

U-MAC: a proactive and adaptive UWB medium access control protocol

Raja Jurdak^{1*}, Pierre Baldi² and Cristina Videira Lopes¹

¹*Donald Bren School of Information and Computer Sciences, Center for Pervasive Communications and Computing, Irvine, CA, USA*

²*California Institute for Telecommunications and Information Technology Cal-(IT)², Irvine, CA, USA*

Summary

Ultra wide band (UWB) technology has received increasing recognition in recent years for its potential applications beyond radar technology to communication networks. UWB is a spread spectrum technology that requires careful coordination among communicating nodes to jointly control link power and transmission rates. Here, we present ultra wide band MAC (U-MAC), an adaptive medium access control (MAC) protocol for UWB in which nodes periodically declare their current state, so that neighbors can proactively assign power and rate values for new links locally in order to optimize global network performance. Simulations comparing U-MAC to the reactive approach confirm that U-MAC lowers link setup latency and control overhead, doubles the throughput and adapts better to high network loads. Simulations also reveal that the basic form of U-MAC favors nodes that are closer to the receiver. As a result, we also introduce novel mechanisms that control the radius around a receiver within which nodes can have fair access to it. We show through simulations the effect of the mechanisms on the tradeoff between network throughput and fair access. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: ultra wide band; medium access control; proactive protocols; quality of service; ad hoc networks; sensor networks

1. Introduction

Ultra wide band (UWB) technology is a spread spectrum technology that is based on the modulation of short nanosecond low power pulses. This technology has been used for radar applications for over half a century. In recent years, UWB has received increasing recognition for its applicability to short to medium range communication networks because of desirable features such as high data rates, low power consumption, precise ranging capabilities, resistance to multipath fading, and penetration of dense objects. UWB is

currently a candidate technology for short range high transfer rate applications such as the simultaneous transfer of multiple video streams in a wireless personal area network (WPAN) [1]. It is also being considered for medium range sensor networks with lower transfer rates [2].

A central problem in UWB networks is the joint optimization of transmission power and transmission rates for active links. The joint rate and power assignment problem in UWB involves complex tradeoffs between fair rate assignment, network efficiency, and quality of service (QoS). A high power link may

*Correspondence to: Raja Jurdak, 444 Information and Computer Science, University of California, Irvine, CA 92697, USA.

†E-mail: rjurdak@ics.uci.edu

achieve high transmission rates, but it also causes high interference which limits the rate available to neighboring links. On the other hand, a low power link promotes fair access to the wireless medium but yields lower transmission rates. Thus, nodes must collaboratively determine the optimal rate and power values for new links in the network.

Another important consideration for rate and power assignments in UWB networks is fair access among nodes at different distances from a common receiver, especially in ad hoc networks with dynamic topologies. An inherent property of wireless communications is that transmission rates drop with increasing distance. This effect is even more pronounced in UWB communications, where the strong correlation of transmission rates with multi-user interference levels further increases the impact of relative distances on transmission rates. This impact may be unsuitable for applications that require all the network nodes to have fair access to the medium regardless of distance. The fair access requirement in such applications imposes an additional constraint on the choice of transmission power.

So far, UWB research has been primarily confined to the investigation of the behavior of the physical layer [3,4]. Research at higher layers of the network stack has been somewhat limited. Previous research proposals for higher layers considered an underlying radio frequency (RF) physical layer, so most of these proposals are not suitable for UWB networks. For instance, the existing wireless medium access control (MAC) protocols [5] for RF networks do not meet the need of UWB networks for joint rate and power assignment, hence the need for new UWB MAC protocols. There have been recent attempts to develop mechanisms for UWB networks at the MAC layer. The work in Reference [6] proposes a simple reactive multiple access protocol that defines the handshaking procedure to establish a new link. The work in References [7,8] discusses a protocol that uses periodic state updates to allow nodes to jointly assign rate and power assignment values locally at each node.

Here, we propose a new MAC protocol ultra wide band MAC (U-MAC) that jointly assigns rate and power values in UWB networks, and reduces the control messaging and latency required for link establishment. To ensure collaboration among nodes, U-MAC requires nodes to periodically announce their state information in hello messages, so that any node can locally select appropriate rate and power values for a link request without polling

neighbors. In U-MAC, the hello message period adapts to the stability of network state, to avoid sending frequent updates unnecessarily. U-MAC also provides a mechanism to adjust the radius around a receiver within which all nodes get fair access to the receiver. Furthermore, the protocol's framework supports the future integration of multi-hop links [9], which limits the impact of distance on fairness and the internode interference.

