

DCC-MAC: A Decentralized MAC Protocol for 802.15.4a-like UWB Mobile Ad-Hoc Networks Based on Dynamic Channel Coding

Jean-Yves Le Boudec, Ruben Merz, Božidar Radunović, Jörg Widmer

EPFL, CH-1015 Lausanne, Switzerland

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

ABSTRACT

We present a joint PHY/MAC architecture (DCC-MAC) for 802.15.4a-like networks based on PPM-UWB. Unlike traditional approaches it fully utilizes the specific nature of UWB to achieve high rates at low protocol complexity. It is the first MAC protocol that adapts the channel code (and thus the bit rate) to interference from concurrent transmissions instead of enforcing exclusion. In order to avoid a complex mutual exclusion protocol at the MAC layer, we propose an interference mitigation scheme. The scheme is based on a modification of the physical layer that cancels much of the interfering energy, in particular from nearby interferers. We further use dynamic channel coding to combat the remaining interference. Sources constantly adjust their channel codes to the level of interference and send incremental redundancy as required. Contention between sources sending to the same destination is solved by a “private MAC” protocol that involves only the nodes that want to talk to the same destination. The private MAC does not use any common channel; this avoids the issues of hidden and exposed terminals altogether. We show by simulation that our MAC protocol fully satisfies the application requirements of 802.15.4a in terms of link lengths, rates and mobility. We further show that it achieves a significant increase in network throughput, compared to traditional MAC protocols like 802.15.4, that are separated from the physical layer.

Keywords: Ultra-wide band, 802.15.4a, Medium Access Control, Dynamic Channel Coding

1. INTRODUCTION

1.1. Low Power UWB Networks

Emerging pervasive networks assume the deployment of large numbers of wireless nodes, embedded in everyday life ob-

The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

jects. An example of such a trend is IEEE 802.15.4 standard [4] that covers a large range of applications, varying from industrial monitoring, automation and control to connecting PC peripherals and home appliances. Another standard for low-power wireless networks with extended application requirements and an alternative physical layer is currently under discussion within the framework of the IEEE 802.15.4a study group. A new application request [1] requires long distant and low-power links that have ranging capabilities and are robust against interference and mobility.

Ultra-wide band (UWB) is a radio technology that has a potential to satisfy these requirements, and is currently under the consideration of the 802.15.4a study group. UWB is characterized by a broad radio spectrum; more precisely, according to Federal Communications Commission regulations, UWB has a bandwidth that is larger than 20% of the carrier frequency or a bandwidth equal to or greater than 500Mhz. Due to a high frequency, such a system has a very high time resolution, thus excellent ranging capabilities.

The 802.15.4a standard targets low-power networks. In order to increase the lifetime of a network, it is important to keep the radiated energy low. Also, an 802.15.4a network is envisaged to have a very large number of nodes. 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. We are interested in a *very low power* physical layer meaning that the radiated energy per node does not exceed $1\mu\text{W}$ ($= -30\text{dBm}$). With the currently planned technology, it is possible with such very low-power to achieve rates of several Mbps per source at distances to the order of tens of meters and up to one hundred meters (Section 4). These rates are reduced when several nearby UWB sources transmit concurrently. In this paper we describe a fully decentralized protocol that avoids much of the rate reduction by the design where a MAC is joined with the physical layer.

Existing MAC protocols for UWB [10, 3, 5, 13] are either based on mutual exclusion (no other communication

is possible within the same collision domain) or on a combination of power control and mutual exclusion. Exclusion is enforced either with a collision management protocol (CSMA/CA or a variant of it [11]), or with a time division scheme (e.g. TDMA), or with a combination of both [10, 3]. A CSMA/CA is used for the MAC protocol in 802.15.4 [4]. It is decentralized, but it has an overhead in backoffs and beaconing. The latter two are proposals for UWB physical layers to be used with the 802.15.3 MAC protocol. The protocols are centralized and rely on a node called a piconet coordinator. It is possible to improve spatial reuse and to reduce interference by controlling the transmitting power of the nodes [5, 13]. CA/CDMA [14] is a power control protocol where the underlying physical layer is based on code division multiple access (CDMA). Here, an interference margin is taken into account when determining the transmitting power. This allows for a limited amount of concurrent transmissions instead of always having to enforce exclusion. All of these proposals either have a fixed rate or allow the users to choose between a very small number of fixed rates.

