A MAC protocol for UWB 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

Abstract

We propose a MAC protocol for very low radiated power (1 micro-watt) ultra-wide band (UWB) mobile networks. Some specifics of UWB, compared to narrowband are that it is optimal to send short pulses and that carrier sensing is impossible. First, we exploit this former property and propose an interference mitigation scheme, at the physical layer, which greatly reduces, but does not entirely cancel, the impact of interference. Second, we analyze how the optimal MAC protocol should be designed, and find that it should not use mutual exclusion (as is commonly done by random access or TDMA protocols) but, in contrast, should allow interference to occur and adapt to it. With the optimal MAC protocol, competing sources are able to send concurrently, causing rate reductions instead of collisions. Third, we design a MAC protocol accordingly. It is made of two components: "Dynamic Channel Coding" and "Private MAC". Through the use of the former component, we can allow sources to send simultaneously at the maximum power permitted by hardware and regulation constraints. Sources then adapt to interference by dynamically adjusting their channel codes (thus their bit rates). Such an approach is suggested by information theoretical results; it sharply contrasts with the traditional, alternative interference management method that would control transmission power instead of channel code. We show by numerical analysis and detailed simulations in ns-2 that our approach is indeed superior. The latter component (private MAC) solves the contention between sources competing for the same destination; it is required because nodes are assumed to be able to receive from only one source at a time. We solve the problem of absence of carrier sensing by a combination of invitation and request, and the use of predictable time hopping sequences. No common channel is used; this avoids the issues of hidden and exposed terminals altogether. Our MAC protocol integrates both components in a single design, and is entirely distributed. It is fully implemented in ns-2. We show by simulation that we achieve a significant increase in network throughput, compared to MAC protocols for UWB that use the traditional approaches of power management or mutual exclusion. Our work shows that it is not optimal to simply port to UWB the design of existing MAC protocols, which were invented for narrowband systems. Our distributed MAC protocol also appears to be the first of its kind to apply channel code adaptation as an alternative to mutual exclusion and power management.

I. INTRODUCTION

A. *Very Low Power Ultra-Wide Band for Ad-Hoc Networks*

The wireless interconnection of all sorts of electronic devices is becoming widespread. However, 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 [7], 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 μW per sender. We are interested in the use of *very low power* UWB, by which we mean that the radiated energy per node does not exceed 1μW $(= -30$dBm). 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 III). 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.

We are interested in the interconnection of active devices such as portable music or video devices. In this context, we want to design a MAC that maximizes throughput, subject to power constraints. We also want our protocol to have reasonable fairness properties.

B. *The Design Space for MAC Protocols*

The existing wireless MAC protocols discussed at the IEEE assume that simultaneous transmissions result in transmission errors and thus employ *mutual exclusion* mechanisms to avoid them. Mutual exclusion is enforced either with a collision management protocol (CSMA/CA or a variant of it [18]), with a time division scheme (allocating time slots with a reservation protocol), or with a combination of both [14], [34]. 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].

A second ingredient used in the design of MAC protocols is *power control*. It is used for example in wireless technologies that have the ability to allow and intelligently manage interference, and thus may accept several transmissions at the same time. This is for example the case with Code Division Multiple Access (CDMA). In a purely synchronous setting (cellular base station), CDMA networks manage multi-user interference primarily by means of power control. In asynchronous settings (ad-hoc networks), CDMA uses both power control and a mutual exclusion protocol [22].

A third, largely unexploited, dimension is *dynamic channel coding*. In [26], 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 [26] 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 [36] that power control does not provide significant gains when dynamic channel coding is used, as long as the goal is to maximize throughput, subject to power constraints. A performance analysis of ad-hoc networks in [24] shows that the optimal MAC layer should use full power when it sends, thus confirming the results in [36].

These existing results tell us that, for our performance objective, we should design a MAC protocol by solely considering the ingredients of *mutual exclusion* and *channel coding*, thus eliminating power control. We verify in our performance analysis that our MAC outperforms a MAC design that uses power control. As we show later, it turns out that the optimal MAC design for our case (once we introduce interference mitigation) does not use mutual exclusion either, thus is entirely based on dynamic channel coding.

Note that we are interested in maximizing throughput subject to very low power constraints. This differs from the requirement to maximize battery lifetime, subject to minimum activity requirements. The latter is typical of sensor networks – a design requirement that we do not address in this paper.

C. Interference Mitigation

It was shown in [38], [35] that the 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 [40]. 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 spread spectrum techniques, but there are some important differences (see Section III-D).

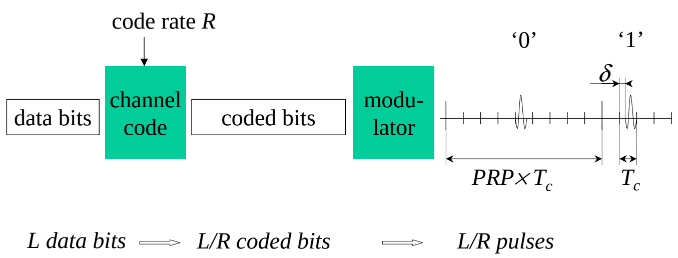

Fig. 1. Model of Ultra Wide Band used in this paper

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 is corrupted.

The frame size PRP must be above some minimum 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 III-C).

We exploit this feature to propose a simple interference mitigation method. Our goal is to remove as much interference as possible, at the receiver, in the physical layer. When an interference occurs, the decoder may wrongly decode the signal, because its correlator is fooled by false values. We propose to replace false bits by *erasures*, as suggested in [31]. Assume a source X interferes with a transmission from some node S to a node D. We propose to use a demodulator at D that detects when the 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 declared lost, which is much better than corrupting the decoder with wrong values. The probability that X causes an erasure is of the order of $\frac{0.5}{\mathrm{PRP}}$, i.e. 0.18% (Section IV). With a capacity approaching code, a small erasure probability translates into an equally small reduction of the rate, of the same order. Our interference mitigation scheme based on erasures is described in detail in Section IV. Using analytical methods, we also show that it clearly improves performance in all cases, regardless of which strategy is used for the MAC layer.

