

Interference Management and Rate Adaptation in OFDM-based UWB Networks

Raed T. Al-Zubi and Marwan Krunz
 Department of Electrical and Computer Engineering
 University of Arizona, Tucson, AZ 85721
 E-mail: {alzubi, krunz}@ece.arizona.edu

Abstract—Ultra-wideband (UWB) communications has emerged as a promising technology for high data rate wireless personal area networks (WPANs). Several proposals for UWB-based WPANs have been made. One widely popular proposal was standardized by ECMA-368, and is based on OFDM. In this paper, we address one of the important aspects that impact the performance of this standard, namely interference management between different uncooperative beacon groups that operate simultaneously over the same area. We first propose an interference management distributed reservation protocol (IM-DRP) for OFDM-based UWB communications. IM-DRP aims at improving the throughput of an UWB WPAN by reducing interference between uncooperative beacon groups. We then integrate IM-DRP into the design of a rate adaptation strategy that exploits the multi-rate capability of OFDM-based UWB systems. Besides maintaining a target packet error rate, our proposed strategy attempts to reduce the required reservation time over a link, hence allowing more links to be simultaneously activated. This improves the overall network throughput. Simulations are used to demonstrate the performance gain of our proposed schemes.

Index Terms—OFDM-based UWB WPANs, interference management, multi-rate capability, rate adaptation schemes .

1 INTRODUCTION

In 2002, the FCC issued its First Report and Order, which permitted the deployment of ultra-wideband (UWB) devices [7] over the spectrum from 3.1 GHz to 10.6 GHz. This report motivated researchers to exploit the interesting features of UWB for various applications, including short-range wireless communications, wireless sensor networks, imaging and radar systems, and precision location tracking systems. In general, UWB techniques can be categorized into two classes: single-band or impulse UWB (I-UWB), which is based on sending ultra-short pulses without frequency carriers, and multi-band or multi-carrier UWB (MC-UWB), which is based on sending data using multiple simultaneous OFDM sub-carriers. In this paper, we focus on the latter class, as standardized by ECMA-368 [6]. Industry advocates of this class formed an organization called the Multi-band OFDM Alliance (MBOA) [16]. This alliance eventually evolved into a large industrial alliance known as WiMedia, which defines, certifies, and supports wireless technology for multimedia applications. WiMedia members approached and asked the European Computer Manufacturers Association (ECMA) to consider WiMedia specifications for OFDM-based UWB and undertake the system standardization process. OFDM-based UWB has since emerged as a promising technology for high data rate WPANs.

An UWB WPAN uses the concept of a *beacon group* of a node, which consists of the neighbors of that node

that operate on the same *time-frequency code* (TFC) (see Fig. 1). The union of overlapped beacon groups is called an *extended beacon group*. Nodes in an extended beacon group are time-synchronized, and they use a specific TFC for data transmission. For simplicity, in this paper, we use the term *virtual network* (VN) to refer to a group of nodes that are time-synchronized and use the same TFC. Using different TFCs for different VNs (ECMA-368 defines 10 TFCs) facilitates channel reuse over the same area.

The ECMA-368 standard supports both random access and time-based reservations. The latter mode, known as the *distributed reservation protocol* (DRP), is particularly suited for real-time (voice and video) traffic. According to this protocol, devices that want to communicate with each other reserve their required medium access slots (MASs) from the available MASs that are not reserved by their neighbors in the same VN.




Fig. 1: Concepts of beacon group and extended beacon group in ECMA-368. A solid circle around a node indicates the beacon group of that node, whereas the union of these beacon groups (dashed line) indicates an extended beacon group.

One of the important performance metrics in the design of an OFDM-based UWB system is network throughput. Because different VNs use TFCs that are not completely orthogonal (as explained later, there can be up to 1/3 overlap between two TFCs) and because the operations of these VNs are not coordinated in time, they can interfere with each other and impact the system throughput. The DRP of the ECMA-368 standard does not discuss how devices can account for the reservations of neighboring devices that operate over different TFCs.

Motivated by the above, in this paper we first propose an interference management distributed reservation protocol (IM-DRP) for multi-band OFDM-based UWB communications. IM-DRP addresses the limitation of the current DRP and provides a mechanism to reduce the interference from transmissions made on different TFCs. This results in an overall improvement in network throughput. IM-DRP exploits the ability of an UWB device in one VN to scan all TFCs and discover MASs reserved by the device's neighbors in other VNs. Based on this scanning, the transmitter and the receiver of a new reservation can agree on the time slots to reserve with minimum to no interference. The proposed IM-DRP is easy to implement and does not require significant changes in the standard.

One of the important features of the OFDM-based UWB technology is its multi-rate capability. Specifically, an UWB transmission can use one of eight data rates, ranging from 53.3 Mbps to 480 Mbps. Rate adaptation has great impact on network performance. To illustrate, consider the establishment of a new session in an UWB network. The two end nodes must first determine the required number of MASs to be reserved. This number clearly depends on the traffic demand and the transmission rate. The higher the transmission rate, the smaller the number of required MASs, and hence the higher the number of MASs that are available for other prospective reservations. This leads to an overall improvement in network throughput. Furthermore, the smaller the number of MASs that are allocated to a given reservation, the lower the probability of time-overlapped reservations from different TFCs, which in turn increases the throughput of each VN. At the same time, a higher transmission rate leads to a higher packet error rate (PER) over a link. This places an upper limit on the transmission rate that can be used to support a target PER.

In this paper, we propose a rate adaptation scheme that exploits the multi-rate capability of an OFDM-based UWB system. According to this scheme, for a given link, rate adaptation is done at two levels. The first level takes place at the MAS time scale during the lifetime of a session. This type of rate adaptation aims at satisfying a target PER. The second level of rate adaptation is done at a coarser time scale (e.g., session admission times), and aims at minimizing the number of required MASs. This results in more admitted requests and higher network throughput. Moreover, reserving fewer MASs reduces

the probability of time-overlapped reservations, hence decreasing the interference between these reservations. For the second level of rate adaptation, we study four variants. The first one, called the interference-based (IB) scheme, uses the proposed IM-DRP protocol to determine on a slot-by-slot basis the maximum allowable transmission rate that satisfies a target PER. The second one, called the channel-gain-based (CGB) scheme, uses one transmission rate for all reserved MASs. This rate is determined based on the channel gain between the transmitter and the receiver. The third scheme is called reservation-size-based (RSB) scheme. Unlike IB and CGB, which are designed for peer-to-peer single-hop transmissions, RSB supports multi-hop transmissions in order to reduce the overall required number of reserved MASs. Note that from a control standpoint, the network setup considered in this paper allows any pair of nodes in the network to communicate directly if they use the smallest transmission rate. Under RSB, a direct transmission may be replaced with a multi-hop transmission. Finally, our fourth variant is the fixed-rate (FR) scheme, in which the transmission rate is fixed and multi-hop transmissions are allowed. FR mainly serves as a reference point for evaluating the RSB scheme and also to show the impact of multi-hop routing when rate adaptation is not implemented. Our rate adaptation schemes can be integrated with both the proposed IM-DRP and the standardized DRP.

The rest of the paper is organized as follows. Section 2 presents some related work. Section 3 overviews multi-band OFDM-based UWB. Section 4 describes the proposed IM-DRP protocol. Section 5 presents our rate adaptation schemes. In Section 6, we use simulations to evaluate our designs. Finally, concluding remarks are drawn in Section 7.

2 RELATED WORK

In OFDM-based UWB systems, three types of interference are of interest: interference from narrow-band signals (NB interference) [19] [9] [8], interference from devices in the same VN, which is handled through time-based reservation, and interference from UWB devices in other VNs. In this paper, we focus on the third type of interference. ECMA-368 handles this type of interference by defining different logical channels (TFCs) that can be used by different VNs. However, these TFCs are not completely orthogonal. Little work has been done in the literature to account for this type of interference. In [1], we outlined the design of an interference management scheme that handles interference between different VNs, assuming no rate adaptation. In [18], interference is controlled by combining the interleaver (which increases the effective coding rate) with the OFDM symbol repeater (which reduces the probability of symbol collision). The authors in [13] used the approach in [21], which was proposed for receive antenna arrays, for interference suppression in OFDM-based UWB systems. Both [18]

and [13] consider the interference during an ongoing transmission. In contrast, our proposed IM-DRP protocol considers and manages interference at the reservation time.

