as they are outside a well-defined exclusion region around a destination, and to adapt the channel code to this interference; in contrast, interference from inside the exclusion region should be combatted. We use pulse position modulation, as is commonly done with UWB. The very low power assumption implies that pulses are infrequent. This property, together with a demodulation scheme that cancels much of the interfering energy, allows us to mitigate interference from nodes in the exclusion region. This scheme is entirely local to a destination and involves no protocol action. In addition, sources constantly adjust their channel codes (thus their bit rates) and send incremental redundancy as required. Contention between sources sending to the same destination is solved by a “private MAC” protocol, i.e., a set of interactions that involve only the nodes that want to talk to the same destination. Carrier sensing is not used, as it is practically impossible with UWB. We show by simulation that we achieve a significant increase in network throughput, compared to traditional MAC protocols that are separated from the physical layer.

1 Introduction

1.1 Very Low Power Ultra-Wide Band for Ad-Hoc Networks

Emerging pervasive networks assume the deployment of large numbers of wireless nodes, embedded in everyday life objects. For such networks to become accepted, it is important that the level of radiated energy per node be kept very small; otherwise, environmental and health concerns will surface. Ultra-wide band (UWB) is an emerging radio technology for wireless networks, which has the potential for satisfying this requirement. UWB is characterized by an extremely broad use of the radio spectrum; more precisely, according to FCC, UWB has a bandwidth that is larger than 20% of the center frequency or a bandwidth equal to or greater than 500MHz. The radiated power per node depends on technological choices; it is of the order of 0.1 mW to less than 1 μ W per sender. We are interested in *very low power* UWB, by which we mean that the radiated energy per node does not exceed 1 μ W (= -30dBm). With currently planned technology, it is possible with such very low power to achieve rates of 1 to 18 Mb/s per source at distances on the order of tens of meters (Section 4). This is because UWB is particularly robust to channel impairments such as multipath fading. Of course, these rate values are reduced when several near UWB sources transmit concurrently. It is precisely the goal of this paper to design a Medium Access Control (MAC) protocol for very low power UWB that avoids much of the rate reduction. We achieve this goal by designing a MAC protocol that is joint with the physical layer.

1.2 Optimal MAC Protocol and Exclusion Region

The existing wireless MAC protocols discussed at the IEEE assume that simultaneous transmissions with the same channel code result into transmission errors and thus employ temporal exclusion mechanisms to avoid them. Exclusion is enforced either with a collision management protocol (CSMA/CA or a variant of it [15]), with a time division scheme (allocating time slots with a reservation protocol), or with a combination of both [12, 27]. 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 [10].

In contrast, wireless technologies that have the ability to allow and intelligently manage interference may accept several transmissions at the same time. This is for example the case for code division multiple access. In a purely synchronous setting (central base station), CDMA manages access by controlling power. In asynchronous settings (Ad-hoc networks), CDMA uses both power control and an exclusion mechanism [19].

Another largely unexploited dimension is channel coding. In [23], sources send to a central 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. By adapting the channel code, a source also changes its bit rate. A striking feature of the model in [23] is that sources send at full power – there is no power control and the authors show that this is optimal (in their centralized scenario, and with specific assumptions on optimal coding). Similarly, it is shown in [29] that power control does not provide significant gains when dynamic channel coding is used. In practice, dynamic channel coding can be implemented with incremental redundancy codes (Section 5).

It is possible to combine exclusion, power control and dynamic channel coding. A performance analysis in [21, 22] indicates that the optimal MAC layer should use full power when it sends, thus confirming the results in [29]. 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 a destination, and to adapt the channel code to these interferences; in contrast, interferences from inside the exclusion region should be combatted. The main reason is the non-linearity of the achievable bit rate as a function of the signal to noise and interference ratio: exclusion mechanisms divide time and rate linearly, whereas interferences reduce the rate less than linearly outside the exclusion region. We compute in Section 4 the exclusion region for our model; we find that it is not large but is not negligible either (it can be up to 2 meters).

These existing results tell us that the optimal MAC design for our case should (1) allow sources to send at the maximum power allowed by their budget (2) allow interference outside the exclusion region and (3) combat interference from the exclusion region.

A similar goal is implicitly achieved by CA/CDMA [19], which combines CDMA and an exclusion protocol. We go beyond [19] in two ways: First, we use dynamic channel coding instead of power control, since the theoretical findings mentioned above indicate that this is optimal. Second, we exploit specific features of very low power UWB that allow us to considerably simplify the exclusion problem, as we now explain.

1.3 Specificity of Very Low Power UWB

It was shown in [31, 28] that optimal wide-band signaling consists of sending short pulses. There are currently several proposals for the UWB physical layer; the model we use in this paper is based on Win-Scholtz’s proposal [33], which appears to be very close to one of the current IEEE 802.15.3a proposal [12]. It uses pulse position modulation (PPM) and a coherent receiver. Time is slotted in chips of very short duration T_c (0.2 ns in our model); chips are organized in frames of length PRP chips (Figure 1). PRP stands for “Pulse Repetition Period”. A node transmits one pulse in one chip per frame, and uses a pseudo-random *Time Hopping Sequence* (THS) to determine in which chip to transmit. The achievable capacity for one user is maximum if interfering pulses appear to be independent. This is achieved if, as

we assume, source-destination pairs use independent, pseudo random Time Hopping Sequences, and if sources are not synchronized with each other. This has some similarity with code division multiple access (CDMA) and spreading code techniques, but there are some important differences (see Section 3.4).

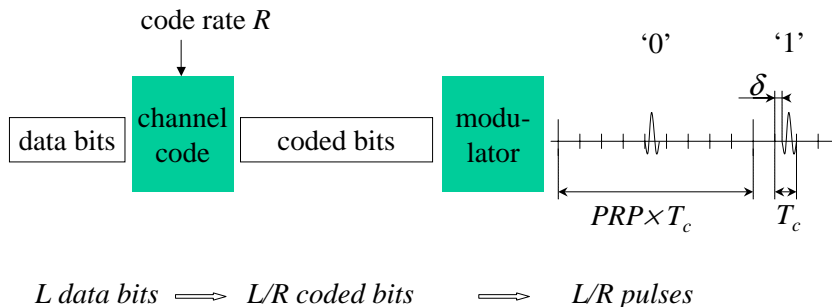


Figure 1: Model of Ultra Wide Band used in this paper

In addition to a Time Hopping Sequence, a source also uses a *Channel Code* to translate data bits into encoded bits that are in turn encoded as pulses by the *Modulator* (Figure 1). If the channel code is well chosen, it is able to deliver data bits with a small error probability even if some percentage of pulses suffer from interference.

The frame size PRP must be above some specified value (10 to 100) to avoid energy peaks in the frequency domain – a requirement imposed by regulation to avoid interference with other, non-UWB systems. Further, the radiated power depends linearly on PRP. Our very low power constraint imposes that PRP be large (the value we consider in this paper is $PRP = 280$, see Section 3.3).

