Multiband Media Access Control in Impulse-Based UWB Ad Hoc Networks

Ioannis Broustis, *Student Member*, *IEEE*, Srikanth V. Krishnamurthy, *Member*, *IEEE*, Michalis Faloutsos, *Member*, *IEEE*, Mart Molle, *Member*, *IEEE*, and Jeffrey R. Foerster, *Member*, *IEEE*

Abstract—We propose a MAC protocol for use in multihop wireless networks that deploy an underlying UWB (Ultra Wide Band)-based physical layer. We consider a multiband approach to better utilize the available spectrum, where each transmitter sends longer pulses in one of many narrower frequency bands. The motivation comes from the observation that, in the absence of a sophisticated equalizer, the size of a slot for transmitting a UWB pulse is typically dictated by the *delay spread* of the channel. Therefore, using a wider frequency band to shorten the transmission time for each pulse does not increase the data rate in proportion to the available bandwidth. Our approach allows data transmissions to be contiguous and practically interference free, and, thus, highly efficient. For practicality, we ensure the conformance of our approach to FCC-imposed emission limits. We evaluate our approach via extensive simulations, and our results demonstrate the significant advantages of our approach over single-band solutions: The throughput increases significantly and the number of collisions decreases considerably. Finally, we analyze the behavior of our MAC protocol in a single-hop setting in terms of its efficiency in utilizing the multiple bands.

Index Terms—UWB, short-range communications, medium access control, ad hoc networks.

1 Introduction

Ultra Wide Band (UWB) is a novel wireless short-range technology which has been the focus of a lot of interest [6], [10], [13], [14], [16], [17], [23]. Our objective in this effort is to design a MAC protocol that fully utilizes the capabilities of UWB communications in an ad hoc network setting. The use of impulse-based UWB in military ad hoc networks is especially attractive given that a low probability of signal detection by an adversary is a desirable property. While physical layer technologies on UWB communications have been developed to some extent [14], MAC and higher layer technologies that enable the use of UWB in ad hoc networks are yet to mature [6]. The unique properties of UWB pose challenges to the design of a MAC protocol and require the MAC layer to be synergetic with the underlying physical layer. We present three of these practical challenges which motivate our multiband approach.

First, with impulse-based UWB, pulses are subject to multipath *delay spread*, due to which multiple time-shifted copies of each transmitted pulse appear at the receiver. This delay spread causes *intersymbol interference* (ISI), wherein the delayed copies of one pulse interfere with subsequent pulses [3]. In indoor settings, the magnitude of this delay spread is of the order of tens of nanoseconds. The use of sophisticated equalization to combat ISI adds considerable hardware complexity to the transceivers and increases the synchronization overhead. In fact, UWB communications already require a long acquisition time for nodes to be synchronized prior to communications [18], which becomes longer due to the training sequence overheads required with equalizers.[1] One can also reduce ISI by ensuring that the spacing between the received pulses is larger than the delay spread; thus, the delayed copies of one pulse will not interfere with the next pulse.[2] With this approach, as opposed to the width of a pulse, the *interpulse spacing* constrains the throughput of the channel. Thus, in this case, *a smaller bandwidth channel, which requires an elongated pulse duration, can yield a throughput comparable to that of a wider band, which allows a much shorter pulse duration for a fixed equalizer complexity*. Hence, we note that we can partition the UWB spectrum into multiple comparatively narrow frequency bands that are mutually orthogonal and can be used simultaneously, and, thus, use the available spectrum more efficiently.

The second motivating observation stems from the absence of carrier sensing capabilities in UWB. With impulse-based UWB, data is transmitted in the form of pulses[3] and there is no contiguous carrier, although these pulses are possibly modulated by means of a high frequency signal (referred to as the *pseudocarrier*). Thus,

- *I. Broustis, S.V. Krishnamurthy, M. Faloutsos, and M. Molle are with the Department of Computer Science and Engineering, University of California, Riverside, Engineering BU2, Room 351, Riverside, CA 92521. E-mail: {broustis, krish, michalis, mart}@cs.ucr.edu.*
- *J.R. Foerster is with Intel Corporation, Architecture Labs, 2111 NE 25th Avenue, Hillsboro, OR 97124-5961. E-mail: jeffrey.r.foerster@intel.com.*

Manuscript received 30 Jan. 2006; revised 24 June 2006; accepted 18 July 2006; published online 15 Feb. 2007.
For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-0035-0106.
Digital Object Identifier no. XXX

1. We wish to point out here that the WiMedia Alliance supports an OFDM (Orthogonal Frequency Division Multiplexing)-based specification [5] for UWB; the motivation for dividing the available spectrum into multiple bands is to overcome the need for complex equalization. OFDM, however, first requires complex signal processing in terms of complex inverse fourier transform computations. Second, a MAC protocol for use with OFDM for UWB-based ad hoc networks has yet to emerge.

2. For a given average power constraint, the peak power constraint also imposes restrictions on the pulse repetition frequency (PRF), as we will discuss later.

3. Recent developments with OFDM and Multicarrier CDMA use carrier-based methods; the trade-offs between the use of impulse-based UWB and OFDM-based UWB are discussed in [14].

the commonly used protocols that rely on carrier sensing are not necessarily applicable with UWB. In addition, the very limited number of UWB-based MAC protocols that have been proposed previously are based on arbitration via time-hopping on a single channel. But, time-hopped sequences with a *short* spacing between the time-hops can lead to collisions, while *long* durations between time-hops can lead to excessive delays and low efficiency. Thus, the second key objective of our design is to reduce collisions to the extent possible without resorting to long time-hopping sequences.

The third motivation for our multiband approach is the associated flexibility in spectrum use and the interoperability with other networks. With multiband operations, UWB communications can coexist with other networks (such as IEEE 802.11a-based networks), a definite requirement in urban, disaster recovery, and military settings. For example, in the presence of an IEEE 802.11a network, the multiband system can avoid using the bands centered at 5.2 GHz, 5.3 GHz, or 5.775 GHz. We also wish to point out that the multiband transceiver circuit remains simple [20], i.e., the cost, power, and integration concerns are similar to those in a single-band system.

Thus motivated, we propose and develop a novel multiband MAC protocol for use with UWB-based ad hoc networks. To the best of our knowledge, this is the first multiband MAC protocol that is synergetic with UWB communications and is designed for use in ad hoc networks. The key concept of our design is the use of different bands for control and data transmission (the separation is not pure, as we will see later). Simply put, two nodes first use a control channel to facilitate a rendezvous in another band for a data exchange. The first advantage of the approach is that, since all the nodes share a common unreserved channel only for short control messages, the contention on the shared channel is limited. Second, once a pair of nodes agrees to communicate on a data band, the communication can be continuous (no need for the use of time-hopping sequences), and, thus, it is highly efficient. This efficiency is also enhanced by the fact that, once a communication is established in a data band, our protocol practically eliminates the possibility of collisions of transmissions of large data packets. Extensive simulations indicate that the throughput of our scheme is significantly higher compared to a single-band approach that combats delay spread by increasing the spacing between pulse transmissions. In addition, the number of pulse level collisions also drops dramatically.

We perform analytical assessments of our protocol to understand the efficiency with which bands are utilized. We use a single hop setting, which lends itself to analysis and yet provides significant insight into the transient behavior of our protocol. The study provides an estimate on the efficiency of the protocol in terms of utilizing the multiple bands. We validate our analytical results with complementary simulation experiments.

Finally, note that we design our protocol to conform to the requirements of the Federal Communications Commission (FCC) to ensure the practical relevance of this work. More specifically, we adhere to the FCC-specified [1] average and peak emission power levels. FCC requires that the effective isotropic radiated power (EIRP) be no higher than -41.25 dBm/Mhz.

We wish to point out here that, while, conceptually, the division of the bandwidth into a multiplicity of simultaneously usable bands is similar to a frequency division multiple access approach, the challenge is in dynamically provisioning access to these bands in an ad hoc networked setting. Previous work on multiband access in ad hoc networks has received some attention [36]; however, the approaches are based on carrier sensing (not possible in impulse-based UWB systems) and are built on top of the IEEE 802.11 MAC standard. We also wish to point out that, given the delay spread, it is impossible to design an *equivalent* time-division-based single-band approach, wherein pulses can be packed closer to each other via the use of a larger bandwidth (as mentioned earlier).

We organize the paper as follows: In Section 2, we provide the relevant background on UWB communications and discuss the physical layer dependencies. In Section 3, we provide a detailed description of our protocol. In Section 4, we present our simulation framework and results, and deliberate on the observations. We analyze our scheme in a single hop ad hoc network to obtain a more in-depth understanding in Section 5. Related work is discussed in Section 6. Finally, Section 7 concludes the paper.

2 Physical Layer Dependencies

In this section, we discuss the UWB physical layer and highlight its impact on the design of our protocol. Detailed descriptions of some of the aspects of UWB communications can be found in [1], [13] and [14].

Facilitating Multiband Impulse-Based UWB Communications. UWB communications, as per the specifications of the FCC, use the spectrum from 3.1 GHz to 10.6 GHz [1]. FCC imposes that UWB signals span at least 500 MHz of absolute bandwidth or occupy a fractional bandwidth of $W/f_c \geq 20\%$, where W is the transmission bandwidth and f_c is the frequency at the center of the band [13]. UWB systems have traditionally achieved these high bandwidths by using pulses that are of very short time duration; we refer to these as *impulse-based* UWB systems. A typical UWB pulse belongs to the family of *Gaussian* shaped doublets [13], [14].

Multiband modulation facilitates the division of the 7.5 GHz of spectrum into multiple smaller frequency bands, each of which conforms to the aforementioned FCC specifications. With impulse-based UWB, the pulse shape determines the distribution of energy in the frequency domain and therefore allows for the separation and, thus, the simultaneous use of the bands. Depending on the spectrum of operation, the Gaussian pulse is modulated by a set of carriers that belong to the particular band. This center frequency component is typically referred to as the *pseudocarrier*. Note here that these high frequency modulating signals are simply used to shape the pulse and are *not* used to reflect encoded bit information as in traditional modulation methods (such as frequency shift keying or FSK [2]). We also wish to point out that the center frequency components of the different bands must be separated

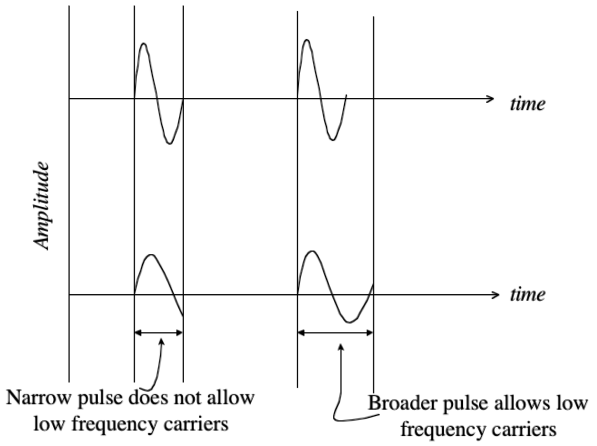

Fig. 1. The use of multiple bands requires elongated pulse durations.

sufficiently in the frequency domain to avoid interband interference effects. A detailed discussion of pulse shaping can be found in [14]. In our simulations, we use a simple Gaussian-shaped pulse and assume that appropriate modulating signals can be employed (as shown in [14]) to divide the bandwidth into multiple bands, each of which is 500 MHz wide. The pseudocarrier of the highest band is 10.35 GHz and that of the lowest band is 3.35 GHz. The pulse-width will depend on the occupied bandwidth. Since the occupied bandwidth in the multiband case is smaller than a single band case (wherein the entire allocated spectrum is used), the pulse-width would be longer with multiband impulse-based UWB, as shown in Fig. 1. Note that the elongation of the pulse width is a necessity with a multiband approach.