Although the ECMA-368 standard allows devices to support multiple transmission rates, it does not say how these rates should be selected. In [10], the authors proposed a rate adaptation scheme for OFDM-based UWB systems. According to this scheme, the transmission rate is adapted on per superframe basis according to the measured distance between the transmitter and the receiver, and the average received signal strength indicator (RSSI). A fast link-feedback mechanism for rate adaptation in OFDM-based UWB systems was proposed in [20]. According to this mechanism, the receiver uses Block acknowledgements (B-ACK) to recommend a transmission rate to be used in the upcoming reserved MASs. The main goal of rate adaptation in [10] and [20] is to satisfy a required PER during the ongoing transmission. The measured SINR in some MASs are used to select the transmission rate in other MASs. As mentioned before, time-overlapped reservations from different TFCs can occur in an unpredictable manner, resulting in interference that varies from one MAS to another, even within the same reservation. Therefore, it is inaccurate to estimate the SINR for one MAS based on the measured SINR in another MAS, as this may lead to the selection of a too high transmission rate that does not meet the required PER. In our work, we address this issue by assigning the transmission rate in a MAS based on the measured SINR of that same MAS.

3 MULTI-BAND OFDM-BASED UWB SYSTEMS

3.1 Physical Layer

According to ECMA-368, the spectrum between 3.1 GHz and 10.6 GHz (which is the spectrum assigned by the FCC for unlicensed UWB use) is divided into 14 frequency bands, each with a bandwidth of 528 MHz (see Fig. 2). These bands are grouped into six *band groups*. Each of the first four groups consists of three bands, whereas the fifth band group consists of two bands. The sixth band group consists of three bands (one band from the third band group and two bands from the fourth band group). TFCs are used to spread the transmitted symbols over one to three frequency bands in each band group, where 312.5 nsec are spent in each frequency band to send one symbol. Fig. 3 [6] illustrates the hopping sequence for the first TFC in band group 1. ECMA-368 uses three types of TFCs: Time-Frequency Interleaving (TFI), in which data are sent over three bands, two-band TFI (TFI2), where data are sent over two bands, and Fixed Frequency Interleaving (FFI), where data are sent over one band. As shown in Fig. 4, different types of TFCs are defined for the various band groups:

- Four TFIs, three TFI2s, and three FFIs for each of the first four band groups.
- Two FFIs and one TFI2 for the fifth band group.
- Three TFIs for the sixth band group.

TFCs are not completely orthogonal (as shown in Fig. 4, there can be up to 1/3 overlap between any two TFCs). Accordingly, interference between two neighboring links that operate on different TFCs may not be fully suppressed.




Fig. 2: Band group allocation in ECMA-368 [6].




Fig. 3: Realization of a TFC in ECMA-368.

3.2 MAC Layer

According to the ECMA-368 standard, time is divided into 65.536 msec intervals, called *superframes*. Each superframe is further divided into 256 MASs. The superframe consists of two parts: a beacon period (BP) and a data transfer period (DTP), as shown in Fig. 5. The BP is divided into beacon slots, the duration of each is 85 μ second. Each device chooses a beacon slot, transmits a beacon in this slot, and listens to its neighbors' beacons in all other beacon slots of each superframe. Beacons are very important for fast device discovery and synchronization between devices in the same extended beacon group [4]. They also contain important *information elements* (IEs) to control, manage, and facilitate distributed protocol operation. In the DRP mode, a device that wants to communicate with another device must reserve the required number of MASs from MASs that are not reserved by its neighbors in the same VN. The device announces its reservation via DRP IE that is transmitted in beacon frames. The reservation gives the device exclusive access to the medium during its allocated period.

TFC Number	Hopping Sequence					
1	1	2	3	1	2	3
2	1	3	2	1	3	2
3	1	1	2	2	3	3
4	1	1	3	3	2	2
5	1	1	1	1	1	1
6	2	2	2	2	2	2
7	3	3	3	3	3	3
8	1	2	1	2	1	2
9	1	3	1	3	1	3
10	2	3	2	3	2	3

(a) band group 1.

TFC Number	Hopping Sequence					
5	13	13	13	13	13	13
6	14	14	14	14	14	14
8	13	14	13	14	13	14

(b) band group 5.

TFC Number	Hopping Sequence					
1	9	10	11	9	10	11
2	9	11	10	9	11	10
3	9	9	10	10	11	11

(c) band group 6.

Fig. 4: Time frequency codes (TFCs) in ECMA-368.

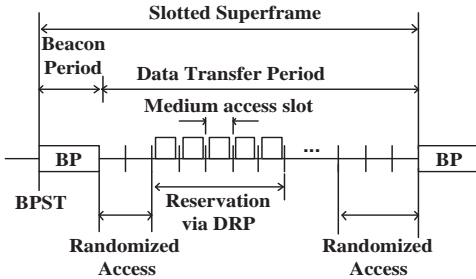


Fig. 5: Superframe structure in ECMA-368.

3.3 Interference Management

ECMA-368 addresses two types of interference: interference between devices in the same VN and interference between devices from different VNs. The first type of interference is controlled through time-based reservations. Interference from devices of different VNs is one of the key factors that determine the performance of the system. The impact of this interference is captured using the signal-to-interference ratio (SIR), given by [2]:

$$\text{SIR} = \frac{S}{I} \frac{W_{\text{eff}}}{R} \quad (1)$$

where S/I is the signal-to-interference power ratio, R is the transmission rate at the physical layer, and W_{eff} is the effective bandwidth of the transmitted signal. This W_{eff} is

given by

$$W_{\text{eff}} = \frac{N_b N_d}{T_{\text{sym}}} \quad (2)$$

where N_b is the number of bands in the TFC, N_d is the number of data sub-carriers, and T_{sym} is the symbol duration. The second term ($\frac{W_{\text{eff}}}{R}$) is the bandwidth expansion factor, and is related to the processing gain available to suppress inter-group interference. ECMA-368 specifies several approaches to increase SIR using bandwidth expansion, including coding, time-domain spreading, frequency-domain spreading, and time-frequency interleaving.

4 PROPOSED INTERFERENCE MANAGEMENT PROTOCOL

4.1 Motivation

As discussed in Section 3, ECMA-368 defines 10 channels (TFCs) per band group for data transmission. The average overlap between any two TFCs of type TFI or TFI2 is $\frac{1}{3}$. In other words, interference between different geographically co-located VNs is not fully suppressed. Furthermore, within one VN, the DRP grants exclusive channel access to the transmitter and receiver of a reservation. It does not give devices the ability to consider the reservations of other VNs that operate on different TFCs. To address this limitation, we propose the IM-DRP protocol, which provides a mechanism to manage interference between different reservations made on different channels. In a nutshell, for each reservation, IM-DRP tries to identify time slots that have low interference on both ends of a reservation (i.e., the transmitter and receiver). Both ends negotiate to identify the best slots for their reservation.

4.2 IM-DRP

IM-DRP follows two main steps. First, before making a new reservation, both the transmitter and receiver of the prospective session execute the standard DRP to determine available MASs that are not reserved by neighboring devices in the same VN. Second, the required MASs for the new reservation are selected from these available MASs such that the interference between the new reservation and already established reservations on other TFCs is minimized. The selection mechanism in the second step is the key novelty of IM-DRP. This mechanism requires the transmitter and receiver of the new reservation to scan all TFCs and identify the reserved MASs of potentially interfering neighboring devices. As shown in Fig. 4(b), for each of the first 4 band groups, there are 10 different TFCs, with potentially 10 mutually interfering VNs. Note that since there is no overlap between TFCs 5, 6, and 7, devices that want to establish a reservation on these channels scan only the first 4 and last 3 TFCs.

