Concurrent and Parallel Transmissions are Optimal for Low Data-Rate IR-UWB Networks

Jean-Yves Le Boudec, Ruben Merz

EPFL, School of Computer and Communication Sciences

{jean-yves.leboudec,ruben.merz}@epfl.ch

(Invited Paper)

Abstract—

   The Internet of Things, emerging pervasive and sensor networks are low data-rate wireless networks with, *a priori*, no specific topology and no fixed infrastructure. Their primary requirements are twofold: First, low power consumption and, due to environmental concerns, low emitted power. Second, robustness to poor propagation environments and multi-user interference. Impulse-radio ultra-wide band (IR-UWB) physical layers have the potential to satisfy these requirements. Because the features of IR-UWB physical layers differ from narrow-band physical layers, the design rules of IR-UWB networks are likely to be different than for narrow-band wireless networks. Indeed, to optimally use the resources available, it is crucial for the network layers to take into account and take advantage of the underlying physical layer. Therefore, we are interested in the design of IR-UWB networks in a low data-rate, self-organized, and multi-hop context. We concentrate on the medium access control (MAC) layer and the physical layer. In the case of low data-rate IR-UWB networks, the optimal design is to allow for parallel and concurrent transmissions at the MAC layer. Interference is managed with rate adaptation, no power control and an interference mitigation scheme at the physical layer. A protocol that implements the optimal design and allows for parallel transmissions outperforms protocols that use exclusion or power control.

I. INTRODUCTION

   An impulse-radio (IR) physical layer is a spread-spectrum physical layer [1]. It makes use of ultra-short duration pulses that can yield ultra-wide bandwidth signals [2]. A pulse duration of two nanoseconds yields a bandwidth of roughly 500 MHz. The pulses are sent infrequently, with a characteristic duty cycle lower than one percent in the case of low data-rate systems. For instance, a data-rate of 1 Mbit/s is obtained with a binary modulation and a duration of 1000 ns between the transmission of two pulses.

   Impulse-radio ultra-wide band (IR-UWB) physical layers exhibit several distinctive features. The large bandwidth of UWB radios, typically on the order of the gigahertz, allows for the resolution of multipath components [3]. This property, combined with the use of a proper radio receiver, offers a great resistance to multipath fading that usually plagues narrow-band radios. The wide bandwidth also allows for multiple-access and provides robustness to interference. The large number of degrees of freedom available can be shared by

several communications. In practice, time-hopping [2] can provide multiple-access to an IR-UWB physical layer. In a low data-rate setting, it allows *a priori* for many asynchronous and concurrent transmissions with few interferences between simultaneous transmissions. Another advantage of IR-UWB radios is high precision ranging [4], with a potential for centimeter accuracy in indoor environments. Because of their very wide bandwidth that would overlap with the bandwidth of existing systems, there are stringent radio spectrum regulations already in effect in several countries (see [5] for the USA and [6] for the European Union). Consequently, UWB systems are also characterized by extremely low power spectral densities.

   An IR-UWB physical layer might provide both robust communication and ranging capabilities for dense and low data-rate wireless network scenarios. In fact, it appears to be an ideal candidate for pervasive and sensor networks, and for Internet of Things applications.

   There are several challenges facing the designers and implementers of these networks. One challenge is robustness to interference. These networks might range from a few dozen nodes, to large-scale networks composed of hundreds of nodes with multi-hop topologies. They have no global synchronization. In this case, a tight control and coordination of the network is unrealistic and interference inevitably occurs. In addition, they might have dense topologies with a potentially high level of interference between nodes. Besides, with the increasing deployment of wireless networks, uncontrolled interference becomes problematic. Uncontrolled interference typically occurs due to several independent networks functioning in close vicinity of each other. The networks must also be robust to poor radio propagation because they might be operated in hostile environments with poor radio propagation properties, for instance heavy multipath in indoor environments. Another challenge is low power consumption, because nodes might operate with batteries or energy harvesting techniques. Furthermore, for environmental and health concerns, as well as for coexistence with other wireless technologies, it is important that the level of radiated power per node be kept very low.