1.2. A Rate-Adaptive MAC Protocol with Interference Mitigation

A largely unexploited dimension is to let the rate vary with the level of interference. A mathematical analysis of an optimal MAC design for UWB, including exclusion, power control, and rate adaptation is given in [16, 17]. It is proven in [16] that the optimal MAC layer should not use power control but should send at full power whenever it sends. Furthermore, results from [16] show that 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 it is optimal to adapt the channel code to these interferences; in contrast, interference from inside the exclusion region should be combated. The main reason is the non-linearity of the achievable bit rate as a function of the signal to interference and noise ratio (SINR): exclusion mechanisms divide time and rate linearly, whereas interference reduces the rate less than linearly outside the exclusion region. These results indicate that in our case the optimal MAC design should (1) allow sources to send at maximum power, (2) allow interference outside the exclusion region and (3) combat interference inside the exclusion region.

Instead of enforcing exclusion within the exclusion region, we propose a different form of interference management called *interference mitigation*. To this end, we modified the existing PPM UWB [19]. Our UWB physical layer is also based on transmitting short pulses. At a receiver, interference is most harmful when pulses from a close-by interferer collide with those of the sender. The interference mitigation scheme discussed in Section 3.2 is based on detecting and removing pulses that have significantly higher energy

than the energy of the signal previously received from the sender. In contrast with exclusion-based mechanisms, this interference mitigation scheme does not require coordination between senders to determine who is allowed to send next, which significantly simplifies the design of the MAC layer. It is, however, still necessary to enforce exclusion between sources that send to the same destination, since we assume that a destination can receive only from one source at a time. This is solved by means of a *private MAC* protocol, described in Section 4.2. Although the interference from inside exclusion regions can be taken care of by interference mitigation, it remains for the sources to adapt their rate to the level of interference on the channel. To this end, we use “dynamic channel codes” with incremental redundancy as discussed in Section 4.1. To our knowledge, this is the first MAC protocol for ad-hoc networks that uses dynamic channel coding.

Our design simplifies the problem of multiple access for UWB. It 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). Problems such as hidden or exposed nodes naturally disappear. There is no need for a separate channel for global control. Simulation results (Section 5) show a significant increase in throughput vs. traditional MAC protocol design.

Our MAC protocol is well suited for mobility, which is another requirement of the 802.15.4a working group. A source constantly tracks and adapts to channel variations, hence there is almost no loss when one of the nodes is mobile. Also, since our MAC protocol is fully distributed and independent of positions of surrounding nodes, changes in the network topology does not affect the performance. We demonstrate these findings by simulations in Section 5.

We show by simulations that DCC-MAC, due to its tight integration with the physical layer, is better than the other existing MAC protocols, including 802.15.4, for various network scenarios. In addition, our protocol is simple, has much less overhead, and is less power consuming than all the other proposed protocols.

2. SPECIFICS OF UWB

There are several proposals for a UWB physical layer [19, 12, 10, 3]. It was shown in [18] that the optimal wide-band signaling consists of sending short infrequent pulses. Consequently, our model is based on the widely used proposal of [19]. Time is slotted in chips of very short duration T_c ; chips are organized in frames of the length of *PRP* chips. PRP stands for “Pulse Repetition Period”. A node transmits one pulse in one chip per frame. In a chip, binary pulse

position modulation (2-PPM) is used. Since in our model $T_c = 0.2$ ns, the bandwidth is roughly 5 GHz. In addition, a pseudo-random *Time Hopping Sequence* (THS) specifies in which chip to transmit. THSs permit different sources to share the channel, i.e. source-destination pairs use independent, pseudo random, uniformly distributed THS.¹ Further, they avoid energy peaks in the frequency domain – a requirement imposed by regulation. The minimum value for the PRP is on the order of 10 to 100. The average radiated power P_{rad} depends linearly on PRP $P_{rad} = \frac{P_{pulse}}{PRP \cdot T_c}$ where $P_{pulse} = 0.28$ mW [10]. Hence, to achieve $P_{rad} = 1 \mu\text{W}$, we use $PRP = 280$. Note that despite the rather large PRP, the maximum rate is equal to $\frac{1}{T_c \cdot PRP} = 18$ Mb/s.