We exploit this feature to construct a simple exclusion method. We use an erasure method, inspired by the ideas in [24]. Before explaining this, we need to say more about the exclusion region. Assume a source X interferes with a transmission from some node S to a node D . For X to be in the exclusion region around D , it must be that the interference caused by X at D is much larger than the signal from S received at D . Otherwise X is outside the exclusion region and it is optimal to allow X to interfere with S and to let S pick an appropriate channel code. Now we can come back to our point. We use a demodulator at D that detects when a received energy is much higher than the intended received power from S and declares an erasure. A possible pulse from S received at about the same time is lost by the erasure. The probability that X causes an erasure is of the order of $\frac{0.5}{PRP}$, i.e. 0.18% (Section 4). With a dynamic channel code as we use, a small erasure probability translates into a small reduction of the rate, of the same order. Compare to a temporal exclusion protocol (as in 802.11): the effect here would be to share the rate between X and S , which means that, in the best case (neglecting protocol overhead) both X and S have a reduction in rate equal to 50% (versus of a fraction of a percent). When there are n interfering sources in the exclusion region, the rate reduction grows less than linearly with n (since several interferences may overlap and because of the law of large numbers); the exact value

depends on the coding scheme and the spatial distribution of nodes, but in all cases we find reductions that are extremely small, much less than if temporal exclusion scheme would be used. Our interference mitigation scheme based on erasures is described in detail in Section 4.

Using erasures as a way to mitigate interference is only possible when PRP is large. In the hypothetical limit where PRP would be close to 1, erasures would destroy the entire signal sent by S and the method would not apply. Thus, this method is specific for very low power UWB. We found that our method continues to perform well for smaller values of PRP (down to 100), i.e. for higher power, but a detailed investigation remains for further research.

1.4 A MAC Protocol Based on Dynamic Channel Coding and Interference Mitigation

The above discussion suggests the following MAC. Since interference from inside exclusion regions can be taken care of by erasures at the demodulator, all that remains is for sources to adapt their channel codes (hence bit rates) to the level of interference on the channel. To this end, we use “dynamic channel coding”, with incremental redundancy. A source picks a channel code according to the protocol described in Section 5.1, receives feedback from the destination and, if needed, sends incremental redundancy for the destination to be able to decode. Contrary to some other protocols, when a source sends, it always sends at the maximum power allowed by its budget.

There remains, however, the need to support exclusion between sources that send to the same destination, because we assume that a node can receive from only one source at a time. We solve this problem by means of a *private MAC* protocol, described in Section 5.2. The private MAC concerns only nodes that have a common destination. Thus, our design moves the complexity of the MAC protocol away from global exclusion between competing sources (a difficult problem) to channel coding (a private affair between a source and a destination) and a collection of independent private MAC protocol instances (one instance per destination).

1.5 Goal and Paper Contributions

Our goal is to design a MAC protocol for a mobile UWB ad-hoc network with very low radiated power. We want to maximize rate subject to power dissipation constraints. We focus on a network offering a single class of service and leave service differentiation for further study. Optimal flow control and routing specific to this MAC protocol are also outside the scope of the paper. Section 3 gives our assumptions in detail.

Our design draws upon conclusions from mathematical results on optimal MAC available in the literature. Our design uses three parts: an interference mitigation procedure at the demodulator (Section 4),

a dynamic channel coding with incremental redundancy, and a private MAC (Section 5). We call our protocol Dynamic Channel Coding MAC (DCC MAC).

To our knowledge, this is the first MAC protocol for ad-hoc networks that uses dynamic channel coding. Contrary to other protocols, it does not adjust the power, but the channel code.

Our design simplifies the problem of multiple access for UWB. Problems like hidden or exposed nodes naturally disappear (Section 5.2). There is no need for a separate channel for global control, as in [19], which saves system capacity. We show in Section 6 that we achieve a significant increase in network throughput at no additional energy cost, compared to protocols that do not use dynamic channel coding. Further, we solve the problem of the impossibility of carrier sensing (there is no carrier in UWB) by a combination of receiver based and invitation based choice of THSs. Also, because the source constantly adapts to the varying channel, mobility is well supported.

We implemented the protocol in the network simulator ns-2 [2]. This required us to redesign the physical layer support in ns-2, in order to account for interferences that can vary during packet transmissions. For bit error rates and transmission rates, our ns-2 implementation uses interpolation from lookup tables that we created by extensive offline Matlab experiments.

2 Related Work

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

In [19] the authors propose a CDMA [20] based MAC protocol for wireless Ad-Hoc networks called Controlled Access CDMA (CA/CDMA). Based on power control and fixed rate transmissions, it permits interfering nodes to transmit concurrently, assuming they do not destroy any ongoing transmission in their vicinity. A node can know the maximum power at which it can transmit by overhearing CTS packets transmitted by neighbors. The protocol requires the use of a separate control channel to disseminate RTS/CTS type of packets and to estimate the path loss between a source and destination pair.

Both [5] and [13] propose distributed multiple-access scheme for UWB mobile Ad-Hoc networks. The physical layers use PPM with THS to allow multiple access. However, they use a fixed channel code.

In [5] a distributed control admission function is based on the evaluation of the interference generated by each potential new link over active links, thus is essentially controlling power and not channel code. The private MAC access problem is not mentioned. Another power controlled UWB MAC protocol is described in [16].

In [13] there is no power and no channel code control. Node to node communications are initiated by

destinations as in [26, 30]. All nodes share a common signaling channel by using a common THS. A destination broadcasts an invitation on the common THS when it is ready to receive. The invitation is followed by a contention period where potential sources have to compete for starting a communication.

The IEEE 802.15 Task Group 3a is currently reviewing proposals for alternate physical layer for the IEEE 802.15.3 MAC[1]. It is based on the concept of piconets. Each member of a piconet share the same physical channel. Each piconet is controlled by a piconet coordinator which grants access to members of its piconet on a TDMA basis. In [4] a UWB physical layer based on multi-band OFDM presents a completely different paradigm of UWB modulation. In [12] the physical layer is based on PPM and the use of THS's. In this case, all members of the piconet use the same THS. In both, fixed rate channel coding is used. Note that our work is different in that we are designing a MAC specific to UWB and to very low power, and do a joint Physical Layer and MAC design. Also, our MAC is fully distributed.

Finally, in [14], a rate-adaptive MAC protocol is introduced. Rate is controlled by adjusting power, contrary to our approach, based on channel code control. Moreover, our range of available rate is much larger.

With a UWB PPM physical layer, a source and destination pair have to decide on a common THS. With receiver based scheme, the THS of the destination is used. A major issue with this scheme is that communication from different sources would collide at the destination. However, the destination has only to listen to its THS. On the contrary, with transmitter based scheme, the source uses his THS to communicate with a destination. Although this permits to avoid collisions at the destination, it forces the destination to listen to the whole sets of THS. Note that hybrid combinations of the two schemes are possible[25]. Another solution is a receiver initiated scheme as in [30].

3 System Assumptions

3.1 Model of the Physical Layer

The physical layer at the transmitter contains a channel encoder followed by a modulator (see Figure 1). The channel encoder adds redundancy to an incoming block of data bits to produce a *block of coded bits*. It is also in the channel encoder that the appropriate *encoding rate* is selected. Let L be the length of the data block in bits and R the encoding data rate currently used by the channel encoder (data bits per coded bits). The length of the coded block is $\frac{L}{R}$ bits. The modulator transforms the coded block into a form suitable for transmission on the physical medium. Similarly, at the receiver, the two symmetric components are a demodulator followed by the channel decoder. The demodulator transforms the continuous received signal from the physical layer into blocks of $\frac{L}{R}$ noisy coded bits.