   IR-UWB physical layers have the potential to address these challenges because of their wide bandwidth, resistance to multipath propagation, robustness to interference, and limited transmission power. Although the issue of low power consumption remains open, there is evidence it is possible [7]. In fact, an IR-UWB physical layer has been chosen for the IEEE 802.15.4a [8] amendment to IEEE 802.15.4 [9], a standard that

targets low data-rate wireless networks with extensive battery life and very low complexity. UWB is also attractive in high data-rate settings (e.g. Wireless USB): the wide bandwidth is used by only one source to pack as many bits as possible.

From a network design point of view, the challenges associated with low data-rate wireless networks, such as pervasive and sensor networks, are cross-layer. They concern not only the choice and design of an appropriate physical layer, but also the design of the upper layers. Moreover, the properties of UWB physical layers are very different than those of narrow-band physical layers. Consequently, the design rules and the architecture of a network composed of IR-UWB nodes are likely to be fundamentally different than those for narrow-band wireless networks.

We are interested in the design and architecture of the medium access control (MAC) and physical layers for IR-UWB networks, in a low data-rate, self-organized, and multi-hop setting. Two of the main tasks of the MAC layer are to manage interference and multiple-access to the physical layer. Existing wireless MAC protocols for narrow-band or code division multiple-access (CDMA) physical layers mostly employ mutual exclusion schemes (CSMA or TDMA), or power control, or a combination of both. With mutual exclusion, interference is simply prevented. Mutual exclusion schemes are used because of the assumption that simultaneous transmissions result in transmission errors. In contrast, thanks to the robustness of IR-UWB physical layers, it might well be that an exclusion scheme is not necessary. Then, allowing for and intelligently managing interference might actually provide a better utilization of the resources and prove to be more efficient. For instance, power control is a well-known way to manage interference. But there are also less commonly exploited possibilities for interference management. In particular, rate adaptation where the rate is adapted to the level of interference. Hence, we want to understand the design principles to manage multiple-access and interference in IR-UWB networks.

For low data-rate IR-UWB networks, the optimal network organization consists of four design principles: (1) to allow for concurrent and parallel transmissions, (2) to use an interference mitigation scheme at the physical layer, and (3) to manage interference without power control but with rate adaptation (see Section III-A). The fourth design principle is related to packet detection and timing acquisition: (4) a private acquisition preamble for each destination must be used. For complexity reasons, the network might be designed with a single and common acquisition preamble for the entire network. However, this entails a negative impact on the throughput and creates hidden terminal effects. On the contrary using a private acquisition preamble per destination leads to a higher and stable throughput compared to a common acquisition preamble (see Section III-B). Clearly, the optimal network design for low data-rate IR-UWB networks contrasts strongly with narrow-band wireless networks where exclusion schemes and power control protocols are necessary. Protocols that implement all [10] or parts [11] of the above design principles outperform protocols that use exclusion or power control.

In this paper, we do not address routing or interference management techniques for narrow-band interference (see [12], [13], [14] and the references therein). The remainder of this paper is organized as follows. We give the necessary material on IR-UWB networks in Section II, in particular on IR-UWB physical layers, the characterization of interference and interference mitigation. We address the optimal design of IR-UWB networks in Section III and how to implement it in a practical protocol in Section IV. We discuss the results and conclude the paper in Section V.

II. IR-UWB NETWORKS

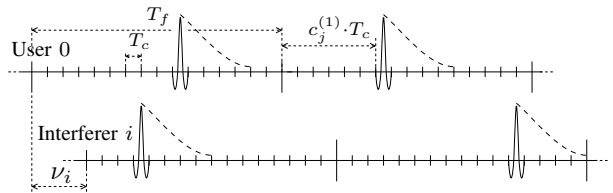### A. IR-UWB Physical Layers

Fig. 1. Impulse-radio UWB physical layer with time-hopping and an asynchronous interferer: $c_j^{(i)}$ denotes the time-hopping sequence of user i and ν_i is the delay between interferer i and the user of interest (user 0). The dashed curve following each pulse represents multipath propagation. The time-hopping positions for user 0 are 9 and 6 and for user i are 3 and 12.