Prior to the modulation, a source also uses a channel code to translate data bits into encoded bits that are in turn modulated as pulses. 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. In our case, we need variable rate channel codes. We use the so-called rate-compatible punctured convolutional (RCPC) codes [9, 15], in particular those from [7]. A variable encoding rate is achieved by puncturing [9], where a high-rate code is created from a low-rate code by removing coded bits from the lowest rate block of coded bits. Let $R_0 = 1 > R_1 > R_2 > \dots > R_N$ be the set of rates offered by the channel code (where R_N is the so-called “mother code” with the lowest rate). These codes allow for an easy implementation of incremental redundancy, since a block of coded bits with rate R_{n+1} is a subset of the block of coded bits with rate R_n .

At the receiver, a demodulator transforms the received pulses into symbols that are fed to the channel decoder that attempts to recover the data bits that were sent. Finally, we also use an interleaver [15]. 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. PERFORMANCE ANALYSIS OF THE OPTIMAL MAC

3.1. Exclusion Region

A simple performance analysis of the UWB physical layer described in Section 2, is presented in [19, 5, 6]. In these analyses the interference caused by concurrently transmitting nodes is approximated with a Gaussian white noise. Therefore, the total noise is the sum of the background noise and all the interfering signals, and is also Gaussian. This approximation is very coarse, but it is used to make the analy-

sis tractable. The analyses also assume that repetition coding is used to recover from possible errors. Under these assumptions it is shown that a rate is a linear function of signal-to-noise ratio at the receiver.

A mathematical analysis of an optimal MAC design for networks with a linear rate function is done in [16]. The results of the analysis show that a node should adapt the rate and not the transmitting power. When sending, it should always send with the maximum power, and it should adapt the rate according to the level of the interference at the receiver. Also, in the optimal schedule, there is a well defined exclusion region around a receiver. Nodes within the exclusion region of a receiving node should not transmit concurrently, whereas the nodes outside can transmit in parallel, regardless of the amount of interference they produce. Finally, the size of the exclusion region depends only on the transmitting power limit of a source, and not the link length nor on the positions of other nodes. The lower the transmission power is, the smaller the exclusion region is. An approximate formula for calculating the size of the exclusion region is given in [16]; using physical layer parameters from Section 2, we see that the optimal exclusion region has the radius of approximately 1 m.

Our physical layer differs in two aspects from the analytical models used in [19, 5, 6, 16]. As already mentioned, the interference from concurrent transmissions is not Gaussian. It is shown in [6] that a non-Gaussian interference yields to a significantly worse performance of a receiver than a Gaussian interference of the same power. The second difference is that the analytical models use a simple and non-efficient repetition coding, again to facilitate calculations. We use more advanced RCPC codes, as explained in Section 2. Although the above analysis gives an approximate size of the optimal exclusion region, it is not clear how sensitive this finding is to the inaccuracies of the model.

In order to verify the above result, we consider a cylindrical scenario from Figure 2. We consider cylinders with 4 and 32 nodes, and fix $l = 24$ m. We vary the distance d to adjacent interfering nodes. We consider two transmission scenarios. The first is without mutual exclusion, where all senders send at the same time. The second scenario is with mutual exclusion, where every second link transmits; the links adjacent to the transmitting link, one on the left and one on the right, are silent. We verified numerically that this is the optimal mutual exclusion schedule for cylindrical example.

As can be seen from Figure 3, when $d \leq 1$ m, having all nodes send at the same time is not optimal. This means that the radius of the exclusion region is approximately 1m. If an interfering node is closer than 1 m to a receiving node, the interference will significantly deteriorate the performance of the receiver and an exclusion is necessary.

¹Note that this is not equivalent to traditional CDMA and spreading code techniques and that the orthogonality of THSs is not required.