Knowing the locations of reservations on other TFCs is not enough to design a cooperative mechanism for interference management. Therefore, IM-DRP requires the transmitter and receiver of a new reservation to evaluate, during each available MAS, the level of interference they may inflict on devices in other TFCs, and vice versa. To accomplish this, devices have to distinguish between interference due to data transmission and interference due to ACK transmission. Because a data packet is much larger in size than an ACK packet and the ACK mechanism is optional in the standard, interference caused by a data transmission has a greater effect on performance. To illustrate, consider the situation in Fig. 6. Suppose that node A needs to reserve one MAS (for the time being, assume that the transmission rate is fixed and known to all nodes). Suppose that nodes B, C, D, and E operate on a different channel (TFC) than A's, and that there are only two MASs available to A (i.e., two MASs that are not reserved by any node in A's beacon group). One of these MASs is reserved by node C and the other one is reserved by node D. Now, if node A is a transmitter, then its best choice to avoid interfering with D is to reserve the MAS that is reserved by node C. However, to avoid interfering with C, node A must select the MAS that is also reserved by node D. Another fact that can be used in the selection process is that the receiver can evaluate the SIR in each MAS. Referring to the previous example, if node A is a receiver and all its neighbors are transmitters, then its best choice is the MAS during which node A achieves the highest received SIR.

Inspired by the above, IM-DRP uses the following selection mechanism to achieve cooperative interference management:

- The sender s and the receiver t of a prospective reservation individually classify unused MASs ("unused" is taken with respect to the superframe of the VN that the two nodes belong to) according to the measured interference into four sets, denoted by T_1 through T_4 for the sender s , and R_1 through R_4 for the receiver t . T_1 consists of unused MASs that are not reserved by any of s 's neighbors that operate on others TFCs (interference-free MASs). T_2 consists of the MASs that are reserved only by transmitting neighbors of s operating on other TFCs. T_3 consists of the MASs that are reserved only by receiving neighbors of s operating on other TFCs. Finally, T_4 consists of unused MASs that are reserved by both types of neighbors (transmitters and receivers). R_1 consists of the MASs that are not reserved by any neighboring receiver of t from other TFCs and during which the calculated SIR is greater than a predefined SIR threshold (SIR_{th}). R_2 consists of the MASs that are reserved only by t 's neighbors from other VNs that are receivers and during which the SIR is greater than SIR_{th} . R_3 consists of available MASs that are not reserved by any receiving neighbor of t from other VNs and during which the SIR

is less than SIR_{th} . Finally, R_4 consists of available MASs that are reserved by out-of-group receiving neighbors of t and during which the SIR is less than SIR_{th} .

- After the above classification is done, the receiver of the reservation provides the sets R_1 , R_2 , R_3 , and R_4 to the sender using a beacon or control frame.
- The sender sorts the unused MASs from best (MASs with low interference) to worst (MASs with high interference), as follows: $\{\{R_1 \cap T_1\}, \{R_1 \cap T_2\}, \{R_2 \cap T_1\}, \{R_2 \cap T_2\}, \{R_1 \cap T_3\}, \{R_2 \cap T_3\}, \{R_1 \cap T_4\}, \{R_2 \cap T_4\}, \{R_3 \cap T_1\}, \{R_3 \cap T_2\}, \{R_4 \cap T_1\}, \{R_4 \cap T_2\}, \{R_3 \cap T_3\}, \{R_4 \cap T_3\}, \{R_3 \cap T_4\}, \{R_4 \cap T_4\}\}$.
- The sender sequentially selects the required number of MASs for the new reservation from the above sorted set.
- Finally, both the sender and receiver use beacon frames to announce the new reservation so that it will be considered by other devices in the same VN that wish to establish new reservations.




Fig. 6: Interference due to data and ACK transmissions.

4.3 Compatibility and Protocol Overhead

To support the IM-DRP protocol, devices must be able to scan (for beacon frames) all the TFCs defined by the ECMA-368 standard, a process that is already supported by the standard. The scanning process incurs no more than 10 superframes of delay (≈ 0.65 second) before each new reservation. Since TFCs 5, 6, and 7 are orthogonal to each other, any new reservation that uses one of these TFCs will only need to scan the first 4 TFCs as well as the last 3 TFCs. In this case, the scanning delay is at most 7 superframes (≈ 0.45 second). Expectedly, the efficiency of IM-DRP increases with the session length.

5 RATE ADAPTATION

In general, rate adaptation has two aspects. The first one deals with selecting the highest transmission rate that satisfies a target PER. In SNR-based rate adaptation, rate selection is performed by the receiver, which is generally better informed of the channel state than the transmitter. The second aspect of rate adaptation is related to the feedback mechanism that the receiver uses to inform the transmitter of the selected rate. ECMA-368 deals only with this aspect, using as a feedback mechanism an information element (IE) in a beacon frame. According to ECMA-368, beacon frames are transmitted at the lowest rate (53.3 Mbps) so as to achieve the farthest transmission range. Non-beacon frames (control, command, and

data frames) can be transmitted at any of the following rates: 53.3, 80, 106.7, 160, 200, 320, 400, and 480 Mbps. These rates are achieved using time-domain spreading, frequency-domain spreading, and forward error correction (FEC).

Several rate selection techniques were proposed in the literature, including SampleRate [3], AMRR [11], and Onoe [14]. SampleRate selects the transmission rate by monitoring the success or failure of periodically sent probe messages. It tries to predict the transmission rate that provides the maximum link throughput. AMRR utilizes an estimate of the number of frame retransmissions to select the best transmission rate. A sampling period is used to change the value of the transmission rate. The length of this period is adjusted depending on the transmission status of probe packets. Onoe increases a credit value for each successful transmission, and on reaching a certain threshold, the current transmission rate is increased to the next higher rate. Similarly, the credit is reduced for each failed transmission, and the current transmission rate is reduced to the next lower rate when the credit goes below a certain threshold.

Both SampleRate and AMRR require extra bandwidth to send probe messages. Onoe is less sensitive to individual packet failures, and it adapts slowly to short-term variations in channel conditions. In contrast, in our work, the transmission rate is adapted based on the received SINR. This SINR-based technique is easy to implement (according to ECMA-368, nodes are supposed to measure the received SINR), it does not require additional bandwidth, and it adapts quickly to channel variations.

As explained in Fig. 7, our rate adaptation strategy is executed at two time scales: at the MAS level within each superframe (fine time scale) and at the session level (coarse time scale). Notice that adapting the transmission rate during a given superframe does not necessarily require measuring channel conditions in the same superframe. In a WPAN, channel conditions are unlikely to change dramatically between two successive superframes. Therefore, the channel state measured during a given reservation (i.e., a sequence of reserved MASs) in a superframe can be adequately used to adapt the transmission rate for the same reservation in the subsequent superframe [10].

At the fine time scale, rate adaptation aims at adapting the transmission rate per each reserved MAS during the ongoing session in order to maintain a target PER. The transmission rate for each reserved MAS in the upcoming superframe is updated according to the received SIR in the corresponding MAS of the current superframe. The measured SIR is affected by channel dynamics caused by mobility and interference from nodes in other VNs. In fact, it was shown in [12] that the performance of an OFDM-based UWB system under time-varying channel conditions is not much worse than such performance under a static channel, as long as node mobility is limited to 10 meters/second. Furthermore, by using IM-DRP, interference from other VNs is reduced.

Recall that IM-DRP requires the sender and receiver of a new reservation to evaluate, during each available MAS, the level of interference they would inflict on devices in other TFCs, and vice versa. Therefore, the determined transmission rates in the current superframe are reliable enough to be used for the next superframe. This type of rate adaptation requires the receiver to convey back to the sender the proper transmission rate to be used. We envision three options to carry such feedback: beacon frames, control frames, or block acknowledgement frames [20]. Without loss of generality, in the rest of this paper, we assume that beacon frames are used to convey rate information.

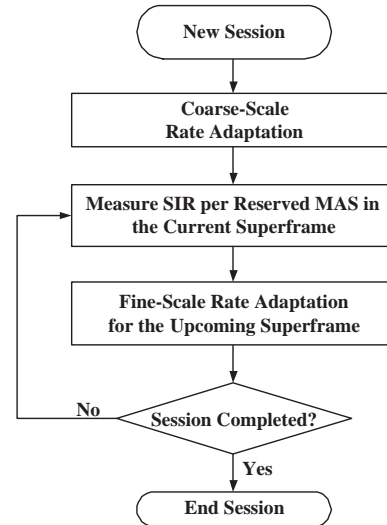


Fig. 7: Overview of the proposed rate adaptation strategy for OFDM-based UWB systems.