The classic IR-UWB physical layer model [2] is illustrated in Figure 1 and explained in the following. Time is divided into frames of duration T_f and there is one pulse transmitted per frame. As the pulses are sent infrequently, several transmitters can share the medium concurrently. However, the transmission time of each pulse is randomized to avoid catastrophic collisions [2]. Hence, a frame is further subdivided into N_c non-overlapping *chips*; for each frame, these chips define the possible locations for the transmission of a pulse. To avoid inter-symbol interference (ISI) due to the multipath propagation channel, a guard time can reduce the number of effective available positions by N_g chips to $N_c - N_g$. A so-called time-hopping sequence (THS) of integers in $[0, N_c - N_g - 1]$ indicates which position to choose in each frame for the transmission of a pulse. Information can be transmitted thanks to pulse position modulation (PPM) or amplitude modulation.

One distinctive characteristic of UWB systems is their multipath resolvability. Multipath occurs due to reflection, refraction and scattering of radio waves by the surrounding environment. The transmitted signal reaches the receiver by more than one path. In a narrow-band system, fading occurs because the multipath components of the signal combine to produce a distorted version of the transmitted signal. In a UWB system, a direct consequence of the short duration of the pulse is that multiple paths may be separately identified [3].

B. Multiple-Access and Interference Characterization in IR-UWB Networks

The multiple-access capability of IR-UWB physical layers stems from time-hopping. From a networking point of view, the IR-UWB physical layer can be seen as a multi-channel

physical layer by considering each THS as a particular channel. Despite the infrequent transmission of the pulses and time-hopping, these channels are quasi-orthogonal and multi-user interference (MUI) can still occur whenever two IR-UWB signals overlap. Signal collisions can happen mainly because the THSs are not perfectly orthogonal and due to asynchronous signal transmissions. The multipath propagation channel also increases the occurrence of signal collisions.

The density of the MUI exhibits an impulsive shape and heavy tail characteristics [15], [16]. Indeed, the statistics of the MUI can generally not be approximated with a Gaussian distribution [17]. Nevertheless, if several users transmit concurrently, only occasional signal collisions will occur between the concurrent signals. Unlike narrow-band systems, the collisions of packets from different IR-UWB transmitters do not fully destroy the underlying radio signals, especially in low data-rate scenarios. The IR-UWB physical layer is relatively robust to interference stemming from concurrent transmissions

However, there are two additional important factors concerning interference. The first is the choice of an adequate receiver. For example, we show in [18] that the robustness of IR-UWB to MUI and thus the possibility of allowing for parallel transmissions is completely annihilated by the use of a simple energy detection receiver. A second factor is the near-far effect: interferers close to the receiver can have a much higher signal strength compared to the user of interest. The near-far effect in IR systems with time-hopping is less severe than in CDMA systems [2]. In fact, the near-far effect is only a factor when a strong pulse overlaps with a weak pulse. Still, it can drastically affect the performance in IR-UWB networks [10]. In CDMA networks, the near-far effect is solved by power control. But as we will see in Section III, power control is not optimal in IR-UWB networks. Furthermore, power control requires a large amount of coordination among nodes in the network, which entails a cost on the performance because of signaling overhead.

C. Interference Mitigation Mechanisms for IR-UWB Physical Layers

In IR-UWB networks, interference mitigation schemes at the physical layer can take advantage of the structure of IR signals to reduce the effect of interference from concurrent transmitters with little cost on the performance, especially in near-far cases (e.g. [19], [20], [10], [21], [22]).

Interference mitigation schemes for IR-UWB physical layers are *not* multi-user detectors. They have a much lower complexity and do not require estimating the signal characteristics from the interferers. Actually, they originate from work on the design of optimal receivers for the detection of signals in non-Gaussian interference: With an impulsive interference model, [23] suggests applying a non-linear function on the received signal prior to demodulation. In [24] a very simple thresholding operation is used as a possible non-linearity. A similar principle is used by many interference mitigation schemes for IR-UWB. The next section demonstrates that allowing for concurrent transmissions with an interference mitigation scheme at the physical layer is optimal for low data-rate IR-UWB networks.

Of course, one of the simplest ways to mitigate interference could be to prevent it, for example by coordinating access to the physical layer. But, as we already explained, this cannot prevent MUI due to uncontrolled activities in neighboring networks (e.g., several IEEE 802.15.4a piconets running in parallel). Interference must be taken into account already in the design of the physical layer.