Encoding Information: Pulse Position Modulation. The modulation scheme that we use is a commonly studied scheme called Pulse Position Modulation or PPM [14].[4] We also assume the use of a rate 1/3 convolutional code [3]; thus, the information in each bit is encoded into three *encoded bits*. Each pulse represents an encoded bit. The information in the encoded bit is determined by the *position* of the pulse within what we call a **chip time** T_c. If the pulse occupies the first part of the chip-time, it represents an encoded bit value of "0"; else, an encoded bit value of "1" is implied. We assume that a Viterbi decoder is deployed at the receiver [2] to enable the soft-decision decoding of the received information.

Time-Hopping. Time-hopping has been used in previous approaches for sharing a single impulse-based UWB frequency band among multiple users [6], [7]. In mobile ad hoc networks, imposing fixed TDMA-like schedules is difficult due to the fact that nodes could be mobile. A completely random access scheme is 1) unlikely to provide high throughput and 2) requires nodes to acquire synchronization at arbitrary unpredictable time instants. Time-hopping is a form of spread spectrum communications specifically designed for impulse-based systems. In single-band approaches, nodes transmit as per pseudorandom time-schedules. The pseudorandom nature of time-hopping provides a reasonable level of robustness to collisions; however, at low loads with time-hopping, the spectrum is not used efficiently (as we show with our simulations). To the best of our knowledge, all protocols designed thus far, for impulse-based UWB ad hoc networks, use time-hopping as the basic means of providing multiple-access. In contrast, with our approach, we use time-hopping only in the control band and not in the data bands, as we explain later. In time hopping, a fixed number of chip-times are aggregated to form a **sequence frame**. The duration of each sequence frame is T_f, and, thus, the number of chip-times per sequence frame is T_f/T_c. Each transmitter sends a pulse in only one of the chip-times in each sequence frame. The specific chip-time is determined by the node's time-hopping sequence (THS), generated as per a pseudorandom number (PN) code. The distribution of PN codes (for making a node's THS known to its neighbors) has been the topic of a few efforts [21], [22]. In our work, we assume that the PN code is a function of a node's identifier (possibly the MAC layer address). The generators of these PN code sequences are initialized at system setup. Nodes periodically use out-of-band techniques to announce the state of their PN code generators.[5] The technique is similar to the proposal in [21].

Time-hopping sequences may be either sender-based or receiver-based. In receiver-based time hopping, a receiver expects to receive a pulse only in one of the chip-times in a sequence frame. In the sender-based case, the transmitter sends pulses based on its own THS. The sender-based strategy is robust; however the receiver has to be synchronized with all of its potential transmitters. The receiver-based approach is much simpler; however, one could encounter collisions between the pulses from different transmitters directed toward the same receiver. Protocols could use both approaches, as in [6].

Note that it is extremely difficult to guarantee that time hopping sequences of nodes are orthogonal to each other. To satisfy this requirement, the number of chips within a sequence frame would have to be *larger* than the number of nodes in the network, i.e., each node would have to have a dedicated chip-time for duration. This would result in extremely long sequence frames and has two consequences: 1) the utilization of the available spectrum is heavily affected and 2) longer delays are entailed, especially at low loads. Continuing the above discussion, the average spacing between successive transmissions as per the THS will have an effect on the achieved performance. With shorter spacing between the time-hops, the pulses could be sent at a faster rate;[6] however, there is a higher possibility of collisions. With longer spacing, the possibility of collisions is reduced; however, large delays could be incurred. With our scheme, as mentioned earlier, time hopping is only used for the transfer of short control messages; since these messages are infrequent and fairly short in duration (low load), the probability of experiencing collisions remains low even with a relatively short spacing between the time-hops.

4. We wish to point out here that multiband systems may be based on direct sequence CDMA modulation or OFDM. For this work, we use PPM modulation. While no CDMA or OFDM-based MAC layer solutions have been completely developed for ad hoc networks, we discuss relevant related work in a later section.

5. These announcements are made in special frames that we refer to as *Availability frames*. We discuss this in Section 3.

6. The FCC regulations impose a limit on the pulse repetition frequency, as will be discussed later.

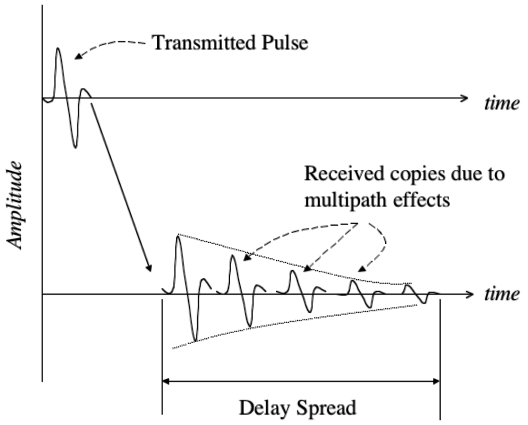

Fig. 2. The multipath delay spread phenomenon.

Channel Impairments and Effects. We next discuss the effects of the wireless channel on UWB communications and the associated impact on our MAC protocol design. A signal typically experiences three types of channel impairments: pathloss, shadowing, and multipath effects. The pathloss factor is given by Frii's law [13] and is $\alpha = \left(\frac{c}{4\cdot\pi\cdot d_{ij}\cdot f_c}\right)^2$, where c is the speed of light, f_c is the center frequency of the band, and d_{ij} is the distance between the transmitter and the receiver. Note that the above equation depicts the observed effects *on average*, and does not imply that *each* transmitted signal experiences the *same* level of attenuation. (For wideband systems, the geometric mean of the upper and lower frequency limits of the pulse band is more accurate than using the center frequency in the Friis equation, but the center frequency is sufficient for this study). Furthermore, at a given distance d_{ij}, higher frequencies will experience *higher levels* of attenuation than lower frequencies. Shadowing is ignored since short range transmissions ($\approx$ 10 meters with UWB) do not experience shadow fading [19].

UWB transmissions (high data rates) will experience *multipath delay spread*. A transmitted UWB pulse, radiated using an isotropic antenna, will take multiple paths (as a consequence of reflections from various objects) and results in multiple time-shifted copies at the receiver. Each received copy may have a different amplitude, phase, and delay. Beyond a certain delay threshold, called the delay spread of the channel, the signal amplitudes may be considered negligible. This multipath phenomenon is depicted in Fig. 2. For indoor environments, measurements have shown that the delay spread is of the order of tens of nanoseconds [19]. If the time-spacing between the UWB pulses is smaller than the delay spread of the channel, copies of a transmitted encoded bit interfere with the subsequent encoded bits. This is called *intersymbol interference*, or ISI for short. Equalizers could be used to combat ISI [2]. The higher the ISI, the higher the complexity and sophistication of the required equalizer. Equalizers also require the transmission of a training sequence prior to information communication. This can be expensive in terms of the overhead consumed. With UWB transmissions, a preamble is needed to allow for the sender and receiver to synchronize prior to communications. By acquisition, we mean that the receiver learns how to recognize the presence of a pulse train in the presence of thermal or other noise factors. The aforementioned acquisition preamble is considered expensive in terms of overhead [18]. The deployment of a sophisticated equalizer will further increase the overhead costs incurred with UWB.

Another strategy for combatting ISI would be to use direct sequence CDMA in conjunction with a Rake receiver. However, the long codes with CDMA could still incur capacity penalties. Furthermore, with CDMA, the sender and receiver require code synchronization in addition to the acquisition and this would incur a further cost in terms of overhead.

The alternative that we explore in this work is to separate the pulses by at least the delay spread of the channel. Thus, the time-spacing between the pulses is chosen to be at least 30 ns[7] (delay spreads in indoor environments [19]). We recognize that, by doing so, the pulse width could be increased to some extent since this is unlikely to interfere with future encoded bits. Increasing the pulse width allows for the use of lower bandwidths and, thus, facilitates the use of multiple frequency bands, as discussed earlier.

Conformance with FCC Regulations. The FCC regulations limit the effective isotropic radiated power (EIRP) to -41.25 dBm/MHz (Part 15 of the regulation) [1], [13]; the power used on average per bit cannot exceed this imposed limit. Let us denote the transmit power by P_T dBm/MHz, the received SNR at a distance d by SNR_d dBm/MHz, and the center frequency in the band used by f_c. Let the power spectral density of the thermal noise be N_o dBm/MHz. Then, the signal to noise ratio is given below [13]:

$$SNR_R = P_T - N_o - N_f - 20\log\left(\frac{4\pi f_c}{c}\right) - 20\log d + 10\log BT_c. \quad (1)$$

In the above equation, N_f refers to the noise figure of the receiver, B is the bandwidth of the UWB pulse, and T_c is the time-spacing between pulses. The last factor in the above equation is typically refered to as the pulse processing gain, since a UWB pulse can increase its transmission power when it is "on" and still meet the FCC average transmit power limits which simply averages the "on" and "off" periods over a short time duration (on the order of a millisecond). As long as the peak power limits are not exceeded, this pulse processing gain can be fully realized. For example, if P_T is -41.25 dBm/MHz, SNR_R is set to 3 dB,[8] the thermal noise density is $N_o = -114$ dBm/MHz [13], $N_f = 7$ dB, $B = 500$ MHz, $T_c = 60$ nsec, and the highest center frequency is $f_c = 10.35$ GHz, then the maximum theoretical free-space range could be 17.3 meters. However, this does not take into account non-line-of-sight propagation and possible shadowing or obstructions by people or other objects. Since this study is focused on short-range ad hoc networks, we simply assume a maximum range of 7 meters for all possible bands, which is achieved by reducing the power for each band in order to maintain the target SNR at a maximum range of 7 meters. For example, the transmit power for the band with a center frequency of

7. Note that this translates to having a *chip-time* of 60 ns.
8. SNR = 3 dB.

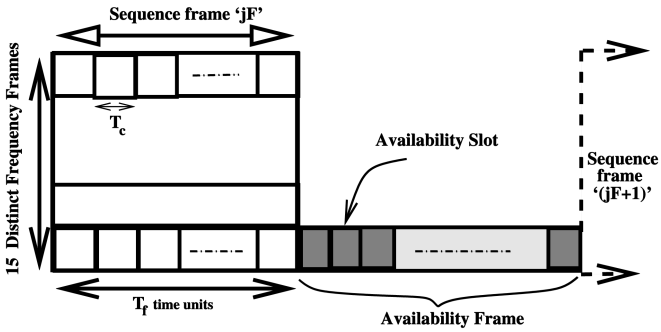

Fig. 3. The frame structure with our protocol.

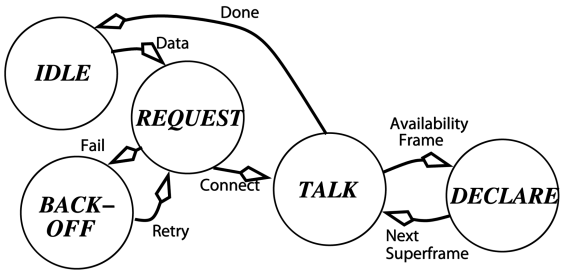

Fig. 4. Depiction of protocol operations.