The second level of rate adaptation is performed at the beginning of a session. Such coarse adaptation aims at reducing the required number of MASs to be reserved for a new session. As a result, more links in the network can be established and the probability of time-overlapped links decreases, leading to higher network throughput. For such session-level rate adaptation, we propose the following variant schemes.

5.1 Channel-gain-based (CGB) Rate Control Scheme

As mentioned before, the PER over a link is a function of the transmission rate and the channel gain between the sender and the receiver. In our work, a target PER is given. The CGB scheme requires that at the time of a new reservation, the sender and the receiver should choose the maximum transmission rate that achieves the target PER, based on the channel gain (or the measured SIR at the receiver). Table 1 shows the required SIR for various transmission rates at a PER of 8%. These values are extracted from the BER vs. SIR curves in [17].

According to the ECMA-368 standard, an UWB device should have the ability to measure the SIR of the received signal [6]. In the CGB scheme, the receiver

Transmission rate (Mbps)	Minimum Receiver Sensitivity (dBm)	Required SIR (dB)
53.3	-80.8	5.3
80	-78.9	5.9
106.7	-77.8	6.1
160	-75.9	6.9
200	-74.5	7.9
320	-72.8	8.5
400	-71.5	9.7
480	-70.4	11.2


TABLE 1: Minimum receiver sensitivity and required SIR for various transmission rates (PER = 8%).

measures the SIR of the received beacon frame, and uses it along with the minimum SIR values associated with various transmission rates to determine maximum transmission rate to be used.

5.2 Interference-based (IB) Rate Control Scheme

The CGB scheme chooses one transmission rate per session based on the measured SIR during the beaconing period. It does not consider the interference during available MASs. Such simplicity comes at the expense of possible degradation in performance, as shown in Fig. 8. In this figure, two nodes, A and B, wish to establish a new session. Assume that after they have applied the CGB scheme, they found that the maximum allowable transmission rate is 160 Mbps and the corresponding number of required MASs at this rate is three. Suppose that there are five available MASs. Based on the measured SIR during these MASs, the corresponding maximum transmission rates are shown in the figure. Under the standard DRP, nodes A and B will reserve the first three available MASs (5, 6, and 20) and transmit at 160 Mbps during these MASs. This will lead to an unacceptably high PER during the first two MASs because 160 Mbps is higher than the maximum allowable transmission rate during these MASs. However, if instead the two nodes apply IM-DRP, they will first sort the available MASs as follows: 20, 21, 5, 6, and 22. Next, the nodes reserve the required number of MASs. In this example, the nodes reserve the first three MASs from the arranged MASs (20, 21 and 5) and use 160 Mbps during these MASs. This will lead to an unacceptably high PER during only MAS number 5. From this example, it is clear that CGB may use a transmission rate greater than the maximum allowable transmission rate in some reserved MASs.

To address this problem, we propose the interference-based (IB) rate control scheme. The main difference between CGB and IB is that the latter considers the interference at the MAS resolution. As discussed before, IM-DRP allows devices to estimate at the reservation instant the SIR during each available MAS. The IB scheme exploits this feature to determine the maximum allowable transmission rate per available MAS. Referring to Fig. 8, under the IB scheme, nodes will first reserve MASs 20 and 21. Then, they will reserve the fewest additional MASs that satisfy the required demand, such that



MAS 5: maximum achievable R = 80 Mbps

MAS 6: maximum achievable R = 80 Mbps

MAS 20: maximum achievable R = 160 Mbps

MAS 21: maximum achievable R = 160 Mbps

MAS 22: maximum achievable R = 53.3 Mbps

Fig. 8: Example illustrating the IB rate control scheme.

during each reserved MAS, nodes will use the maximum allowable transmission rate for that MAS. Clearly, the IB scheme may require reserving more MASs than the CGB scheme and this may reduce the probability for other reservations to be established. However, the maximum allowable transmission rate during each reserved MAS will be used (i.e., PER requirement is satisfied during each reserved MAS). In that sense, we expect IB to perform better than CGB.

Any rate control scheme requires feedback from the receiving device. In the CGB scheme, the receiver can use the link feedback IE to suggest the proper transmission rate to be used by the sender. However, for the IB scheme, this approach cannot be used because this IE allows for only one transmission rate per superframe. Instead of a link feedback IE, we suggest using a control or data frame (defined by ECMA-368) as a feedback for transmission rate adaptation if it is not possible to use the beacon frame, i.e., the maximum beacon length is not enough. The control or data frame would be sent during the data transfer period of a superframe, and would contain the maximum allowable transmission rate for each selected MAS.

5.3 Reservation-size-based (RSB) Rate Control Scheme

In CGB and IB, rate adaptation is optimized for a single hop. In contrast, the RSB scheme takes into consideration the possibility of multi-hop transmissions. This scheme replaces a direct transmission with a multi-hop transmission if the channel occupancy time, i.e., the required number of MASs per superframe for the reservation, under multi-hop transmission is less than that of direct transmission. As a result, this scheme achieves higher end-to-end throughput by allowing more sessions to be established. In RSB, each device measures the SIR of each beacon frame it receives from its neighbors, and announces these SIR values in a beacon frame. This step allows each node to construct a *neighborhood graph*,

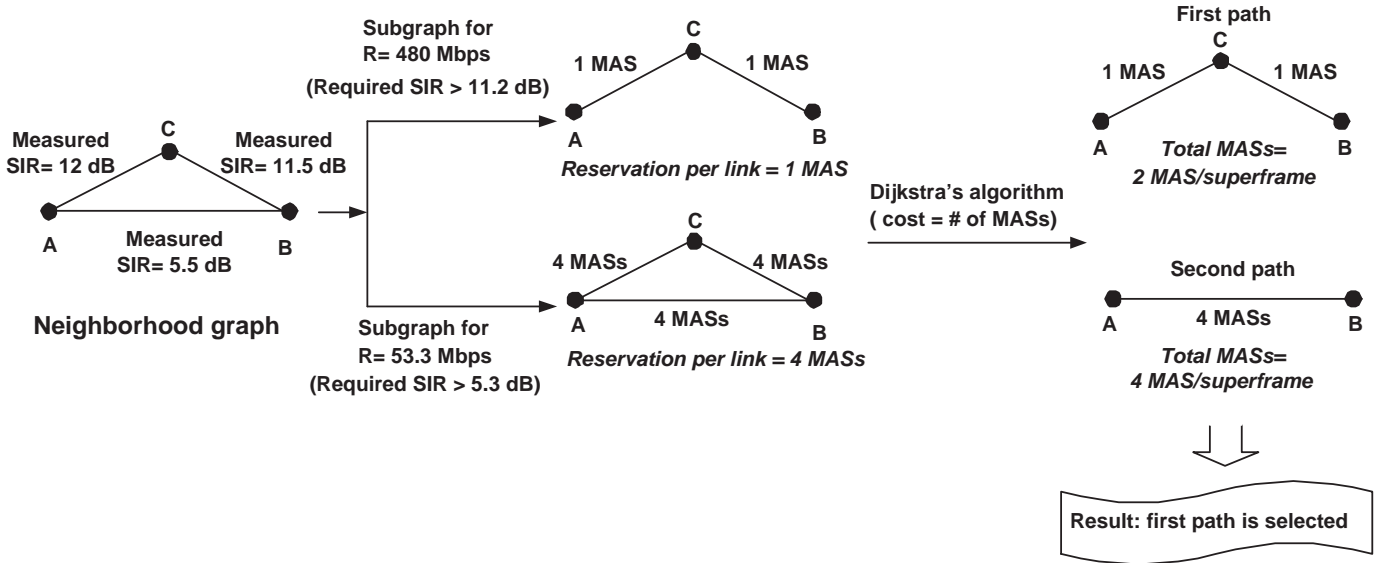


Fig. 9: Example illustrating the RSB scheme for two transmission rates (480 Mbps and 53.3 Mbps).

which reflects the network topology under the smallest transmission rate (which is used by beacon frames). Each link in this graph is assigned a weight value that represents the measured SIR over that link. From the neighborhood graph, the transmitter extracts the sub-graphs that correspond to individual transmission rates (53.3 to 480 Mbps). Specifically, the sub-graph associated with a given transmission rate consists of links whose measured SIRs are greater than the SIR threshold of that transmission rate. For each sub-graph, the source device then assigns one weight for all the links in that graph, which is the channel time in MASs for the required reservation. Note that this channel time depends on the used transmission rate. For each weighted sub-graph, the source subsequently runs Dijkstra's algorithm [5] to find the least-cost path to the destination on that sub-graph, where the cost is the channel time for the required reservation. In here, the total cost of a path is the sum of the link costs along that path. Finally, among the min-cost paths of various sub-graphs, the transmitter selects the one that results in the least channel time for the required reservation.