Within the emerging UWB MAC framework, the main novel contributions of this work are:

- The introduction of adaptive periods for hello messages in an UWB network so that control message overhead is minimized.
- A comparative assessment with the reactive approach regarding control overhead, link setup latency, network throughput, and adaptability.
- The development of mechanisms that promote fair access among nodes in an UWB network.

The remainder of the paper is organized as follows. Section 2 reviews the fundamental concepts in UWB radio, and provides the framework upon which we design U-MAC. Section 3 introduces U-MAC and explains the mechanisms and features that characterize this protocol. Section 4 presents the simulation results for U-MAC. Section 5 discusses U-MAC in light of existing literature and future research directions.

2. UWB Network Principles

2.1. UWB Principles

Recently, UWB radio has received increasing recognition for its applicability to multi-user wireless communication networks. UWB radio relies on periodic sequences of short sub-nanosecond pulses (referred to as monocycles) for data transmission. The short duration of UWB pulses yields a low power spectrally wide signal. In a single sender/receiver environment, a common modulation technique used with UWB radio is pulse position modulation (PPM), which encodes symbols by shifting the monocycles according to the following expression:

$$s(t) = \sum_{i=-\infty}^{\infty} \sum_{j=0}^{N_s-1} g(t - jT_i - b_i\tau) \quad (1)$$

where $s(t)$ is the transmitted signal, $g(t)$ is the pulse, and T_f is the frame time. N_s is the number of pulses that encode each symbol, and the sequence b_i encodes the information bits. Win and Scholtz [4] have proposed a multi-user access scheme for UWB using time-hopping (TH) codes. TH codes accommodate multiple users by further shifting the pulse $g(t)$ according to one of many chipping codes. Consequently, UWB radio has the potential for supporting multiple users within the same frequency and spatial channel.

In a centralized UWB network, the base station could send periodic beacons to allow nodes to stay synchronized. Global synchronization among nodes in an ad hoc UWB network requires excessive signaling overhead, which is a waste of valuable battery resources in mobile nodes. Consequently, it is more realistic to assume that only each sender and receiver that shares a link is synchronized. This assumption may lead to collisions of some monocycles among different links.

The impact of monocycle collisions due to the lack of global synchronization is reduced by sending multiple pulses for each symbol at the source as a form of forward error correction, so that collisions contribute only to mutual interference. Provided that proper quality margins are set, collisions only reduce signal quality and do not affect correct reception of data at the receiver.[‡]

The binary bit rate of an M -ary PPM UWB signal is given by the following expression [10]:

$$R_b = \frac{1}{T_f N_s} \log_2 M \quad (2)$$

Both T_f and M are difficult to modify in an UWB system. Changing M for different transmissions is undesirable for communication systems since it leads to processing overhead. Similarly, modifying T_f for each transmission increases the complexity of the hardware design of the system. Thus, the simplest way to adjust R_b is to vary the value of N_s . The only requirement for allowing different N_s values is that the receiver of each link must integrate the correct number of pulses for each symbol received on that link. A protocol that is adaptive to network behavior should vary N_s based on the interference levels in the network. More specifically, high interference levels

increase the probability of pulse collisions, which require more pulses per symbol.

We adopt the framework of Cuomo *et al.* [6] for an UWB radio resource sharing model that assumes continuous values for R_b . The framework considers that a new link request arrives when there are N pairs of communicating UWB terminals, with each pair consisting of one transmitter and one receiver. Each pair of sender and receiver are synchronized to the TH code of their common link, and both background and UWB noise impact the SNR of UWB links. Consequently, the SNR at the receiver of the i th link is:

$$\text{SNR}_i = \frac{P_i}{R_i \text{PL}_{ii}(\eta_i + T_f \sigma^2 \sum_{k=1, k \neq i}^N P_k g_{ki})} \quad (3)$$

where P_i is the power of the i th transmitter, R_i is the binary bit rate of the i th link, η_i is the background noise energy plus interference from non-UWB sources at receiver i , PL_{ij} is the path loss from the i th transmitter to the j th receiver, g_{ki} is the path gain from the k th transmitter to the i th receiver, and σ^2 is a parameter depending on the shape of the monocycle. Common values for the above parameters are [4]: $T_f = 100$ ns; $\sigma^2 = 1.9966 \times 10^{-3}$; and $\eta = 2.568 \times 10^{-21} \text{V}^2\text{s}$; with a pulse duration of 0.75 ns.

