Concurrent Transmissions in IR-UWB Networks: an Experimental Validation

James Colli-Vignarelli[†*] Jérôme Vernez[#] Ruben Merz[‡]
Catherine Dehollain[†] Stephan Robert[#] Jean-Yves Le Boudec[*]
† EPFL, School of Engineering, ∗ EPFL, School of Computer and Communication Sciences,
‡ Deutsche Telekom Laboratories, TU Berlin, # HEIG-VD
{james.colli-vignarelli,catherine.dehollain}@epfl.ch, {ruben.merz}@telekom.de,
{jean-yves.leboudec}@epfl.ch, {jerome.vernez,stephan.robert}@heig-vd.ch

Abstract—

We demonstrate and validate that concurrent and parallel transmissions are feasible for low data-rate impulse-radio ultra-wide band (IR-UWB) physical layers. The optimal organization for a low data-rate IR-UWB network is to allow for concurrent transmissions at the link layer, and to use interference mitigation techniques at the physical layer. Indeed, low data-rate IR-UWB physical layers can suffer from multi-user interference (MUI), especially in near-far scenarios. Therefore, we perform an experimental validation with a packet-based, low data-rate IR-UWB physical layer testbed. Our results clearly demonstrate that concurrent IR-UWB transmissions are possible. This shows that completely uncoordinated low data-rate IR-UWB networks are feasible. We also demonstrate that, in the presence of MUI, an interference mitigation scheme at the physical layer is indeed necessary. In particular, because it is the first component for the proper reception of a packet, we focus on packet detection and timing acquisition. We show that a traditional scheme is not robust against multi-user interference and prevents concurrent transmissions. On the contrary, a scheme designed to take MUI into account, even with a very simple interference mitigation scheme, allows for concurrent transmissions, even in strong near-far scenarios.

I. INTRODUCTION

An impulse-radio ultra-wide band (IR-UWB) physical layer makes use of ultra-short duration pulses to produce extremely wide bandwidth signals [1]. IR-UWB physical layers exhibit several distinctive features. The large bandwidth of UWB radios, typically on the order of a gigahertz, allows for the resolution of the multipath components of the propagation channel [2]. This property, combined with the use of a proper radio receiver, offers a great resistance to multipath fading that usually plagues narrowband 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 provides multiple-access capabilities [1]. In a *low data-rate* setting, time-hopping allows *a priori* for many

asynchronous and concurrent transmissions with few interferences between simultaneous transmissions. However, as we explain later in this paper, concurrent transmissions require the presence of interference mitigation functionalities at the physical layer. Another advantage of IR-UWB radios is high precision ranging [3], with a potential for centimeter accuracy in indoor environments. Hence, IR-UWB physical layers may provide both robust communication and ranging capabilities for dense and low data-rate wireless network scenarios.

The properties of UWB physical layers are very different than those of narrowband physical layers. Consequently, the design rules and the architecture of a network composed of IR-UWB nodes are fundamentally different than those for narrowband wireless networks [4], [5], [6]. Narrowband networks need exclusion protocols and power control. On the contrary, for low data-rate IR-UWB networks, the optimal network organization is not to use any exclusion protocols, or power control [4] (see also [7] regarding the absence of power control). Rather, concurrent transmissions should be allowed at the link layer, rate adaptation should be used for interference management and an interference mitigation scheme (see [8], [9] and the references therein) should be used at the physical layer. Indeed, even in a low data-rate setting, an IR-UWB physical layer can suffer from multi-user interference (MUI), especially in near-far scenarios [5]. Essentially, allowing for a bit of complexity at the physical layer allows for a completely uncoordinated network design [6]. Protocols that implement all [5] or parts [10] of the above design principles outperform protocols that use exclusion or power control.

