In Proceedings of Ad-Hoc Now 2004 (3rd International Conference on AD-HOC Networks & Wireless), July 22-24, 2004, Vancouver, British Columbia, Canada. © Springer-Verlag, Lecture Notes in Computer Sciences
Available at `http://lcawww.epfl.ch/Publications/Merz/MerzLBWR04.pdf`

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

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

EPFL, School of Computer and Communication Sciences
CH-1015 Lausanne, Switzerland
`{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. Based on these theoretical findings, we design a fully distributed rate-adaptive MAC protocol for ultra-wide band (UWB) where sources constantly adapt their channel code (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. Through simulation we show that we achieve a significant increase in network throughput compared to traditional MAC proposals.

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. 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 *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. These rate values are reduced when several nearby UWB sources transmit concurrently. Our protocol avoids much of the rate reduction through a joint design of MAC and physical layer.

It was shown in [1] that the optimal wide-band signaling consists of sending short infrequent pulses. Consequently, our physical layer is based on the widely used proposal of [2]. It is a multiple access physical layer using Time-Hopping Sequences (THS) with a chip time $T_c = 0.2$ ns (which corresponds roughly to 5 GHz) and Pulse Repetition

---

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

Period $PRP = 280$. Hence, the average radiated power $P_{rad} = P_{pulse}/(PRP \cdot T_c) = 1\ \mu$W where $P_{pulse} = 0.28$ mW [3].[1] Additionally, we use the so-called rate-compatible punctured convolutional (RCPC) channel codes [4, 5] which provide a variable encoding rate. In particular we use the RCPC codes from [5].

Existing MAC protocols [6–10] 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. All of these proposals either have a fixed rate or allow the users to choose between a very small number of fixed rates. 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 [11]. It is proven that the optimal MAC layer should not use power control but should send at full power whenever it sends. Furthermore, it is optimal to allow interfering sources to transmit simultaneously, as long as they are outside a well-defined *exclusion region* [11] around a destination. In contrast, interference from inside the exclusion region should be combatted. We base our MAC layer design on these findings.

Instead of enforcing exclusion within the exclusion region (a difficult problem), we propose a different form of interference management called *interference mitigation*. At a receiver, interference is most harmful when pulses from a close-by interferer collide with those of the sender. Even though the probability of collisions is fairly low (below 1% with PRP = 280 and one interferer) collisions with strong interferers do cause a significant rate decrease. Inspired by the work of [12], we use a *threshold* demodulator at the receivers that detects when the received energy is larger than some threshold (i.e. when high energy pulses from one or more strong interferers collide with pulses from a source). In such a case, the chip is skipped and an erasure is declared. The loss incurred by those erasures can mostly be recovered by our channel codes and therefore translates into a small reduction of the rate. In addition, this technique reduces the size of the exclusion region to a negligible value.

Hence, what remains is (1) adapt the rate to the varying channel conditions and (2) enforce exclusion between sources that simultaneously send to the same destination. This is solved by means of *dynamic channel coding* and a *private MAC* protocol (DCC-MAC) respectively. Our design moves the complexity of the MAC protocol away from global exclusion between competing sources (a difficult problem) to the combination of 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. Simulation results (Section 3) show a significant increase in throughput compared to traditional MAC protocol designs.

2 Joint PHY/MAC Protocol for UWB

2.1 Dynamic Channel Coding and Incremental Redundancy

To make the best use of the channel, the rate needs to be constantly adapted to the highest rate that still allows successful reception of the data packet at the receiver.

---

[1] Note that the maximum achievable rate is still equal to $\frac{1}{T_c \cdot PRP} = 18$ Mb/s.

In our case, a variable encoding rate is achieved by puncturing [4], 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 $R_0 = 1 > R_1 > R_2 > \ldots > R_N$ be the set of rates offered by the channel code. The rate compatibility feature [4] implies that a block of coded bits with rate R_n is a subset of the block of coded bits with rate R_{n+1}. Hence these codes permit the use of incremental redundancy (IR) using a typical hybrid-ARQ protocol as explained below.

ARQ procedure To communicate with a destination D, a source S has to perform the following steps:

- S adds a CRC to the packet and encodes it with the lowest rate R_N.
- S punctures the encoded data to obtain the desired code rate R_i 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 (ACK) is sent back to S. Otherwise, a negative acknowledgement (NACK) is sent.
- As long as S receives NACKs, further packets of punctured bits (IR Data) 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 below).