2.2. UWB Traffic Classes

We consider two traffic classes for UWB networks, in accordance with the specifications of the European Whyless Project [6,11,12], to address the requirements of different application types: (1) reserved bandwidth (RB) and (2) dynamic bandwidth (DB).

The RB traffic class is geared towards continuous, real-time, or multimedia traffic, since it requires quality guarantees prior to establishing a link. The continuous nature of traffic that exploits the RB traffic class requires that the link rate remains constant throughout the lifetime of the link.

DB traffic does not offer any rate guarantees and is thus suitable for best-effort data, such as internet traffic. As the name implies, the rate of a link is dynamic and elastic, and depends on the number of other active DB links and on interference levels in the network. For instance, if the traffic load in a network is low, then individual DB links may use higher rates.

In short, the goal of RB traffic is to offer a certain QoS for the sender under varying network conditions.

[‡]This concept is similar to increasing the processing gain in CDMA systems.

The goal of DB traffic is to provide adaptive and efficient overall network behavior for asynchronous data, and to ensure a constant interference level of bursty traffic by modifying rates of all DB channels dynamically.

3. U-MAC Protocol

3.1. Problem Definition

U-MAC addresses the joint rate and power assignment problem for UWB links for both RB and DB traffic. In general, each node in the network is the receiver for a certain number of communication links. Based on the quality requirements of its currently active links, the node can tolerate a finite amount of additional interference, referred to as maximum sustainable interference (MSI) [13]. The MSI at each node must be efficiently and fairly divided between RB and DB traffic and among links of each traffic class.

Initially, each node's resources are split evenly between RB and DB traffic. As a node starts receiving link requests, the MSI portion allocated to each traffic type can adapt to the relative number of links in each traffic class. In general, at any point in time, each node allocates a portion λ of its MSI to DB traffic, and the remaining portion $(1 - \lambda)$ to RB traffic, where λ is less than one. Subsection 3.6 provides a more detailed discussion of an MSI allocation technique that avoids starvation.

From the transmitting node's perspective, the challenge is selecting the rates and power levels for new links that adhere to the MSI states of its neighbors. The generic relation between link quality, transmission rate, and transmission power is:

$$\text{Quality} \propto \left(\frac{\text{Power}}{\text{Rate} \times \text{Noise}} \right) \quad (4)$$

During the lifetime of the link, new communication links may cause the noise to increase, which subsequently causes quality degradation of the link. Avoiding quality degradation can be achieved in several ways:

1. Increasing transmission power.
2. Decreasing transmission rate.
3. Providing a quality margin above the minimum quality requirement initially so that when new links are set up, the link can tolerate additional interference.

Although increasing the transmission power maintains link quality and transmission rate, it degrades the quality of neighboring links which may require additional power or rate adjustments in the network. Furthermore, the FCC has imposed tight limits on UWB emissions [14], so increasing transmission power to maintain quality is impractical. Alternatively, the link transmission rates can be lowered to maintain quality in response to increasing interference. This option does not require reconfiguration of neighboring links, but it leads to inefficient use of the medium and may cause quality violations if the link carries RB traffic. Thus, U-MAC allows reducing link transmission rates only for DB traffic. Finally, providing quality margins avoids both rate and power adjustments of any active links, which suits RB links. The drawback of quality margins is that they also lead to less than optimal medium utilization.

U-MAC adopts signal-to-noise ratio (SNR) as the main link signal quality metric. The SNR of a new link must have some margin above the minimum acceptable SNR for the link. In U-MAC, the parameter μ determines the size of SNR quality margin of links (see Equations (7) and (9) below). The value of μ could be static or adaptive to the spatial distribution of nodes, the traffic load, and the lifetime of the link. The rest of the discussion assumes that source nodes set the SNR margin μ statically for simplicity.

3.2. Protocol Overview

Figure 1 illustrates our vision of the complete U-MAC protocol. Here, we present a protocol that covers both RB and DB traffic for a single hop distributed topology. The discussion for a distributed topology could easily apply to the centralized or hybrid topology, since the distributed case is inherently more complex in nature. Furthermore, a single hop topology could be extended to a multi-hop topology through the use of a global cost function to enable multi-hop links [9].

