

A Rate-Adaptive MAC Protocol for Low-Power Ultra-Wide Band Ad-hoc Networks

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

Swiss Federal Institute of Technology - Lausanne (EPFL)
School of Computer and Communication Sciences
CH-1015 Lausanne, Switzerland

Tel: +41 21 693 5258 Fax: +41 21 693 6610

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

Abstract

Recent theoretical results show that it is optimal to allow interfering sources to transmit simultaneously as long as they are outside a well-defined exclusion region around a destination and to adapt the rate to interference. In contrast, interference from inside the exclusion region needs to be controlled. Furthermore, when a source transmits, it should do so at maximum power. Based on these theoretical findings, we propose a cross-layer design of MAC and physical layer and present a fully distributed, rate adaptive MAC protocol for ultra-wide band (UWB) where sources constantly adapt their channel codes (and thus their rate) to the level of interference experienced at the destination. To mitigate the interference of sources inside the exclusion region, we propose a specific demodulation scheme that cancels most of the interfering energy. This scheme is entirely local to a destination and involves no protocol action. In contrast, existing MAC designs for UWB are either based on mutual exclusion, where no other communication is possible within the same collision domain, or based on power control. Through simulation we show that we achieve a significant increase in network throughput compared to traditional MAC proposals. Overall, we find that existing proposals do not fully take advantage of the specific nature of an ultra-wide band physical layer.

Topic: Architectures of Ad-Hoc Networks

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

*contact author

1 Introduction

Emerging pervasive networks assume the deployment of large numbers of wireless nodes, embedded in everyday life objects. For such networks to become accepted, it is important that the level of radiated energy per node be kept very small; otherwise, environmental and health concerns will surface. Ultra-wide band (UWB) is a radio technology for wireless networks, which has the potential to satisfy this requirement. 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. The radiated power per node depends on technological choices; it is of the order of 0.1 mW to less than $1 \mu\text{W}$ per sender. We are interested in *very low power* UWB, by which we mean that the radiated energy per node does not exceed $1\mu\text{W}$ ($= -30\text{dBm}$). 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 3). These rate values are reduced when several nearby UWB sources transmit concurrently. In this paper we describe a protocol that avoids much of the rate reduction by designing a MAC that is joint with the physical layer.

Existing protocols for UWB [6, 1, 2, 9] 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 [7]), or with a time division scheme (e.g. TDMA), or with a combination of both [6, 1]. The latter two are proposals for UWB physical layers to be used with the 802.15.3 MAC protocol recently discussed at the IEEE. The protocols are centralized and rely on a so-called piconet coordinator. It is possible to improve spatial reuse and to reduce interference by controlling the transmit power of the nodes [2, 9]. CA/CDMA [10] 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 transmit 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.

1.1 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 including exclusion, power control, and rate adaptation is given in [12, 13]. It is

proven in [12] that the optimal MAC layer should not use power control but should send at full power whenever it sends. Furthermore, results from [12] 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 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*. Our UWB physical layer is 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 3.1 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 to 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 3.3. Since interference from inside exclusion regions can be taken care of by interference mitigation, it remains for 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 3.2. 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 like hidden or exposed nodes naturally disappear. There is no need for a separate channel for global control. Also, because the source constantly adapts to the varying channel, mobility is well supported. Simulation results (Section 4) show a significant increase in throughput compared to traditional MAC protocol design.

2 Specifics of UWB

There are several proposals for a UWB physical layer [15, 8, 6, 1]. It was shown in [14] that the optimal wide-band signaling consists of sending short infrequent pulses. Consequently, our model is based on the widely used proposal of [15]. Time is slotted in chips of very short duration T_c ; chips are organized in frames of length 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 $inourmodelT_c = 0.2$ 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 [6]. Hence, to achieve $P_{rad} = 1\mu$ 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.




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

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 [5, 11], in particular those from [3]. A variable encoding rate is achieved by puncturing [5], 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 which attempts to recover the data bits that were sent. Finally, we also use an interleaver [11]. An interleaver pseudo-randomly interleaves bits at the output of the channel encoder; de-interleaving is done at the input of the

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

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 Joint PHY/MAC Protocol for UWB

As described in the introduction, interference should either be mitigated if the interferer is within the exclusion region or should be allowed to exist. Furthermore, when sending it is optimal to use as much power as allowed by the budget. Besides interference mitigation, the function of the proposed *Dynamic Channel Coding MAC* protocol (DCC-MAC) 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 and the interference mitigation scheme are described in the rest of this section.

3.1 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%).

We exploit this feature and the fact that the received power from a nearby² interferer is usually much larger than the power received by the source to greatly reduce the effect of interference. This *interference mitigation* mechanism is inspired by the work in [8]. With interference mitigation, the size of the exclusion region is reduced to a negligible value, which in turn greatly simplifies our protocol design.