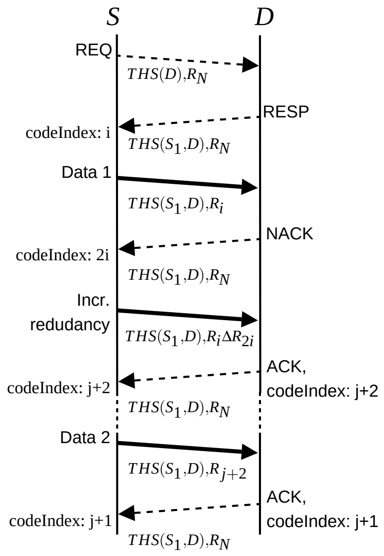

Fig. 1. Dynamic Channel Coding

Note that 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.

Rate selection 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. If the destination can decode the data packet successfully, it will estimate which higher rate $R_j, j \leq N$ would still allow to decode the data. Decoding of the data packet with rate R_N is performed by step-wise traversal of the trellis of the Viterbi decoder [13]. 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_i', with $R_i' > R_i$. Senders maintain a cache of channel codes for a number of destinations. If a sender does not communicate with a receiver for a certain amount of time, the corresponding cache entry times out and the sender bootstraps the code selection procedure with code R_N as described previously.

In summary (Figure 1), the algorithm for the selection of code is as follows. Remember that a large code index means a small rate.

- S keeps in a variable `codeIndex` the value of the next code index to use. Initially or after an idle period, $\texttt{codeIndex} = N$.
- When D sees that a packet is sent but cannot decode it, it sends a NACK to S.
- When S receives a NACK, it sets `codeIndex` to $\min(2*\texttt{codeIndex}, N)$.
- When D 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 S with $\texttt{codeIndex} = i$ in the cache receives an ACK with $\texttt{codeIndex} = i'$, if $i' < i$ then S sets `codeIndex` to $i+1$, else it sets `codeIndex` to i'.
- When S 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.

The heuristic used by a destination of increasing the optimal rate index by 2 is used to cope with the fact that the channel may vary until the next data transmission.

2.2 Private MAC

The goal of the private MAC protocol is to enforce that several senders cannot communicate simultaneously with one destination.

We cannot use traditional carrier sensing scheme since carrier sensing is not possible with our UWB physical layer. We solve the problem by a combination of receiver-based and invitation-based selection of THSs.

Contention for a destination uses the *public* THS of the destination, but an established communication uses the *private* THS to a source-destination pair. The public THS of the user with MAC address S is the THS produced by a pseudo-random number generator (PRNG) with seed S. The private THS of users S and D is the THS produced by the PRNG with a seed equal to the binary representation of the concatenation of S and D.

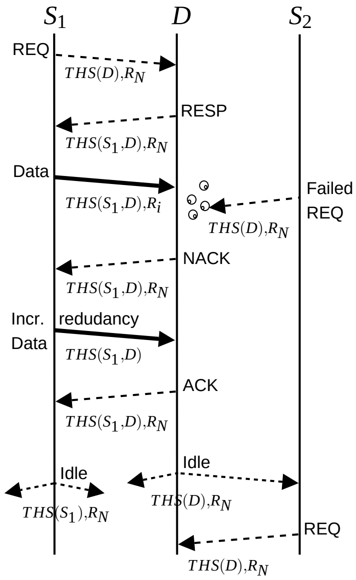

Fig. 2. Private MAC

As shown in Figure 2, a successful data transmission consists of a transmission request (REQ) by the sender S_1, a response (RESP) by the receiver D, the actual data packet, and an acknowledgement (ACK). Assume S_1 has data to transmit to D, and no other node is sending data to D. When D is idle, it listens on its own public THS. As soon as S_1 wants to communicate with D, it sends a REQ on D's public THS using the lowest possible rate R_N. D answers with a RESP using the private THS of the pair S_1-D coded with rate R_N. This response contains the channel code $R_i \geq R_N$ to be used for subsequent data packets dictated by the channel code assignment procedure. When S_1 receives the reply, it starts transmitting the data packet on the THS private to S_1 and D. After the transmission, S_1 listens for an ACK sent by D on the private THS with rate R_N. If a negative ACK is received, S_1 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 $R_i > R_j \geq R_N$.

If no feedback is received, S_1 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 public THS to inform other nodes that they are idle.

Assume now that S_2 wishes to communicate with D while D is receiving a packet from S_1. It sends out a REQ on D's public THS; this may create some interference but will usually not disrupt the private communication between S_1 and D since it is on a different THS. S_2 then switches to D's public THS and listens for the idle signal. When it hears the idle signal, it waits for a random small backoff time. If the timer expires without the node overhearing a REQ from another node, a REQ is sent. Otherwise, the node defers transmission and pauses the backoff timer until it hears the idle signal again.