$f_c = 3.35$ GHz will be 9.8 dB lower than the transmit power for the band with a center frequency of 10.35 GHz for the same range. As a result, the appropriate powers are set to ensure that the range is identical, irrespective of the band being used. Clearly, the powers used will be lower than the FCC imposed limit, but this may be a very desirable property for military systems, for example, in order to reduce the probability of intercept. To summarize, with the settings as above, we conform to the FCC imposed restrictions on the EIRP.

In addition to the imposed restriction on EIRP, the FCC also imposes a limit on the peak power that can be used for UWB transmissions. As specified in [1], if the average power limit is met and the frequency of pulse transmissions is higher than 1 MHz, the peak power limitation is also met. With our scheme, since the maximum distance between the pulses is 60 ns (the chip-time), the frequency is 16.67 MHz. This implies that our scheme inherently conforms to the peak power constraint.

Coding and Higher Layer Abstractions. In addition to the rate 1/3 convolution code (mentioned earlier), we employ a *repetition code* of 2 for the control messages. With the repetition code, the output of the convolutional encoder is repeated twice. Again, as alluded to earlier, the control messages are transmitted as per a time-hopping pattern and the code helps alleviate the effects of collisions between pulses. In our simulations, we assume the presence of the convolutional encoder and decoder and do not implement them. Instead, we use the bit error rate of 10^{-7} and discard bits at this rate.

Time Synchronization. Our approach requires the division of time into frames, which implies that communicating nodes must be synchronized in time. This requirement is not unresonable because of the following reasons: First, with our MAC frame structure (to be described), nodes rendezvous on a periodic basis in what we call "availability frames." The nodes can adjust their clocks based on when this frame appears. It is possible that the nodes have different views of time as long as they are aware of the clocks of each of the neighbors that they communicate with. Furthermore, several time-synchronization methods have been proposed for ad hoc networks [25], [26]. If GPS is available, it might be used to provide clock synchronization. Finally, we assume that nodes are equipped with accurate clocks (as with current technologies, such as Kernco laser-based Atomic Clocks) [37], which lose a second in approximately 10,000 years. We also assume that appropriate (fairly short) guard bands will be employed. For the synchronization between transmitter and receiver, the transmitter sends a preamble just before the transmission of request.

3 The Multiband MAC Protocol

In this section, we present our multiband MAC protocol. The key idea is to have a communicating pair of nodes exchange data over a private band as opposed to a single common band. We do not use time-hopping and, thus, avoid its disadvantages (discussed earlier) in the private bands. We first give a brief overview of the basic concepts and the operation of the protocol.

The Multiple Bands. We divide the available frequency bandwidth into B bands. $B-1$ of these bands are used for data transmissions and are referred to as *data bands*. The remaining band is used for request control packets only; we call it the *Request Band* or *Req-Band*; the first band is assigned to be the Req-band. As discussed earlier, if each band is of bandwidth 500 MHz, $B = 15$. The protocol is designed based on the physical separation of the available UWB bandwidth of 7.5 GHz into multiple bands (as discussed in the previous section), each of which spans 500 MHz of the spectrum.[9]

The Frame Structure. Across all the bands, time is broken into *superframes*, which are separated by smaller *availability* frames. All data and control communication takes place during superframes. The availability frame is used to indicate whether each band will be busy or not in the next superframe (explained later). The availability frames alleviate the possibility of collisions of *data transmissions* in the superframes. Note that each superframe consists of F sequence frames, each of which in turn consists of T_f/T_c chip-times. Fig. 3 depicts the frame structure; the availability frame is sandwiched between the last sequence frame of the jth superframe and the first sequence frame of the $(j+1)$st superframe.

Protocol Operations in a Nutshell. We first provide a high level overview of our protocol. The protocol implementation at each node can be represented by a finite state machine, as shown in Fig. 4. Initially, a node is in the IDLE state. When data needs to be sent, it enters the REQUEST state as shown. In this state, it attempts to *initiate* a request to the appropriate receiver. Toward this, it transmits the request as per the *receive* THS of the receiver in the Req-band. If this request

9. Note that FCC specifications require that each UWB band is at least 500 MHz in bandwidth.

were to succeed, the node enters the TALK state, switches to a data band (the rules for choosing a band will be discussed later), and attempts to establish a connection with the receiver. If the request were to fail, it enters the BACK-OFF state and tries again at a later time. Upon a successful request, the node sends data. In addition, it periodically announces (by transitioning to the DECLARE state), by means of the availability frame, that the specific data band being used is *occupied*. This precludes other nodes from claiming the particular band and causing collisions. After the data transfer, the node returns to the IDLE state.

Detailed Protocol Descriptions. Next, we discuss the protocol in greater detail and, in particular, the nuances of protocol operations in each of the aforementioned states.

Request Initiation. Upon having data (either its own or data that it has to forward) to send to a neighbor, a node will first have to send a request to the receiver. Our design mandates that transmissions are to be initiated at the beginning of each superframe, i.e., right after the availability frame. The availability frame (as we will discuss later) reflects the *occupancy* of each of the data bands. By the above design mandate, we ensure that nodes have *up-to-date* information on which of the data bands are occupied prior to initiating new transmissions. This would prevent these new transmissions from colliding with previously initiated data transfers that might be in progress. Thus, if a node (whose queue was empty at the end of a particular availability frame) generates packets for transmission in the middle of the following superframe, it is precluded from initiating a transmission *before* the end of the upcoming availability frame. At these allowed times, in order to initiate a request, the sender sends a *REQ packet* in the Req-Band as per the THS of the receiver. The REQ packet identifies the particular band that the sender has *chosen* for the data exchange. After transmitting the REQ packet, the sender switches to the indicated data band and awaits a response from the receiver. Note that the above operations occur in the REQUEST state discussed earlier.

Acknowledgment of the Request. If the REQ packet is correctly received, the receiver will switch to the specific data band indicated in the REQ packet and will send a *RACK (Request Acknowledgment)* packet to the originating sender. If the RACK packet is successfully received by the sender, it completes a successful handshake and the sender can then begin the data transfer.

Transfer of Data. The reception of the RACK asserts that the band is almost surely free for exclusive use for data transfer. In the chosen band, nodes (now in the TALK state) transmit data in *consecutive chip-times* instead of using time hopping. As discussed in the previous section, the spacing between the pulses is at most 60 ns and we ensure that the FCC emission regulations are met. Upon the successful reception of a complete data packet, the receiver sends a *DACK (Data Acknowledgment)* packet to the sender. Even if collisions are completely eliminated, it is possible that other noise factors (thermal noise) can corrupt the data packet. If the receiver is unable to correctly decode the packet, it does not issue a DACK back to the sender. The sender would then reattempt to transmit the data packet up to a fixed number of times, after which the packet is dropped.

The Availability Frame and the DECLARE State. As mentioned earlier, superframes are interspersed with the so-called availability frames. During the much smaller availability frame, data communications stop temporarily so that nodes currently occupying a data band can signal their intention to continue using it during the next superframe. This signaling takes place in the Req-Band (we could have chosen any band, since availability frames are exclusively used for signaling availability and no data transfers occur during these frames). The availability frame is divided into time intervals that are different in size from those in the sequence frames. The number of these intervals corresponds to the number of data bands. We call these intervals *availability slots*. Communicating nodes "saturate" the availability slot that corresponds to the data band that they intend using in the next superframe. As an example, if a communicating pair is using data band j, where $2 \leq j \leq B$, the pair would transmit in the jth slot of the availability frame. The sender saturates the first half of the availability slot and the receiver the second half. This is done to ensure that the neighbors of both the sender and the receiver are made aware that the corresponding band is occupied. Nodes in search of an available band listen to the availability frame and select an unused band for their upcoming data transfers. Note that, due to the consecutive transmission of pulses during the availability frame, nodes are able to detect (or sense) the pulses. The size of each availability slot is chosen so as to accommodate an adequate number of pulses to facilitate acquisition and to combat noise effects. The availability frame corresponds to the DECLARE state discussed earlier.

Choosing a band for communication. Initially, each sender selects a band randomly from the set of free bands, as indicated by the availability frame. The following mechanisms are incorporated to further reduce the possibility of collisions due to multiple new senders choosing the same band:

1. **Persistent Band Selection.** Nodes keep a history of bands that they successfully used in the past. The random choice process is biased with time such that the nodes would prefer to reuse these previous successfully used bands. In the long run, this can further reduce the possibility of two (or more) senders selecting the same band. This can be particularly helpful when traffic is bursty and the same sender nodes are active for repetitive interspersed busy and idle periods. Note that, in most practical networks, traffic and communication patterns are indeed bursty [24].
2. **Availability Eavesdropping.** Nodes can determine the bands that are being used, even when they do not have packets to send, by means of the availability frame. Thus, they can keep track of bands that are consistently occupied (as per persistent band selection). A band that is often busy is more likely to be used in the future. Thus, new senders can avoid the use of these bands.

In fact, the two mechanisms discussed above enable a self-organizing behavior, where groups of nodes that have disjoint periods of activity end up having the same *preferred*

bands. A trivial example is the case of two nodes sharing two bands: Persistence and eavesdropping can lead to each node using a band without conflicts. This can be generalized to a case with any number of nodes and bands.

Failure of the Request Process and the BACK-OFF state: There are three cases where the receiver does not reply successfully to the sender with a RACK: 1) there were more than one REQs that collided, 2) the receiver is busy, and 3) two or more pairs of communicating nodes attempt to use the same data band. To elaborate on case 1, if two nodes (or more) transmit their REQs to a common receiver at the same time, a collision will occur. In this case, the two senders after the REQ transmission will switch to their own selected data bands and will wait for a response from the common receiver. As a result of the collision of the REQ packets, they do not receive a response. The sender nodes wait for a specified time interval in their selected bands and, at the end of this period, they conclude that a collision has occurred. They will then initiate back-off timers and, at the end of their back-offs, reattempt to initiate the request.

We employ a simple additive back-off scheme[10] for retransmission attempts after a failure. Upon experiencing a collision, a sender chooses, with a uniform probability, one of the M subsequent superframes to reattempt its request. The number M is given by $M = N + xL$, where x is the number of consecutive failures and N and L are system parameters that define the aggressiveness of the back-off policy. We impose a *maximum limit* on the number of retransmission attempts x, after which the packet is dropped.

To elaborate on case 2, if the receiver is busy in another data band either sending or receiving data, it does not receive the REQ packet. The sender will, as in the previous case, transmit the REQ packet and await the RACK packet in the data band of its choice. Clearly, in this case, no RACK packet is forthcoming. The sender cannot distinguish this case from case 1, in which a collision occurs. Therefore, it enters the BACK-OFF state as discussed earlier and reattempts a request at a later time.

In case 3, if two or more pairs of nodes select the same band, their transmissions may collide in that data band. The problem is exacerbated when the number of sender nodes is much larger than the number of bands. This problem is alleviated to a large extent by our policy of initiating new transmissions only at the beginning of a superframe. Thus, when two pairs of nodes choose the same band, their RACK packets collide. The nodes would infer that a collision has occurred and retract to reattempt a reservation. Note that the collision is quickly and efficiently detected.