Fig. 9 shows an example of RSB's operation for two transmission rates (53.3 Mbps and 480 Mbps). In this example, there are two paths from node A to B: the two-hop path A-C-B and the direct path A-B. According to the load demand and the maximum allowable transmission rate for each path, the required number of MASs per hop is one for the first path and four for the second path. Therefore, the total number of required MASs per superframe for the first path (two MASs) is less than that for the second path (four MASs). According to the RSB scheme, the first path will be selected.

It should be noted that if the RSB scheme outputs the direct path, then both CGB and RSB schemes work similarly. The two schemes will transmit over this path

at the maximum allowable transmission rate. Therefore, we can say that the CGB scheme is a special case of the RSB scheme. Both schemes aim at reducing the required channel time per superframe for each session, which effectively improves the network throughput by allowing more sessions to be established in the network.

5.4 Fixed Rate (FR) Scheme

To facilitate comparison with various rate adaptation schemes, we also consider a fixed rate scheme in which all nodes use the same transmission rate. Recall that the transmission rate affects network connectivity; the higher the rate, the sparser is the network topology. In fact, enforcing a relatively high transmission rate on all links may lead to a disconnected topology, especially when sessions are limited to single-hop. To temper the effect of fixing the transmission rate on network connectivity, in the FR scheme we allow for multi-hop paths. To find the path between a source and its destination, the FR scheme follows the same procedure as in the RSB scheme, but for one transmission rate. Note that even though multi-hop transmissions are allowed in FR, network connectivity is not always guaranteed.

6 PERFORMANCE EVALUATION

In this section, we study the performance of the proposed IM-DRP and contrast it with the standard DRP. Our evaluation includes the non-adaptive version of IM-DRP as well as its various rate-adaptive versions. Our results are based on simulation experiments conducted using Matlab. The determination of collisions and interference is done according to the physical (SINR) model. The simulations account for the IM-DRP scanning overhead, time-frequency interleaving (defined by ECMA-368 to achieve bandwidth expansion), the framing structure

(i.e., number of MASs per superframe, beacon period, MAS duration), transmission powers and rates, receiver sensitivity, and MAC policies (e.g., limits and locations of reservations). Table 2 depicts the default parameter values used in our simulations. These values correspond to realistic hardware settings [2] [6].

Our main performance metrics are the effective network throughput (i.e., goodput) and the PER. Because of the TDMA structure of the underlying system, the end-to-end delay-jitter experienced by a flow (i.e., the maximum allowable deviation from the nominal inter-packet time) cannot exceed the duration of one superframe (~ 65 msec), which meets the jitter requirement of typical streaming applications.

Transmission rate (non-adaptive case)	200 Mbps
Transmission rate (adaptive case)	53.3-480 Mbps
Average transmission power	-10.3 dBm
Transmitter antenna gain	0 dBi
Receiver antenna gain	0 dBi
Path loss factor	2
Average noise power per bit	-85 dBm
Receiver noise figure	6.6 dB
Required E_b/N_0	7.9 dB
Hardware-related loss	2.5 dB
Data packet size	1 KB

TABLE 2: Parameters used in the simulation.

6.1 IM-DRP without Rate Adaptation

We consider a multi-band OFDM-UWB based WPAN, where N nodes are equally split among 10 TFCs and are uniformly placed in a square of length 10 meters. In each TFC, nodes are randomly paired such that the source and destination of a reservation are neighbors (according to the minimum transmission rate). For a given source-destination pair, the session length is randomly selected from the range $[0, T_{\max}]$, where T_{\max} is a controllable parameter. Once the session terminates, a new session is immediately initiated with a newly selected duration. For each session, the traffic load of a reservation (in bits/superframe) is uniformly chosen from 1 to 6 Mbps.

In Fig. 10 and 11, we show both metrics as functions of N for two mean session lengths. Both the proposed IM-DRP and the standard DRP are similarly affected by network density. For long sessions, IM-DRP has a significantly higher network throughput than the standard DRP (54% on average). It also gives a much lower PER, as shown in Fig. 11. This improvement is attributed to the cooperative nature of IM-DRP.

Fig. 12 and 13 depict the network throughput and the PER versus the session length for $N = 40$ and $N = 60$. Fig. 12 reveals that the throughput gain of IM-DRP decreases as the session length decreases. This is attributed to the scanning overhead of IM-DRP that precedes the establishment of new sessions. Typically, DRP is used to support isochronous services (e.g., video streaming), which are generally associated with long

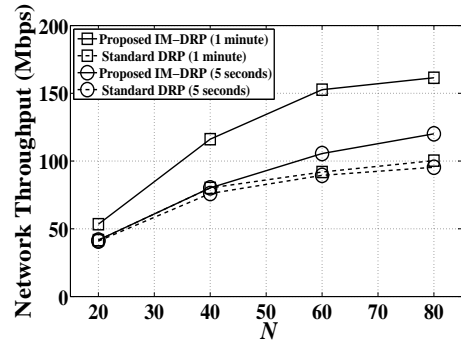


Fig. 10: Network throughput versus N for two session lengths (no rate adaptation).

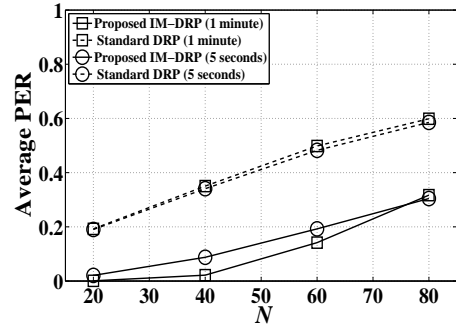


Fig. 11: Average PER versus N for two session lengths (no rate adaptation).

sessions. For instance, streaming content from a DVD would require the reservation to last for at least several minutes. Fig. 13 shows that IM-DRP significantly reduces the PER under various session lengths. The figure also shows that, for very short sessions (less than 1 second on average), the PER increases with the session length. This is because very short sessions have a high probability to end during the scanning process of other reservations. Note that during the scanning process, the nodes of the upcoming reservation do not send data, so they do not interfere with their neighbors. For longer sessions (average length greater than 1 second), the PER under IM-DRP decreases with the session length. This is attributed to the fact that as the session length decreases, the number of simultaneous scanning attempts increases. Therefore, the probability that the durations of these scanning attempts are overlapped will increase. Nodes from different VNs that have time-overlapped scanning processes cannot know the reservations of each other because their reservations are not determined yet. As a result, these nodes cannot coordinate their reservations, leading to higher interference, i.e., higher PER for shorter sessions.

In Fig. 14, we study the throughput gain and PER reduction achieved by IM-DRP relative to DRP, both as functions of the percentage of short sessions (δ), where the average length of a session is 1.5 seconds for a short session and 1 minute for a long session. In this experiment, we set $N = 60$. This figure shows that the throughput gain and the PER reduction decrease as δ

increases, which supports the trends in Fig. 12 and 13.




Fig. 12: Network throughput versus session length (no rate adaptation).

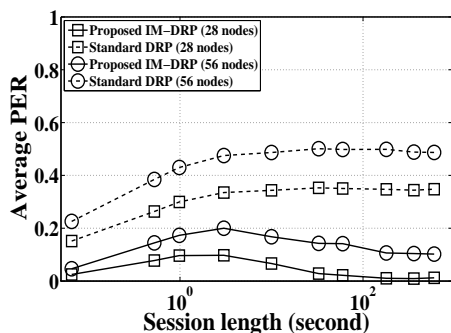


Fig. 13: Average PER versus session length (no rate adaptation).