These are then fed to the the channel decoder that uses them to attempt to recover the transmitted data block.

3.2 Variable Encoding with Incremental Redundancy

As mentioned in Section 1, we use variable rate channel coding. In addition, we require that coding offers *incremental redundancy*, defined as follows. Assume a system offers a choices of rates R_{low} and R_{high} . Assume a source has L data bits to send, and uses the high rate, i.e. sends a block of $\frac{L}{R_{high}}$ coded bits. Also assume that the destination is not able to decode at rate R_{high} because the quality of the channel required the lower rate R_{low} . Thus the source did a wrong choice and has to send more information to the destination. We say that the system of codes offers incremental redundancy if it is sufficient for the source to send $\frac{L}{R_{low}} - \frac{L}{R_{high}}$ additional bits to repair the error. The redundancy that must be sent is what had been saved by using the higher rate, not more. In other words, none of the already transmitted coded bits are thrown away but are used to improve decoding. Incremental redundancy is useful in our framework because we never know in advance the channel condition.

We use Rate Compatible Punctured Convolutional codes (RCPC codes) [11, 9]. These are convolutional codes [32, 20] providing variable encoding rate as well as incremental redundancy. Variable encoding rate is given by puncturing. This is the process of creating a high-rate code from a low-rate code simply by removing (i.e. puncturing) coded bits from the lowest rate block of coded bits. The advantage of such a scheme is that only one encoder/decoder pair is needed to generate blocks of coded bits of any rate. Incremental redundancy is given by the rate compatibility feature of RCPC codes. Let $R_0 = 1 > R_1 > R_2 > \dots > R_N$ be the set of rates offered by our channel coding scheme. For a given block of data bits, rate compatibility means that a block of coded bits with rate R_{n+1} is a subset of the block of coded bits with rate R_n . That is, the coded bits of the block of rate R_{n+1} are contained in the block of rate R_n . In other words, it is possible to obtain any block of coded bits by removing appropriate bits in the block of coded bits produced with the lowest rate code R_N (the so called “mother” code).

We use the particular RCPC codes from [9]. We have 31 possible rates:

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

Rate 1 is achieved by sending the block uncoded data bits. Rates $8/31$ to $8/9$ are achieved by puncturing the mother code of rate $8/32$. The rates $1/5$ to $1/10$ are obtained by *nesting* additional coded bits to the coded bits encoded with the mother code. Obviously, these rates are still rate-compatible with the higher ones.

Finally, we also use an interleaver [20]. An interleaver pseudo-randomly interleaves bits at the output of the channel encoder; de-interleaving is done at the input of the channel decoder. The goal is to make the

noise added by the channel independent from coded bit to coded bit, which improves the performance of the channel decoder.

3.3 Modulation and Power

Theoretical results on the wide-band communication channel [28, 31] show that the optimal signaling consists of sending very short pulses of maximum allowed energy. We assume that the pulses generated by our physical layer have a width T_p , and the peak power $P_{peak} = E_{peak}/T_p = 0.28$ mW, the limits allowed by regulations and hardware constraints [12].

As previously described, the modulator uses binary Pulse–Position Modulation(PPM) [33, 7] to transmit a block of coded bits. A coded bit equal to 0 is sent as a pulse transmitted at the beginning of the chip, whereas for a bit of 1 the pulse is offset by a fixed δ , smaller than the chip time $T_c = 0.2$ ns¹. Note that the simple repetition coding scheme of [33] is replaced by the sophisticated channel coding scheme that we use.

An active source sends one pulse per frame; in order to avoid peaks in the power spectrum a pulse must not be sent in the same chip in every frame. Hence, a Time Hopping Sequence (THS) is used to determine in which chip to transmit. The number of chips per frame PRP is dictated by hardware constraint as well as rate and multi–user interference (MUI) consideration. A commonly found value is $PRP = 10 T_c$. Higher values of PRP reduce the bit rate and the radiated power, but decrease the level of MUI.

The radiated power P_{rad} is defined as the average power during transmission; it is $P_{rad} = \frac{P_{peak}}{PRP T_c}$. We are interested in a low power UWB network, with a radiated power $P_{rad} = 1\mu$ W.

As the radiated power depends linearly on $1/PRP$, our goal is achieved by taking $PRP = 280$. With this value, the maximum rate available to one source is $(PRP T_c)^{-1} = 18$ Mb/s. We show in Section 4 that this rate goes down to 6Mb/s when the distance between source and destination is 30 meters. In Section 6 we find average throughputs (including the overhead of the MAC protocol) per source of several Mb/s.

3.4 Time Hopping Sequences

A THS is a periodic sequence $[x_1, x_2, \dots, x_n]$ where x_i is the position of the chip to be used for transmission in the i th frame. Though deterministic, the sequence x_i has to look like a sequence of independent, random variables with each x_i uniformly distributed over the range of chip positions $[0, PRP - 1]$. Also, different THSs used by different sources should appear to be independent.

¹We have that $\delta < T_p < T_c$

We use the following method to satisfy these requirements. Every user has an identical pseudo-random number generator (PRNG) and a unique identifier (its MAC address). Communication uses either public or private THSs. The public THS of user with MAC address A , called $\text{THS}(A)$, is the THS produced by the PRNG with seed = A . The private THS of users A and B , called $\text{THS}(AB)$ is the THS produced by the PRNG with seed = the number whose binary representation is the concatenation of A and B . The private MAC protocol (Section 5.2) governs which THS is used. Note that a node can always compute the THS used by a potential source. There is no protocol for THS distribution.

The THS is reset at its seed value for every packet transmission. In addition, there is a predefined THS used for broadcast.

THSs share similarities with the spreading codes used with CDMA techniques, but they are not the same: the determination of time hopping sequences is almost trivial, whereas finding good spreading codes for CDMA is difficult. In particular, it is well known that due to asynchronicity, a lot of attention has to be given to ensure that CDMA sequences have a very low cross-correlation[20]. The difference stems from the fact that, with CDMA, a node sends a signal in every chip, whereas with PPM, pulses are sent only in one chip out of PRP, in average.

3.5 Synchronization

Synchronization of the physical layer is required only between a source and a destination (for unicast), and is performed at the destination only. This is similar to IEEE 802.3 and 802.11. There is no global synchronization.

The general method for achieving synchronization [20] relies on the presence of a synchronization preamble at the beginning of a packet. This synchronization preamble is known by the destination of the packet. In order for the destination to synchronize it needs to detect the synchronization preamble. When this preamble is detected by the destination, the demodulation of the received signal starts. When established, synchronization must be maintained.

In [17], the authors propose an efficient method for detecting synchronization patterns over UWB channels. The time required for achieving synchronization between completely de-synchronized systems is of the order of tens of μs whereas the time to refresh is of the order of $10\mu\text{s}$.

Hence, we assume that synchronization can be maintained over the whole duration of a packet, and can be re-established for each data packet.

We assume that a node is able to listen both on the broadcast THS and on one other THS. This means that it is able to acquire synchronization on either of the THSs. However, a node can receive on one THS at a time.