Enabling Receptions while in Back-Off. While a sender node's back-off counter is counting down, the node switches to the Req-Band. In this band, it resumes receptions as per its THS while its back-off counter ticks down. If the node successfully receives a preamble for a REQ packet (from any of its neighbors), the node temporarily freezes its back-off timer and switches to the band specified in the REQ packet to attempt a reception from the originator of the REQ packet. Upon the completion of the reception, or upon the detection of a collision, the node under discussion would switch back to the Req-Band and resume the countdown of its back-off timer. Without this, nodes that are attempting to contact the sender under discussion would be forced to back off, while the node is counting down. This would degrade the efficiency of the system, since a node can be blocked waiting to send to another blocked node forming a chain or even a cycle.

Finally, while awaiting RACK, Data, or DACK packets, a node will wait only up to preset time limits (system parameters). If the expected packet is not received within this time limit, the node assumes that the communicaton has failed. The sender would then attempt to resend a request after an appropriately chosen back-off time.

Multihop Communications: Coping with Hidden Terminals. The hidden terminal problem is already alleviated to a great extent since the transmitter and the receiver both send messages in the availabilty band to indicate the occupancy of the band on which they currently communicate. This ensures, to a large extent, that neither the neighbors of the sender nor those of the receiver claim the same band. However, note that, after transmitting their REQ packets, the nodes rendezvous in the chosen data band. Since the REQ packets are sent according to THSs, the transgression of communicating pairs onto data bands is not synchronized. Thus, it may happen that two pairs choose the same bands but move to that band for the rendezvous at different times. Now, if, after such a rendezvous between a given sender and a receiver, a neighbor of the receiver, hidden from the sender, switches to the same band and initiates a new message transfer, a collision would occur at the receiver. In order to avoid such effects, we require that 1) when a pair of nodes switches to a new band for data transfer, they wait for a duration of T_n nanoseconds (during which they listen to other possible communications on the band) prior to completing their handshake and beginning the data transfer, and 2) receivers send short *occupancy* indication messages in the band on which they are on with a periodicity of T_n nanoseconds. This will further reduce the possibility of collisions due to hidden terminals. In fact, our simulations suggest that the two schemes together practically eliminate collisions.

4 Simulation Results

We present the evaluation of our idea through simulations using a C++ simulator that we have developed by extending a previous simulation effort [10]. Our focus is on the performance at the MAC layer. Thus, we assume that data is injected at the MAC layer and the transmissions of a node are intended for a neighbor. However, we wish to clarify that nodes are distributed over a region of interest for multihop operations; thus, MAC layer effects, such as the presence of hidden terminals, are accounted for in our simulations. In our simulations, we use assumptions and conventions that are widely used in UWB studies and try to incorporate as many realistic details [6], [10] as possible. Some of our simulation assumptions were alluded to in Section 2.

Comparisons. We compare our scheme with a single-band approach in order to demonstrate the benefits of our

10. Our simulations suggest that this simple scheme is very effective in resolving collisions.

multiband scheme. In a nutshell, the single-band approach is based on using a single band with time-hopping as the basic means of access. We choose this, given that there do not exist MAC protocols that are based on other basic multiple access methods, for ad hoc networks. We do not assume the presence of an equalizer and, hence, the pulses are spaced apart as in the multiband approach. One might think that the single band approach is disadvantaged to a large extent; while this is true in some sense and it is intuitively clear that the multiband approach can yield a significant increase in the achievable throughput, especially when the number of users is small (and, thus, the bottleneck is not the reservation channel), the comparison quantifies the achievable gains. Furthermore, we provide some sample results, wherein we eliminate some of the collision effects in the single-band approach (the approach we take for doing this is discussed below); this provides a fairer comparison of the two approaches.

The Single-Band Approach. With the single-band approach, data and control packets use the entire 7.5 GHz bandwidth (whereas up to $B-1$ simultaneous users can transmit data packets on different bands during the same superframe in our multiband scheme). The approach is *loosely* based on the approach in [6]. Initially, the nodes exchange the control messages (as with our protocol) to establish a handshake. If the handshake is successful, the nodes switch to a unique THS on which they communicate. However, note that the bandwidth is shared among the plurality of users and the data transmissions will also have to compete with the transmission of the control information. This would put the single-band approach at a distinct disadvantage, especially at low loads; at these low loads, with the multiband approach, transmissions will be practically collision-free, whereas collisions would be higher with the single-band approach. In order to avoid giving our scheme an unfair advantage, we provide a version of the single-band approach where we *magically* eliminate the effects of pulse collisions on the reception of data packets; when the communicating nodes switch to the unique predetermined THS (mentioned above) to exchange data packets,[11] they communicate collision-free. Note that this assumption now shifts the unfair advantage to the single-band case, since many more than $B-1$ simultaneous data transfers could be supported if the requests get through.[12] One can envision this to be akin to using a perfect equalizer, which is calibrated during the reception of a request packet, to eliminate the ISI during the reception of the following data packet. Note, however, that with both the *collision-free* version of the single-band *and* the multiband approaches, pulse collisions may occur during the initial handshake wherein a request is transmitted (we refer to this as request transmissions in Section 3) as per the receiver's THS. In our plots, we label the more realistic single-band approach as simply *single-band*; we label the collision free version of the approach as CF-Single-Band.

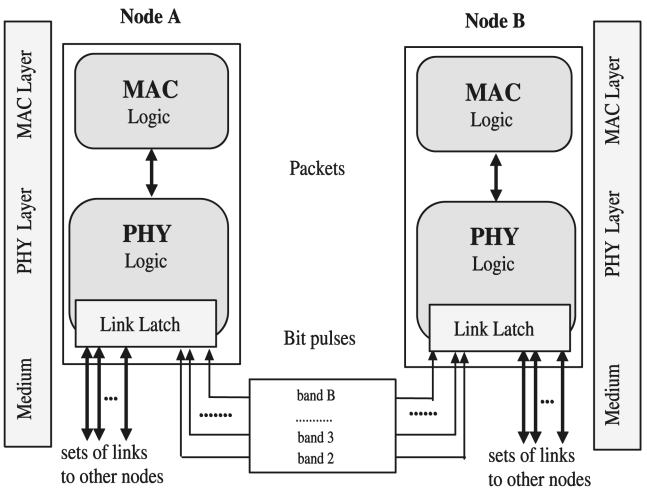

Fig. 5. Simulation implementation platform.

Simulator Implementation Details. In our implementation, the physical layer consists of a number, m, of sets of virtual links, as shown in Fig. 5. This number is equal to the number of bands; each set of links has a separate buffer and connects a node with its neighbors.[13] As a result, a node has m links with a neighbor node, each representing a different band. The MAC layer of the transmitter delivers the packet to the appropriate link of the appropriate band. The physical layer component converts the bits to pulses, which will be transmitted through this link. The channel characteristics, discussed earlier in Section 2, are applied and distort the transmission. The receiver picks each pulse, decodes a set of pulses that form a bit if possible, and stores the bit in a buffer. A bit may be discarded either due to a collision (elaborated below) or due to its being corrupted by thermal noise as discussed in Section 2. When a set of bits that form a packet have been received correctly, the packet is reconstructed and delivered to the receiver's MAC layer. The arrival of two or more pulses, simultaneously from different links of the same band, denotes a collision.

Simulation Scenarios.

Network Layout. The nodes are mobile and form an ad hoc network. We vary the number of nodes from 6 to 30. We restrict the nodes in a $30\ \mathrm{m} \times 30\ \mathrm{m}$ square region and this is seen to maintain an average node degree larger than 3. As mentioned in Section 2, the maximum range of a transmitter is considered to be 7 meters. The total number of bands in the multiband system is 15, as mentioned earlier. A transmitter always selects a receiver randomly from within its transmission range.

Frame Structures. Every sequence frame consists of six T_c frames (chip-times). The duration of the superframe is set to 11,200 chip-times, which is approximately equivalent to a successful packet exchange including the control overhead. We divide each availability frame into $B-1=14$ availability slots, each of which is 33 T_c units of time in duration.

11. A similar single-band scheme is described in [6].

12. However, we wish to point out that there may be other single band approaches that simply have pulse boundaries commensurate with the actual duration of the pulse as opposed to the delay spread. This may lead to a higher number of pulse collisions; however, at the same time, higher levels of redundancy may be employed, given that more pulse transmissions are potentially possible. A systematic evaluation to determine the best possible single band approach is beyond the scope of this work.

13. Note that two neighbors may have more than just their common neighbors.

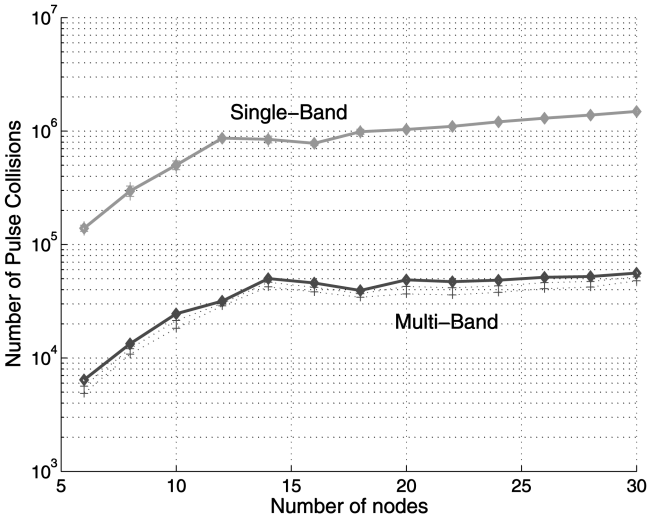

Fig. 6. Number of pulse collisions. Logarithmic scale is used.

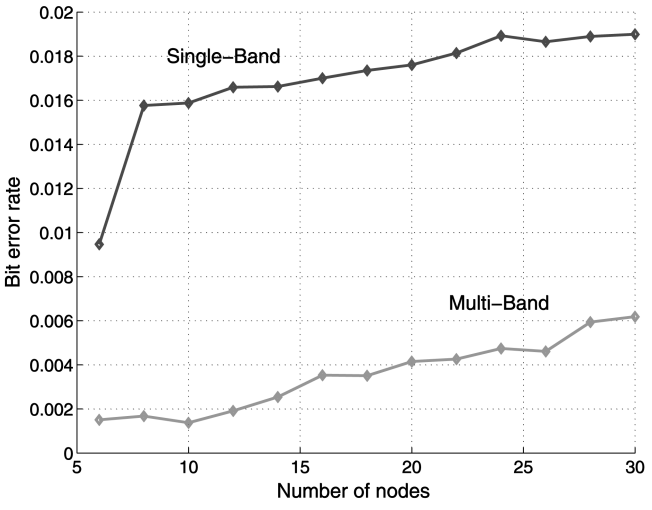

Fig. 7. Bit error rate in the network. The multiband scheme outperforms.

This duration is sufficient for neighbor nodes to detect the pulses and correctly infer that the corresponding data band is occupied.

Traffic Characteristics. We use both CBR *(Constant Bit Rate)* and bursty Poisson traffic in our simulation experiments. In the experiments shown, the packets are of size 250 bytes and, with CBR, are transmitted once every 40 msec unless otherwise stated. Each control packet is 120 bits long, in accordance with the control packets used with other wireless protocols, such as with the IEEE 802.11 MAC protocol [11]. The 120 bits of the request packet correspond to the transmission of 4,320 pulses. This also includes the synchronization preamble. Even though the per-source CBR rate is low, with a multiplicity of sources, the load on the network fits in with the use of UWB.