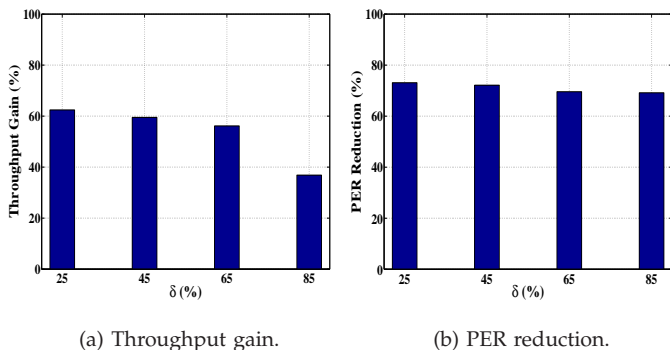


Fig. 14: Throughput gain and PER reduction achieved by IM-DRP relative to DRP (no rate adaptation).

6.2 IM-DRP with Rate Adaptation

We now evaluate the performance under various rate adaptation schemes. Except for the IB scheme, all the rate adaptation schemes are evaluated under both IM-DRP and the standard DRP. As discussed in Section 5, the IB scheme involves obtaining the interference during each unused MAS (i.e., candidate MASs for reservation) in order to determine the maximum allowable transmission rate in each reserved MAS. Therefore, this scheme is only applicable to IM-DRP.

The random waypoint model is used for mobility, with a typical human walking speed that is uniformly distributed between 0 and 2 meters/second. The pause time between changes in direction and speed is exponentially distributed with a mean of 5 seconds. Unless stated otherwise, the total number of nodes is 40 nodes. Nodes are split among 10 TFCs, and are uniformly placed within a disk area of diameter 17 meters. This size is chosen to ensure that any source-destination pair can communicate directly, even if they choose to use a multi-hop path. It is also representative of realistic deployment scenarios (e.g., indoor office, indoor residential, body-area networks [15]). In each TFC, the source-destination pairs are randomly selected. For each pair, the average session length is 15 seconds. Once a session terminates, a new session is immediately initiated. For each session, the traffic load (in bits/superframe) of a reservation is a controllable parameter.

To evaluate and compare the performance of the proposed rate control schemes, we mainly focus on five performance metrics: network throughput, PER, blocking rate, *deficiency*, and network connectivity. To explain the metrics of blocking rate and deficiency, we should clarify the procedure of establishing a session between two nodes. Nodes start by checking the available channel time, i.e., unreserved MASs in the superframe. If adequate channel time is available to support the given traffic load (in bits/superframe), nodes proceed with their reservation. On the other hand, if there are no available MASs, the request is blocked. However, if some MASs are available, but are not enough for the entire traffic load, the admission decision depends on the type of application. In our simulation, we consider two types of applications: elastic and non-elastic applications. If the application is “elastic”, the session will be established using the available channel time, and the unsatisfied traffic load is captured via the deficiency metric. On the other hand, if the application is non-elastic, the request will be blocked. According to the above discussed procedure, we calculate the blocking rate as the ratio between the number of blocked sessions and the total number of generated sessions. Deficiency is calculated as the ratio between the unsatisfied load and the total offered load for elastic traffic.

Note that energy consumption is impacted by the average PER. High PER means more retransmissions, which leads to higher energy consumption. Moreover, since we do not consider transmission power control in this work, energy consumption increases with the path length, i.e., number of hops between the source and the destination. Therefore, expectedly the energy consumption in the RSB and FR schemes is higher than that in the IB and CGB schemes, as the former schemes involve multi-hop transmissions (see Fig. 15 and 16).

Fig. 17 depicts the performance under elastic traffic of the IB, CGB, RSB. We also include the performance of the rate adaptation schemes in [10] and [20]. This figure shows that the IB scheme achieves high throughput

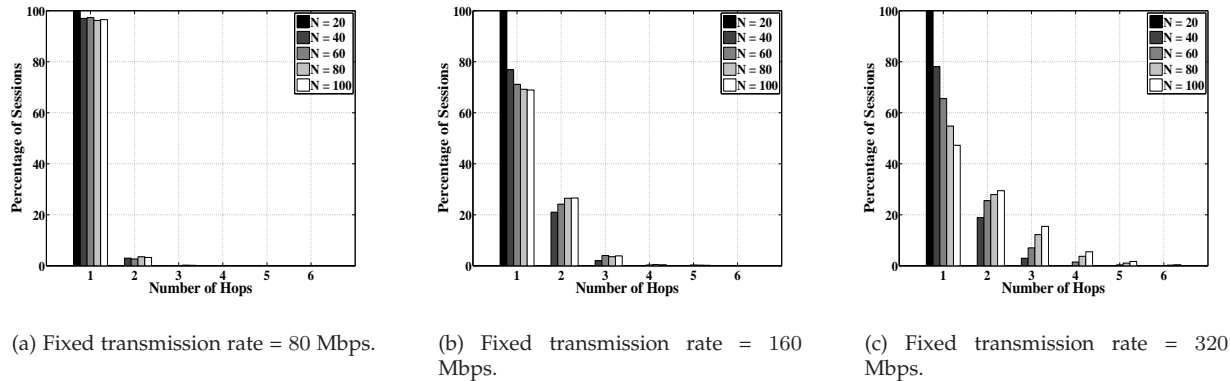


Fig. 16: Histogram of path length in the FR scheme.

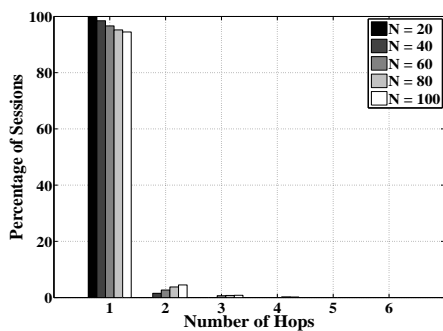


Fig. 15: Histogram of path length in the RSB scheme.

and low PER relative to the other schemes, but it has higher deficiency and blocking rate. This is because the IB scheme sometimes uses low transmission rates during some reserved MASs to achieve low PER, leading to higher number of required MASs for a target reservation. Accordingly, this reduces the probability for other reservations to be established, hence increasing the deficiency and the blocking rate. The performance of IB relative to CGB and RSB can be clearly seen at high traffic loads. As shown in Fig. 17, the performance of CGB and RSB in terms of network throughput and average PER deteriorates with the traffic load. This is attributed to the fact both schemes do not adapt the transmission rate on a MAS-by-MAS basis, leading to a higher transmission rate than the maximum allowable one in some reserved MASs, and hence higher PER during these MASs. Fig. 17 shows that our proposed schemes (IB, CGB, and RSB) achieve better performance than the schemes in [10] and [20]. Table 3 summarizes the average performance gain achieved by the IB scheme over CGB, RSB, and the schemes in [10] and [20].

Another result shown in Fig. 17 is that, at low traffic loads (e.g., 2-4 Mbps), RSB achieves slightly lower throughput than the CGB scheme. This is because, in the RSB scheme, the reservation process over a multi-hop path starts from the source and proceeds sequentially to subsequent hops such that any reservation over a hop cannot be done until the reservation over the preceding

Gain (%)	CGB (IM-DRP)	RSB (IM-DRP)	CGB (DRP)	RSB (DRP)	[10]	[20]
Throughput Increase	23	27	64	73	154	86
PER Reduction	-43	-38	-62	-60	-72	-65

TABLE 3: Average performance gain achieved by the IB scheme over other schemes (elastic traffic).

hop is announced. This sequential process incurs a delay before the source starts its transmission. Moreover, in the case of the RSB scheme with IM-DRP, additional delay is incurred due to the IM-DRP scanning overhead. To illustrate, suppose that two nodes A and B want to establish a path between them through node C. Assume that node C already has a path to another node D. In this case, C suspends its transmission to D during the IM-DRP scanning phase, which is required to establish the path between A and B.

Note that the reservation size increases with the traffic load, leading to a higher probability of time-overlapped reservations. This in turn increases the average PER under the CGR scheme. On the other hand, the RSB scheme aims at reducing the reservation size, which reduces the likelihood of time-overlapped reservations, hence achieving lower PER compared to the CGB scheme. However, due to the previously mentioned sequential reservation process and IM-DRP scanning delay, the throughput achieved by RSB at high loads ends up being comparable to that of CGB, as shown in Fig. 17.