Yet, the practical feasibility of such a design has to be experimentally tested and validated. In fact, all the results previously mentioned were obtained through theoretical analysis and numerical simulations. Therefore, the contribution of this paper is an experimental validation with a packet-based IR-UWB hardware testbed that shows concurrent IR-UWB transmissions are possible in a low data-rate setting. In particular, we focus on packet detection and timing acquisition: We show experimentally for timing acquisition that a traditional scheme is not robust enough against multi-user interference, and prevents concurrent transmissions. Whereas, a scheme designed to take MUI into account [11], even with a very simple interference mitigation scheme, allows for

concurrent transmissions, even in strong near-far scenarios, and outperforms a traditional scheme. Indeed, in the presence of MUI, an interference mitigation scheme at the physical layer is necessary.

Our results do not demonstrate a completely working uncoordinated IR-UWB network because our testbed consists of one receiver and several transmitters. Also, we implement a robust packet detection and timing acquisition, but not yet a robust demodulation and decoding. In fact, our results in Section III show that a robust demodulation and decoding is definitely required too. This is scheduled for future work. However, our results demonstrate functional, parallel and asynchronous IR-UWB transmissions and the feasibility of completely uncoordinated detection and timing acquisition.

Note that the primary objective of our experimental hardware testbed is extensibility, modularity and flexibility. We want the ability to easily exchange components in the RF chain. We want full access and control of the parameters of the physical layer, and the possibility to easily implement and test new algorithms at the receiver. For instance, either by programming an FPGA for real-time processing, or by capturing signal traces to use offline with an algorithm implemented in a high-level programming language. Challenges such as low power consumption or integration are definitely important but are not primary objectives. This is why most of the RF elements are either off-the-shelf components or are built with discrete-components.

The remainder of this paper is organized as follows: In Section II, we describe the experimental system model of our testbed and the algorithms used along with some important assumptions and limitations of our hardware. We present the results of our experimental performance evaluation in Section III. We conclude the paper in Section IV.

II. EXPERIMENTAL SYSTEM MODEL AND ASSUMPTIONS

Our experimental setup comprises several transmitters and one receiver (we use two transmitters in Section III for our experiments).
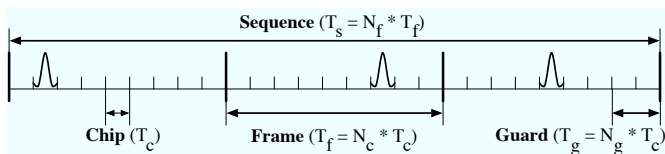

Fig. 1. Classic IR-UWB signal [1] and its parameters: T_c is the duration of a chip, $T_f = N_c T_c$ is the duration of a frame and $T_s = N_f T_f$ is the duration of a sequence. $T_g = N_g T_c$ is guard time used to prevent ISI.

The physical layer signal generated by each transmitter is a classic IR-UWB signal with time-hopping (TH) as in [1] (see Figure 1): Time is divided into frames of duration T_f and there is one pulse transmitted per frame. Because 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 [1]. 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 reduces the number of effective available positions by N_g chips to $N_c - N_g$.

A so-called pseudo-random time-hopping sequence (THS) of integers uniformly distributed in $[0, N_c - N_g - 1]$ indicates which position to choose in each frame for the transmission of a pulse. Hence, each transmitter has its own THS, which is independently generated. Information is transmitted thanks to binary pulse position modulation (BPPM), where the position of a pulse carrying a one is shifted by a duration T_m and the position of a pulse carrying a zero is left unchanged. More formally, the baseband IR-UWB signal with BPPM of the n-th transmitter is

$$s^{(n)}(t) = \sum_i p(t - iT_f - c_i^{(n)}T_c - d_i^{(n)}T_m) \qquad (1)$$

where $p(t)$ is a pulse, $c_i^{(n)}$ is an element of the THS of transmitter n and $d_i^{(n)} \in \{0, 1\}$ is an information-bearing bit.

In practice, the pulses of the IR-UWB signal generated by our transmitters are simply square pulses of duration T_p up-converted at f_c. For time-hopping, for complexity reasons, we do not generate a continuous stream of time-hopping positions. Instead, for each transmitter, we generate a sequence of length N_f of time-hopping positions and use this sequence repeatedly i.e. $c_i^{(n)} = c_{i+N_f}^{(n)}$. The aggregation of N_f frames is a sequence.