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) but a detailed investigation remains for further research.

D. The Theoretically Optimal MAC Protocol

Before designing a MAC protocol in detail, we first determine what is the theoretical optimal design, neglecting protocol overhead. The challenge is to find the optimal combination of mutual exclusion and dynamic channel coding. In Section V we consider simple analytical models with small and large interferers. We find that the optimal strategy, in our setting, is always to use dynamic channel coding, and never mutual exclusion. In the case of a few large interferers, this result is not surprising, since interference mitigation cancels most of the interference. With mutual exclusion, both source and interferer share the access to the channel, which reduces the rate for each by up to 50%, far more than the rate reduction caused by erasures. In contrast, when there are many small interferences, this result may look surprising. It comes from the non-linearity of the rate reduction caused by interferences.

This result suggests the design of a MAC protocol that radically differs from existing, exclusion based, protocols: we should let sources send concurrently, and arrange the coding such that this causes rate reductions instead of collisions.

E. A MAC Protocol Based on Dynamic Channel Coding and Interference Mitigation

Our second step is to design a protocol that follows the design suggested by the analysis. Since it is optimal to allow interference, 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 sources picks a channel code according to the protocol described in Section VI-A, 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 VI-B. 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).

F. 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.

First, we exploit the specifics of UWB and propose an interference mitigation scheme at the physical layer, which greatly reduces, but does not entirely cancel, the impact of interference. Second, we analyze how the optimal MAC protocol should be designed, and find that it should not use mutual exclusion (as is commonly done by random access or TDMA protocols) but, in contrast, should allow interference to occur and adapt to it. Third, we design a MAC protocol accordingly. We show by numerical analysis and detailed simulations in ns-2 that our approach is indeed superior.

A unique feature of our design is the tight coupling with the physical layer. Problems like hidden or exposed nodes naturally disappear; as a result, the performance remains good even for multi-hop TCP scenarios (Section VII-B). There is no need for a separate channel for global control, as in [22], which saves system capacity. We show in Section VII 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.

To our knowledge, this is the first MAC protocol for ad-hoc networks that uses dynamic channel coding as an alternative to mutual exclusion or power control. Our work also shows that it is not optimal to simply port to UWB the design of existing MAC protocols, which were invented for narrowband systems.

Finally, the fact that the optimal MAC protocol is based on dynamic channel coding, without mutual exclusion, is specific to our setting. However, our idea of arranging the physical layer and the MAC protocol such that collisions may be replaced by rate reduction is quite general; the optimal MAC protocol in other settings is likely to be a combination of dynamic channel coding and mutual exclusion – but we leave such a study outside the scope of this paper.

The rest of the paper is organized as follows. Section II describes related work. Section III gives our assumptions in detail. Section IV describes our first contribution, interference mitigation. Section V analyzes the MAC design space and concludes and determines that mutual exclusion is not optimal. Section VI describe our MAC protocol. Section VII describes our simulation results.

II. 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 [22] the authors propose a CDMA [23] 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 [15] 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. A similar power controlled UWB MAC protocol is described in [19].

In [15] there is no power and no channel code control. Node to node communications are initiated by destinations as in [33], [37]. 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 access to the destination.

The IEEE 802.15 Task Group 3a is currently reviewing proposals for an alternate physical layer for the IEEE 802.15.3 MAC [1]. The MAC is based on the concept of piconets. All members 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 [3] a UWB physical layer based on multi–band OFDM presents a completely different paradigm of UWB modulation. In [14] the physical layer is based on PPM and the use of THSs. 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, by taking a systematic design method, we find that the optimal MAC should not be based on allowing interference; in contrast, the MAC protocols in [1], [3], [14] are based on mutual exclusion. We verify by simulation in Section VII that a mutual exclusion MAC provides much less throughput with the same power than our proposal.

Rate adaptation protocols were already proposed for Aloha [27] or CSMA/CA, narrowband networks [30], [16], [29]. They differ from our proposal in that rate adaptation is performed only to track the state of the channel, excluding interference from competing sources in the same network. Such interferences are either treated as collisions, or avoided by the protocol. In contrast, we use rate adaptation as a mechanism to support multiple access. Furthermore, with UWB, it is not possible to use the channel measurement methods used in these references. Also, we use a set of channel codes which has a much larger range, precisely in order to address rate reduction during interference.

Other proposals [12] consider the combination of multiple access and rate adaptation, albeit by means of adjusting power (in order to reduce interference). This is justified in settings where energy conservation is the primary objective; in contrast, we focus on maximizing the utility of the system, subject to power constraints. This is a different objective, although both methods result in energy savings. In our setting, power control is not at all optimal, as we show in the Section VII.

With a UWB PPM physical layer, source and destination 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 collide at the destination. However, the destination has only to listen to its own THS. In contrast, with a 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 set of THSs. Note that hybrid combinations of the two schemes are possible [32]. Another solution is a receiver-initiated scheme as proposed in [37].

III. System Assumptions

A. 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 ($R \leq 1$). The modulator transforms the coded block into a form suitable for transmission on the physical medium. Similarly, at the receiver, the two corresponding 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.

B. Variable Rate Encoding with Incremental Redundancy

As mentioned in Section I, we use variable rate channel coding. In addition, we require that coding offers *incremental redundancy*, defined as follows. Assume a system offers a choice 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) [13], [9]. These are convolutional codes [39], [23] providing variable encoding rate as well as incremental redundancy. Variable encoding rate is given by puncturing. This process creates 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 $R_0 = 1 > R_1 > R_2 > \ldots > R_N$ 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 (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 RCPC codes from [9]. We have 31 possible rates:

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

Rate 1 is achieved by sending an uncoded block of data bits. Rates $8/31$ to $8/9$ are achieved by puncturing the mother code of rate $1/4$ i.e. $8/32$. 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 [23]. 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.

C. Modulation and Power

Theoretical results on the wide-band communication channel [35], [38] 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 [14].