Fig. 1. Model of Ultra Wide Band used in this paper

It is a realistic scenario to have one or more interfering nodes at 1m from the receiver. Still, an implementation of exclusion regions is tedious, requires a complex MAC protocol and a large overhead, which is clearly undesirable in a low-power network. Instead, we propose an interference mitigation technique, presented in the following section, that will alter the physical layer and further decrease the size of the exclusion region. This will render mutual exclusion no longer necessary even for $d \leq 1\text{m}$, and will greatly simplify the MAC layer.

3.2. Interference Mitigation

In our context, the cause of interference is twofold. There is thermal noise and there are collisions between a pulse from a source and pulses from one or several interferers. The probability of collision depends on the PRP, where a large PRP implies a low probability of collision (with $\text{PRP} = 280$ and one interferer, it is below 1%).

It might seem that with such a low probability of collision, the effect of the interference is very small and can be neglected. However, if an interferer is very close, within the radius of the exclusion region, the energy of an interfering pulse is going to be much larger than the energy of the received signal. A typical decoder decodes a large sequences of received symbols together. A colliding pulse will thus impact the decision on a large number of bits, with the probability of error being larger as the energy of the colliding pulse increases. Therefore, even with a low probability of collision of less than 1%, the impact of interference is huge when the interferer is in the exclusion region, as predicted by analysis in Section 3.1, and verified by simulations on Figure 3.

In order to combat the interference coming from the nodes within the exclusion region, we exploit the fact that the received power of a pulse from an interferer within an exclusion region is much larger than the power of a pulse received from the source to greatly reduce the effect of interference. This *interference mitigation* mechanism is inspired by the work in [12].

Assume a source S communicates with a destination D in the presence of a nearby interferer X . If a pulse from X collides with a pulse from S , the received energy is much higher than the intended received power from S . We use a *threshold* demodulator at D that detects when the received

energy is larger than some threshold \mathcal{B} . In this case, we skip the chip and declare an erasure. With our choice of PRP, the probability that an erasure occurs is very low. The loss incurred by those erasures can mostly be recovered by our channel codes, unlike losses from collisions which are difficult to compensate. The resulting rate reduction is much smaller than if we do not use interference mitigation and let the channel code attempt to recover the errors created by a pulse collision. A small erasure probability translates into a small reduction of the rate. We simulate the physical layer in Matlab and evaluate the performance of our interference mitigation in the scenario of Figure 2. The results are depicted in Figure 3.




Fig. 2. Multiple interferers scenario: n nodes are symmetrically distributed on the edges of a cylinder. Every second link is inverted such that each destination is close to an interfering source. l is distance between a source and a destination, d is the distance between a destination and the adjacent interfering source. Number of links is $n/2$ (on the figure $n = 6$)

With the interference mitigation, we change the performance of our physical layer. The effect of the mitigated interference is highly non-linear, and even further from the model of [16] than our initial physical layer. Thus we cannot use




Fig. 3. Rate achieved in the multiple interference scenario, for our physical layer (with and without interference mitigation) and mutual exclusion versus distance d . Link lengths $l = 24$, $n = 4$ and 32.

the analysis from [16], and we turn again to the cylindrical example from Figure 2 to estimate the size of the exclusion region. We see from Figure 3 that for the case with interference mitigation, the achieved rates never fall below the rates obtained by mutual exclusion. Thus, interference mitigation renders the size of the exclusion region negligible and significantly increases the achieved throughput in the case of a close interferer.

Note that the optimal value of the threshold \mathcal{B} depends on both the power of the interferer and the white noise. A too large \mathcal{B} is equivalent to the case without erasures, whereas a small \mathcal{B} will declare too many erasures. Our goal is to set \mathcal{B} such that the erasures are declared only due to collisions, and not due to the white noise. We found by simulation that a suitable value is $\mathcal{B} = 3\mathcal{N} + \mathcal{A}$ where \mathcal{N} is the average white noise power and \mathcal{A} is the estimated received signal power. Using erasures as a way to mitigate interference is suitable when PRP is rather large. However 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 is for further research, as well as the determination of the optimal value for the threshold \mathcal{B} .