Our experimental system is packet based (see Figure 2). Prior to a payload of length L_{pay} pulse (or equivalently L_{pay} bit because of the binary modulation), there is a preamble of length L_{pre} followed by a so-called preamble delimiter of length L_{del}. The preamble delimiter is used to detect the beginning of the payload once timing acquisition is achieved. Notice that thanks to time-hopping, each transmitter has its own distinct preamble *and* preamble delimiter. Indeed, the preamble delimiter uses another THS than preamble. Hence, a receiver can detect and acquire the timing of the packet from a given transmitter while another transmitter is active.

L_{pre} (32*N_f)	L_{del} (8*N_f)	L_{pay} (127 bytes)
Preamble	Preamble delimiter	Payload (Binary Pulse Position Modulation)

Fig. 2. The structure of the packet sent by the transmitter of interest is loosely based on the IEEE 802.15a standard [12]. It contains three parts: (1) a preamble for packet detection and timing acquisition, (2) a preamble delimiter is used to detect the beginning of the payload and (3) the payload. N_f is the length of a sequence.

Our system is completely asynchronous: There is no global synchronization between the transmitters and the receiver. Hence, at the receiver side, the first operation is packet detection and timing acquisition [13]. We implement two possible algorithms: (1) a conventional correlation based algorithm [14] and (2) a multi-user (MUI) robust algorithm [11]. Both algorithms are correlation based: the received signal is continuously correlated with a template. For the conventional algorithm, the correlation Φ of the received signal with the template is the sum of several individual correlations with a

pulse template corresponding to a squared impulse i.e.

$$\Phi = \sum_{i=0}^{L_T-1} \Phi_i \qquad (2)$$

where Φ_i is the i-th individual correlation result and L_T is the length of the template. However, the algorithm in [11] uses a thresholding-based interference mitigation mechanism (see [8], [9] and the references therein): instead of summing the individual correlation, the algorithm in [11] applies a hard decision on the output of each individual correlation, i.e. (2) becomes

$$\Phi = \sum_{i=0}^{L_T-1} 1_{\{\Phi_i \geq \nu\}} \qquad (3)$$

where ν is the hard decision threshold and $1_{\{x\}}$ is the indicator function. This thresholding operation prevents spurious strong correlation outputs, essentially due to near-far interferers, which may completely blind the presence of a valid transmitter. As shown by simulations in [11] and experimentally in Section III, the robust algorithms allows for successful timing acquisitions in the presence of MUI.

A detailed description of both algorithms lies outside the scope of this article. Still, in our implementation, both algorithms bypass coarse timing acquisition and perform directly a fine timing acquisition (they are performed in the digital domain). Once the fine timing acquisition is achieved, a verification phase follows. The initial timing acquisition is declared successful if $\Phi > \sigma_1$. The verification phase is declared successful if for N_v subsequent correlations, we have $\Phi > \sigma_2$. If the verification phase is successful, the receiver begins to search for the preamble delimiter. If successful, the receiver can start the demodulation of the payload.

For the demodulation, we currently perform energy detection: The decision rule compares the energy contained at the two possible locations for a zero or a one [14]. The energy at each position is gathered over a duration T_{int}. We use energy detection because of its simplicity and also because we can avoid any channel estimation. At this moment, we do not have any error correction code. In practice, after band-pass filtering and a low-noise amplification, our receiver performs an IQ down-conversion. Even though we perform energy detection, we use a mixer for the demodulation for extensibility reasons: We want the ability to implement coherent processing in addition to energy detection. The signal after down-conversion is fed to an analog-to-digital converter (ADC) running at f_S. The ADC is coupled to an FPGA. The samples are then directly moved into a large DRAM on the FPGA and stored. The remaining operations, packet detection and timing acquisition and demodulation, are all performed in the digital domain. For instance, for energy detection, the integration over a duration T_{int} is replaced by a summation of N_{int} samples where $N_{int} = T_{int} \cdot f_S$. These operations are also all performed offline after the signal trace contained in the DRAM is offloaded to a PC. Typically, our ADC is running at $f_S = 2$ GS/s with an 8 bits resolution and we have 512 Mbyte of DRAM. Hence, we can store a signal trace of about 0.256 seconds. In our setting, this is sufficient to capture one or more packets and process them offline. We could

use the FPGA to process the signal in real-time. However, an offline processing, although time consuming, allows for a greater flexibility. Real-time processing is scheduled for future work.