D. Packet Detection and Timing Acquisition

Another crucial aspect of IR-UWB networks is packet detection and timing acquisition. As the network is packet-based and there is no global synchronization, the first step towards the correct reception of a packet is packet detection and timing acquisition. For a particular destination, it consists of detecting the packets that are intended for itself and finding the time reference of the source. Only then can the destination recover the payload by demodulating the received signal. Notice that even if their exists a global synchronization in a network, packet *detection* may still be necessary. Also, because of the wide bandwidth and low emitted power, carrier-sensing is sorely feasible with IR-UWB physical layers

Packet detection and timing acquisition is performed on a per packet basis and typically relies on the presence of a so-called acquisition preamble at the beginning of each packet. This acquisition preamble can be generated using a THS, amplitude modulation, or a combination of both. The exact functioning of a packet detection and timing acquisition algorithm is out of the scope of this article. But, how this preamble is chosen is a network design issue and may have quite an impact on the network performance. A simple design choice of the network is to use a common acquisition preamble for the whole network. A second design choice is to use an acquisition preamble that is private to each destination The next section will discuss the effect of this choice on the throughput of IR-UWB networks.

III. DESIGN OF IR-UWB NETWORKS

A. Interference Management and Multiple-Access in IR-UWB Networks

1) Optimal Design of IR-UWB Networks: The work in [25] addresses the optimal design of IR-UWB networks. The findings are the following: While receiving, a node should maintain an exclusion region around itself; nodes inside the exclusion region should remain silent during the reception, whereas nodes outside of this region can transmit in parallel. The size of the exclusion region depends only on the power constraints of the source of the transmission, and not on the length of the link or the positions of other nodes. Additionally, the receiver and its corresponding sender should adapt the rate of the transmission according to the amount of noise and interference at the receiver. Finally, when a node is transmitting, it should do so with maximum power. An earlier result [26] also demonstrates that nodes should transmit at maximum power.

These findings indicate that the optimal MAC protocol in an IR-UWB network should be, *a priori*, a combination of rate adaptation and mutual exclusion. An exclusion protocol

or power control are not the optimal strategies in this setting. The results in [25] are optimal in terms of rate efficiency. In terms of energy efficiency, they are not optimal. But there is strong evidence in [27] that they are close to optimal. Although we do not address routing in this article, results in [25] for static networks show that minimum energy routing is optimal, both from an energy and rate performance viewpoint.

The size of the exclusion region around destinations is a critical parameter. An exclusion region around a particular node is *active* whenever this node is receiving. The size of the exclusion region should be adapted to the parameters of the network. If this size is large, then a protocol is necessary to ensure that sources inside an active exclusion region are prevented from sending. If it is rather small, then the overhead of the exclusion protocol is likely to be balanced by the increased spatial reuse in case no exclusion is enforced. If the size of the exclusion region is negligible, then no exclusion protocol is necessary and it becomes optimal to allow for concurrent and parallel transmissions. The optimal MAC protocol in this case requires only rate adaptation for interference management. For multiple-access nodes can decide to transmit whenever they have some packet ready to transmit.

For low data-rate and low power IR-UWB networks, it might be that the size of the exclusion region is small enough that no exclusion protocol is necessary. Computing the size of the exclusion region is a highly difficult problem [25]. But, numerical simulations can be used instead [10], [27].

2) There is no Exclusion Region With Interference Mitigation in a Low Data-Rate IR-UWB Network: It might seem obvious that the size of the exclusion region becomes negligible for low data-rate IR-UWB networks. However, it is not. In fact, numerical simulations in [10] using topologies with strong near-far scenarios and in [27] with a set of random topologies indicate that that even in low data-rate conditions (with a rate around 1 Mbit/s), the size of the exclusion region might be non negligible. The reason is that even in low data-rate settings, the IR-UWB physical layer can be sensitive to near-far effects (Section II-B). But, further results from [10], [27] show that with an interference mitigation scheme at the physical layer, the size of the exclusion region becomes negligible. Actually, even with interference mitigation, the activity of one near-far user can have a severe effect on the rate achievable by other users. But the rate reduction is still less that what would be lost by an exclusion protocol [10].

These results have been recently confirmed in [28], where the authors compute lower and upper bounds on the achievable rate with an IR-UWB receiver able to mitigate interference at the physical layer. The authors use this result to demonstrate that a non-coordinated MAC exhibits better performance than more complex coordinated solutions. This result clearly indicates that most of the complexity should be invested in a receiver design (to mitigate MUI, as in [21], [22]) instead of intricate MAC or signaling protocols.