As previously described, the modulator uses binary Pulse–Position Modulation (PPM) [40], [6] 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 $T_c = 0.2$ ns[1]. This implies that the bandwidth of the generated signal is about 5 Ghz which corresponds to values proposed for future UWB devices [3]. Note that the simple repetition coding scheme of [40] 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) considerations. Commonly found values for PRP are on the order of tens to few hundreds of 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 IV-B that this rate goes down to 8 Mb/s when the distance between source and destination is 25 meters. In Section VII we find average throughputs (including the overhead of the MAC protocol) per source of several Mb/s.

[1]We have that $\delta < T_p < T_c$

D. *Time Hopping Sequences*

A THS is a periodic sequence $[x_1,\ x_2, \ldots,\ x_n]$ where x_i is the position of the chip to be used for transmission in the ith 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.

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 (receiver-based) or private THSs. The public THS of user with MAC address A, called THS(A), is the THS produced by the PRNG with seed = A. The private THS of users A and B, called THS(AB) is the THS produced by the PRNG with a seed equal to the number whose binary representation is the concatenation of A and B. The private MAC protocol (Section VI-B) 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: as previously explained, 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 must have a very low cross-correlation [23]. They can not be computed on the fly and an assignment protocol becomes necessary [32]. In our case, since a node can always compute the THS used by a potential source, no assignment protocol is required. 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 every PRP, in average. As such, CDMA signals from two different users transmitting concurrently always interfere, whereas two concurrent PPM signals interfere only when the chips in the two THS coincide. Moreover, this effect is attenuated by the asynchronicity between two transmitting sources. Hence, generating the THSs as periodic sequences of independent uniform random variables over the range of chip positions $[0, PRP - 1]$ is sufficient, since we operate in the low spectral efficiency regime and PRP is high.

E. *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 [23] relies on the presence of a synchronization preamble at the beginning of a packet. This synchronization preamble is known to the destination of the packet. To synchronize, the destination needs to detect the synchronization preamble. Upon detection, the demodulation of the received signal starts. When established, synchronization must be maintained. In our setting, the preamble is modulated with the THS [28].

In [21], the authors propose an efficient method for detecting synchronization patterns over UWB channels. The time required to achieve synchronization between completely de-synchronized systems is on the order of tens of μs whereas the time to refresh it is of the order of 10μs.

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 simultaneously on the broadcast THS, on its own THS, and on a private THS. This means that it is able to acquire synchronization on either of the THSs. However, a node can receive on only one THS at a time.

F. 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 a 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.

IV. INTERFERENCE MITIGATION

A. Mitigation of Interference by Erasures

With the UWB physical model of [40] described in Section III, it is clear that an interfering user, reasonably close to a destination, will significantly lower the transmission rate. We introduce interference mitigation as a mechanism to reduce this effect. Our proposal is motivated by [31], and is based on the concept of erasures.

When a source sends data to its destination the transmitted pulses are distorted with noise. Moreover, if an interferer chooses the same chip to send data, the transmitted pulses will be additionally corrupted by interferer pulses. In order to transform the received signal at the destination into a sequence of noisy coded bits, a correlator receiver is used [40]. The noisy coded bits correspond to the output of the correlator. These are then fed to the channel decoder.

For a given chip, the amplitude of a noisy coded bit depends roughly on the total energy present in the chip. If no collision occurs, the noisy coded bits correspond to a slightly distorted version of the their respective coded bits. The same happens if a collision occurs with a distant interferer. In this case, the decoder is powerful enough to recover the transmitted bits.

However, a collision with a very strong interferer will result in a highly distorted version of the coded bit. This will pollute the decoding process to a large extend, and will likely cause several subsequent decoding errors.

Therefore, the idea is to remove noisy coded bits that are the result of a collision with strong interferers and to replace them with erasures. In other words, when the received signal energy in a chip is larger than some threshold B, the noisy coded bit is replaced by an erasure.

This is analogous to the simple information theoretic result that states that the capacity of a binary channel that flips bits with probability p is much less than that of the channel that transforms bits into erasures with the same probability p (however, our channel is not a simple binary channel, so this analogy is only indicative).

Due to the time-hopping sequences, the probability of a symbol level collision (i.e. that both a source and an interferer choose the same chip to transmit a bit) is low. Hence only a few percentage of erasures are produced and the channel decoder can recover from them.

It has been shown through indoor channel measurements that variations in the received signal power are typically caused by shadowing rather than fast fading [31]. Hence, a receiver can track the received strength of the signal during several chips, and estimate its average over time. In the same manner, the receiver can estimate the white noise by tracking it in idle chips.

If in a given chip the received signal is by far larger than expected, the receiver concludes that a collision occurred in the chip. More precisely, if for a given threshold B, the received signal amplitude is outside of the interval $[-B, +B]$, we declare an erasure and disregard the chip. An erasure can happen both in the case of a collision with a high power interfering pulse, or in the case of a high white noise sample.

Our goal is to set B such that the erasures are declared only due to collisions, and not due to the white noise. The optimal value of the threshold B depends both of the power of the interferer and of the white noise. On the one hand, a too large B is equivalent to the case without erasures. On the other hand, a small B will declare too many erasures. We set $B = E[|A_s|] + kE[|A_n|]$, where $E[|A_s|]$ is the estimate of the average absolute amplitude of the signal and $E[|A_n|]$ is the estimate of the average absolute amplitude of the noise. We use $k = 5$, and we verify by simulations that it offers good performance. The optimal choice of B remains to be further analyzed.

B. Analysis of Interference Mitigation

We evaluate the performance of interference mitigation on simple, but tractable examples. We need to compute the rate achievable by a source: a classical method is to use the Win-Scholtz UWB physical model as presented in [40]. However, this model has two major drawbacks. First, interference from other users is assumed to be 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], [8], and the achieved performance is lower than the one the predicted by the model. Second, it assumes simple repetition codes, whereas we use more powerful codes [13], [9] which have higher performance. Therefore we use the more elaborate model with non-Gaussian interference defined in [6], [8], which analyzes the signal pulse by pulse.

We analyze several symmetric scenarios, with varying signal and interference intensities as well as the number of interferers, as shown on Figure 2. As a special case with 4 nodes (2 links), we get the so-called near-far scenario. The scenarios are chosen to be tractable (due to the symmetry), while incorporating near-far effects.