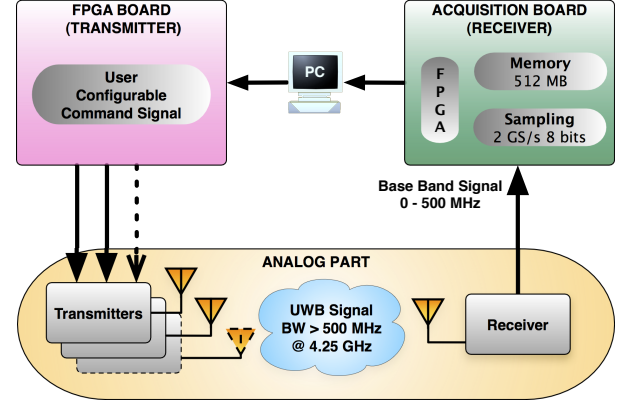An overview of our experimental testbed with its characteristics is shown in Figure 3.

Fig. 3. Overview of the testbed: an FPGA board controls the transmitters with the desired signal characteristics. The receiver amplifies, filters and down-converts the received signal to the baseband. The acquisition board of the receiver samples the baseband signal and stores this signal trace in DRAM on another FPGA. The signal trace is later offloaded to a PC for offline processing.

III. PERFORMANCE EVALUATION

A. Measurements Settings and Scenarios

We consider three different scenarios for our experimental performance evaluation. First, a single user scenario (scenario A, *single user*, Figure 4(a)): one transmitter located at distance L of the receiver. We use Scenario A to validate our implementation and to obtain reference results. For the second scenario, (scenario B, *equal power*, Figure 4(b)), we have an additional interfering transmitter also at distance L from the receiver. Hence, the received power from the transmitter of interest and from the interfering transmitter is roughly identical. For the third scenario (scenario C, *near-far*, Figure 4(c)), the additional interfering transmitter is positioned in a near-far situation at a distance $L_{nf} < L$ from the receiver. Hence, the received power from the interfering transmitter is much higher than the power received from the transmitter of interest. This is a typical sensor networks or ad hoc network situation. We perform all experiments with line-of-sight (LOS) propagation. We consider $L = 1, 3, 6$ meters and $L_{nf} = 0.2$ meter. Remember that thanks to time-hopping, each transmitter has its own distinct preamble and preamble delimiter (Section II). Hence, a receiver can detect and acquire the timing of the packet from a given transmitter while another transmitter is active.

Our main performance metric is the percentage of achieved packet detection and timing acquisition (A_{PDTA}). A packet detection and timing acquisition is declared `achieved` if a preamble delimiter is detected. However, this does not ensure that packet detection and timing acquisition is absolutely successful. For this, we would need a time reference for the received packet in order to compare with the timing acquired
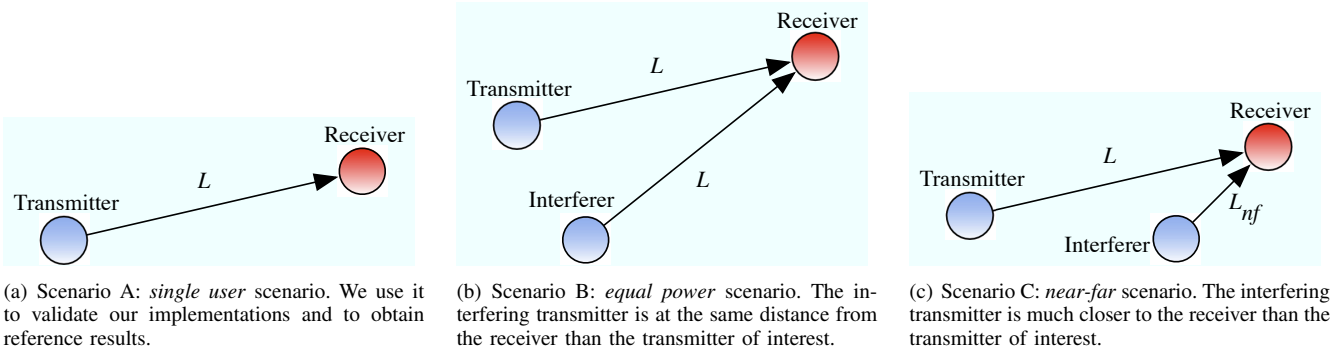
Fig. 4. The three topologies used for our experiments: a single user scenario, an equal power scenario and a near-far interference scenario. L and L_{nf}, where $L_{nf} < L$, are distances between the transmitters and the receiver.