Pulse Collisions and Bit Errors. As stated earlier, we assume the use of a rate 1/3 convolutional encoder for the data bands. In the control band or in the case of the single-band system, we use a repetition code of 2, i.e., each encoded bit is repeated twice. Thus, in these bands, six pulses form a bit. A pulse collision occurs when two or more pulses arrive during the same T_c period in the same band. A bit is received in error, when all of the pulses that make up the bit collide or if it is corrupted due to thermal noise.

Mobility Model. We assume that nodes move as per a Brownian motion mobility model. Each node chooses a new position that differs from its current position by at most 10 cm in a randomly chosen direction, once every 6 milliseconds.

Back-Off Policy. With our back-off algorithm (discussed in Section 3) for packet retries, we set the initial back-off to a randomly chosen value between 0 and 5 superframes. After each retry, the maximum value increases by 2, until it reaches a maximum of 15. We have varied these values and the results obtained demonstrate behavioral traits that are similar to those considered in our sample set presented here. The packet is discarded if, after 15 attempts, a node is unable to deliver it to its intended neighbor.

Providing for Consecutive Packet Transmissions. With any given reservation, we allow a transmitter to send two consecutive packets to its receiver. This would, in some sense, amortize the preamble and request costs over a larger transmission. We restrict this number to two to prevent the dominance of a channel by a single communicating pair. The overall simulation time is 15 million chip-times T_c.

Performance Metrics. We evaluate the performance of our scheme by measuring the number of pulse collisions, the bit error rate, the overall number of transmitted data packets during the simulation, the average packet delay, and the utilization of a band. We define the metrics in detail when we discuss the results. In a few of the following graphs, we include 95 percent confidence intervals.

Results. Due to space constraints, we only present a sample set of results. We first present results when CBR traffic is considered.

In Fig. 6, we plot the total number of pulse collisions for each approach as a function of the number of nodes in the network. For this set of experiments, we use the Single-Band approach for comparisons; in effect, we do this since we wish to compare the *actual collision rates* with the two schemes. We observe that our protocol decreases the number of pulse collisions by an order of magnitude as compared with the single-band approach. This is expected since, with our protocol, data packets are transmitted practically free of collisions, since they are exchanged on an exclusively reserved data band. In contrast, in the single-band case, packets suffer frequent collisions due to overlaps between nodes' THSs.

In Fig. 7, we plot the bit error rate averaged over the observations from all the nodes in the network as a function of the number of nodes. We observe a much higher (more than 4 times) bit error rate in the single-band system, which is, again, a direct result of collisions of data packets.

Next, we report the observed average packet delay in the network. The packet delay is the duration between the instance that a packet arrives to the MAC layer queue of a node until the instance that it is completely reconstructed at its destination. With the multiband approach, this delay accounts for retransmissions that may occur due to the failure of the packet transfer due to the packet being corrupted or collided with. We consider the CF-Single-Band

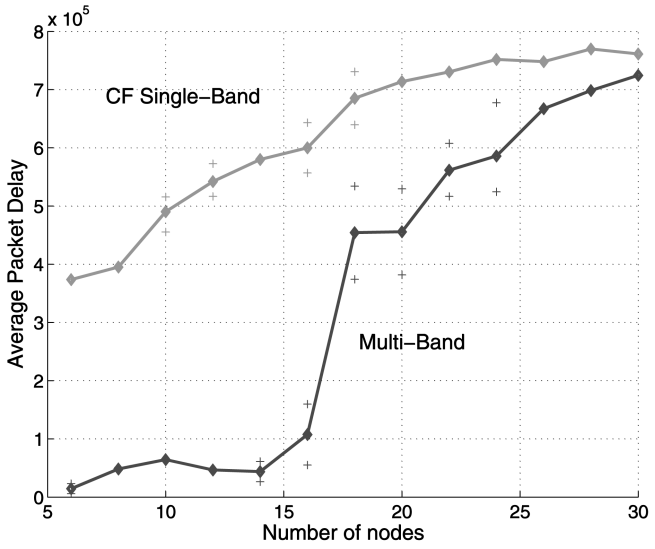

Fig. 8. Average packet delay in the network, in pulse slots.

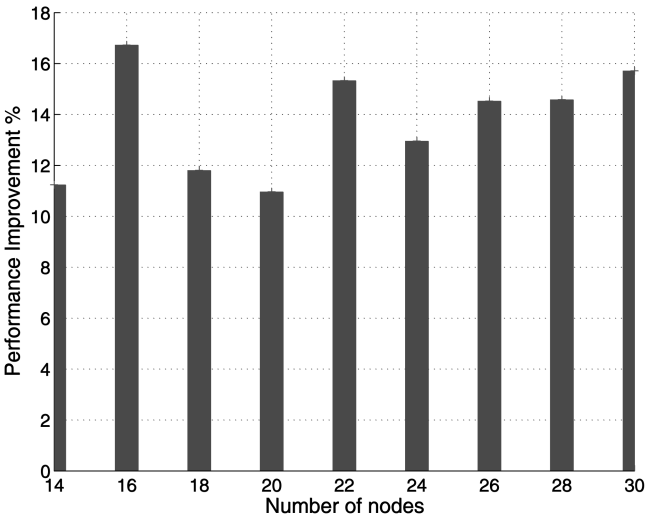

Fig. 9. Performance improvement in terms of throughput.

approach for comparisons, i.e., the data transmissions are expected to be collision-free if the request handshake is successful with the single-band approach as described earlier. In Fig. 8 we plot the average packet delay as a function of the number of nodes in the network with CBR traffic.

In our protocol, packet delays are lower by a factor of six as compared with the delay incurred with the CF-Single-Band scheme for low network densities (i.e., when the network consists of 15 to 16 nodes). With more nodes, the multiband delay rapidly increases and approaches the delay that is observed with the CF-Single-Band case when there are approximately 30 nodes in the network. Note that this behavior with a large number of nodes is an artifact of the system reaching its capacity. Recall that, as the number of nodes increases, the network load increases as well in this experiment.

We also measured the total number of transmitted data packets for the duration of the simulation with CBR traffic (Fig. 9). The improvement is calculated as

$$\text{Improvement} = ((T_{Multiband} - T_{Singleband})/T_{Singleband}) * 100\%.$$

Note from Fig. 9 that the network throughput in terms of transmitted packets is higher with the multiband scheme. Our protocol performs better by as much as 16.72 percent, a significant increase at these higher capacities. Furthermore, notice that this improvement is over the *unfairly* advantaged CF-Single-Band system, which magically eliminates the effects of pulse collisions on the reception of data packets and, hence, can support any number of simultaneous data transmissions as long as their request handshake was successful.

In all of the previous examples, we assume that the load increases with the number of users. We perform experiments to demonstrate the benefits of our scheme with high loads when the number of users in the network is small. In such cases, communicating pairs can be allocated exclusive bands in the multiband approach. With the single-band approach, however, throughput is much lower due to collisions. For this experiment, we assume that packets are generated at each node at a constant rate of one every 1.4 milliseconds. We observe that, now, the achieved throughput is more than an order of magnitude better than with the single-band approach (shown in Fig. 10).

We wish to point out here that we do not present results with extremely high loads. This is because, under these conditions, there will be a very high rate of collisions in the Req-band in both the multiband and single-band systems. As a result, the throughput is driven to very low (and, therefore, uninteresting) values.

Our final simulation experiment examines the band occupancy with the multiband approach. The objective of this experiment is to determine if the traffic load is uniformly distributed across the data bands or if some bands are preferentially used with our policy. The motivation for this study is that, if some of the bands are hardly ever used, one might consider the usage of additional control bands to improve efficiency. For facilitating understanding, in this simulation, we assume a clique with ad hoc

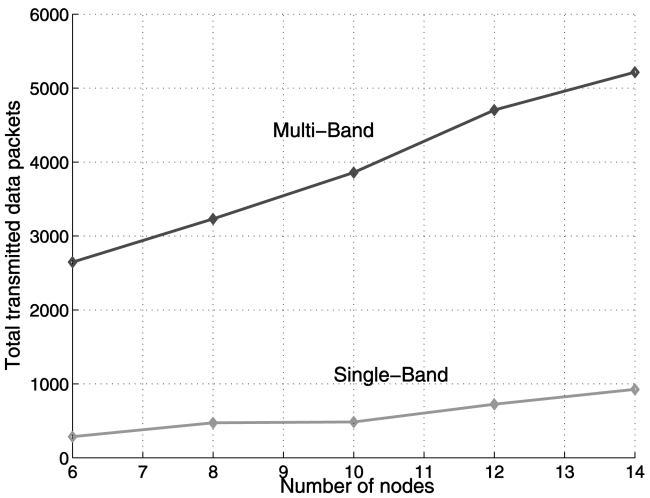

Fig. 10. Throughput improvements at heavy loads with a small number of users in the network.

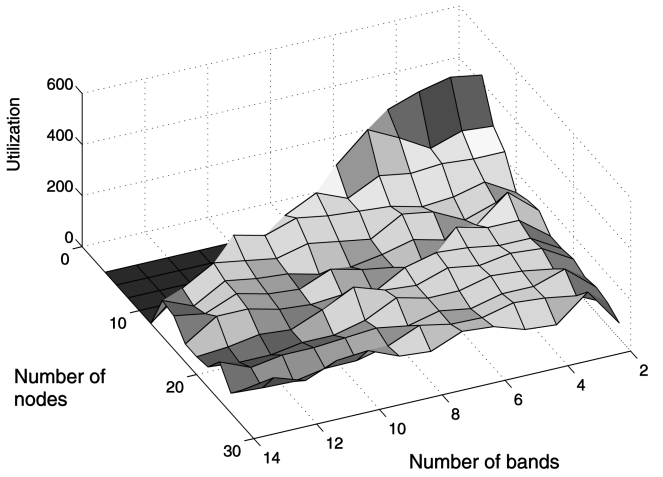

Fig. 11. Use of bands for successful transmissions as a function of nodes.

communications and we perform measurements in the cases wherein the network is very heavily loaded and the individual nodes always have packets to send. It is in this regime that we wish to observe band occupancy for the motivating reason specified above.

Initially, we assume that the band that a node chooses is determined via a simple mapping function from its identifier (ID) (a simple hash function is assumed). Later, the preferential band selection process discussed in Section 3 is adopted. In Fig. 11, we plot the number of times each band is chosen (utilization) as a function of the number of nodes in the network. When the number of nodes exceeds the number of bands, there is a tendency for all bands to be used uniformly in the long run.

The highest utilization values occur with small numbers of nodes, where no collisions in data bands occur. When all bands are used, the maximum utilization is observed for 12 nodes and decreases as the number of nodes increases. Notice that, when the number of nodes is small, not all bands are utilized. Thus, with small scale deployments, one may consider the use of more bands for dissemination of control information to improve performance. We will consider such possibilities in future work.