In Fig. 18, we compare the performance of the RSB and FR schemes. For FR, two (fixed) transmission rates are separately used: 80 and 320 Mbps. As shown in this figure, RSB achieves higher throughput and lower PER compared with the FR scheme. This is because in the RSB scheme the transmission rate is adapted. In the FR scheme, a change in channel conditions during the lifetime of a session will have a strong impact on network performance. In fact, the performance of the FR scheme depends on the transmission rate, network

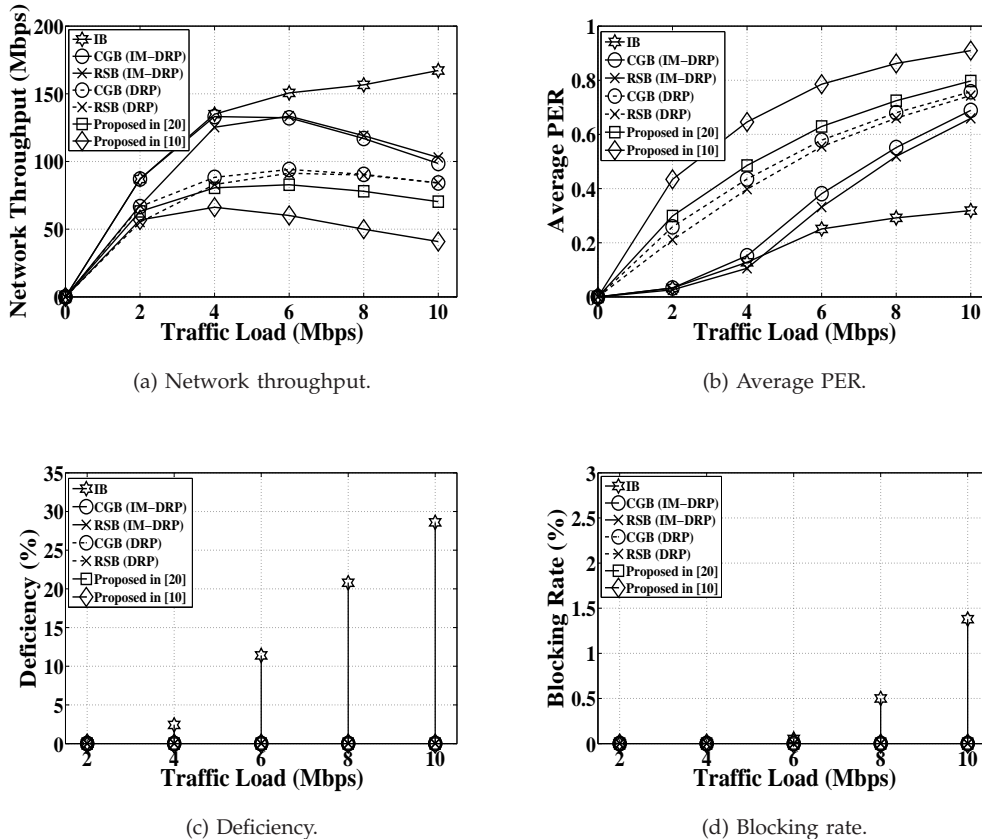


Fig. 17: Performance of various rate-adaptation approaches versus load (elastic traffic).

density, and traffic load. Collectively, these factors determine network connectivity and reservation size. In general, it is difficult to find a good compromise between various performance metrics under a fixed transmission rate strategy. For instance, at high traffic loads, using a low transmission rate results in low throughput and high PER, as shown in Fig. 18. On the other hand, network connectivity improves when the transmission rate is decreased, as shown in Fig. 19. Furthermore, the path length (i.e., number of hops) between a source and its destination decreases with a reduction in the transmission rate, as shown in Fig. 16, leading to lower energy consumption.

Finally, in Fig. 20, we study the performance under non-elastic traffic. In general, our simulation results show that in both cases, i.e., elastic and non-elastic applications, the schemes achieve comparable throughput and PER. However, in the non-elastic case, the IB scheme at heavy traffic loads has a higher blocking rate than in the elastic case. This is expected since in the non-elastic case, the session is established only if the channel time for the entire required load is available. This constraint becomes more difficult to achieve at heavy traffic load, leading to a higher blocking rate.

7 CONCLUSIONS AND FUTURE WORK

OFDM-based UWB systems, as standardized by ECMA-368, is a very promising technology for high-speed WPANs. In this paper, we proposed an interference management distributed reservation protocol (IM-DRP) for such systems. IM-DRP aims at improving the throughput of an UWB WPAN by managing interference between uncooperative beacon groups that operate simultaneously over the same area. The design of IM-DRP lends itself to a practical rate adaptation strategy, which has not been defined in the ECMA-368 standard. Motivated by this feature, we further proposed a novel rate adaptation strategy for OFDM-based UWB systems. We integrated this strategy into our proposed IM-DRP. We proposed four rate-control variants that exploit the multi-rate capability of the OFDM-based UWB system. Besides maintaining a target packet error rate, our proposed strategy attempts to reduce the reservation time over a link, hence allowing more links to be established. Our simulation results showed that IM-DRP significantly improves network throughput and reduces PER relative to the standard DRP. Furthermore, the integration of the proposed rate adaptation strategy into IM-DRP achieves an efficient MAS utilization; this integration allowed for establishing more reservations, and at the same time achieved high overall network throughput and low PER. One of our proposed rate-control schemes, which

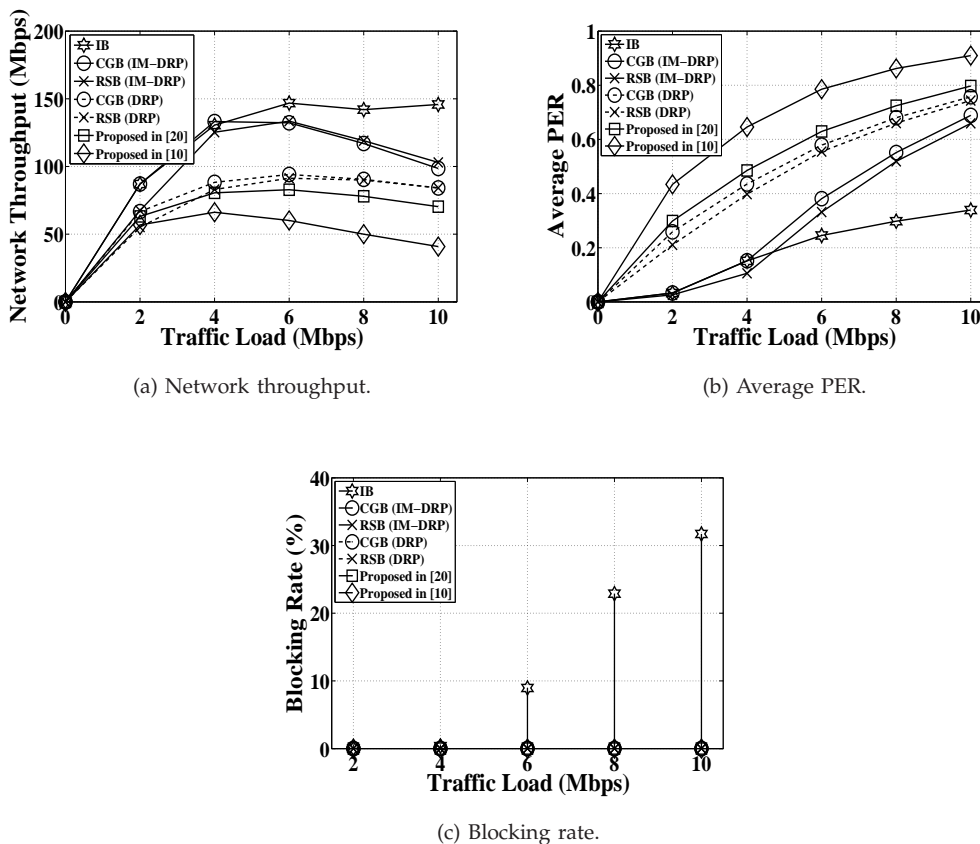


Fig. 20: Performance of various rate-adaptation approaches versus load (non-elastic traffic).

is called IB, achieved about 154% (86%) improvement in throughput and about 72% (65%) reduction in PER relative to the rate adaptation techniques in [10] ([20]).