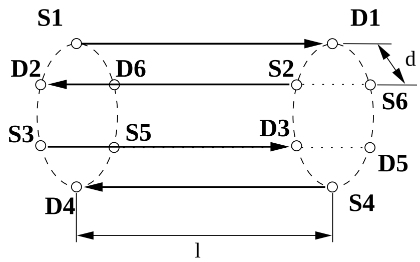

Fig. 2. Multiple interferers scenario: n nodes are symmetrically distributed on the edges of a cylinder. Corresponding peers are located on the adjacent disks. Every second link is inverted such that each destination is close to an interfering source. The distance between a source and a destination is the length of the cylinder l, and the distance between a destination and the adjacent interfering source is d. Number of links is n (on the figure $n = 6$). The case of $n = 2$ represents the near-far scenario.

For every value of the parameter set of the scenario, we test several possible codes. For each code we select random data to be transmitted and random interference. The achieved rate is calculated 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 [40] is lower than at the input. Each interferer resynchronizes with its own destination for every packet (Section III-E). This makes the synchronization offset very variable. We model it as a uniform random variable.

We implemented this model in Matlab and solved it numerically. The results are shown on the top curves in Figure 3, and can be summarized as follows.

Near Far Scenario: (Figure 3 (a)) When the interferer is close to the receiver, it is more likely to be canceled, and interference to be decreased. Therefore, a benefit of interference mitigation is more evident for small d. For example, for $l = 12$m and $d = 0.5$m if the interference mitigation is not used, there are no codes powerful enough to cope with the interference, and the rate achieved is zero. On the other hand, when the interferer is far away, it is very unlikely for an erasure to occur, hence the performances with and without interference mitigations are the same.

There is an additional effect of desynchronization which leads to smaller rates than expected in the case of small d. Namely, an extremely close interferer may be desynchronized and have an interference that is close to, but below the bound B. This in turn leads to slightly worse performance than in the case of an erasure.

Multiple Large Interferences. The case of large interferences corresponds to small values of d in Figure 3 (b) and (c). We see that the interference mitigation improves performance even in the case of a large number of strong interferences. This is because, as we explained in Section III-C, we have a rather large PRP. 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. Even for a very large number of interferers (64), which we hardly find in practice, the rate drops at most by 20%.

Many Small Interferences. The case of small interferences corresponds to large values of d in Figure 3 (b) and (c). In contrast, here, interference mitigation does not improve the rate. However, as we discuss in the next section, this occurs in the region where it is optimal to allow interference, rather than combat it.

V. Optimal MAC Protocol Design

In this section we first determine what is the theoretical optimal design, neglecting protocol overhead. Based on the discussion in Section I, the issue is to find the optimal combination of mutual exclusion and dynamic channel coding with concurrent transmission. A similar problem is studied in [24], where the authors model the network as a complex optimization problem and the function to maximize is the sum of the logarithms of the rates. This corresponds to searching for the proportionally fair allocation, an objective justified for example by [25]. In general, the authors find that the problem is not solvable, even numerically, for arbitrary networks of reasonable size. Therefore, in this section, we restrict our analysis to a few simple cases, which address the impact of large and small interferers. We verify by simulation in Section VII that the findings hold for arbitrary cases, and also when all protocol overhead is included. In spite of these restrictions, this analytical phase is essential, because it gives us a clear guideline for protocol design.

We analyze the same case of the cylindric network as in Section IV, using the same modeling method, but now consider two strategies

- *Mutual Exclusion*: only one node sends at a time. This assumes an underlying scheduling protocol (collision avoidance, polling or TDMA), the overhead of which is not modeled in this section.
- *All-at-Once*: all sources transmit concurrently.

In both cases, senders use the maximum power when transmitting, and adapt codes, hence rates, according to the signal-to-noise level at the receiver. Here too, due to symmetry, the optimization gives the same rates to all links. In the Mutual Exclusion case, this means each link transmits $1/n$ of the time. Interference mitigation is used with the All-at-Once strategy; with Mutual Exclusion, the performance is the same whether interference mitigation is used or not, since sources alternate and do not interfere.

We considered the same three examples as in Section IV: near-far scenario (Figure 3 (a)), multiple large interferences (Figure 3, for small values of d) and many small interferences (Figure 3, for large values of d).

We see that the All-at-Once strategy always performs better than Mutual Exclusion. As expected, the difference in performance of All-at-Once and Mutual Exclusion increases with the interference. All-at-Once is better than Mutual Exclusion even in the scenario with a large number of nearby interferers. This suggests that the optimal protocol is to let transmissions occur concurrently, and adapt to them.

Note that if we would not use interference mitigation, the results would be different: there is a level of interference beyond which mutual exclusion is optimal; in such case, the optimal protocol should implement mutual exclusions when interference is high, and allow parallel transmissions when interference is low. Also, for $l = 12$m and $d = 0.5$m there are no channel codes powerful enough to cope with the interference, if interference mitigation is not used.