We have seen in this section that the optimal MAC protocol for the low power UWB system described in Section 2 needs to enforce mutual exclusion for sources and destinations that are less than 1 m away. An implementation of mutual exclusion requires a complex protocol, which becomes even more difficult in the case of UWB since it is not possible to perform carrier sensing. We modify the physical layer instead; we proposed an interference mitigation technique that decreases the destructive impact of the interference. We verify that with the modified physical layer the size of the exclusion region is reduced to a negligible value, which in turn greatly simplifies our protocol design.

4. DCC-MAC: JOINT PHY/MAC PROTOCOL FOR UWB

DCC-MAC is a joint PHY/MAC protocol and it consists of two parts. The first part is dynamic channel coding (DCC), and it adapts the rate to the interference and the channel condition. The second part is a private MAC part that resolves contentions of several senders for a single receiver.

4.1. Dynamic Channel Coding and Incremental Redundancy

When a source decides to send packets, it transmits regardless of the other users sending in parallel, and using the maximum available transmitting power. Channel conditions at the receiver depend on several factors: the ongoing transmissions in the destinations neighborhood, short and long-term channel fadings, and mobility. To make the best use of

the channel, the channel code needs to be 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 slightly more powerful code than required) to avoid losing a packet when channel conditions deteriorate. We use a typical hybrid-ARQ protocol. 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. The 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 [9] 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 the decoding is successful, an acknowledgement is sent back to S . Otherwise, a negative acknowledgement (NACK) is sent.
- 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 4.2).
- If the receiver cannot even detect reception of data it cannot send a NACK. In this case the sender will time out and retry communication with a more powerful code.

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 a step-wise traversal of the trellis of the Viterbi decoder [15]. 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, code R_i can be eliminated. Because of the rate compatibility feature of RCPC codes, this allows for the elimination of 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. This rate is sent back to the sender in the acknowledgement.

The same calculations are performed for all successive data transmissions, albeit based on the current code instead of the lowest rate code. If 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 a bootstrap, as soon as the packet can be decoded. 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.

4.2. Private MAC

The goal of the private MAC protocol is solely to enforce that several senders cannot communicate simultaneously with one destination. This is traditionally solved by a carrier sensing scheme. However, carrier sensing 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). We solve the problem by a combination of a receiver-based and an invitation-based selection of THSs. Contention for a destination uses the permanent THS of the destination, but an established communication uses the THS private to a source-destination pair.

Recall that a THS is a periodic sequence that specifies which chip position to use for transmission in each frame. We use the following method to generate the THSs. 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 a user with MAC address S is the THS produced by the PRNG with seed S . The private THS of users S and D is the THS produced by the PRNG with the number whose binary representation is the concatenation of S and D as seed.²

As shown in Figure 4, 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 S has data to transmit to a node D , and no other node is sending data to D . The idle node D listens on its own THS. When node S wants to communicate with D , it sends a transmission request (REQ) on D 's THS. The channel code uses the lowest possible rate R_N , so that all nodes within reach that want to talk to D may overhear it. D answers with a response packet (RESP) using the THS private to S and D also coded with rate R_N . This response contains the channel code R_i to be used for data

²Note that a node can always compute the THS used by a potential source. There is no protocol for THS distribution.

³This scheme differs from RTS/CTS as used in 802.11 in that it only reserves a per-destination collision domain.

packets dictated by the channel code assignment procedure discussed in Section 4.1. When S receives the reply, it starts with the transmission of the data packet on the code private to S and D . After the transmission, S listens for an acknowledgement sent by D on the private THS with the smallest rate channel code R_N . If a negative acknowledgement is received, S sends incremental redundancy until a positive ACK is received (which marks the end of the packet transmission). Together with the previous data, this results in a code of rate R_j with $j > i$. If no feedback is received, the sender S retries transmission after a random backoff, up to a certain retry limit. After a transmission (either successful or unsuccessful), both sender and receiver issue a (short) idle signal each on their own THS to inform other nodes that they are idle.

Assume now that a node S' wishes to communicate with D while D is receiving a packet from S . It sends out a request on D 's THS; this may create some interference but will usually not disrupt the private communication between S and D since it is on a different THS. S' then switches to D 's THS and listens for the idle signal. When it hears the idle signal, it waits for a random, small backoff time and transmits the request again. If the timer expires without the node overhearing another transmission request, a request is sent. Otherwise, the node defers transmission and pauses the backoff timer until it hears the idle signal again.