3 Simulations

We implemented DCC-MAC in the network simulator ns-2.[2] We compared our protocol to a CSMA/CA-like exclusion-based protocol as well as a power control protocol. The exclusion based protocol is similar to the 802.11 MAC layer. The transmit power is fixed but dynamic channel coding is used. The power-control protocol is based on the CA/CDMA protocol [10] and uses a fixed channel coding.

Note that *all* MAC layers have the same UWB physical layer in common as well as the same maximum power limit. We consider a network with 32 parallel sender-receiver pairs. The distance between sender and receiver varies from 1m to 40m. Simulation results are depicted in Figure 3 along with the result for a single sender-receiver connection without any interference. We also performed similar simulations for random topologies. Due to space constraints we omit the simulation results but again, DCC-MAC performance is far superior to the performance of power control and the exclusion-based protocols.

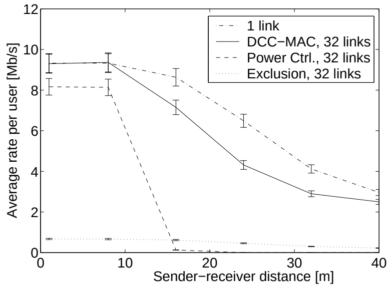

Fig. 3. Throughput vs. sender-receiver distance in the parallel link scenario

---

[2] We redesigned the physical layer support in ns-2 to account for varying interference over the course of a packet transmission. Bit error rates and transmission rates are obtained by interpolation from lookup tables created by offline Matlab experiments. The channel model is the one from [14]

4 Conclusion

In this paper, we presented a MAC protocol for low-power UWB ad-hoc networks. To the best of our knowledge, this the first MAC protocol based on dynamic channel coding where the rate is constantly adapted to the level of interference. Its design is closely coupled with the physical layer. It is based on the assumptions that all nodes have simple receivers and transmitters (single user decoding, one transmitter per node) and all have the same PRP. No common channel is necessary. Also, due to the constant adaptation to varying channel condition, mobility is well supported

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. For future research we plan to investigate optimal policies for channel code adaptation and under which conditions it is better to directly send the data packets instead of having a request/response packet exchange beforehand (i.e., when the additional overhead of the REQ).

References

1. Verdu, S.: Spectral efficiency in the wideband regime. IEEE Transactions on Information Theory **48** (2002) 1319–1343
2. Win, M.Z., Scholtz, R.A.: Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications. IEEE Transactions on Communications **48** (2000) 679–691
3. Hélal, D., Rouzet, P.: ST Microelectronics Proposal for IEEE 802.15.3a Alternate PHY. IEEE 802.15.3a / document 139r5 (2003)
4. Hagenauer, J.: Rate-compatible punctured convolutional codes (RCPC codes) and their applications. IEEE Transactions on Communications **36** (1988) 389–400
5. Frenger, P., Orten, P., Ottosson, T., Svensson, A.: Rate-compatible convolutional codes for multirate DS-CDMA systems. IEEE Transactions on Communications **47** (1999) 828–836
6. IEEE: 802.15 WPAN high rate alternative PHY task group 3a (TG3a). (http://www.ieee802.org/15/pub/TG3a.html)
7. Cuomo, F., Martello, C., Baiocchi, A., Fabrizio, C.: Radio resource sharing for ad hoc networking with UWB. IEEE Journal on Selected Areas in Communications **20** (2002) 1722–1732
8. Hicham, A., Souilmi, Y., Bonnet, C.: Self-balanced receiver-oriented MAC for ultra-wide band mobile ad hoc networks. In: IWUWBS'03. (2003)
9. Kolenchery, S., Townsend, J., Freebersyser, J.: A novel impulse radio network for tactical military wireless communications. In: IEEE MILCOM'98. (1998) 59–65
10. Muqattash, A., Marwan, K.: CDMA–based MAC protocol for wireless ad hoc networks. In: Proceedings of MOBIHOC'03. (2003) 153–164
11. Radunovic, B., Le Boudec, J.Y.: Optimal power control, scheduling and routing in UWB networks. (Accepted for publication in IEEE Journal on Selected Areas in Communications)
12. Knopp, R., Souilmi, Y.: Achievable rates for UWB peer-to-peer networks. In: International Zurich Seminar on Communications. (2004)
13. Proakis, J.G.: Digital Communications. 4th edn. McGraw–Hill, New York,NY (2001)
14. Ghassemzadeh, S., Tarokh, V.: UWB path loss characterization in residential environments. In: IEEE Radio Frequency Integrated Circuits (RFIC) Symposium. (2003) 501–504