One of the design requirements of our proposed schemes is compatibility with the ECMA-368 standard. For a real implementation of our schemes, UWB devices must have the ability to measure the SINR, exchange reservation information, and scan (for beacon frames) all the TFCs defined by the standard. All these functions are already supported by the ECMA-368 standard. Devices must also be able to select the SINR thresholds associated with various transmission rates to achieve a target PER. These thresholds are usually determined using offline measurements obtained for specific setups [10] [20].

Our future work will address the case of dynamic reservations where devices modify their reservations on the fly. In such a scenario, scanning prior to establishing a reservation is not enough by itself, so a dynamic approach is needed.

ACKNOWLEDGEMENTS

The authors would like to thank Dr. Alaa Muqattash for providing many valuable comments during the early stage of this research.

REFERENCES

[1] R. Al-Zubi, M. Krunz, and A. Muqattash. Interference management distributed reservation protocol for OFDM-Based UWB

communications. In *Proceedings of the IEEE GLOBECOM Conference*, December 2008.

[2] A. Batra, J. Balakrishnan, G. R. Aiello, J. R. Foester, and A. Dabak. Design of a multiband OFDM system for realistic UWB channel environments. *IEEE Transactions on Microwave Theory and Techniques*, 52(9):2123–2138, September 2004.

[3] J. C. Bicket. Bit-rate selection in wireless networks. In *Master's thesis, Massachusetts Institute of Technology*, 2005.

[4] C. Chou, J. del Prado Pavon, and S. Shankar. Mobility support enhancements for the WiMedia UWB MAC protocol. In *Proceedings of the IEEE BROADNETS Conference*, October 2005.

[5] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

[6] European Computer Manufacturers Association (ECMA). *ECMA-368 (3rd Edition): High Rate Wideband PHY and MAC Standard*, December 2008.

[7] FCC. *First Report and Order: In the Matter of Revision of Part 15 of the Commissions Rules Regarding Ultra-Wideband Transmission Systems*. FCC 02-48, April 2002.

[8] D. Gerakoulis and P. Salmi. An interference suppressing OFDM system for ultra wide bandwidth radio channels. In *Proceedings of the IEEE UWBST Conference*, May 2002.

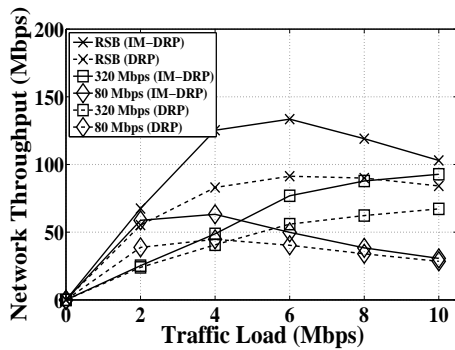
[9] Z. Irahauten, A. Yarovoy, G. Janssen, H. Nikoogar, and L. Ligthart. Suppression of noise and narrowband interference in UWB indoor channel measurements. In *Proceedings of the IEEE ICUWB Conference*, September 2005.

[10] J. Kim and J. Huh. Rate adaptation scheme for slot reservation in WiMedia MAC. In *Proceedings of the IEEE ICCE Conference*, January 2007.

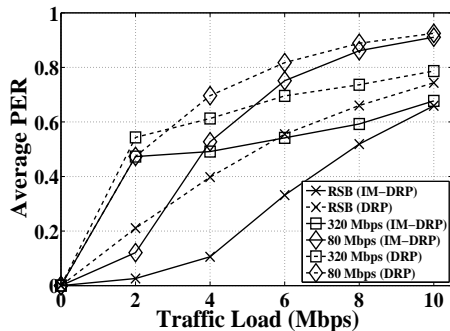
[11] M. Lamage, M. H. Manshaei, and T. Turletti. IEEE 802.11 rate adaptation: A practical approach. In *Proceedings of the ACM MSWiM Conference*, October 2004.

[12] N. Laurenti and P. Toniolo. Performance of the multi-band OFDM UWB system with time-varying channels. In *Proceedings of the IEEE WPMC Conference*, September 2004.

[13] Y. G. Li, A. F. Molisch, and J. Zhang. Practical approaches to



(a) Network throughput.



(b) Average PER.

Fig. 18: Performance of the RSB and FR schemes versus load (elastic traffic).




Fig. 19: Blocking rate due to unreachable destination (disconnected topology) in the FR scheme versus number of nodes.

channel estimation and interference suppression for OFDM-based UWB communications. *IEEE Transactions on Wireless Communications*, 5(9):2317–2320, September 2006.

- [14] MadWifi driver documentation, Onoe Rate Control, <http://www.madwifi.org>.
- [15] A. Molisch, K. Balakrishnan, C. Chong, S. Emami, A. Fort, J. Karedal, J. Kunisch, H. Schantz, U. Schuster, and K. Siwiak. IEEE 802.15.4a channel model - final report. *IEEE 802.15 Working Group for WPANs, Technical report IEEE P802.15-04/0662*, November 2004.
- [16] Multi-Band OFDM Alliance (MBOA). *Multi-Band OFDM Physical Layer Proposal for IEEE Task Group 3a*, September 2004.
- [17] S. Nowak, O. Hundt, and R. Kays. Joint efficiency and performance enhancement of multiband OFDM ultra-wideband (WiMedia) systems by application of LDPC codes. In *Proceedings of the IEEE ISCE Symposium*, April 2008.

- [18] S. Park, G. Shor, and Y. S. Kim. Interference resilient transmission scheme for multiband OFDM system in UWB channels. In *Proceedings of the IEEE ISCAS Symposium*, May 2004.
- [19] D. C. Popescu and P. Yaddanapudi. Narrowband interference avoidance in OFDM-based UWB communication systems. *IEEE Transactions on Communications*, 55(9):1667–1673, September 2007.
- [20] K. Schoo and J. Tervonen. Prompt rate adaptation in WiMedia considering time variant UWB channels. *Proceedings of the IEEE PIMRC Symposium*, September 2007.
- [21] J. Winters. Signal acquisition and tracking with adaptive arrays in wireless systems. In *Proceedings of the IEEE VTC Conference*, May 1993.



Raed T. Al-Zubi received B.S. and M.S. degrees in electrical engineering from Jordan University, Amman, 2002 and 2005, respectively. He is currently a research assistant working toward a Ph.D. degree at the advanced networking Lab., Department of Electrical and Computer Engineering, University of Arizona, Arizona, USA. His current research interests include system architecture and communication protocol designs for wireless networks with emphasis on cross-layer designs for ultra-wideband (UWB)

networks. Other interests include biometrics, pattern recognition, and image processing.



Marwan Krunz is a professor of electrical and computer engineering at the University of Arizona. He directs the wireless and networking group in the ECE Department. He is also the UA site director for Connection One, a joint NSF/state/industry IUCRC cooperative center that focuses on RF and wireless communication systems and networks. At present, the center consists of five participating universities and 17+ industrial affiliates. Dr. Krunz received his Ph.D. degree in electrical engineering from Michigan

State University in 1995. He joined the University of Arizona in January 1997, after a brief postdoctoral stint at the University of Maryland, College Park. He previously held visiting research positions at INRIA, HP Labs, University of Paris VI, and US West (now Qwest) Advanced Technologies. His research interests lie in the fields of computer networking and wireless communications. His current research is focused on cognitive radios and SDRs; distributed radio resource management in wireless networks; channel access and protocol design; MIMO and smart-antenna systems; UWB-based personal area networks; energy management and clustering in sensor networks; media streaming; QoS routing; and fault monitoring/detection in optical networks. He has published more than 150 journal articles and refereed conference papers, and is a co-inventor on US patents. M. Krunz is a recipient of the National Science Foundation CAREER Award (1998). He currently serves on the editorial boards for the *IEEE Transactions on Mobile Computing* and the *Computer Communications Journal*. He previously served on the editorial board for the *IEEE/ACM Transactions on Networking* (2001–2008). He was a guest co-editor for special issues in *IEEE Micro* and *IEEE Communications magazines*. He served as a technical program chair for various international conferences, including the IEEE WoW-MoM 2006, the IEEE SECON 2005, the IEEE INFOCOM 2004, and the 9th Hot Interconnects Symposium (2001). He has served and continues to serve on the executive and technical program committees of many international conferences and on the panels of several NSF directorates. He gave several tutorials and participated in various panels at premier wireless networking conferences. He is a consultant for a number of companies in the telecommunications sector.