5. 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 a UWB physical layer, as well as new 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 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 [2].⁴

5.1. Goals

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

⁴The implementation forms a basis for the simulation of multi-access based physical layers. It can easily be extended to support for example the simulation of CDMA in ns-2.




Fig. 4. Channel coding and private MAC

We first verify that the performance of our MAC protocol complies with 802.15.4a requirements [1]. More precisely, we simulate networks with different link lengths, rates, interferers' distances and mobility levels, and compare to [1].

Next, through the comparison with other MAC proposals, we verify the two basic assumptions of our design described in Section 3: 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 power control protocol uses variable power, has a fixed channel coding, and allows interference (to some degree). The exclusion-based protocol uses maximum power, has a fixed channel coding, and controls interference through exclusion.

We analyze the number of successful data packet transmissions per node over time. From this we calculate our main performance metric, the average throughput achieved by all nodes. The throughput thus takes into account the loss in bit rate due to channel coding and the overhead due to the transmission of control packets.

To compare the fairness of the MAC protocols (i.e., the distribution of rates achieved by the nodes), we use Jain's fairness index: $F(x) = (\sum x_i)^2 / (n \sum x_i^2)$. Since for our protocol, the fairness index was very close to 1 (in the range of 0.99 – 0.999) in all of the simulations, we leave out fairness graphs for reasons of brevity.

In all the simulations we use UDP traffic. We also ran the simulations with TCP traffic and we obtained very similar results. To isolate the effect of mobility on the MAC protocol (and not investigate the performance of a particular routing protocol under mobility instead), we do not use multi-hop routing.

5.2. 802.15.4a Requirements

5.2.1. Link Lengths

We first inspect what rates can be achieved for various link distances. We consider a single link whose length we vary, and we plot what the minimum code index (that is, the maximum rate) we can use, given the link size. We use the UWB propagation model described in [8]. Depending on whether the path is line-of-sight (LOS) or not (NLOS), the path loss

exponent is 1.7 or 3.5. The model further includes a random component that captures building-specific differences in propagation.

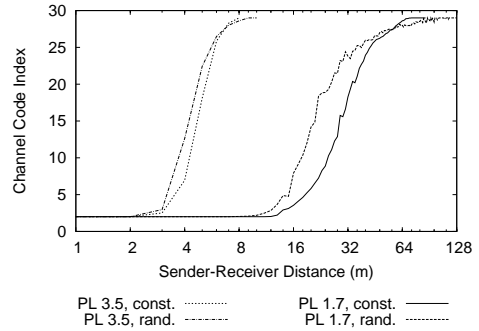


Fig. 5. UWB signal propagation for LOS and NLOS over distances of 1m-100m. The lower the SINR due to the sender-receiver distance, the higher the index of the chosen channel code. Non-line-of-sight communication is possible for up to 10m, line-of-sight communication for up to 60m.

Figure 5 shows the average channel code index for a single pair of nodes communicating over a given distance. The highest code corresponds to the minimum rate of 1.8 Mbps. LOS communication with a path loss exponent of 1.7 achieves distances of more than 100m, as required by [1]. At the same time, the rate is above 1 Mbps. In contrast, with a range of ca. 10m, indoor NLOS communication is more or less restricted to a single room. Additionally, more powerful codes can be constructed to further increase links' lengths at the expense of further decreasing rates.

When the channel varies from packet to packet ("rand."), more powerful channel codes are used compared to the same simulations without the random component ("const."). The more powerful channel codes result in an approximately 20% lower data rate compared to the codes used for the constant channel.

At the same time, these more powerful codes provide additional protection against signal degradation due to interference. Similarly, signal power drops off quickly for the NLOS channel and nodes have to be placed rather close together for perceptible interference.

Therefore, the simulations in the following sections are per-

formed with a LOS channel without channel variations. This allows for an easier analysis of the impact of interference from other nodes.