(a) Scenario A: *single user* scenario. We use it to validate our implementations and to obtain reference results.

(b) Scenario B: *equal power* scenario. The interfering transmitter is at the same distance from the receiver than the transmitter of interest.

(c) Scenario C: *near-far* scenario. The interfering transmitter is much closer to the receiver than the transmitter of interest.

by our timing acquisition algorithm. However, for each packet where timing acquisition is achieved, we can additionally verify the BER of the payload. We observe during trial runs that if timing acquisition is incorrect, the BER is generally above 0.3 (even with MUI). Hence, for each packet where timing acquisition is achieved, we always compute the BER of this packet. If the BER is above a verification threshold $\gamma = 0.3$, we declare that timing acquisition failed.

In addition, we consider two other performance metrics: the packet error rate (PER) and the bit error rate (BER). For the PER a packet is considered corrupted when at least one bit is erroneous. The BER is computed only with payloads where timing acquisition was achieved. The PER reports both timing acquisition errors and erroneous bits in the payload.

In the equal power and near-far scenario, the transmitter of interest sends packets. However, the interfering transmitter sends a continuous signal. This ensures that we evaluate the performance in a worst case scenario. To obtain statistically meaningful results we transmit 10,000 packets per experiment run. For each experiment run, we always run both the MUI robust algorithm and the conventional algorithm on the same trace.

As explained in Section II, for each packet transmission, we first store the received signal in the DRAM of the FPGA. We then transfer the signal trace to a PC for offline processing. We perform all the processing with Matlab where we extensively use external functions written in C. Still, a typical 10,000 packet experiment run lasts around 20 hours. Also, before each experiment run, we always carefully recalibrate the hardware: We verify and adjust the power spectral density of our signal to make sure it is FCC compliant and we adjust the hard decision thresholds; We set the range of the ADC to avoid any clipping on the strongest received signal. These last two operations, threshold setting and range calibration, are normally done automatically by estimating the noise variance and by performing automatic gain control (AGC) at the receiver. This is left for future work.

The parameters of the physical layer (see Section II) are assigned the following values: $N_c = 128$ chips, $N_g = 8$ chips and $N_f = 8$ frames. With $T_c = 6$ ns (Appendix A), this corresponds to a rate around 1.3 Mbit/s. For the ADC, $f_s = 2$ GS/s. For the packet parameters (see Figure 2), we use $L_{pre} = 32*N_f$, $L_{del} = 8*N_f$, $L_{pay} = 127$. With these values, the duration of packet is about 1 ms. The timing acquisition algorithms use a template with a pulse width of 0.5 ns. The initial timing acquisition and verification phase thresholds are $\sigma_1 = 0.625$ and $\sigma_2 = 0.75$. These values are the optimal results found in [11]. For the demodulation, $N_{int} = 40$. A summary of theses values is given in Table I.