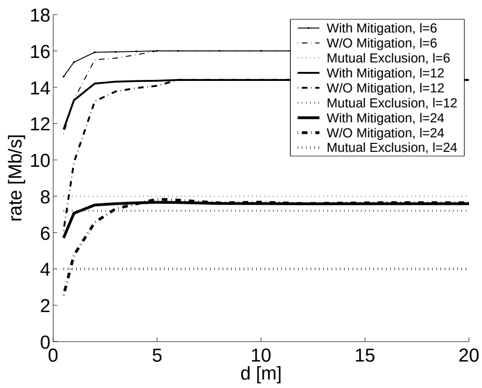

(a) $n = 2$

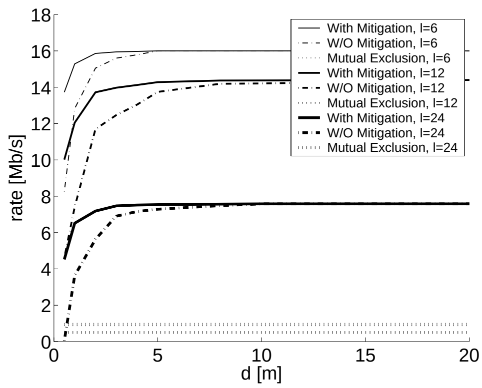

(b) $n = 16$

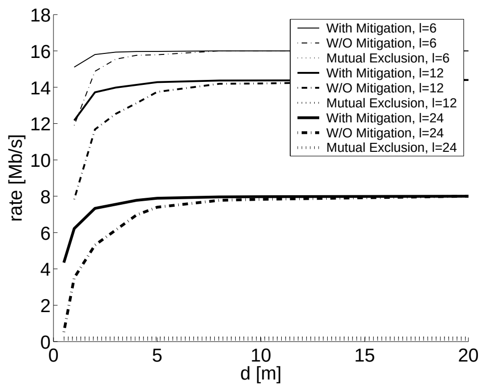

(c) $n = 64$

Fig. 3. Rate achieved in the multiple interference scenario,, for All-at-once and mutual exclusion, and various link lengths l, versus distance to d from the interferer to the destination, for various numbers n of interfering links.

VI. PROTOCOL DESCRIPTION

As discussed in Section V, interference should always be allowed. Furthermore, it is optimal to use maximum power when sending. Thus, the function of our DCC-MAC protocol is 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. The two functions are integrated in one single protocol, but for the ease of presentation, we introduce them one by one.

A. Dynamic Channel Coding and Incremental Redundancy

While channel coding and Automatic Retransmission reQuest (ARQ) (hybrid–ARQ) [20] 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.

Incremental Redundancy: 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. In addition, the packet contains a short header with the packet length, encoded at the lowest rate code. This ensures that D is able to detect that a packet transmission attempt did take place even when decoding of the entire packet fails. In such a case, D sends a negative acknowledgement (NACK) (Figure 4).
- 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 VI-B).

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.

Adaptation of Code Rate: The choice of which code to use exploits the feature of our codes that a destination that can decode can also determine the highest rate code that could have been used.

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 [39], [20]. 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 with $R_j > R_i$. 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 rate R_{i+2} if the highest possible code rate is R_i performs sufficiently well. The code R_{i+2} is indicated to the sender in the 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.

If packet transmission is unsuccessful on the first attempt (thus a NACK was sent), the receiver determines the highest possible rate in the same way as during bootstrap, as soon as the packet can be decoded. On receiving a NACK, the source doubles the rate index of the code, i.e. switches to a much lower bit rate.

The sender determines the code to use as follows: in case the code indicated by the receiver is of higher rate than the current code (i.e., the safety margin at the source was larger than necessary), the sender does not directly switch to the new channel code but decreases its code index by one. Otherwise, if a lower rate 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]

In summary, the algorithm for the selection of code is as follows (see also Figure 4). Remember that a large code index means a small bit rate.

- A source keeps in a variable `codeIndex` the value of the next code index to use. Initially or after an idle period, `codeIndex`$= N$.
- When a destination sees that a packet is sent but cannot decode it, it sends a NACK to source.
- When a source receives a NACK, it sets `codeIndex` to $\min(2*\texttt{codeIndex}, N)$.
- When a destination can decode, it computes the smallest code index, say j, that could have been used, and returns a `codeIndex` attribute in the ACK equal to $j + 2$.
- When a source with `codeIndex`$= i$ in the cache receives an ACK with `codeIndex`$= i'$, if $i' < i$ then source sets `codeIndex` to $i + 1$, else it sets `codeIndex` to i'.
- When a source times out on a packet it sent, it sets the corresponding `codeIndex` to $\min(2*\texttt{codeIndex}, N)$ and resumes a transmission request/reply exchange.

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

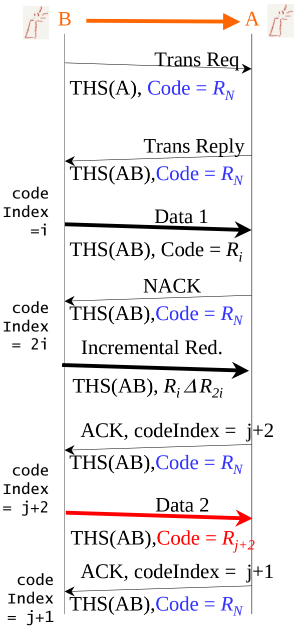

Fig. 4. Dynamic Channel Coding: text above an arrow refers to the content of the message, text below gives the message's THS and channel code.

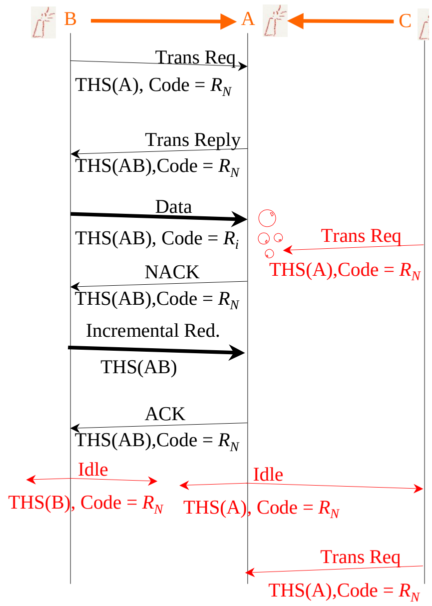

Fig. 5. Private MAC: B sends to A, C attempts to send to A and is deferred. No collision occurs because B and C use different Time Hopping Sequences.