5 Performance Analysis

To supplement the simulation results described in the previous section, we have developed a simple analytical model to provide additional insight into the dynamic behavior of our protocol. We perform the analysis for a simplistic case wherein the ad hoc nodes are all within the communication range of one another, i.e., we consider a *clique* topology. The analysis becomes much more complex and intractable in a multihop setting. The analysis, although simple, does provide us with some insights with regard to the utilization of bands with the multiband approach. In particular, it provides insight with regard to whether or not bands are all efficiently utilized and helps us identify the primary reasons for inefficiencies in utilizing the bands. We proceed with the analysis and highlight these insights at the end of the section. Finally, we validate the results from the analysis by additional simulation experiments.

Our primary objective is to find an expression for the probability that a node in the REQUEST state succeeds in establishing a session with an available receiver, on a free band, during the next superframe. Given our *clique* assumption, we can model our protocol as an embedded Markov chain [27]. We will determine the possible states and, later, use a fluid flow approach to analyze system behavior; our procedure is described below.

Since a node in the REQUEST state can only initiate the transmission of a REQ control message immediately following an availability frame, we focus our attention on the beginning instance J_{start} of each availability frame. In this case, we can represent the state of the Markov chain by the pair (N, K), where M is the total number of nodes, $N \leq M$ is the number of nodes currently in the REQUEST state, B is the total number of available data bands, and $K \leq B$ is the number of currently occupied data bands.[14] Note that, if K data bands are occupied, then $2K$ nodes must be in the TALK state and, hence, actively engaged in some session as a transmitter or a receiver.

Let us now find the probability that a particular node, say A, successfully establishes a new session during the current superframe, given that the state of the system at the start of the availability frame is (N, K) and node A is in the REQUEST state. More precisely, we see that, in order for node A to be successful, we require:

1. $K < B$, i.e., there are some free bands available that could be allocated to a new session. In this case, node A will randomly select one of the $B - K$ available bands for inclusion in its REQ message, say, band b_A.
2. Node A's intended receiver, R_A, must be a valid target. In particular, if A is not aware of the complete state of the system, then it might direct its REQ message to a target node that will not be able to hear it, i.e., to one of the other $N - 1$ nodes in the REQUEST state that are busy sending their own REQ or to one of the $2K$ nodes currently in the TALK state that are already engaged in another session. Let $Pr[valid]$ be the probability that node A directs its REQ message to a node currently in the IDLE or BACKOFF state and, hence, picks a valid target. Assuming that node A randomly picks any node except itself as its target, we have immediately that
$$Pr[valid|N, K] = \frac{M - N - 2K}{M - 1}. \quad (2)$$
3. None of the other nodes in the REQUEST state have selected the same valid receiver R_A as our target node A. Otherwise, they will all use the same time-hopping code to send their REQs, causing a collision from which R_A will be unable to successfully identify any valid reservation. Let $Pr[unique|valid, N, K]$ be

14. We will assume that $M > B$ to avoid trivializing the MAC protocol by allowing each communicating pair to be assigned a unique band all the time.

the probability that node A successfully delivers its reservation request to node R_A, given that R_A is a valid target. Since R_A is valid, it cannot be one of the $N-1$ other nodes in the REQUEST state and, hence, it would have been included in each of their respective lists of $M-1$ possible targets. Thus,

$$Pr[unique|valid, N, K] = \left(1 - \frac{1}{M-1}\right)^{N-1}. \quad (3)$$

4. No other REQ message that was successfully delivered to a different target can specify the use of band b_A. Otherwise, node A's session will fail because a collision will occur in band b_A. Let $Pr[private|unique, valid, N, K]$ be the probability that no other active node that successfully delivered its reservation to a valid receiver picked band b_A, given that node A has successfully delivered its reservation to node R_A for using band b_A.

In order to compute $Pr[private|unique, valid, N, K]$, we will first calculate the conditional probability

$$Pr[private|unique, valid, V, N, K]$$

that no other successful REQs have specified band b_A, given that node A has successfully delivered its reservation to node R_A for using band b_A and that the total number of valid REQs is V. In that case, there are $B-K-1$ other band choices available (besides A's selection of band b_A), and there are $V-1$ other transmitters making valid band selections. Therefore, we will have:

$$Pr[private|unique, valid, V, N, K] = \left(1 - \frac{1}{B-K}\right)^{V-1}. \quad (4)$$

In order to derive the unconditional probability $Pr[private|unique, valid, N, K]$, it remains to calculate the probability $Pr[V|unique, valid, N, K]$ that there are exactly V valid reservations, including the one by node A, given that node A has successfully delivered its reservation to node R_A for using band b_A. Each of the N active users may pick among the other $M-1$ potential receivers. Thus, the probability $Pr[V|unique, valid, N, K]$ can be expressed as the ratio $\mathcal{N}_V/\mathcal{D}_V$, where the numerator, $\mathcal{N}_V$, represents the number of ways for N users to pick targets, such that exactly V REQs, including A's, are valid, and the denominator, $\mathcal{D}_V$, represents the total number of ways in which N users may select target receivers.

We begin by computing $\mathcal{D}_V$. Node A picks one of the $M-N-2K$ available valid targets, and every other node is required to select a target different from that chosen by A. Thus, the other nodes have $M-2$ choices (to exclude the target of node A and the node itself). Hence,

$$\mathcal{D}_V = (M-N-2K)\cdot(M-2)^{N-1}. \quad (5)$$

We next compute the numerator $\mathcal{N}_V$ as follows. There are $\binom{M-N-2K}{V}$ ways in which we can select the V valid targets. One of the successful reservations must be from node A, and there are $\binom{N-1}{V-1}$ ways to select the other successful nodes. Moreover, there are $V!$ distinct ways to match success nodes to their respective targets, giving a total of

$$S_V = \binom{M-N-2K}{V} \cdot \binom{N-1}{V-1} \cdot V!. \quad (6)$$

In addition, we must find F_V, the number of different ways in which all of the other $N-V$ attempts fail because their REQ packets either collide or are directed to one of the $N+2K$ invalid targets. To simplify the problem, we first condition on t the number of valid receivers that were targeted by the failed attempts, where $0 \le t \le t_{max} = \min\{M-N-2K-V, (N-V)/2\}$ and the second term in the minimizaton is included because each of the t unsuccessful valid targets must have been selected by at least two unsuccessful active nodes. The conditional value of F_V, given t, may be calculated as follows: First, there are $C_1 = \binom{M-N-2K-V}{t}$ different ways to select the t unsuccessful valid target nodes. We need to pick $2t$ active nodes to serve as "spoiler nodes" to force a REQ collision at each of these targets, which leaves the remaining $N-V-2t$ active nodes free to select any receiver that will not increase the number of requests successfully completed. There are $C_2 = \binom{N-V}{2t}$ ways to select the spoiler nodes, which can then be assigned to targets in $C_3 = (2t)!/2^t$ different ways. The remaining $(N-V-2t)$ transmitters are free to choose any of the $2K$ occupied nodes, the $(N-1)$ other transmitters, or the t collision targets as their chosen receiver, which can be done in $C_4 = (N-1+2K+t)^{N-V-2t}$ ways. Thus,

$$F_V = \sum_{t=0}^{t_{max}} (C_1 \cdot C_2 \cdot C_3 \cdot C_4) \quad (7)$$

and, hence, $\mathcal{N}_V = S_V \cdot F_V$.

We are now finally ready to calculate the probability that no other active node that successfully delivered its reservation to a valid receiver picked band b_A, given that node A has successfully delivered its reservation to node R_A for using band b_A. Using the law of total probability [27], we uncondition on the number of valid reservations, V, to obtain

$$\begin{aligned} &Pr[private|unique, valid, N, K] = \\ &= \sum_{V=1}^{N} \left(1 - \frac{1}{B-K}\right)^{V-1} \cdot \frac{\mathcal{N}_V}{\mathcal{D}_V}. \end{aligned} \quad (8)$$

Therefore, the probability of a successful connection establishment instance by node A is computed to be

$$\begin{aligned} Pr[success|N, K] = \\ Pr[valid|N, K] \cdot \\ \cdot Pr[unique|valid, N, K] \cdot \\ \cdot Pr[private|unique, valid, N, K]. \end{aligned} \quad (9)$$

We consider the short-term dynamical behavior of the protocol if we force it to continue executing from some arbitrary state (N, K). Following the classical stability analysis for slotted ALOHA by Lam and Kleinrock (see, e.g., [28, pp. 379-385]), and the capacity analysis for the 0.487 tree conflict resolution algorithm by Gallager [30], we

will now evaluate the *average drift* of the process if it is currently in state (N, K). In other words, we approximate its trajectory by a "fluid flow" model that deterministically moves the process from state (N, K) to state (N', K'), which represents the average of all possible transitions leaving that state. This method is quite simple to apply to our system, since (9) is the only complicated formula we need.

First, we consider how K', the value for the number of occupied data bands at the next embedding point, is related to N and K. Each data band that is currently occupied will remain occupied if the reserving nodes transition to the DECLARE state at the end of the current superframe. For simplicity, we assume that the session lengths are geometrically distributed, with probability p_c of completion, before the next availability frame. Thus, the number of occupied data bands at the next embedding point will decrease due to a binomially distributed number of session completions, with a mean decrease of $N \cdot p_c$ bands. Conversely, the number of occupied bands will increase because some of the N nodes currently in the REQUEST state succeed in establishing a new session. Since the probability of success for each of those nodes is given by (9), and the expectation of a sum is the sum of the expectations even when the terms are not independent, we see that

$$K' = K \cdot (1 - p_c) + N \cdot Pr[success|N, K]. \tag{10}$$

Next, we consider N', the value for the number of nodes in the REQUEST state at the next embedding point; by the Markovian property, this is related to N, K, and K'. Referring to Fig. 4 and the associated protocol descriptions, we see that nodes do not just remain in the REQUEST state over a series of superframes until they succeed in establishing a new session. Instead, those N nodes which fail to establish a session transition to the BACKOFF state, where they remain for a random delay before returning to the REQUEST state at the beginning of a superframe. On the other hand, those $M - N - 2K$ nodes that were in the IDLE state at the start of the current superframe will transition to the REQUEST state if they generate new data before the start of the next superframe. For simplicity, in this analysis, we will approximate both the generation of new packets by a node and expiry of its back-off delay by a Poisson process with rate λ per node per superframe, so that

$$N' = \lambda \cdot (M - 2K'). \tag{11}$$

This analysis can be used to construct a two-dimensional drift field diagram to illustrate the stability of the protocol for representative choices of the above parameter values. One example is shown in Fig. 12, where we have set $M = 40$, $B = 14$, $\lambda = 0.25$, and the average session duration to 3.5 superframes, so that $p_c = 2/7$. Other parameter combinations we tested show qualitatively similar behavior. Each line segment in Fig. 12 represents the average state change from a regular grid of starting points, $\{(N, K)|N = 0, 2, \ldots, 20; K = 0, 2, 4, \ldots, 12\}$. From this set of starting points, we plot the trajectories that the points take as the system evolves using (9) and (11). From this diagram, we can see that the process always moves in the direction of an (almost vertical) "attractor line" that is determined by (11) and, in this example, reduces to the form $K' = 2 \cdot (10 - N')$. In addition, the process moves upward when started from small values of K, moves downward when started from large values of K, and seems to converge to the fixed point where $N' = K' = 20/3 \approx 7$ in steady-state.