3.6 Other Assumptions

We focus on a mobile ad-hoc scenario and assume that a node can either send or receive, but cannot do both at the same time. A node can receive from only one source at a time. However, a node can listen to several THSs at the same time. Extension to more powerful nodes (like cellular network base stations) that can do multiple receptions/transmissions at a time is for further study.

4 Interference Mitigation

First, in Section 4.1 we estimate the size of the exclusion region in our specific case. We find that the exclusion region cannot be neglected, thus we need to do something about interference from the exclusion region. We then present in Section 4.2 our interference mitigation scheme which is very simple, efficient and does not require any signaling. In Section 4.3, we verify on theoretical examples that the performance is good – a final evaluation is by simulation of the complete MAC protocol in Section 6.

4.1 Size of The Exclusion Region

The size of the exclusion region can be approximated by analyzing the simple network, depicted on Fig. 2, which we call *the near-far scenario*. More precisely, an estimate of the radius of the exclusion region is the value of the distance d below which it is more advantageous to have a temporal exclusion than to accept interference.

To this end we need to compute the rate achievable by source $S1$: a classical method is to use the Win-Scholtz UWB physical model as presented in [33]. However, this model has two major drawbacks:

- Interference from the other users is assumed Gaussian. This is true only in the case of a large number of equally powerful users. In the case of a small number of users, the interference is less regular [6, 7, 8], and the achieved performance is lower than what the model predicts.
- It assumes simple repetition codes, whereas we use more powerful codes [11, 9] which have higher performance.

Therefore we use a more elaborate model, as follows. We model a non-Gaussian interference as in [6, 7, 8]. For a given link length l and interferer distance d , we test one of the several possible codes. For each code we select a random data to be transmitted, a random interference, and calculate the achieved rate as a function of the bit error rate of the experiment and the rate of the code, using the formulas in [9].

Another important issue in UWB physical layer modeling is synchronization. While a sender is synchronized with a receiver, an interferer is not, hence its contribution at the output of the correlator [33] is lower than at the input. Each interferer resynchronizes with its own destination for every packet (Section 3.5). This makes the desynchronisation jitter very variable. We model it as a uniform random variable.

We implemented this model in Matlab and obtained the following results.




Figure 2: Near-far scenario: Source S1 sends data to the destination D1, and S2 to D2. The distance between a source and a destination is l and the distance between a source and the corresponding interferer is d . For large d it is optimal to have both links transmitting at the same time; for small d , both S1 and S2 send half of the time.

As in [21, 22], we compare two schedules: TDMA and All-at-once. TDMA protocol assumes only one node sends at a time, and All-at-once assumes both send all the time. Senders use the maximum power when transmitting, and adapt codes, hence rates, according to the signal-to-noise level at the receiver. It is easy to see that due to symmetry, the optimization gives the same rates to both links. In the TDMA case, this means each link transmits half of the time. We analyze the performance of the two scheduling for fixed l and variable d . As one expects, TDMA performs better for small d and All-at-once for large d . The estimate of the exclusion region is the value of d at which both schedules perform equally.

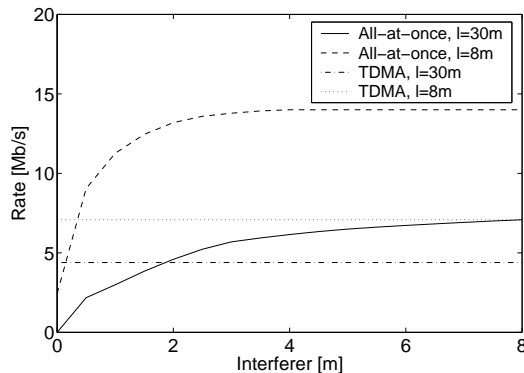


Figure 3: Rates achieved in the near-far scenario, for All-at-once and TDMA schedules, and link lengths $l = 8\text{m}$ and $l = 30\text{m}$. The rate is given on the y-axis, and the distance d from the interferer to the destination is given on the x-axis.

The results of our analysis are on Fig. 3, for two link lengths, $l = 4\text{m}$ and $l = 15\text{m}$. We see that for $l = 8\text{m}$ and for interferer distances smaller than 0.4m , the performance of All-at-once schedule becomes worse than TDMA, and decays fast with decreasing interferer distance. We conclude that the size of the exclusion region in this case is 0.4m . For link length of 15m , the size of the exclusion region is 2m .

This exclusion region is rather large, and a presence of any interferer inside, who would transmit at the same time, would drastically deteriorate the rate of the link.

4.2 Our Interference Mitigation Scheme

It is clear from the above analysis that an interfering user, reasonably close to a destination, will significantly lower the transmission rate. We are interested in finding a mechanism that will enforce exclusions without an explicit signaling such as RTS/CTS. We introduce an interference mitigation scheme on the UWB physical model of [33]. The goal of the scheme is to reduce the interference from nodes in the exclusion region of the destination. Our interference mitigation is motivated by [24], and is based on the concept of erasures.

Suppose that A is the amplitude of the signal, A^I the amplitude of the interference, and the A_n^N is the amplitude of the noise in the n -th chip of a codeword. Let X_n be the symbol a source and let X_n^I be the symbol an interferer is sending in n -th chip (which can be either $+1$ or -1). The bit error rate of the code depends on the signal-to-interference plus noise ratio [33]

$$\text{SINR} = \frac{\sum_{n=1}^N AX_n}{\sum_{n=1}^N A_n^N + A^I X_n^I}. \quad (1)$$

Due to the time-hopping sequences, the probability of a symbol level collision (both a source and an interferer choose the same chip to transmit a bit) is low. However, if an interferer is very close, the interfering signal A_n^I will be very strong. This means that the average total noise at the input of the decoder will be very high, and the performance will be weak.

It has been shown through indoor channel measurements that variations in the received signal power are typically caused by shadowing rather than fast fading [24]. Hence, a receiver can track the received signal's strength during several chips, and have its average estimate over the time. The average estimate corresponds to the level of the useful signal, since the average amplitude of the noise and interference is zero. If in a given chip the received signal is by far larger than expected, the receiver concludes there was a collision at the symbol level. More precisely, if for a given bound B , the received signal's amplitude is outside of the interval $[-B, +B]$, we declare an erasure and disregard the symbol from that chip. An erasure can happen both in the case of a collision with a high power interfering symbol, or in the case of a high white noise sample. The total signal-to-interference ratio in this case is

$$\text{SINR} = \frac{\sum_{n=1}^N AX_n E_n}{\sum_{n=1}^N (A_n^N + A^I X_n^I) E_n}, \quad (2)$$

$$E_n = I(|AX_n + A_n^N + A^I X_n^I| < B) \quad (3)$$

where E_n is the event of the erasure in the n -th chip, and I is the indicator function. The total signal strength will be lower by the fraction of erasures, but the destructive interference will be lower, and in the case of high interference (A^I very large), (2) is larger than (1).