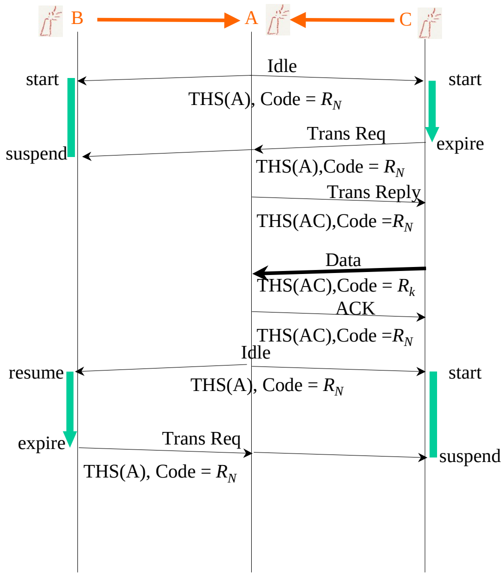

Fig. 6. Private MAC, showing state of backoff timer. B and C compete for sending to A. C initially draws a shorter backoff timer and gains access to A.

B. 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 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 fi x would be to use ALOHA, but its performance is not acceptable except for broadcast traffi c. 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 II. 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 III-D).

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 5). 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 (and only those nodes) 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 VI-A. By doing so, A indicates to B that it is idle. When B receives the reply, it starts with the transmission of the data packet on the

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

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 VI-A). 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.

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. When C does not immediately receive a reply, it switches from the private THS AC to A's THS and listens for A's 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 5).

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 note that this contention window is only used to arbitrate between nodes accessing the *same* 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 it expires without the node overhearing another transmission request, a request is sent. Otherwise, the node defers transmission and pauses the backoff timer until an idle is heard. See Figure 6 for an illustration.

The value of $cw_{\max}$ is initially set to a number of the order of 70μs, i.e. a small multiple of the time it takes for a destination to acquire synchronization (Section III-E). In the rare event that an attempt to access a receiver after the expiry of the backoff timer is unsuccessful (for example because of temporarily getting out of range due to mobility), $cw_{\max}$ 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 a request packet is sent, another timer ($T1$) is computed and immediately started. The value is of the order of the transmission time of a receiver response, a maximum size data packet at the smallest rate, an acknowledgement, and an idle signal. 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 not heard 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 vulnerable period is bounded by the synchronization time (i.e. 10 μs). This is less than the 51.2μs of vulnerable period on Ethernet at 10 Mb/s and is thus very low. Note that we achieve a small vulnerable period 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. However, broadcast packets will not be received by nodes that happen to be sending or receiving at the time of transmission.

A source arbitrates between sending and receiving using a scheduler. In our ns-2 implementation used for this paper, arbitration is based on a simple FIFO scheduler. The performance improvements that can be obtained for example with a fair scheduler remain for further study. Another item that is left for future work is to apply the ideas of CDMA/HDR [4], which consist in picking the best destination at any single point in time.

The hidden terminal problem affects protocols that require that, by listening to the medium, a source is able to detect a collision at the destination. In 802.11, the problem is solved by implementing a RTS/CTS exchange before data transmission, which in turn causes the exposed terminal problem. An exposed terminal is one that is unnecessarily prevented from sending because it hears many CTSs from non-coordinated parts of the network. These problems do not exist with our solution, since 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.

VII. 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 payed to accurately model radio interference of concurrent transmissions. For signal propagation we use a UWB-specific propagation model proposed in [11], which is derived from indoor UWB measurements. Further details of the ns-2 implementation are described in [reference omitted for double blind reviewing].[4]

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 under mobility 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 III-C.

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

Power Control. – The power control MAC is based on the CA/CDMA protocol proposed in [22]. We adjusted the protocol to work together with a UWB physical layer instead of CDMA for which it was originally designed. While our implementation abstracts from some protocol details, it captures the main aspect of adjusting the power instead of the channel code. 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 for 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.

Mutual Exclusion-based MAC with Random Access. (Mutual Excl. (RA)) – All nodes use the same time hopping sequence. Therefore, if a node is transmitting, all other nodes within communication range will receive the packet and cannot send (since a node cannot send and receive at the same time). All nodes but the destination discard the packet. 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. Backoff timers are increased after consecutive unsuccessful retries and reset upon success.

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

While MAC protocol details differ, the principles on which the implemented power control MAC is based are the same as the ones of other power control protocols proposed for UWB, such as [5], [19]. 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.

[4]The 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.

A. 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 I-B: 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.

B. Simulation Results

We first perform a simple simulation to demonstrate the impact of the UWB propagation model [11]. 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 which captures building-specific differences in propagation.

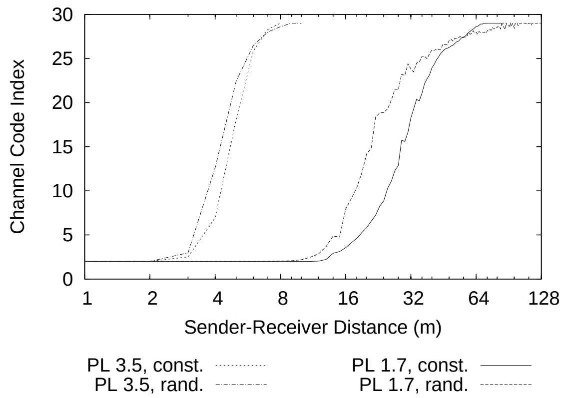

Fig. 7. UWB signal propagation for LOS and NLOS over distances of 1m-100m. The lower the SINR due to the sender-receiver distance, the higher the index of the chosen channel code. Non-line-of-sight communication is possible for up to 10m, line-of-sight communication for up to 60m.

Figure 7 shows the average channel code index for a single pair of nodes communicating over the given distance. With a range of ca. 10m, indoor NLOS communication is more or less restricted to a single room. In contrast, LOS communication with a path loss exponent of 1.7 achieves distances in the range of 50m-60m. 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 performed with a LOS channel without channel variations. This allows for an easier analysis of the impact of interference from other nodes.

Generalized Near-Far Scenario. The near-far scenario we used for the simulations is an "unfolded" two-dimensional version of the one shown in Figure 2, since ns-2 does not allow for three-dimensional simulation topologies. We consider networks with 2 to 32 senders. The distance between sender and receiver varies from 1m to 20m.