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 erasures occur is very low. The loss incurred by those erasures can mostly be recovered by our channel codes. 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

²the distance between the interferer is smaller than the distance between source and destination

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.




Figure 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 $n = 6$)

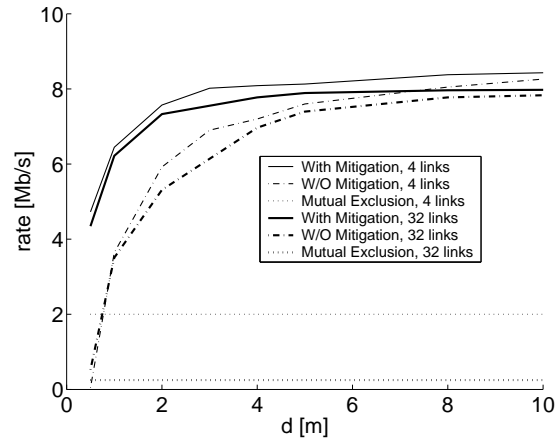


Figure 3: Rate achieved in the multiple interference scenario, for our physical layer (with and without interference mitigation) and mutual exclusion versus distance d . Link lengths $l = 24$, $n = 4$ and 32.

First, we see how the size of the exclusion region can be determined in practice. It corresponds to the values of d where the rate achieved with mutual exclusion becomes larger than the rate obtained when all sources send simultaneously, that is for $d < 1$ m in the case without interference mitigation. 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 case of a close interferer. Finally, the distance between sources and destinations is large ($l = 24$ m). As such maximum rate achieved is equal to 8 Mb/s since it is a function of l .

Note that the optimal value of the bound \mathcal{B} depends on both the power of the interferer and the white noise. On the one hand, a too large \mathcal{B} is equivalent to the case without erasures. On the other hand, 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 $\mathcal{B} = 5\mathcal{N} + \mathcal{A}$ 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} .

3.2 Dynamic Channel Coding and Incremental Redundancy

When interference mitigation is used, channel conditions also depend on the ongoing transmissions in the destinations neighborhood and are therefore more variable than with exclusion between concurrent transmissions. 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 [5] 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 3.3).
- 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 step-wise traversal of the trellis of the Viterbi decoder [11]. 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 to also eliminate all codes $R_j > R_i$. The highest rate code that remains is still powerful enough to decode the packet. Ideally, the more stable the channel conditions, the closer the code used for the next transmission should be to this highest rate code. In practice, we find that the heuristic of using a channel code of R_{i+2} if the highest possible code is R_i performs sufficiently well. 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 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.

3.3 Private MAC

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. 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 receiver-based and 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 and a unique identifier (its MAC address). Communication uses either public or private THSs. The public THS of 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.³

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.⁴ Assume a node S has data to transmit to a

³Note that a node can always compute the THS used by a potential source. There is no protocol for THS distribution.

⁴This scheme differs from RTS/CTS as used in 802.11 in that it only reserves a per-destination collision domain.

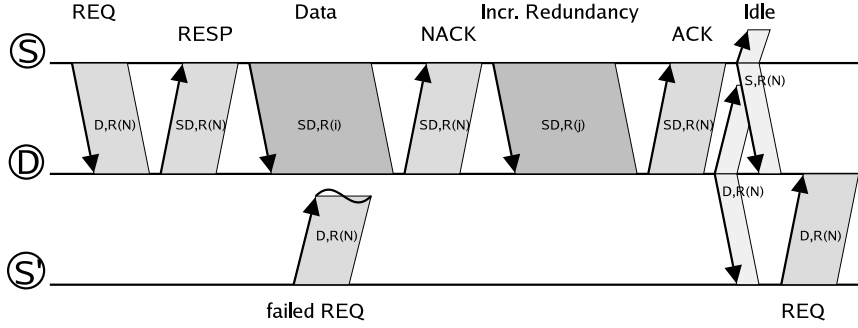


Figure 4: Channel coding and private MAC

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 packets dictated by the channel code assignment procedure discussed in Section 3.2. 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.

4 Simulations

We implemented the DCC-MAC in the network simulator ns-2 to allow an in-depth analysis under more realistic conditions. This required us to redesign the physical layer support in ns-2, in order to account for varying

interference over the course of a packet transmission. For bit error rates and transmission rates, our ns-2 implementation uses interpolation from lookup tables that we created by extensive offline Matlab experiments. We further implemented the UWB channel model described in [4].

We compare our protocol to a traditional exclusion-based protocol similar to CSMA/CA as well as a power control protocol. Through the comparison with these alternatives we verify the two basic assumptions of our design: 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 CSMA-like protocol is similar to the 802.11 MAC layer but is used together with our UWB physical layer. It uses maximum power, dynamic channel coding, and controls interference through exclusion. The power-control protocol is based on the CA/CDMA protocol proposed in [10]. It has fixed channel coding and allows interference up to a certain degree. We define a minimum SINR that is necessary to achieve a given probability of error, plus a safety margin. We then adapt the transmission power of the packet in order to achieve the desired SINR. If this is not possible due to the power limit at the sender, the sender attempts another transmission after a random backoff.