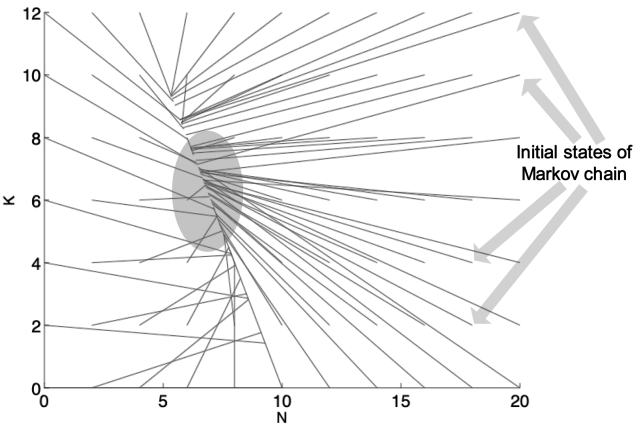

Fig. 12. Fluid diagram for $M = 40$ total nodes, $B = 14$ data bands, and $p_c = 2/7$. The elliptic shaded area corresponds to the convergence region at steady state.

It is interesting to compare the average number of occupied bands at this fixed point (i.e., $\approx$ 7) with an upper bound on the number of data bands that this system could fill with traffic in the best case, where there are no transmission errors and every REQ message is successful. In this case, each node alternates between IDLE and TALK periods. Since the idle period for a particular node A ends as soon as A generates a new packet or another idle node generates a new packet and selects A as its receiver, the average duration of A's idle period will be $\frac{1}{2\lambda} = 2$ superframes. Because of our perfect scheduling assumption, the two nodes will immediately find a free data band and occupy it for an average of $1/p_c = 7/2$ superframes before they return to the IDLE state. Thus, in the best case, each pair of nodes will occupy a single data band for $\frac{7}{2}$ out of every $\frac{11}{2}$ superframes, which represents an average occupancy of $\frac{7}{11}$ data bands per node pair. Since our system can support at most $M/2 = 20$ disjoint node pairs, *the average number of occupied data bands is at most* $\frac{140}{11} \approx 13$*—which is almost twice as large as the fixed point for our protocol!*

The explanation for this large discrepancy between the fixed point and upper bound provides a key insight into the importance of selecting a valid target for the receiver in MAC protocols for this type of ad hoc network. In our case, we have assumed that receivers are randomly selected, which leads to (2). Thus, if many of the other nodes are already in the TALK state, then it will be very difficult for nodes entering the REQUEST state to guess which of its possible receiver choices is in the IDLE or BACK-OFF state. For example, our upper bound scenario states that each node is in the TALK state for an average of $\frac{7}{2}$ out of every $\frac{11}{2}$ superframes, and one of the two nodes in a session must have spent the previous superframe in the REQUEST state to set it up. Thus, a given node in the upper bound system is only a valid target during $\frac{3}{11}$ of the superframes. ***This observation shows why it would be impossible for a MAC protocol that uses random receiver selection to achieve such***

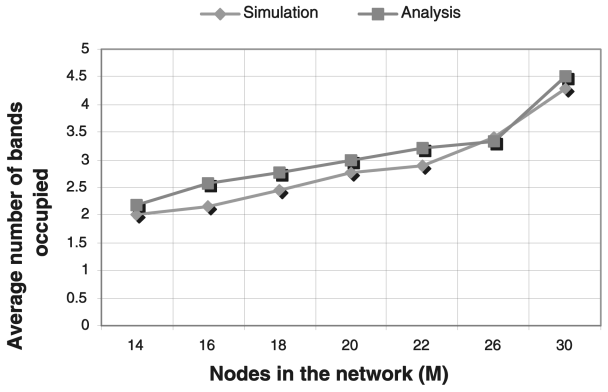

Fig. 13. Comparison between simulations and analysis with regards to the average number of data bands utilized.

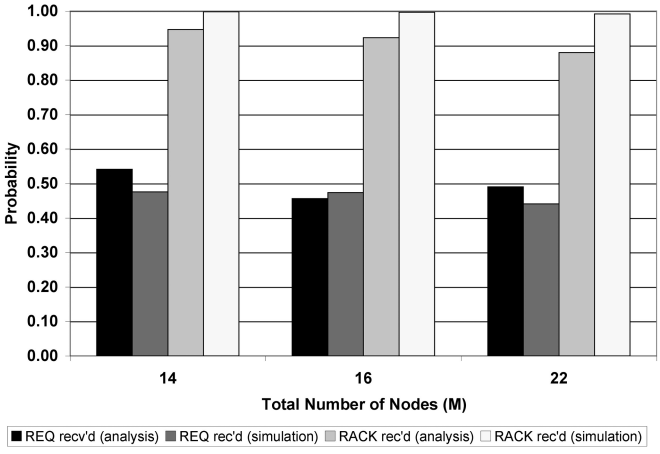

Fig. 14. Comparison of success probabilities for REQ and RACK messages between simulation and analysis.

a high occupancy of the data bands. Conversely, since halving the occupancy of the data band would double the availability of each node to serve as a valid target, we see that random receiver selection is the key factor that defines the fixed point for our protocol.

Note that this observation of the chosen receiver being busy applies not just to the multiband approach, but also to the single-band approach (and other approaches as in [6]). If the node degree increases, a more efficient utilization of the bands may be achieved. While this is beyond the scope of this present work, one modification that may further improve the efficiency of our MAC protocol is the use of an additional band (or a time-period) for allowing busy receivers to announce that they are not accommodating requests. This would reduce the back-offs that are experienced by the senders (they could choose receivers that are free) and thereby further increase the achieved throughput.

To validate our analysis, we perform a set of simulations using the simulator that was described in the previous section. In these simulations, we restrict the nodes to a $3\ \mathrm{m} \cdot 3\ \mathrm{m}$ square region, so as to create a clique topology; this is to reflect the scenarios considered in the analysis. As per the assumptions made in our analysis, packets are generated at each node in accordance to a Poisson process. In Fig. 13, we compare the average number of occupied bands as a function of the total number of nodes, M, using simulation and analysis when we set $\lambda = 0.8$ and $p_c = 0.5$. Each simulation value represents the sample mean from an experiment lasting for 15 million chip intervals. The analytical values represent the steady-state convergence point for the corresponding fluid diagram. We see that the simulation results are very close to those computed analytically for all the considered scenarios. We have repeated the experiments for other values of λ and p; the behaviors observed were similar to those reported here.

In Fig. 14, we present results that quantify the accuracy of the analytical methodology in more detail for population sizes of $M = 14, 16, 22$, for which the corresponding analytical results predicted $N = 2, 3, 3.8$ nodes in the REQUEST state at the convergence point, respectively. First, we compare the probability that a REQ packet is successfully received by its target node. For the simulation, we simply measure the ratio between the total number of RACK messages sent divided by the total number of REQ messages sent during an experiment. For the analysis, we compute $Pr[valid, unique|N, K]$ by multiplying (2) and (3). Second, we compare the probability that a RACK packet is successfully received because there is no data band collision. For the simulation, we measure the ratio between the number of successfully transmitted RACK messages divided by the total number of RACK messages during an experiment. For the analysis, we compute $Pr[private|valid, unique, N, K]$ via (8). The match between the results from the simulations and the analytical computations validates our analysis.

6 Related Work on UWB Networks

There is no prior work on the design of a MAC protocol for multiband UWB-based ad hoc wireless networks. However, there have been some interesting studies on single-band implementations.

Previous Ad Hoc UWB Schemes. Le Boudec et al. [6], [7], [8] propose a scheme that uses dynamic channel coding. The transmitter dynamically varies the code rate upon receiving feedback from the receiver with regard to channel conditions. The scheme uses two types of THS: a receiver-based THS and an invitation-based THS. This invitation-based THS is unique for each communicating pair. After the successful transmission of a request using the receiver-based THS, the pair switches to a *unique* invitation-based THS and uses this THS for the duration of the session. In [12], the authors describe theoretical and practical approaches toward the development of a THS-based MAC protocol for radio resource sharing in UWB ad hoc networks. The effects of multipath and, in particular, delay spread are not addressed in these papers. The authors appear to implicitly assume the presence of a perfect equalizer.

Ding et al. study issues related to channel acquisition [18]. They conclude that existing MAC solutions are unsuitable for UWB networks. Specifically, the authors examine TDMA and CSMA/CA approaches, which have been successful in other environments. In [29], all nodes share a THS and the receiver broadcasts an invitation, as per this sequence. Potential transmitters compete during a contention period to lock on to the receiver. In [32], the

authors propose a full-duplex access scheme for impulse-based UWB networks. The scheme takes advantage of the low duty cycle to maintain physical links among two nodes for the lifetime of their logical link, thereby removing the requirement that the sender and receiver resynchronize for every packet to be exchanged. A theoretical treatment on optimal routing, scheduling, and power control appears in [9]. The authors show that 1) the design of the optimal MAC is independent of the choice of the routing protocol and 2) a minimum energy route is preferable to establishing long hops or invoking direct transmissions.

A MAC protocol for single-transceiver UWB ad hoc networks based on the use of busy signals was proposed in [33]. The key objective of the MAC protocol was to facilitate the detection of collision of UWB signals by using UWB pulses. A single band is used and the protocol is not designed with the intent of maximally exploiting the available spectrum in the presence of delay spread.

WPAN Configurations. Most other studies consider master-slave configurations [4], [31], [10]. The IEEE 805.15.3a task group proposal [4] for media access control is based on the notion of piconets. Each piconet includes a master-coordinator, which assigns resources to slaves. The task group has evaluated numerous proposals for the UWB physical layer. The Multiband OFDM Alliance (MBOA), comprised of over 170 companies, supports a UWB specification that is based on an OFDM approach [5]. Note here that the OFDM approach attempts to reduce the equalizer complexity by dividing the available spectrum into multiple bands. However, the use of OFDM requires 1) frequent complex inverse fast Fourier transform computations [14] and 2) simultaneous receiver synchronization with multiple carriers. Thus, there are trade-offs between the use of OFDM-based UWB and impuse-based UWB. A MAC layer protocol for use with OFDM for ad hoc networks is yet to emerge. Yomo et al. [10] study the interference between distinct WPANs (Wireless Personal Area Networks) that operate in a master-slave configuration. Their studies focus on the interference between co-located WPANs. Various other schemes have been proposed; however, they either address the physical layer only or they assume master-slave configurations.

Reservation-Based MAC Protocols. The use of reservations for arbitrating access to a plurality of orthogonal bands has been considered in wireless and satellite networks [19], [35]. However, the presence of a centralized arbiter (a satellite or base-station) makes allocation much easier as compared to allocation in ad hoc networks. Recently, the use of multiple bands in ad hoc networks that use the IEEE 802.11 MAC protocol has been considered in [36]. However, carrier sensing is possible with IEEE 802.11 and the issues related to MAC access are different from those that arise due to the use of UWB.

7 Conclusions

In this paper, we propose a novel multiband MAC protocol for use with impulse-based UWB ad hoc networks. The design of our protocol is motivated by the following factors: 1) In the absence of a complex equalizer, due to the effects of the multipath delay spread, the entire UWB spectrum cannot be efficiently utilized by a single band approach. 2) Arbitration methods based on the use of time-hopped sequences suffer from inefficiencies due to collisions or large delays. 3) The use of a multiband approach provides an inherent flexibility in operation to coexist with other wireless networks. The approach we present is conjoint with the UWB physical layer and takes into account the regulations imposed by the FCC. We perform extensive simulations to demonstrate that our protocol achieves extremely high throughput and much lower latencies as compared to a single-band approach, wherein no equalizer is available. We provide an analytical model that facilitates the understanding of our approach in terms of its stability and the uniform usage of the different bands.