The optimal value of the bound B depends both of the power of the interferer and of the white noise. On one hand, too large a B is equivalent to the case without erasures. On the other hand, a small B will declare too many erasures, and the decoding will become hard-decoding, which is known to have worse performance than the soft-decoding. Our goal is to set B such that the erasures are declared only due to collisions, and not due to the white noise. In practice, the smallest SINR our codes can operate on are -10 dB. This means that the average amplitude of the noise is approximately 3 times larger than the amplitude of the signal. We set $B = 5A$. We verify that on a link without interference, and with SINR=-10 dB, we will have on average 10% of erasures due to the Gaussian noise, which does not deteriorate much our performance. In most cases we operate on a much higher SINR, and the amount of erasures due to noise from outside our system is negligible.

4.3 Performance Evaluation of Interference Mitigation

We first evaluate the performance of our interference mitigation on the near-far scenario. We vary the distance of the interferer and take random desynchronisation jitters. The results are depicted on the Fig. 4.

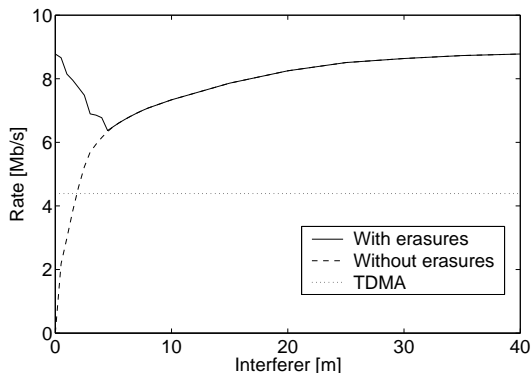


Figure 4: Rates achieved in the near-far scenario with interference mitigation, for All-at-once and TDMA schedules, and the link length $l = 30\text{m}$. The rate is given on the y-axis, and the distance d from the interferer to the destination is given on the x-axis.

In this example, there is only one interferer. Since the value of $PRP = 280$ is rather large, the probability of symbol-level collisions is small. Still, as we have seen on Fig. 3, when we do not apply interference mitigation, the rate is low. From Fig. 4 we see that the interference mitigation successfully eliminates this problem. The All-at-once schedule with interference mitigation always has better performance than the TDMA schedule, and we successfully address interferences from the exclusion region. Furthermore, we see that the performance with interference mitigation is significantly better than without it, even when the interferer is outside of the exclusion region. In the case of $l = 30\text{m}$, this is true for interferer distances up to 4m. For larger distances, the two approaches are equal. The small gap in the rate is due to the asynchronism of the interferer and the receiver, hence the receiver is not always able

to successfully detect and erase symbol level collisions.

We next evaluate the performance of the interference mitigation in the presence of several interferers. We consider a single link, and a variable number of interferers very close to the destination of the link. In order to estimate the effect of the erasures themselves on the code performance, we assume the interferers are very close, hence despite asynchronism, the probability of successful erasures is high.




Figure 5: Rates achieved, with interference mitigation, on a link, as a function of number of cancelled interferers, for link lengths of 5m, 15m and 30m. On the x axis a number of near-by interferers is given (we here assume all these interferers are always successively cancelled), and on the y axis is the rate achieved.

As we explained in Section 3.3, we have a rather large PRP equal to 280. In the presence of one very close interferer, this value of PRP leads to about $P_{\text{erasure}} = 0.18\%$ of symbol level collisions, which is very low. In the case of n interferers, the probability that at least one will collide, given that the source is transmitting, is $1 - (1 - P_{\text{erasure}})^n$, and it grows sub-linearly with the number of interferer. Finally, lower-rate channel codes are much more resilient to erasures than high-rate ones. All this combined lead to an extremely high tolerance of our scheme to erasures, as can be seen on Fig. 5. Even for a very large number of interferers (50), which we hardly find in practice, the rate drops at most by 15%.

5 Protocol Description

As described in Section 4, interference is either mitigated or should be allowed to exist. Furthermore, it is optimal to use as much power as allowed by the budget, when sending. Thus, the function of our DCC MAC protocol is reduced to (1) manage the channel code dynamically in order to adapt to varying interference and other channel conditions and (2) control access from several sources to one same destination. These two functions are described in the rest of this section.

5.1 Dynamic Channel Coding and Incremental Redundancy

While channel coding and Automatic Retransmission reQuest (ARQ) (hybrid-ARQ) [18] are part of most wireless MAC proposals, doing them efficiently is particularly important in an environment where interference is very variable.

To this end, channel coding is constantly adapted to the highest rate code that still allows decoding of the data packet at the receiver. We include a safety margin (i.e., we use a more powerful code than required) to reduce the probability of retransmission when channel conditions deteriorate. If conditions worsen significantly and decoding fails despite the safety margin, additional information is transmitted, until the packet can be decoded or no more redundant information is available and the transmission fails.

Our hybrid-ARQ protocol performs the following steps to transmit a packet from S to D .

- S adds a CRC to the packet content and encodes it with the lowest rate code.
- S then punctures the encoded data (i.e., removes specific bits from it) to obtain the desired code rate and sends the packet. The punctured bits are stored in case the decoding at D fails.
- Upon packet reception, D decodes the data and checks the CRC. If decoding is successful, an acknowledgement is sent back to S . Otherwise, a negative acknowledgement (NACK) is sent (Figure 6).
- As long as S receives NACKs, further packets with punctured bits (each time up to the size of the original packet) are sent, until transmission succeeds or no more punctured bits are available. In the latter case, S may attempt another transmission at a later time (see Section 5.2).

For good performance and a short transmission delay, sending redundant information should rarely be necessary. Hence, it is more important that the transmission succeeds directly without having to send additional punctured bits than using the highest-rate code possible.

When nodes communicate for the first time, it is necessary to bootstrap the code adaptation mechanism. The first data packet is encoded with the most powerful (lowest rate) code R_N . From this, the receiver has to determine the optimum code the sender should use for the next transmission.

Decoding of the data packet with channel code R_N is performed by step-wise traversal of the trellis of the Viterbi decoder [32]. 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 $R_i > R_N$ differs from that of the actual channel code, code R_i can be eliminated. Because of the rate compatibility feature of RCPC codes, this allows to also eliminate all codes $R_j > R_i$. 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 find that the heuristic of using a channel code of R_{i+2} if the highest possible code is R_i performs sufficiently well. The code R_{i+2} is indicated to the sender in the acknowledgement. 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. In case packet transmission was unsuccessful on the first attempt and additional redundancy had to be sent, the receiver can determine the highest possible rate in the same way as during bootstrap, as soon as the packet can be decoded.

The sender determines the code to use as follows: in case the code indicated by the receiver is higher than the current code (i.e., R_{i+2}), the sender does not directly switch to the new channel code but decreases its code index by one. Otherwise, if a more powerful code is necessary, the sender switches directly to this code. The sender maintains a cache of channel codes for a number of receivers. If the sender does not communicate with a receiver for a certain amount of time, the corresponding cache entries time out and the sender bootstraps code selection with code R_N as described above.²

5.2 Private MAC

Overview The goal of the private MAC protocol is to enforce that several senders cannot communicate simultaneously with one destination. This is traditionally solved in our context by a carrier sensing scheme. This is not possible with UWB, as there is no way to tell noise from transmission unless a node actively decodes (there is no carrier to listen to). A simple fix would be to use ALOHA, but its performance is not acceptable except for broadcast traffic. We solve the problem by a combination of receiver-based and invitation-based selection of THSs. This is inspired by the similar mechanisms used in CDMA described in Section 2. Competition to a destination uses the permanent THS of the destination, but an established communication uses the THS private to a source-destination pair (Section 3.4).