3) Implications of the Optimal Design for Practical Protocols: The main implication from the previous section is that the optimal MAC protocol for low data-rate IR-UWB networks should allow for parallel and concurrent transmissions. Interference is managed with rate adaptation, no power control and

an interference mitigation scheme at the physical layer.

The concept of rate adaptation has also been proposed for 802.11 networks [29], but only to track the state of the channel and adapt the signal to noise ratio at the receiver (basically, the distance to the source and the destination). In contrast, we use rate adaptation as support for multiple-access, to manage MUI.

Consequently, the optimal design for low data-rate IR-UWB networks is fully uncoordinated. Prior to the transmission of a packet, no coordination is required. Indeed, no exclusion or power control protocol that requires coordination among nodes is necessary. And rate adaption concerns only the source and the destination, Still, even though nodes can send whenever they have a packet ready to transmit, this does not mean that a protocol is not necessary at the link layer. There are two issues. First, there actually remains some exclusion to implement because generally a node can be engaged exclusively in either the reception or the transmission of a single packet. A protocol is necessary to arbitrate access to the medium between several senders that want to communicate to the same destination. This is not a straightforward task because carrier-sensing is sorely feasible with IR-UWB physical layers. Second, a protocol is required to continuously adapt the transmission rate to noise and interference. In Section IV, we will present DCC-MAC [10]. It is a MAC protocol based on the findings of the previous section, which (1) solves the above issue of multiple-access to a given destination and (2) implements a rate adaptation algorithm that does not rely on channel measurements.

In the context of low complexity implementations, there is also a need for feasible and low-complexity interference mitigation schemes. The work in [22] presents a practical and low-complexity IR-UWB receiver with interference mitigation capabilities.

B. Effect on the Throughput of the Choice of the Acquisition Preamble

The low data-rate IR-UWB networks that we consider are packet based and have no global synchronization. Hence, the first step towards correct packet reception is packet detection and timing acquisition: Before recovering the payload of the packet, the destination must detect the packet on the medium and determine when exactly the payload begins. In fact, even if there exists a global synchronization in the network, packet detection is still necessary. Packet detection and timing acquisition are performed on a per packet basis and typically rely on the presence of an acquisition preamble at the beginning of each packet. How this preamble is chosen is a network design issue and has an impact on the performance of the network.

1) A Network Design Choice: Common or Private Acquisition Preambles: Two extreme possible design choices can be compared in order to evaluate how the choice of the acquisition preamble affects the throughput for IR-UWB networks. First, similar to the IEEE 802.15.4a amendment [8], a simple design choice of the network is to have an identical and common acquisition preamble for the entire network. Second, as for DCC-MAC or [11], another design choice is to have a private acquisition preamble per destination. With private acquisition

preambles, all communications for a given destination use an acquisition preamble (and possibly a time-hopping sequence) that are private to this destination. For example, with DCC-MAC, a source generates the acquisition preamble of its intended destination using the THS of this destination. This THS itself can be easily derived from the MAC address of the destination by using the MAC address as seed of a random number generator. The throughput of a network using a private acquisition preamble is likely to be much higher than the throughput of a network using a common acquisition preamble. Indeed, in a network with a common acquisition preamble, a packet might contend for timing acquisition with packets sent by any node in the entire network. In contrast, with a private acquisition preamble, the contention is reduced to packets transmitted to the same destination.

2) Common Acquisition Preambles Degrade the Throughput: In [30], [31] we show that a private acquisition preamble can yield a throughput gain larger than 100% compared to a common acquisition preamble. For a given source/destination link, the throughput difference grows with the number of concurrent transmitters. Furthermore, ns-2 [32] simulations on multi-hop topologies with TCP flows demonstrate that a network using private acquisition preambles exhibits a stable throughput. On the contrary, using a common acquisition preamble exhibits several performance anomalies. In particular, compounding effects similar to the exposed terminal issue in IEEE 802.11 networks appear: the throughput is severely degraded and complete flow starvation might occur. In addition, the use of a common acquisition preamble can result in very large performance fluctuations.