5.2.2. Resistance to Interference

Another requirement of 802.15.4a is that link rates are robust to a presence of a nearby interferer. In order to verify it, we consider the near-far scenario that is an “unfolded” two-dimensional version of the one shown in Figure 2, since ns-2 does not allow for three-dimensional simulation topologies. We consider networks with 2 to 32 senders. The distance between sender and receiver varies from 1m to 20m.




Fig. 6. Near-far scenario with 2 senders (i.e., 1 interferer) and for different sender-receiver distances (l). We show average rate per user vs. distance to the interferer. Only for very close interferers does DCC-MAC throughput decrease slightly.

We first analyze a simple scenario with only one interfering node at different distances. We further vary the distance between sender and receiver (l). As shown in Figure 6, per node throughput is almost constant and mostly depends on the sender-receiver distance. Only when the interferer is located much closer to the receiver than the actual sender is it necessary to switch to a lower rate code due to the decrease in SINR. Consequently, we observe a slight decrease in throughput for link lengths of 12m, 18m, and 24m. For a sender-receiver distance of 6m, no degradation in rate is observed at all. The difference in throughput compared to the graphs shown in Section 3 can be attributed to the overhead incurred by having an actual MAC protocol (request/response messages, safety margin for code index, etc.). Note that without our MAC protocol, interferers at about 1m away from a receiver would significantly impact the performance, as discussed in Section 3.

Simulations with a varying number of interfering nodes are depicted in Figure 7. The sender-receiver distance is 20m for all of the communicating pairs of nodes. The DCC-MAC is again clearly resistant to interference. There is only a moderate drop in rate from 2300 Kb/s to 1800 Kb/s when we increase the number of nodes from 2 to 16 (i.e., 1 to 15 interferers).




Fig. 7. Near-far scenario with a link length of 20m. We show average rate per user vs. number of (mutually interfering) senders.

5.2.3. Impact of Mobility

Finally, we analyze the impact of mobility on the performance of the network. We consider the random scenario from Section 5.3.2 and let nodes move according to the random way-point model. Node speed varies between 2m/s and 40m/s (which is the maximum speed required by [1]) with 0 pause time.

Comparing the achieved network throughput in the mobile scenario given in Figure 8 with the throughput of the static network, we observe that our MAC protocol is very resilient to mobility. A change of channel conditions due to mobility is compensated by our channel code adaptation mechanism, 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.




Fig. 8. Random mobile network: Surface area of $20m \times 20m$ 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.

5.3. Comparison with The Other Protocols

In the following sections we describe setup and results of a number of simulations to analyze the performance of the DCC-MAC and compare it to other MAC solutions. In particular, we investigate performance in random scenarios and in near-far scenarios, where the receiver is located closer to interferers than the sender. For all of the 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 2.

In the simulations, the following MAC protocols are compared to the DCC-MAC:

Mutual Exclusion with Random Access (RA): This is a simplified implementation of 802.15.4 MAC. 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, TDMA. (Mutual Excl. (TDMA)) – is the ideal mutual exclusion without overhead. It corresponds to the MAC of 802.15.3. We do not actually implement this protocol in ns-2. Instead, we simulate separately transmissions on every link, and obtain the rate when it sends alone. 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.

Power Control – The power control MAC is based on the CA/CDMA protocol proposed in [14]. We adjusted the protocol to work together with a UWB physical layer instead of CDMA for which it was originally designed. Although our implementation abstracts from some protocol details, it captures the main aspect of adjusting the power instead of the channel code. The coding is fixed to the highest-rate channel code that allows communication between the senders and receivers. 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. The safety margin allows a limited amount of future transmissions to overlap with the current transmission. If the transmission is not possible at the required power level due to the maximum power limit at the sender or because it would exceed the interference margin of ongoing transmissions, the sender defers from transmitting and retries after a random backoff.

While MAC protocol details differ, the principles on which the implemented power control MAC is based are the same as those of other power control protocols proposed for UWB, such as [5, 13]. We therefore believe that their performance would be comparable. Similarly, the MAC layer proposed for 802.15.3 can be seen as a combination of TDMA and the exclusion-based random access MAC.

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.