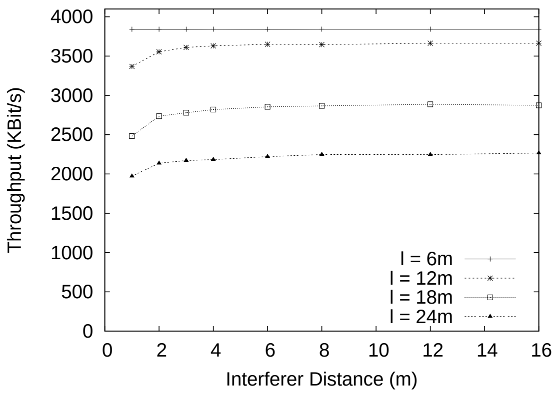

Fig. 8. Near-far scenario with 2 senders (i.e., 1 interferer) and for different sender-receiver distances (l). We show average rate per user vs. distance to the interferer. Only for very close interferers, DCC-MAC throughput decreases slightly.

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 8, 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 IV-B can be attributed to the overhead incurred by having an actual MAC protocol (request/response messages, safety margin for code index, etc.).

Simulations with a varying number of interfering nodes are depicted in Figure 9. The sender-receiver distance is 20m for all of the communicating pairs of nodes. The DCC-MAC clearly outperforms the other MAC solutions. There is only a moderate drop in rate from 2300 Kb/s to 1800 Kb/s when we increase the number of nodes from 2 to 16 (i.e., 1 to 15 interferers). For the other MAC protocols, the drop in rate with an increasing number of senders is more pronounced. Power control comes closest to DCC-MAC performance since it allows for a limited amount of concurrent transmissions. It achieves between 75% and 30% of DCC-MAC's rate. Both exclusion-based protocols, TDMA and random access, have very similar performance which is significantly worse than that of power control or DCC-MAC. The improvement in SINR and the resulting higher channel code rates cannot compensate for the loss in transmission time due to exclusion.

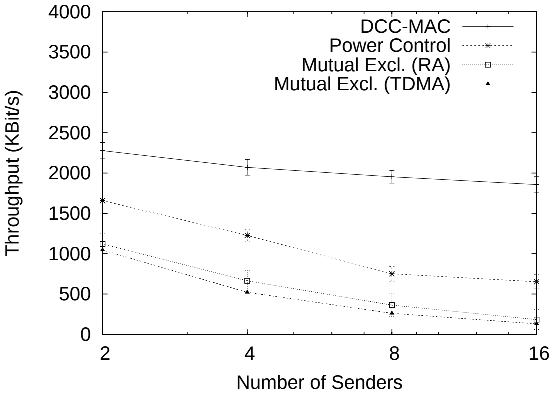

Fig. 9. Near-far scenario with a link length of 20m. We show average rate per user vs. number of (mutually interfering) senders.

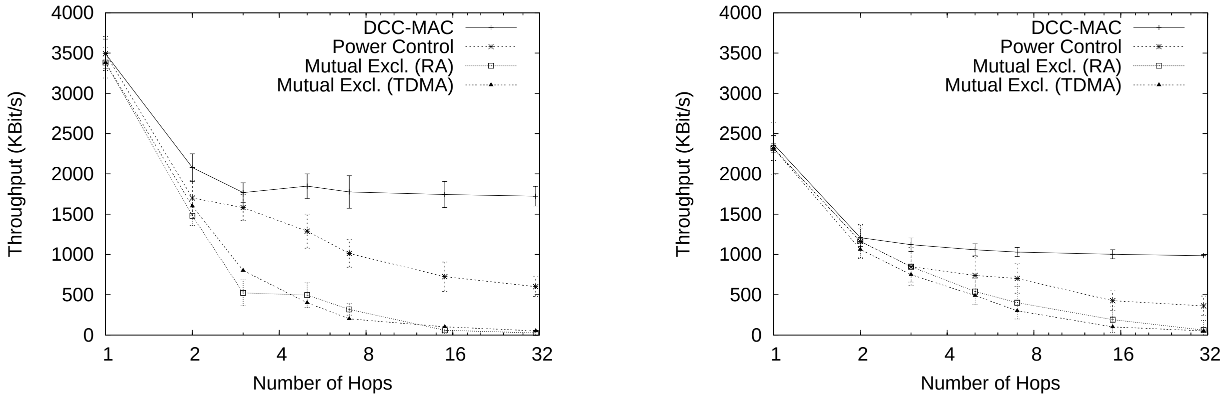

Fig. 10. Throughput on the multi-hop network for UDP (left graph) and TCP (right graph). We show throughput vs. number of hops. There is almost no drop in throughput for the DCC-MAC as the number of hops increases.

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

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

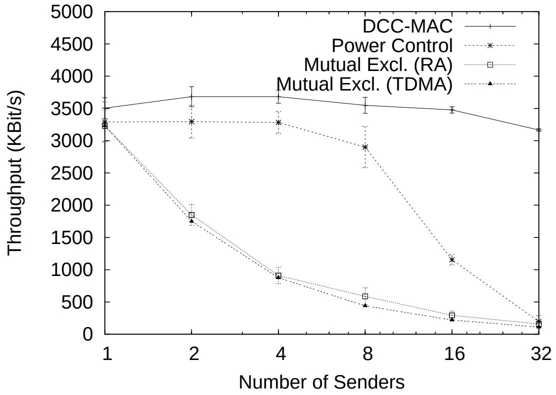

Fig. 11. Random scenario with nodes placed on 20m×20m square. The number of nodes is given on x-axis, and the average rate is given on the y-axes.

for larger numbers of senders due to the dynamic code adaptation that becomes important when the number of nodes (and therefore interference) is high.

Multi-hop Multi-hop forwarding in wireless networks has been extensively studied and was shown to be difficult (see for example [10], [17]). As is usually done, we investigate multi-hop performance of the different MAC protocols using a simple line topology as depicted in Figure 12. Source and destination are at either ends of the line of nodes; intermediate nodes forward packets between them. The distance between nodes is 20m.

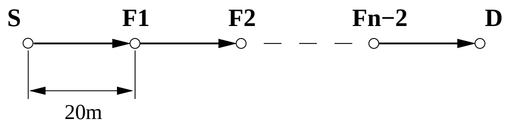

Fig. 12. Nodes equally spaced on a line. The source relays messages over intermediate nodes toward the destination.