For all MAC layers, the same physical layer is used. In particular, we allocate the same maximum power limit for the exclusion-based MAC, the power control MAC, and the DCC-MAC. The simulation topology is similar to the near-far scenario used for the Matlab simulations (Section 3.1). We consider a network with 32 sender-receiver pairs. The distance between sender and receiver varies from 1m to 40m.

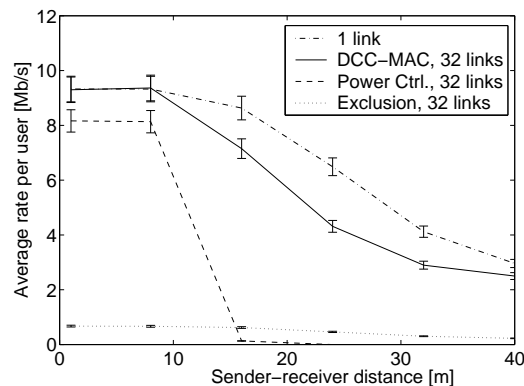


Figure 5: Throughput vs. sender-receiver distance in a multiple interferers scenario with 32 links

Simulation results are depicted in Figure 5. To facilitate comparison, we also give results for a single sender-receiver connection without any interference. For very short distances, the DCC-MAC is equivalent to the optimum rate of a single connection. Over longer distances, it is up to 20% from the optimum, but nevertheless outperforms the other solutions. Over short distances, power control comes close to the DCC-MAC, albeit it

does not quite reach DCC-MAC performance. However, the performance of power control quickly deteriorates as the communication distance increases and the fixed channel code is no longer powerful enough to combat interference and noise. The exclusion-based protocol exhibits the worst performance of all three protocols for short distances since it does not allow parallel transmissions. The improvement in SINR and the resulting higher channel code rates cannot compensate for the loss in transmission time due to exclusion. At longer distances, the exclusion-based protocol outperforms power control protocol due to the adaptive channel coding.

We also performed similar simulations for random topologies. Due to space constraints we omit the simulation results but note that again, DCC-MAC performance is far superior to the performance of power control and the exclusion-based protocols.

5 Conclusion

In this paper, we presented a MAC protocol for UWB ad-hoc networks with very low radiated power. 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 investigated the performance of our design through analysis in Matlab and simulation in ns-2. The results show that it outperforms exclusion-based protocols as well as protocols based on power control. Our protocol works very well for low power UWB, when PRP is large. Our initial results indicate that even for medium values of PRP (around 100) the performance remains similar. For very low PRP, other exclusion mechanisms such as TDMA or CSMA/CA may be required. Further research is needed to clarify this issue.

Thus far, we used PPM modulation at the physical layer. Other, non-coherent modulation schemes are also discussed for UWB [8]. It seems that our MAC protocol would apply with little change to such modulations, but this is also left for further study. Lastly, we focus on a network offering a single class of service. Service differentiation, optimal flow control and routing specific to this MAC protocol are interesting areas of future research.

References

- [1] 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.
- [2] F. Cuomo, C. Martello, A. Baiocchi, and C. Fabrizio. Radio resource sharing for ad hoc networking with UWB. *IEEE Journal on Selected Areas in Communications*, 20(9):1722–1732, December 2002.
- [3] 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.
- [4] S. Ghassemzadeh and V. Tarokh. UWB path loss characterization in residential environments. In *IEEE Radio Frequency Integrated Circuits (RFIC) Symposium*, pages 501–504, June 2003.
- [5] J. Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Transactions on Communications*, 36(4):389–400, April 1988.
- [6] D. Héjal and P. Rouzet. ST Microelectronics Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 139r5, July 2003.
- [7] 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.
- [8] R. Knopp and Y. Souilmi. Achievable rates for UWB peer-to-peer networks. In *International Zurich Seminar on Communications*, February 2004.
- [9] 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.
- [10] A. Muqattash and K. Marwan. CDMA-based MAC protocol for wireless ad hoc networks. In *Proceedings of MOBIHOC'03*, pages 153–164, June 2003.
- [11] J. G. Proakis. *Digital Communications*. McGraw-Hill, New York, NY, 4th edition, 2001.
- [12] 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.
- [13] S. Toumpis and A. Goldsmith. Capacity regions for wireless ad hoc networks. *IEEE/ACM Transactions on Wireless Communications*, 24(5):736–748, May 2003.
- [14] S. Verdú. Spectral efficiency in the wideband regime. *IEEE Transactions on Information Theory*, 48(6):1319–1343, 2002.
- [15] 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.