$N_c = 128$ chips	$N_g = 8$ chips
$N_f = 8$ frames	$T_c = 6$ ns
$f_s = 2$ GS/s	$L_{pre} = 32 * N_f$
$L_{del} = 8 * N_f$	$L_{pay} = 127$ bytes
template pulse width = 0.5 ns	$N_v = 16$
$\sigma_1 = 0.625$	$\sigma_2 = 0.75$
$N_{int} = 40$	$\gamma = 0.3$

TABLE I
PARAMETERS USED FOR THE EXPERIMENTS.
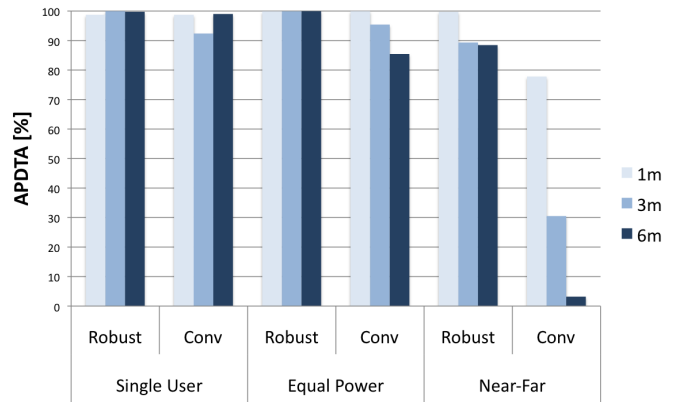
B. Measurements Results

Fig. 5. Measurement results of achieved packet detection and timing acquisition (A_{PDTA}) for the Robust (multi-user robust algorithm) and Conv (conventional correlation based algorithm) algorithms, for our three scenarios: single user, equal power and near-far (see Figure 4). In each case, we consider $L = 1, 3, 6$ meters. For the near-far scenario, $L_{nf} = 0.2$ meter. The robust algorithm outperforms the conventional algorithms, especially in the near-far scenario. The conventional algorithm does not allow for a reliable packet detection and timing acquisition in the presence of MUI.

Figure 5 shows the percentage of achieved packet detection and timing acquisition A_{PDTA}. The MUI robust algorithm

performs better than the conventional algorithm, especially when the distance between the transmitter and the receiver increases. Interestingly, we notice that both algorithms are able to ensure a proper packet detection and timing acquisition in the single user and equal power scenario. However, we display here only the best results obtained with the conventional algorithm. As explained later in Section III-C, the conventional algorithm is extremely sensitive to the surrounding environment and exhibits large variations from one experiment to another. Furthermore, in the near-far scenario, the conventional algorithm yields a much lower A_{PDTA}. It is completely outperformed by the MUI robust algorithm.

PER [%]	Single User		Equal Power		Near-Far	
	ROB	CONV	ROB	CONV	ROB	CONV
1m	1.27	1.27	39.36	39.27	65.98	72.65
3m	0.0	7.62	37.77	40.58	79.47	91.44
6m	0.26	1.01	54.71	61.47	70.43	98.73

TABLE II

MEASUREMENT RESULTS OF PACKET ERROR RATE (PER).

BER [%]	Single User		Equal Power		Near-Far	
	ROB	CONV	ROB	CONV	ROB	CONV
1m	0.0	0.0	0.49	0.48	2.35	2.40
3m	0.0	0.0	0.28	0.27	3.76	4.16
6m	0.004	0.004	0.85	0.85	4.42	3.95

TABLE III

MEASUREMENT RESULTS OF BIT ERROR RATE (BER).

The PER and the BER are reported in Tables II and III. Remember that no error correction or robust demodulation is used. Also, when timing acquisition fails for a given packet, this packet is not used to compute the BER but is taken into account for the PER.

In the single user scenario, we observe a relatively low PER for both methods. In the two other scenarios, the PER is higher due to MUI. As expected, the bit error rate (BER) increases with the distance between the transmitter and the receiver. The PER after the timing acquisition with the MUI robust algorithm can be higher than with the conventional algorithm. This occurs because the MUI robust algorithms allows for timing acquisition in the presence of MUI and consequently allows for more cases where the payload experiences the presence of MUI. Additional observations have been made during trial experiments. For instance, we observed that the BER is lower when N_c is higher and A_{PDTA} is better when N_f is higher for both algorithms.