Acknowledgments

The authors would like to thank the editor, Professor Suresh Singh, and the anonymous reviewers for their constructive suggestions.

References

[1] FCC, "Revision of Part 15 of the Commission's Rules Regarding Ultra-Wideband Transmission Systems," First Report and Order, ET Docket 98-153, 04-2002.
[2] J.G. Proakis, *Digital Communications,* fourth ed. McGraw-Hill, 2001.
[3] R. Gitlin, J.F. Hayes, and S. Weinstein, *Data Communication Principles.* Plenum Press, 1992.
[4] *IEEE 802.15.3 MAC Standard*, 2003.
[5] *Multiband OFDM Physical Layer Proposal for IEEE 802.15.3a,* Sept. 2004, http://www.multibandofdm.org.
[6] J.Y. Le Boudec, R. Merz, B. Radunovic, and J. Widmer, "A MAC Protocol for UWB Very Low Power Mobile Ad-Hoc Networks Based on Dynamic Channel Coding with Interference Mitigation," Technical Report IC/2004/02, Ecole Polytechnique Fédérale de Lausanne, Jan. 2004
[7] J.Y. Le Boudec, R. Merz, B. Radunovic, and J. Widmer, "DCC-MAC: A Decentralized MAC Protocol for 802.15.4a-Like UWB Mobile Ad-Hoc Networks Based on Dynamic Channel Coding," *Broadnets*, 2004.
[8] R. Merz, J.Y. Le Boudec, J. Widmer, and B. Radunovic, "A Rate-Adaptive MAC Protocol for Low-Power Ultra-Wide Band Ad Hoc Networks," *Ad Hoc Now,* 2004.
[9] B. Radunovic and J.Y. LeBoudec, "Optimal Power Control, Scheduling and Routing in UWB Networks," *IEEE J. Selected Areas in Comm.*, vol. 22, no. 7, Sept. 2004.
[10] H. Yomo, P. Popovski, C. Wijting, I.Z. Kovacs, N. Deblauwe, A.F. Baena, and R. Prasad, "Medium Access Techniques in Ultra-Wideband Ad Hoc Networks," *Electronic Trans. Artificial Intelligence*, Sept. 2003.
[11] *ANSI/IEEE 802.11 Standard,* 1999.
[12] F. Cuomo, C. Martello, A. Baiocchi, and C. Fabrizio, "Radio Resource Sharing for Ad Hoc Networking with UWB," *IEEE J. Selected Areas in Comm.*, Dec. 2002.
[13] S. Roy, J.R. Foerster, V.S. Somayazulu, and D.G. Leeper, "Ultrawideband Radio Design: The Promise of High-Speed, Short-Range Wireless Connectivity," *Proc. IEEE*, Feb. 2004.
[14] M. Ghavami, L.B. Michael, and R. Kohno, *Ultra Wideband Signals and Systems in Communication Engineering.* John Wiley and Sons, 2004.
[15] J. Foerster, E. Green, V.S. Somayazulu, and D. Leeper, "Ultra-Wideband Technology for Short—or Medium—Range Wireless Communications," *Intel Technology J.*, 2001
[16] M.Z. Win and R.A. Scholtz, "Impulse Radio: How It Works," *IEEE Comm. Letters,* vol. 2, no. 1, Jan. 1998.
[17] M.Z. Win and R.A. Scholtz, "Ultra-Wide Bandwidth Time-Hopping Spread-Spectrum Impulse Radio for Wireless Multiple Access Communications," *IEEE Trans. Comm.*, Apr. 2000.

[18] J. Ding, L. Zhao, S. Medidi, and K.M. Sivalingam, "MAC Protocols for Ultra-Wide-Band Wireless Networks: Impact of Channel Acquisition Time," *Proc. SPIE-ITCOM,* 2002.
[19] K. Pahlavan and P. Krishnamurthy, *Principles of Wireless Networks: A Unified Approach.* Prentice Hall, 2002.
[20] Discrete Time Communications: "IEEE 802.15.3a 480Mbps Wireless Personal Area Networks: Achieving a Low Complexity Multi-Band Implementation," white paper, 2003.
[21] J.J. Garcia-Luna-Aceves and J. Raju, "Distributed Assignment of Codes for Multihop Packet—Radio Networks," *Proc. IEEE Military Comm. Conf. (MILCOM '97),* 1997.
[22] E.S. Sousa and J.A. Silvester, "Spreading Code Protocols for Distributed Spread-Spectrum Packet Radio Networks," *IEEE Trans. Comm.,* vol. 36, pp. 272-281, 1988.
[23] Texas Instruments, "Multi-Band OFDM Physical Layer Proposal for IEEE 802.15 Task Group 3a," Sept. 2003
[24] M.E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Trans. Networking,* 1997.
[25] K. Romer, "Time Synchronization in Ad Hoc Networks," *Proc. ACM MobiHoc,* 2001.
[26] L. Meier, P. Blum, and L. Thiele, "Internal Synchronization of Drift-Consistent Clocks in Ad-Hoc Sensor Networks," *Proc. ACM MobiHoc,* 2004.
[27] L. Kleinrock, *Queueing Systems Volume 1: Theory.* Wiley, 1975.
[28] L. Kleinrock, *Queueing Systems Volume 2: Computer Applications.* Wiley, 1976.
[29] A. Hicham, Y. Souilmi, and C. Bonnet, "Self-Balanced Receiver-Oriented MAC for Ultra-Wide Band Mobile Ad Hoc Networks," *Ultra Wideband Systems,* June 2003.
[30] R. Gallager, "Conflict Resolution in Random-Access Broadcast Networks," *Proc. AFOSR Workshop Comm. Theory and Applications,* pp. 74-76, Sept. 1978.
[31] F. Cuomo, A. Baiocchi, F. Capriotti, and C. Martello, "Radio Resource Optimisation in an UWB Wireless Access," *Proc. Mobile Comm. Summit,* 2002.
[32] S. Kolenchery, J. Townsend, and J. Freebersyser, "A Novel Impulse Radio Network for Tactical Military Wireless Communications," *Proc. IEEE Military Comm. Conf. (MILCOM '98),* 1998.
[33] N.J. Agusut and D.S. Ha, "An Efficient UWB Radio Architecture for Busy Signal MAC Protocols," *Proc. IEEE Comm. Soc. Conf. Sensor, Mesh and Ad Hoc Comm. and Networks (SECON '04),* 2004.
[34] I. Broustis, S. Krishnamurthy, M. Faloutsos, M. Molle, and J. Foerster, "A Multiband MAC Protocol for Impulse-Based UWB Ad Hoc Networks," *Proc. IEEE Comm. Soc. Conf. Sensor, Mesh and Ad Hoc Comm. and Networks (SECON '05),* 2005.
[35] S. Krishnamurthy, C. Liu, and V. Gupta, "Medium Access Control Protocols for Satellite Communications," *Internetworking and Computing over Satellite Networks,* Y. Zhang, ed. Kluwer, 2003.
[36] J. So and N. Vaidya, "Multi-Channel MAC for Ad Hoc Networks: Handling Multi-Channel Hidden Terminals Using a Single Transceiver," *Proc. ACM MobiHoc,* 2004.
[37] http://www.onr.navy.mil/media/article.asp?ID=62&css=printer, 2007.

Ioannis Broustis received the diploma in electronics and computer engineering from the Technical University of Crete in 2003 and the MSc degree from the University of California, Riverside, in 2005. Currently, he is a PhD candidate at the University of California in Riverside. His research interests include wireless networking, especially mobile ad hoc networks, concentrating on the medium access control (MAC) layer and the network layer. Mostly, he deals with wireless testbeds, network modeling, simulation, and implementation, especially for wireless cross-layer schemes. He is a student member of the IEEE.

Srikanth V. Krishnamurthy received the PhD degree in electrical and computer engineering from the University of California at San Diego in 1997. He received the BE (Hons) degree in electrical and electronics engineering and the MSc (Hons) degree in physics with distinction from the Birla Institute of Technology and Science, Pilani, India, in 1992 and the master of applied science degree in electrical and computer engineering from Concordia University, Montreal, Canada, in 1994. From 1994 to 1995, he was a graduate research assistant at the Center for Telecommunications Research, Columbia University, New York. From 1998 to 2000, he was a research staff scientist at the Information Sciences Laboratory, HRL Laboratories, LLC, Malibu, California. Currently, he is an assistant professor of computer science at the University of California, Riverside. His research interests span CDMA and TDMA technologies, medium access control protocols for satellite and wireless networks, routing and multicasting in wireless networks, power control, and the use of smart antennas and security in wireless networks. Dr. Krishnamurthy has been a PI or a project lead on projects from various DARPA programs, including the Fault Tolerant Networks program, the Next Generation Internet program and the Small Unit Operations program. He is the recipient of the US National Science Foundation CAREER Award from ANI in 2003. He has also coedited the book *Ad Hoc Networks: Technologies and Protocols* (Springer Verlag, 2004). He has served on the program committees of INFOCOM, MobiHoc, and ICC and is an associate editor for ACM MC2R. He is a member of the IEEE.

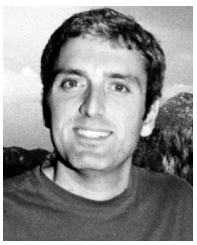

Michalis Faloutsos received the bachelor's degree from the National Technical University of Athens and the MSc and PhD degrees at the University of Toronto. He is currently a faculty member in the Computer Science Department at the University of California, Riverside. His interests include Internet protocols and measurements, multicasting, and cellular and ad hoc networks. With his two brothers, he coauthored a paper on powerlaws of the Internet topology (SIGCOMM '99), which is in the top 15 most cited papers of 1999. His work has been supported by several US National Science Foundation (NSF) and DARPA grants, including the prestigious NSF CAREER award. He is actively involved in the community as a reviewer and a TPC member in many conferences and journals. He is a member of the IEEE.

Mart Molle received the MS and PhD degrees in computer science from the University of California, Los Angeles, in 1978 and 1981, respectively. Between 1981 and 1994, he was a faculty member in the Department of Computer Science at the University of Toronto. In 1994, he joined the University of California, Riverside, where he is currently a professor of computer science and engineering, and served as the department chair from 1999 to 2002. Dr. Molle's research interests include the performance evaluation of protocols for computer networks and of distributed systems. He is particularly interested in efficiently-solvable analytical modeling techniques, fundamental performance limits, applications of queuing theory and scheduling theory to distributed algorithms, and model validation through measurement and simulation. He has served as an editor for the *IEEE/ACM Transactions on Networking* and as a task force chair for the IEEE 802.3 Ethernet Standards Working Group. He is a member of the IEEE.

Jeffrey R. Foerster received the BS, MS, and PhD degrees from the University of California, San Diego, where his thesis focused on adaptive interference suppression techniques for CDMA systems. He joined Intel in August 2000 as a wireless researcher with Intel Architecture Labs in Hillsboro, Oregon. He is currently focusing on future short and medium-range wireless technologies, including Ultra-wideband (UWB) technology and related regulations, system design, and performance analysis. Prior to joining Intel, he worked on Broadband Wireless Access (BWA) systems and standards (IEEE 802.tpd16). He is a member of the IEEE.