5.3.1. Near-far Scenario

We again consider an “unfolded” two-dimensional version of the near-far scenario shown in Figure 2. The sender-receiver distance is 20m for all of the communicating pairs of nodes. We vary the number of interfering nodes. The results are depicted in Figure 7. The DCC-MAC clearly outperforms the other MAC solutions. There is only a moderate drop in rate from 2300 Kb/s to 1800 Kb/s when we increase the number of nodes from 2 to 16 (i.e., 1 to 15 interferers). For the other MAC protocols, the drop in rate with an increasing number of senders is more pronounced. Power control comes closest to the DCC-MAC performance since it allows for a limited amount of concurrent transmissions. It achieves between 75% and 30% of DCC-MAC’s rate. Both exclusion-based protocols, TDMA and random access, have very similar performances that are significantly worse than that of power control or DCC-MAC. The improvement in SINR and the resulting higher channel code rates cannot compensate for the loss in transmission time due to exclusion.

5.3.2. Random Scenario

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




Fig. 9. Random scenario with nodes placed on 20m×20m square. The number of nodes is given on x-axis, and the average rate is given on the y-axis.

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

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

6. CONCLUSION

In this paper, we have presented a MAC protocol for UWB ad-hoc networks with very low radiated power that fits in the design framework of IEEE 802.15.4a. The design of the MAC protocol is closely coupled with the physical layer. It is based on the assumptions 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 PRP. We have also modified the PPM UWB physical layer and introduced the interference mitigation technique.

We have investigated the performance of our design through analysis in Matlab and simulation in ns-2. The results show that it outperforms the standard 802.15.4 exclusion-based protocol, TDMA-based 802.15.3 MAC, as well as protocols based on power control. Our protocol works very well for low-power UWB, when PRP is large. Furthermore, a network based on the UWB physical layer and DCC-MAC fully satisfies the application requests posed by IEEE 802.15.4a.

Our initial results indicate that even for medium values of PRP (around 100) the performance remains similar. For very low PRP, other exclusion mechanisms such as TDMA or CSMA/CA may be required. Further research is needed to clarify this issue.

Thus far, we have used PPM modulation at the physical layer. Other, non-coherent modulation schemes are also discussed for UWB [12]. It seems that our MAC protocol would apply with little change to such modulations, but this is also left for further study.

7. REFERENCES

- [1] IEEE 802.15.4a approved technical requirements document, may 04 - 04/198r2, available at <http://www.ieee.org>.
- [2] A rate-adaptive mac protocol for low-power ultra-wide band ad-hoc networks, <http://icapeople.epfl.ch/widmer/uwb/index.html>.
- [3] 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.
- [4] E. Callaway, et. al. Home networking with ieee 802.15.4: a developing standard for low-rate wireless personal area networks. *IEEE Communications Magazine*, 40(8):70 – 77, August 2002.
- [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 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.
- [7] 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.
- [8] S. Ghassemzadeh and V. Tarokh. UWB path loss characterization in residential environments. In *RFIC Symposium*, pages 501–504, June 2003.
- [9] J. Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Transactions on Communications*, 36(4):389–400, April 1988.
- [10] D. Héjal and P. Rouzet. ST Microelectronics Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 139r5, July 2003.
- [11] 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.
- [12] R. Knopp and Y. Souilmi. Achievable rates for UWB peer-to-peer networks. In *International Zurich Seminar on Communications*, February 2004.
- [13] 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.
- [14] A. Muqattash and K. Marwan. CDMA-based MAC protocol for wireless ad hoc networks. In *Proceedings of MOBIHOC'03*, pages 153–164, June 2003.
- [15] J. G. Proakis. *Digital Communications*. McGraw-Hill, New York, NY, 4th edition, 2001.
- [16] B. Radunovic and J. Y. Le Boudec. Optimal power control, scheduling and routing in UWB networks. *To appear in IEEE Journal on Selected Areas in Communications*, December 2004.
- [17] S. Toumpis and A. Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE/ACM Transactions on Wireless Communications*, 24(5):736–748, May 2003.
- [18] S. Verdú. Spectral efficiency in the wideband regime. *IEEE Transactions on Information Theory*, 48(6):1319–1343, 2002.
- [19] 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.