The results of simulations are shown in Figure 10. In general, TCP throughput is lower than UDP throughput since TCP data packets compete with acknowledgements traveling on the return path. The most apparent drop in throughput occurs when the number of hops increases from one to two. Since the intermediate node in a 2-hop topology can either send or receive, this necessarily halves the throughput for all of the protocols. What is striking is, that the DCC-MAC is able to maintain this rate when the number of hops increases beyond 2. For UDP, there is a small drop in throughput from 2 hops to 3 hops and from there on the rate remains constant. Similarly for TCP, the decrease in throughput is on the order of a few percent. For power control, there are a number of schedules that allow concurrent transmission over at least of few of the hops; throughput is therefore inbetween that of the DCC-MAC and the exclusion-based protocols.

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

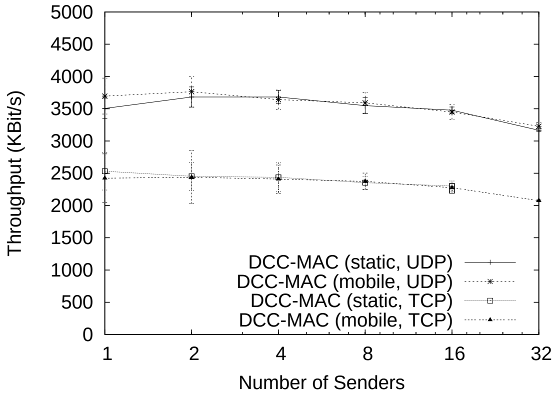

Fig. 13. Random mobile network: Surface area of 20m×20m and mobility according to the random waypoint model. On the x-axis is given the number of nodes in the network and on the y-axis is the average rate per user.

The same holds true when TCP is used instead of UDP. Even for TCP, the variations in channel code caused by mobility are not sufficiently high to result in a perceptible decrease in throughput.

VIII. 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, i.e., when PRP is large. Our initial results indicate that even for medium values of PRP (around 100) the performance remains similar. For very low PRP, interference mitigation is not possible. Exclusion mechanisms such as TDMA, 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 have considered the provision of a single best effort service class. Differentiated services could be supported by having different PRPs – a smaller PRP would provide a larger access opportunity. The details are for further study.

We used PPM modulation. Other, non coherent modulation schemes are also discussed for UWB [31]. It seems that our MAC protocol would apply with little change to such modulations, but this is also for further study.

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

REFERENCES

[1] IEEE 802.15.3 MAC standard, available at http://www.ieee.org.

[2] The network simulator ns-2, http://www.isi.edu/nsnam/ns/.

[3] A. Batra and al. Texas Intruments/Intel and al. Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 267r6, September 2003.

[4] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi. CDMA/HDR: A bandwidth efficient high speed wireless data service for nomadic users. *IEEE Communications Magazine*, 38(7):70–77, July 2000.

[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] Federal Communications Commission. FCC 02–48, First Report and Order, February 14th 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*, 2003.

[11] S. Ghassemzadeh and V. Tarokh. UWB path loss characterization in residential environments. In *RFIC Symposium*, pages 501–504, June 2003.

[12] P. Gupta and P. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications*, 1998.

[13] J. Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Transactions on Communications*, 36(4):389–400, April 1988.

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

[15] A. Hicham, Y. Souilmi, and C. Bonnet. Self-balanced receiver-oriented MAC for ultra-wide band mobile ad hoc networks. In *International Workshop on Ultra Wideband Systems*, June 2003.

[16] G. Holland, N. H. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *Proc. ACM/IEEE MOBICOM*, pages 236–251, 2001.

[17] H.-Y. Hsieh and R. Sivakumar. IEEE 802.11 over multi-hop wireless networks: Problems and new perspectives. In *IEEE VTC-Fall*, Vancouver, Canada, Sept. 2002.

[18] 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.

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

[20] S. Lin and D. Costello. *Error Control Coding: Fundamentals and Applications*. Prentice–Hall, Inc., Englewood Cliffs, NJ, 1983.

[21] I. Maravic, M. Vetterli, and K. Ramchandran. Channel estimation and synchronization with sub-nyquist sampling and application to ultra-wideband systems. Invited paper, ISCAS 2004, 2003.

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

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

[24] B. Radunovic and J. Y. Le Boudec. Optimal power control, scheduling and routing in UWB networks. *IEEE Journal on Selected Areas in Communications*, 2004.

[25] B. Radunovic and J.-Y. Le Boudec. Rate performance objectives of multi-hop wireless networks. In *IEEE INFOCOM*, March 2004.

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

[27] R. Ramanathan and M. Steenstrup. Hierarchically-organized, multihop mobile wireless networks for quality-of-service support. *Mobile Networks and Applications*, 3(1):101–119, 1998.

[28] L. Reggiani and G. Maggio. A reduced-complexity acquisition algorithm for uwb impulse radio. In *UWBST*, 2003.

[29] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic media access for multirate ad hoc networks. In *8th Annual Int. Conf. on Mobile Computing and Networking*, 2002.

[30] Y. Seok, J. Park, and Y. Choi. Multi-rate aware routing protocol for mobile ad hoc networks. In *IEEE VTC-Fall*, 2003.

[31] Y. Souilmi, R. Knopp, and G. Caire. Coding strategies for UWB interference-limited peer-to-peer networks. In *WiOpt'03*, March 2003.

[32] 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.

[33] 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.

[34] M. Teener and R. Gvozdanovic. FDDI-II operation and architectures. *Proceedings of IEEE Conf. on Local Computer Networks*, pages 49–61, October 1989.

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

[36] S. Toumpis and A. Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE/ACM Transactions on Wireless Communications*, 24(5):736–748, May 2003.

[37] 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.

[38] S. Verdu. Spectral effi ciency in the wideband regime. *IEEE Transactions on Information Theory*, 48(6):1319–1343, 2002.

[39] A. Viterbi. Convolutional codes and their performance in communication systems. *IEEE Transactions on Communications*, 19(5):751–772, October 1971.

[40] 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.