C. Validity of Our Results

To assess the validity of our results, two main factors have to be considered. First, our non-perfect hardware and the manual tuning of the thresholds. We have observed that the results obtained with the conventional algorithm show large variations from one experiment to another. In fact, the conventional algorithm is very sensitive to the surrounding environment. The fact that the laboratory where measurements

are performed is not fully immune to electromagnetic radiations has to be considered. We believe, however, that the main reasons for these variations comes from our hardware. The pulse amplitude is not always as constant as it should be, especially because of the imperfections of the oscillator. It is is very dependent on the power supply stability and on temperature variations. Multi-paths also strongly affect the amplitude of the signal. We did notice for instance that the variation of the signal amplitude has a relatively big impact on the conventional algorithm performance. Nevertheless, despite an imperfect hardware, the MUI robust algorithm exhibits a higher stability and outperforms the conventional algorithm.

Second, the choice of the hard decision threshold ν (see Section II) potentially influences the results. The hard decision threshold of both algorithms is separately set, depending on the received power. In practice, we have to decrease the threshold as a function of the distance to obtain the best results. For the scenarios and distances of our measurement sets, the MUI robust algorithm hard decision threshold only requires slight tuning. Whereas, with the conventional algorithm, the hard decision threshold has to be precisely adjusted for each value of the link distance in the first two scenarios (single user and equal power). In the near-far scenario, it is constant because the interference is also constant in power. It is tuned manually and set after a certain number of trials. Therefore, this does not allow for a perfect comparison between the two timing acquisition algorithms. However the overall trend of the results obtained is validated by the multiple measurement runs performed. The MUI robust algorithm is much more resistant than the conventional algorithm, even without MUI.

IV. CONCLUSION

We have experimentally demonstrated that concurrent transmissions are feasible in low data-rate IR-UWB networks. We have also shown that interference mitigation techniques are indeed necessary at the physical layer. In particular, because it is the first component for the reception of a packet, we have focused on packet detection and timing acquisition. We have shown that a traditional scheme is not robust against multi-user interference, and prevents concurrent transmissions. On the contrary, a scheme designed to take MUI into account, even with a very simple interference mitigation scheme, allows for concurrent transmissions, even in strong near-far scenarios.

For future work, we will add real-time processing capabilities on the FPGA board. We will also implement channel coding, a robust demodulation scheme for the payload and add an automatic gain control (AGC) algorithm jointly with an automated way to set the hard decision thresholds.
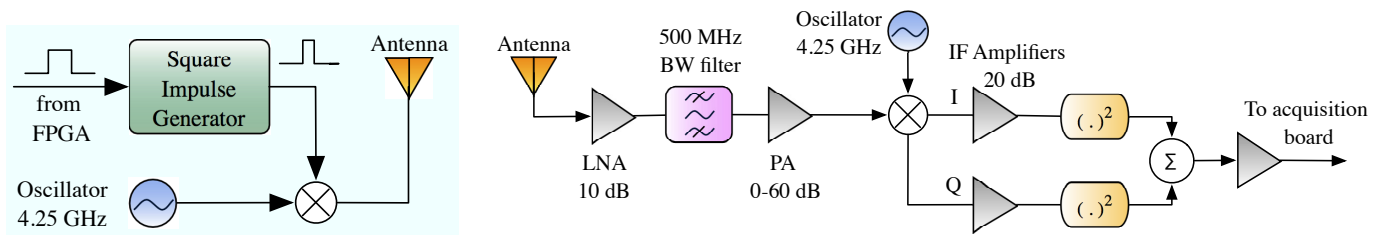
(a) The impulse generator generates a square impulse of 2 ns. This signal drives the mixer that switches on and off the oscillator sine wave and produce the IR-UWB signal.

(b) The received signal is amplified and filtered before being down-converted by an IQ-mixer. The mixer has its local oscillator frequency at 4.25 GHz. The baseband signal is amplified and sent to the FPGA for sampling.

Fig. 6. Overview of the RF part our experimental testbed.

APPENDIX
OVERVIEW OF THE IR-UWB TESTBED

In this section, we give a brief overview of our IR-UWB testbed (see Figure 3). On the transmitter side, we can generate several concurrent IR-UWB signals with a bandwidth of at least 500 MHz. For this purpose, an FPGA creates several command signals that each drive an IR-UWB transmitter. Each command signal is a simple low voltage digitally controlled impedance (LVDCI) signal where rising edges correspond to a pulse transmission. The maximum clock frequency of our FPGA is 166 MHz. This translates to a minimum chip duration T_c of 6 ns. We have the ability to fully control each parameter of the transmitted signal (see Equation (1)) and they can be configured by the computer through registers.