3) Implications of Private Acquisition Preambles: With a private acquisition preamble comes the cost of learning the acquisition preamble of the destination. Hence the throughput increase must be sufficiently large in order to alleviate the associated cost. Learning the acquisition preamble is equivalent to obtaining a unique identifier for a node. If TCP/IP is running, nodes have a unique identifier that must be known by neighboring nodes and the acquisition preamble can be easily generated from this unique identifier. In IEEE 802.15.4 networks, nodes have a unique EUI-64 identifier that could also be used to generate an acquisition preamble. Besides, with digital hardware implementations, the use of private acquisition preambles is essentially not more costly than a single common preamble [7]. A node does not need to listen to more than a few preambles [10]: generally, its own preamble, the one from the destination and a preamble for the broadcast address. Hence, the cost of using private acquisition preamble is very low. From a network design point of view, using private acquisition preambles inherently creates one contention domain per destination.

IV. DCC-MAC: A MAC LAYER THAT ALLOWS FOR CONCURRENT TRANSMISSIONS

This section is not meant to be a complete description of DCC-MAC but summarizes the issues associated with the implementation of the optimal design principles of Section III in a practical protocol. As explained previously in Section III-A3, a protocol is still necessary at the link layer.
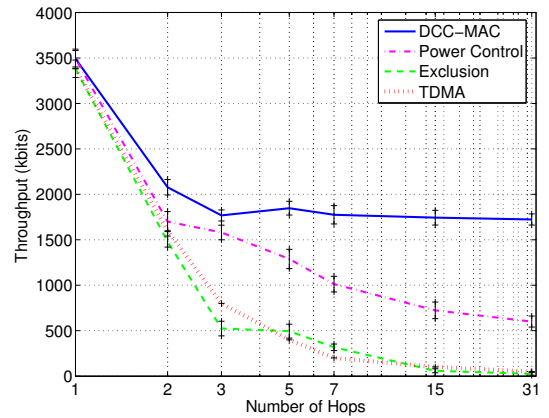

Fig. 2. Average throughput of DCC-MAC vs. number of hops on a multi-hop line topology with the source and destination on each extremity of the line. Traffic is UDP. We also show the throughput of a power control protocol similar to [26], of a mutual exclusions protocol similar to 802.11 and of a TDMA protocol. All MAC use the same IR-UWB physical layer. There is almost no drop in throughput for DCC-MAC as the number of hops increases. More performance results can be found in [10].

For rate adaptation, a protocol is required. To avoid the problem of signal to interference and noise ratio (SINR) measurements, the rate is adapted with an additive increase/multiplicative decrease scheme based on whether a packet is properly received or not. For each packet, feedback to tell the source to increase or decrease the rate is piggybacked in the acknowledgment to the sender. There also remains some exclusion to implement because a node can be engaged exclusively in either the reception or the transmission of a single packet. In addition there is the challenge of absence of carrier-sensing. Contention for the same destination is enforced by the so called "private MAC". It uses private acquisition preambles and private THSs and is an elaborate signaling protocol that alternates between direct access to a receiver and an invitation-based scheme. With DCC-MAC, sources simply send spontaneously. If this fails, they use a backoff mechanism similar to the IEEE 802.11 MAC. However, there is no request message, since potential collisions usually do not result in packet loss. But, sources send an idle message at the end of each packet transmission. This allows sources, where packet transmission failed, to wait for this idle message before attempting the retransmission. With this combination, the behavior in case of contention is the same as with CSMA, even though carrier sensing is never done. DCC-MAC is fully uncoordinated and does not use any separate channel for control signaling to avoid any global contention possibilities. DCC-MAC is actually quite close to the mandatory MAC protocol in the IEEE 802.15.4a amendment [8] with the following notable exceptions: in IEEE 802.15.4a a common acquisition preamble is used, there is no rate adaptation and no idle packet.