In the rest of this section we explain the details, in the form of protocol walkthroughs.

Successful Transmission A successful data transmission consists of a transmission request by the sender, a response by the receiver, the actual data packet, and an acknowledgement.³ Assume a node B has data to transmit to a node A , and no other node is sending data to A (Figure 6). The idle node A listens on its own THS. When node B wants to communicate with A , it sends a transmission request on A 's THS. The channel code uses the lowest possible rate R_N , so that all nodes within reach that want to talk to A may overhear it. A answers with a reply packet using the THS private to A and B , and the channel code dictated by the channel code assignment procedure of Section 5.1. By doing so, A

²For our simulations we use a cache size of 10 and a timeout interval of 1 second.

³Note that this scheme is different from RTS/CTS in that it only reserves a per-destination collision domain.

indicates to B that it is idle. When B receives the reply, it starts with the transmission of the data packet on the code private to A and B . After the transmission, B listens for a feedback sent by A on the private THS with the smallest rate channel code; depending on the feedback, B may have to send more data until A can decode (Section 5.1). If no feedback is received, the sender B retries transmission after a random backoff explained later, up to a certain retry limit. While the sending of additional information is done during the same communication attempt, a retransmission of the packet requires the sender to perform a new request transmission to the receiver after the backoff.

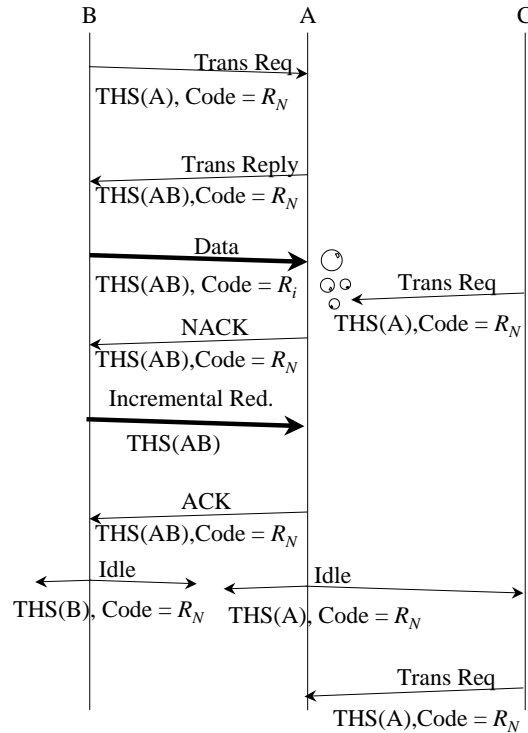


Figure 6: Private MAC: B sends to A , C attempts to send to A and is deferred. No collision occurs because B and C use different Time Hopping Sequences.

Deferred Transmission Assume that a node C wishes to communicate with A while A is receiving a packet from B . It sends out a request on A 's THS; this may create some interference but will usually not disrupt the private communication between B and A . C then switches to A 's THS and listens for the idle signal.

After a transmission (either successful or unsuccessful), both sender and receiver issue a (short) idle signal on their own THS to inform other nodes that they are idle. The idle signal is sent on the smallest rate code. When C hears A 's idle signal, it waits for a random, small backoff time and transmits a request again (Figure 6).

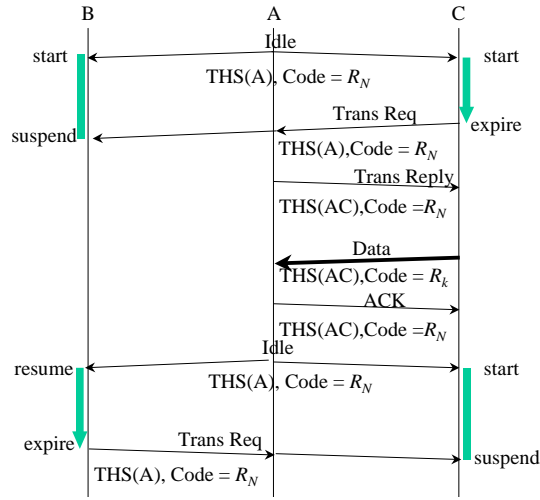


Figure 7: Private MAC, showing state of backoff timer. B and C compete for sending to A . C initially draws a shorter backoff timer and gains access to A .

Race Condition Now we move to the more general scenario with competing senders. Nodes maintain a per-receiver contention window similar to the general contention window used for 802.11 (but is per receiver). The first transmission attempt to a receiver is immediate; if the receiver is idle there is no access delay. If the intended receiver is busy, the transmission attempt fails (without collision). The sender now listens on the receiver's THS and computes a backoff timer to a random value between 0 and the maximum contention window cw_{\max} (but the timer is not started yet). When an idle signal is received, the backoff timer is started; if expires without the node overhearing another transmission request, a request is sent. Else, i.e. if a transmission request is overheard, the node defers transmission, pauses the backoff timer until an idle is heard. See Figure 7 for an illustration.

The value of cw_{\max} is initially set to a number of the order of $70\mu s$, i.e. a small multiple of the time it takes for a destination to acquire synchronization (Section 3.5). Whenever an attempt to access a receiver after the expiry of the backoff timer is unsuccessful, cw_{\max} is doubled and a new backoff time between 0 and cw_{\max} is chosen. Upon successful data transmission, cw_{\max} is reset to its initial value.

When any packet is sent, another timer ($T1$) is computed and immediately started. The value is of the order of the transmission of a largest size packet at the smallest rate, of a receiver response, and an acknowledgement at the smallest rate. This is an upper bound on the time of a successful transmission, including possible incremental redundancies. When a transmission request is sent and not acknowledged, it is most probably because the destination is busy and ignored the request. Normally, this condition ends with the reception of an idle signal as just explained. If, however, $T1$ expires, this may mean that the request was ignored and the destination is not busy, in which case there is no point waiting for an idle signal. Therefore, upon expiration of $T1$, the source waits for a random backoff time as before and a transmission request is re-attempted.

True Collision Assume that A , B and C are idle, and B and C , by chance, send a transmission request to A roughly at the same time. A collision occurs at A only if the synchronization preambles of both requests overlap. Otherwise, A synchronizes to one of the requests, and the other request is an interference. Indeed, though both requests are sent on the same THS, they are not synchronized. Thus the contention window is bounded by the synchronization time (i.e. $10 \mu\text{s}$). This is less than the $51.2\mu\text{s}$ of contention window on Ethernet at 10 Mb/s and is thus very low. Note that we achieve a small contention window without carrier sensing.

When a true collision occurs, the timer $T1$ expires; after expiration of $T1$, the backoff timer is used as after reception of an idle signal.

Other Issues Broadcast is supported by means of well known unique THS and synchronization sequence dedicated to broadcast, and otherwise uses a simple ALOHA. The code is the lowest rate. Note that there is never a collision between broadcast packets and unicast packets, since they use different THSs.

A source arbitrates between sending and receiving by using a fair queuing scheduler. Arbitration is based on a simple fair queuing scheduler in our ns-2 implementation used for this paper. An improvement for further study is to apply the ideas of CDMA/HDR, which consist in picking the best destination at any single point in time.