The FPGA is connected to each IR-UWB transmitter through a Micro DB connector. Hence, our analog IR-UWB transmitter was built to be robust to distortions of the command signal. Its architecture is depicted in Figure 6(a): An integrated PLL sine wave oscillator running at $f_c = 4.25$ GHz is connected to the antenna through a mixer. The mixer behaves like a switch when driven by a square signal. The output of the PLL is amplified before the mixer. An important part of the transmitter is the square impulse generator. Indeed, driven by a possibly degraded FPGA command signal, it must reliably generates short squared impulses of 2 ns to drive the mixer. The duration of the pulse can be finely tuned with a trimmer.

The receiver is a classic direct-conversion circuit (see Figure 6(b)). The signal from the antenna is amplified by a LNA followed by a power amplifier and then down-converted with an IQ mixer driven by a 4.25 GHz sine wave. The signal is additionally band-pass filtered inside the LNA. The design of the antennas is described in [15]. Most of our receiver is built around discrete components with off-the-shelf integrated circuits. Hence, we carefully designed them to avoid any undesired interferences and spurious coupling between components. We made sure to shield components whenever necessary to preserve the integrity of the received signal.

The list of components used for the testbed is in Table IV.

FPGA	Virtex II Pro XC2VP70, Xilinx
Acquisition board	AC-240 (512 MB SDRAM), Acqiris
PLL oscillator	PCA 4260C-LF, Z-Comm
Switching mixer	HMC128, Hittite Corporation
LNA transistor	NESG2030M04, NEC
Power amplifier	MMIC MGA-85563, Agilent
IQ mixer	HMC525, Hittite Corporation
Squaring mixer	ZLW-11, Mini-Circuits
IF amplifier	ZX60-3018G, Mini-Circuits
Duroid substrate	RO4003B, Rogers Corporation

TABLE IV
LIST OF COMPONENTS USED FOR THE TESTBED.

REFERENCES

[1] R. Scholtz, "Multiple access with time-hopping impulse modulation," in *IEEE Military Communications Conference (MILCOM)*, vol. 2, 1993, pp. 447–450 vol.2.

[2] 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.

[3] S. Gezici, Z. Tian, G. Giannakis, H. Kobayashi, A. Molisch, H. V. Poor, and Z. Sahinoglu, "Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks," *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 70–84, July 2005.

[4] 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, September 2004.

[5] 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, pp. 567–580, August 2005.

[6] 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.

[7] 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.

[8] M. Flury and J.-Y. Le Boudec, "Interference mitigation by statistical interference modeling in an impulse radio UWB receiver," in *IEEE International Conference on Ultra-Wideband*, September 2006, pp. 393–398.

[9] J. Mitra and L. Lampe, "Robust detectors for TH IR-UWB systems with multiuser interference," in *Ultra-Wideband, 2007. ICUWB 2007. IEEE International Conference on*, September 2007, pp. 745–750.

[10] 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, October 2005.

[11] A. El Fawal and J.-Y. Le Boudec, "A robust signal detection method for ultra wide band (UWB) networks with uncontrolled interference," *IEEE Trans. Microw. Theory Tech.*, vol. 54, no. 4, pp. 1769–1781, June 2006.

[12] 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 (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks," Jan. 2007.

[13] J. G. Proakis, *Digital Communications*, 4th ed. New York, NY: McGraw–Hill, 2001.

[14] C. Carbonelli and U. Mengali, "Synchronization algorithms for uwb signals," *IEEE Trans. Commun.*, vol. 54, no. 2, pp. 329–338, 2006.

[15] G. Quintero Diaz de Leon and A. Skrivervik, "Analysis of UWB antennas for carrier-based UWP impulse radio," in *European Conference on Antennas and Propagation*, November 2007, pp. 1–4.