A. Parallel Transmissions with Interference Mitigation at the Physical Layer Outperform Exclusion Protocols

The DCC-MAC protocol is fully implemented in ns-2 along with a model of the IR-UWB physical layer [33]. The performance of DCC-MAC has been thoroughly evaluated by simulations in [10] and compared to the performance

obtained with a power control protocol similar to [26], to a mutual exclusions protocol similar to 802.11 and to a TDMA protocol. See Figure 2 for the average throughput on a multi-hop line topology with UDP traffic. The results in Figure 2 and [10] show that an uncoordinated protocol that allows for parallel transmissions, with interference mitigation at the physical layer largely outperforms protocols based on exclusion or power control. Furthermore, DCC-MAC shows a stable performance on multi-hop topologies. In [34], the performance of DCC-MAC was also compared with the UWB^2 protocol in [11]. The UWB^2 protocol uses a global control channel with an RTS/CTS exchange on this channel prior to data transmission. The simulations in [34] clearly exhibited the performance degradation entailed by the use of this control channel compared to DCC-MAC.

V. DISCUSSION AND CONCLUSION

In the case of low data-rate IR-UWB networks, the optimal network organization consists of four design principles: (1) to allow for concurrent and parallel transmissions, (2) to use an interference mitigation scheme at the physical layer, (3) to manage interference without power control but with rate adaptation, and (4) to use a private acquisition preamble and THS per destination. This optimal network design for low data-rate IR-UWB networks contrasts strongly with narrow-band wireless networks where exclusion schemes and power control protocols are necessary. These results are optimal in terms of rate efficiency. In terms of energy efficiency, they are not optimal. But there is strong evidence in [27] that they are close to optimal. Based on the four design principles, we have developed DCC-MAC, a practical uncoordinated MAC protocol for IR-UWB networks where concurrent transmissions are allowed. DCC-MAC clearly outperforms MAC protocols that are based on exclusion or power control.