The hidden terminal problem affect protocols that require that, by listening to the medium, a source be able to detect a collision at the destination. This is avoided here because we do not use carrier sensing. The hidden terminal problem is solved by implemented a RTS/CTS response in 802.11, which in turn causes the exposed terminal problem. An exposed terminal is one that is prevented from sending because it hears many CTSs from non-coordinated parts of the network. The problem does not exist because, with our protocol, because there it is naturally made of a collection of non cooperating private MAC instances. When a node is prevented from transmitting, it must be that it competes with other sources for the same destination, or that the destination is busy sending.

6 Simulations

The feasibility and performance of the proposed MAC layer is analyzed by means of simulation. To this end, the well-known network simulator ns-2 has been significantly extended by incorporating a model for an UWB physical layer and an UWB-specific propagation model as well as MAC layer protocols (including the proposed one). In particular, since interference plays an important role much attention has been paid to accurately model radio interference. Further details of the ns-2 implementation are described in [3].

In the following sections we describe setup and results of a number of simulations to analyze the performance of the DCC-MAC. In particular, we investigate the performance of our protocol under mobility and in near-far scenarios, where the receiver is located closer to interferers than the sender, and we compare it to two exclusion-based MAC protocols.

- **Power Control MAC** - We implement a simplified version of CA/CDMA [19]. While it abstracts from many protocol details, it captures the main aspect of adjusting the power instead of the channel coding. The coding is fixed to the highest-rate channel code. We define a minimum signal-to-interference ratio that is necessary to achieve a given probability of error, plus a safety margin. We then adapt the transmission power of the packet in order to achieve the desired SIR. If this is not possible due to the power limit at the sender, the sender attempts another transmission after a random backoff.
- **CSMA** - All nodes use the same code and if nodes are close enough they are able to overhear transmission from any other node. When one node starts transmitting, other nodes within radio range receive the packet and stay silent during the transmission. All the nodes but the destination discard the packet. The number of collisions is negligible, since synchronization is done on a very small time scale. If a node has a packet to transmit while another node is sending, it does backoff in the manner of IEEE 802.11.

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.

6.1 Goals

Thus far, we analyzed the basic properties of our protocol in very simple scenarios by means of Matlab simulations. The main goal of the ns-2 simulations is to investigate if our protocol works as expected under more realistic network conditions.

Through the comparison with other MAC proposals, we verify the two basic assumptions of our design described in Section 1.2: that a sender should always transmit with maximum power when sending and that a receiver should adjust the channel coding to cope with interference. The CA/CDMA-like protocol uses variable power, has a fixed channel coding, and allows interference. The CSMA-like protocol uses maximum power, uses dynamic channel coding, and controls interference through exclusion. By comparing the two MAC protocols with our proposal, we perform a factorial analysis of the impact of these two assumptions.

6.2 Performance Metrics

We use the total number of packets the network successfully transports (i.e., that can be decoded at the receivers) in 10s as the primary metric in our paper. From this we can compute the total throughput of the network. Dividing the total throughput by the number of links, we obtain the average throughput per flow.

In all of the measurements we simulate UDP traffic. To analyze the effect of our MAC layer design choices in isolation, we use direct links between sources and destinations, hence there is no relaying.

6.3 Numerical Results

We consider three types of network topologies. For each scenario we perform simulations for various parameters of the topologies, and we obtain mean values and confidence intervals for the average per-flow rate.

6.3.1 Generalized Near-Far Scenario

Our generalized near-far scenario is shown in Figure 8. We consider networks with 2 to 32 links. The distance between sender and receiver varies from 1m to 40m.




Figure 8: Generalized near-far scenario: Source S_i sends data to the destination D_i , causing interference to D_{i-1} and D_{i+1} . The distance between adjacent links is 1m, and the distance l of sender and receiver varies.

Simulation results are depicted in Figure 9. To facilitate comparison, we also give results for a single sender-receiver connection without any interference. For very short distances, the DCC-MAC is equivalent to the optimum rate of a single connection. Over longer distances, it is up to 20% from the optimum, but nevertheless outperforms the other solutions. For very long distances, DCC-MAC rates again approach the optimal rate as the low receive power allows to cancel more of the high-power interference.

While the CA/CDMA-like scheme is close to DCC-MAC over short distances, its performance quickly deteriorates as the communication distance increases. We attribute this to the fact that our simple ver-

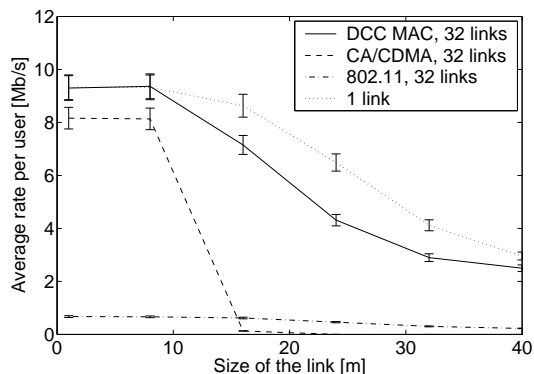


Figure 9: We consider a near-far scenario from Figure 8 with 32 links. The length of the link is given on the x-axis and the average rate per user is given on the y-axis.

sion of the protocol does not include the elaborate calculation of the interference margin as the original CA/CDMA proposal. Exclusion-based IEEE 802.11 exhibits the worst performance of all three protocols since in this configuration it does not allow parallel transmissions. The improvement in SNIR and the resulting higher channel code rates cannot compensate for the loss in transmission time due to exclusion.

6.3.2 Random Scenario

In this scenario, nodes are randomly placed on a square surface of $20\text{m} \times 20\text{m}$. Source-destination pairs are randomly chosen such that each node is either a source or a destination of exactly one link. The number of nodes varies from 2 to 128.

From the simulations results shown in Figure 10, we observe that the rate achieved by the DCC-MAC is comparable to IEEE 802.11, and about 20% better than the CA/CDMA-like scheme for a small number of nodes. For higher numbers of nodes, the performance of both IEEE 802.11 and CA/CDMA drops drastically. The DCC-MAC maintains the same rate due to its interference mitigation capabilities and the dynamic code adaptation that becomes important when the number of nodes (and therefore interference) is high.

6.3.3 Cluster Scenario

In both of the previous scenarios, our scheme successfully copes with interference and its performance is very close to that of a single link. The third scenario is an extreme and the effect of the interference is clearly visible in the performance of the DCC-MAC. We consider a clustered network: nodes are divided in two groups, as shown in Figure 11. The diameter of the cluster is fixed at $d = 8\text{m}$ and we vary the distance between the clusters and the number of nodes in the network.

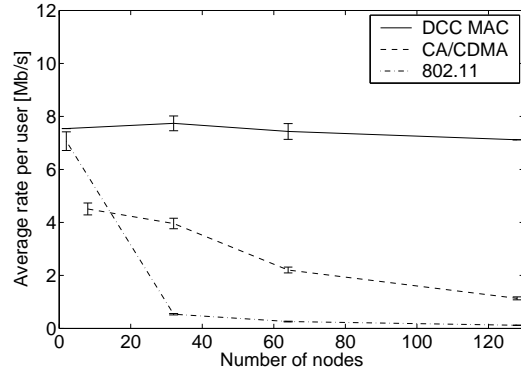


Figure 10: Random scenario with nodes placed on $20\text{m} \times 20\text{m}$ square. The number of nodes is given on x-axis, and the average rate is given on the y-axis.