REFERENCES

[1] R. Pickholtz, D. Schilling, and L. Milstein, "Theory of spread-spectrum communications–a tutorial," *IEEE Trans. Commun.*, vol. 30, no. 5, pp. 855–884, May 1982.
[2] R. Scholtz, "Multiple access with time-hopping impulse modulation," in *IEEE Military Communications Conference (MILCOM)*, vol. 2, 1993, pp. 447–450 vol.2.
[3] M. Z. Win and R. A. Scholtz, "On the robustness of ultra-wide bandwidth signals in dense multipath environments," *IEEE Commun. Lett.*, vol. 2, no. 2, pp. 51–53, 1998.
[4] J.-Y. Lee and R. A. Scholtz, "Ranging in a dense multipath environment using an uwb radio link," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1677–1683, 2002.
[5] Federal Communications Commission, United States, "FCC 02–48, First Report and Order," February 14th 2002.
[6] European Radiocommunications Office (ERO), Electronic Communications Committee, "Document ECC/DEC/(06)04," 2007.
[7] L. Chalard, D. Helal, L. Verbaere, A. Wellig, and J. Zory, "Wireless sensor networks devices: Overview, issues, state-of the art and promising technologies," *ST Journal of Research*, vol. 4, no. 1, April 2007.
[8] IEEE Computer Society, LAN/MAC Standard Committee, "IEEE P802.15.4a/D7 (amendment of IEEE std 802.15.4), part 15.4: Wireless medium access control and physical layer specifications for low-rate WPAN," Jan. 2007.
[9] ——, "IEEE standard for information technology- telecommunications and information exchange between systems- local and metropolitan area networks- specific requirements part 15.4: Wireless medium access control and physical layer specifications for low-rate WPANs," IEEE Std 802.15.4-2006 (Revision of IEEE Std 802.15.4-2003), 2006.
[10] R. Merz, J. Widmer, J.-Y. Le Boudec, and B. Radunovic, "A joint PHY/MAC architecture for low-radiated power TH-UWB wireless ad-hoc networks," *Wireless Commun. and Mobile Comput. J., Special Issue on Ultrawideband (UWB) Communications*, vol. 5, no. 5, Aug. 2005.
[11] M.-G. Di Benedetto, L. Nardis, M. Junk, and G. Giancola, "(UWB)^2: Uncoordinated, wireless, baseborn, medium access control for UWB communication networks," *Mobile Networks and Applications*, vol. 10, no. 5, Oct. 2005.
[12] I. Bergel, E. Fishler, and H. Messer, "Narrowband interference mitigation in impulse radio," *IEEE Trans. Commun.*, vol. 53, no. 8, 2005.
[13] C. da Silva and L. B. Milstein, "The effects of narrowband interference on uwb communication systems with imperfect channel estimation," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 4, pp. 717–723, 2006.
[14] S. M. Mishra, R. W. Brodersen, S. T. Brink, and R. Mahadevappa, "Detect and avoid: an ultra-wideband/WiMAX coexistence mechanism," *IEEE Commun. Mag.*, vol. 45, no. 6, pp. 68–75, 2007.
[15] A. Forouzan, M. Nasiri-Kenari, and J. Salehi, "Performance analysis of time-hopping spread-spectrum multiple-access systems: uncoded and coded schemes," *IEEE Trans. Wireless Commun.*, vol. 1, no. 4, pp. 671–681, Oct. 2002.
[16] Y. Dhibi and T. Kaiser, "On the impulsiveness of multiuser interferences in TH-PPM-UWB systems," *IEEE Trans. Signal Process.*, vol. 54, no. 7, pp. 2853–2857, 2006.
[17] G. Durisi and G. Romano, "On the validity of gaussian approximation to characterize the multiuser capacity of UWB TH PPM," in *IEEE UWBST*, 2002, pp. 157–161.
[18] M. Flury, R. Merz, J.-Y. le Boudec, and J. Zory, "Performance evaluation of an IEEE 802.15.4a physical layer with energy detection and multi-user interference," in *IEEE ICUWB*, 2007.
[19] R. Knopp and Y. Souilmi, "Achievable rates for UWB peer-to-peer networks," in *IZS*, 2004.
[20] W. M. Lovelace and J. K. Townsend, "Threshold discrimination and blanking for large near-far power ratios in UWB networks," *IEEE Trans. Commun.*, vol. 53, no. 9, pp. 1447–1450, 2005.
[21] M. Flury and J.-Y. Le Boudec, "Interference mitigation by statistical interference modeling in an impulse radio UWB receiver," in *IEEE ICUWB*, 2006.
[22] M. Flury, R. Merz, and J.-Y. le Boudec, "An energy detection receiver robust to multi-user interference for IEEE 802.15.4a networks," in *IEEE ICUWB*, 2008.
[23] A. Spaulding and D. Middleton, "Optimum reception in an impulsive interference environment–part I: Coherent detection," *IEEE Trans. Commun.*, vol. 25, no. 9, pp. 910–923, 1977.
[24] K. Vastola, "Threshold detection in narrow-band non-gaussian noise," *IEEE Trans. Commun.*, vol. 32, no. 2, pp. 134–139, 1984.
[25] B. Radunovic and J. Y. Le Boudec, "Optimal power control, scheduling and routing in UWB networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 7, pp. 1252–1270, Sep. 2004.
[26] F. Cuomo, C. Martello, A. Baiocchi, and C. Fabrizio, "Radio resource sharing for ad hoc networking with UWB," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1722–1732, December 2002.
[27] A. El Fawal, J.-Y. Le Boudec, R. Merz, B. Radunovic, J. Widmer, and G. M. Maggio, "Tradeoff analysis of PHY-aware MAC in low-rate, low-power UWB networks," *IEEE Commun. Mag.*, vol. 43, no. 12, 2005.
[28] B. Radunovic, J.-Y. Le Boudec, and R. Knopp, "Optimal PHY and MAC Protocols for Wide-Band Ad-Hoc Networks," in *Forty-Fifth Annual Allerton Conference*, 2007.
[29] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *MOBICOM*, 2001, pp. 236–251.
[30] R. Merz, J.-Y. Le Boudec, and S. Vijayakumaran, "Effect on network performance of common versus private acquisition sequences for impulse radio uwb networks," in *IEEE ICUWB 2006*, 2006.
[31] R. Merz and J.-Y. Le Boudec, "Performance evaluation of impulse radio UWB networks using common or private acquisition preambles," *Submitted to IEEE Transactions on Mobile Computing*, 2008.
[32] "The network simulator ns-2, http://nsnam.isi.edu/nsnam/," Jan. 2008.
[33] R. Merz, J.-Y. Le Boudec, and J. Widmer, "An Architecture for Wireless Simulation in NS-2 Applied to Impulse-Radio Ultra-Wide Band Networks," in *10th Communications and Networking Simulation Symposium*, 2007.
[34] M. Flury, R. Merz, and J.-Y. Le Boudec, "Managing impulsive interference in impulse radio UWB networks," *ST Journal of Research*, vol. 4, no. 1, May 2007.