Figure 11: Cluster scenario: The network is composed of two clusters, each containing half of the nodes. Nodes are randomly placed within the area of the cluster. Source-destination pairs are chosen from different clusters. The diameter of each cluster is d and the distance between clusters is l .

The results are given in Figure 12 and an example with one link is given as a reference. We see that with communication distances of approximately 8m and a high number of nodes in a cluster (16 or 32), the interference significantly decreases the rate of DCC-MAC. Nevertheless, even in this extreme case, the average effective flow rate is above 1Mb/s for communication distances of up to 40m.

6.4 Impact of Mobility

Finally, we analyze the impact of mobility on the performance of the network. We consider the random scenario from Section 6.3.2 and let nodes move according to the random way-point model. Nodes speed varies between 2m/s and 10m/s with 0 pause time. To isolate the effect of mobility on the MAC protocol, we do not use multi-hop routing.

Comparing the achieved network throughput in the mobile scenario given in Figure 13 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, similar to the case of a change of the level of interference. The adaptation is sufficiently fast compared to the node speed to prevent a degradation of the rate.

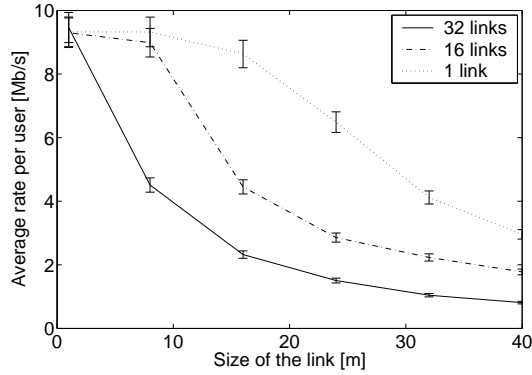


Figure 12: Cluster scenario from Figure 11 with a cluster diameter of $d = 8\text{m}$. We consider the case of networks with 1, 16, and 32 communication pairs. The distance between the clusters is given on the x-axis and the average rate is given on the y-axis.




Figure 13: Random mobile network: Surface area of $20\text{m} \times 20\text{m}$ and mobility according to the random waypoint model. On the x-axis is given the number of nodes in the network and on the y-axis is the average rate per user.

7 Conclusion

We have presented a MAC protocol for very low power UWB that is closely coupled with the physical layer. We have assumed 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, 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. Other exclusion mechanisms such as TDMA, CSMA/CA or OFDM 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 used PPM modulation. Other, non coherent modulation schemes are also discussed for UWB [24]. It seems that our MAC protocol would apply with little change to such modulations, but this is also for

further study.

References

- [1] IEEE 802.15.3 MAC standard, available at <http://www.ieee.org>.
- [2] The network simulator ns-2, <http://www.isi.edu/nsnam/ns/>.
- [3] Ultra-wide band MAC and PHY layer for ns-2, <http://icapeople.epfl.ch/widmer/uwb/ns-2/>.
- [4] A. Batra and al. Texas Instruments/Intel and al. Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 267r6, September 2003.
- [5] 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.
- [6] 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.
- [7] G. Durisi and G. Romano. On the validity of gaussian approximation to characterize the multiuser capacity of UWB TH PPM. In *IEEE Conference on Ultra Wideband Systems and Technologies*, pages 151–161, 2002.
- [8] A. Forouzan, M. Nasiri-Kenari, and J. Salehi. Performance analysis of time-hopping spread-spectrum multiple-access systems: uncoded and coded schemes. *IEEE/ACM Transactions on Wireless Communications*, 1(4):671–681, October 2002.
- [9] 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.
- [10] 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 '03*, March 30, April 3 2003.
- [11] J. Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Transactions on Communications*, 36(4):389–400, April 1988.
- [12] D. H elal and P. Rouzet. ST Microelectronics Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 139r5, July 2003.
- [13] 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 (IWUWBS'03)*, June 2003.

- [14] 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.
- [15] L. Kleinrock and F. A. Tobagi. Packet switching in radio channels: Part i—carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Transactions on Communications*, 23(12):1400–1416, December 1975.
- [16] S. Kolenchery, J. Townsend, and J. Freebersyser. A novel impulse radio network for tactical military wireless communications. In *IEEE Military Communications Conference (MILCOM'98)*, pages 59–65, 1998.
- [17] J. Kusuma, I. Maravic, and M. Vetterli. Sampling with finite rate of innovation: Channel and timing estimation for UWB and GPS. In *IEEE International Conference on Communications (ICC'03)*, pages 3540–3544, 2003.
- [18] S. Lin and D. Costello. *Error Control Coding: Fundamentals and Applications*. Prentice–Hall, Inc., Englewood Cliffs, NJ, 1983.
- [19] A. Muqattash and K. Marwan. CDMA–based MAC protocol for wireless ad hoc networks. In *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC'03)*, pages 153–164, June 2003.
- [20] J. G. Proakis. *Digital Communications*. McGraw–Hill, New York, NY, 4th edition, 2001.
- [21] B. Radunovic and J. Y. Le Boudec. Joint scheduling, power control and routing in symmetric, one-dimensional, multi-hop wireless networks. In *Proceedings of WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, INRIA Sophia-Antipolis, France, March 2003.
- [22] B. Radunovic and J. Y. Le Boudec. Optimal power control, scheduling and routing in UWB networks. Technical Report IC/2003/61, EPFL-DI-ICA, October 2003.
- [23] S. Raj, E. Telatar, and D. Tse. Job scheduling and multiple access. DIMACS Series in Discrete Mathematics and Theoretical Computer Sciences, 2003.
- [24] Y. Souilmi, R. Knopp, and G. Caire. Coding strategies for UWB interference-limited peer-to-peer networks. In *WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, March 2003.
- [25] 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.
- [26] F. Talucci, M. Gerla, and L. Fratta. MACA-BI (MACA by invitation)-a receiver oriented access protocol for wireless multihop networks. In *IEEE PIMRC'97*, pages 435–439, 1997.

- [27] M. Teener and R. Gvozdanovic. FDDI-II operation and architectures. *Proceedings of IEEE Conf. on Local Computer Networks*, pages 49–61, October 1989.
- [28] E. Telatar and D. Tse. Capacity and mutual information of wideband multipath fading channels. *IEEE Transactions on Information Theory*, 46(4):1384–1400, 2000.
- [29] S. Toumpis and A. Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE/ACM Transactions on Wireless Communications*, 24(5):736–748, May 2003.
- [30] A. Tzamaloukas and J. Garcia-Luna-Aceves. A receiver-initiated collision-avoidance protocol for multi-channel networks. In *INFOCOM'01*, volume 1, pages 189–198, 2001.
- [31] S. Verdú. Spectral efficiency in the wideband regime. *IEEE Transactions on Information Theory*, 48(6):1319–1343, 2002.
- [32] A. Viterbi. Convolutional codes and their performance in communication systems. *IEEE Transactions on Communications*, 19(5):751–772, October 1971.
- [33] 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.