The main design goals of U-MAC are to jointly optimize rate and power values in the network to achieve fairness, maximize throughput, and minimize latency and control overhead. To achieve these goals, U-MAC adopts a proactive approach in reporting state information.

Rate and power assignments in U-MAC occur at the source prior to sending any control messages. To enable local assignments at the sender, all nodes periodically update their neighbors with their state information through hello messages. Because frequent hello messages may increase interference in

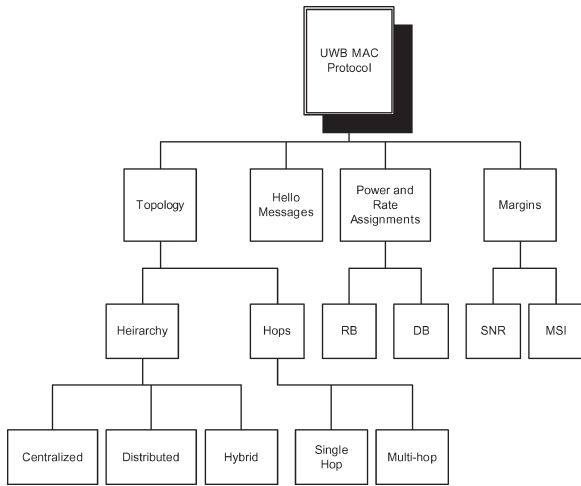


Fig. 1. Protocol overview.

the network, hello message periods always adapt to the stability of network state (see Subsection 3.4). Thus, a highly stable node sends hello messages rarely, while a highly dynamic node frequently updates its neighbors about its state. Every node collects and stores each of its neighbors' most recently advertised state information. Significant state changes at a node also trigger hello messages. Triggered hello messages ensure that each node has a sufficiently up-to-date view of the state of its neighbors.

Figure 2 illustrates the control message exchanges in U-MAC. To set up a new link, a sender *S* first sends a link request in a request to send (RTS) message indicating the rate and power values to the intended

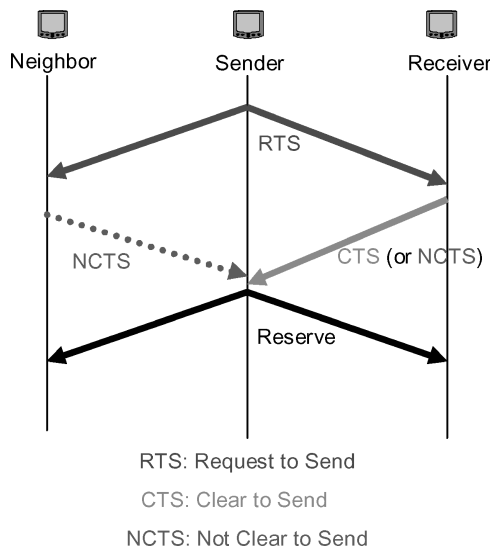


Fig. 2. Control message exchanges.

receiver *R*. Upon receiving an RTS, node *R* and all other neighbors of *S* check whether the requested link is admissible. If so, then *R* notifies *S* with a clear to send (CTS) control message, while other neighbors of *S* refrain from sending any replies if the link parameters are satisfactory. However, if the receiver node *R* or any other neighbor of *S* does not agree with the parameters of the new link, then that neighbor notifies *S* with a not clear to send (NCTS) message that it should reduce either the transmit power or rate or both. After *S* collects all the replies, it declares the duration and parameters of the new link (which may have changed according to neighbor replies) in a reserve message, and immediately sets up the link. In the rest of this section, each subsection discusses the design choices for one main branch of Figure 1.

3.3. Topology

U-MAC supports a hybrid multi-hop topology, which provides a node with flexibility in switching between centralized mode when an access point is available, or ad hoc mode when an access point is not reachable. To determine its current mode of operation, a node monitors a dedicated hello message channel. Whenever it detects any access point hello messages, it switches to centralized mode. The node must hear access point hello messages periodically; otherwise, it switches to distributed mode. The remainder of the discussion in this section focuses on the case of distributed mode in a single hop topology. U-MAC can easily be extended to support multi-hop links through a global cost function that quantifies link costs in order to determine optimal routes [9].

3.4. Hello Messages

U-MAC requires nodes to advertise their local states periodically through hello messages [15], which provides for quick and appropriate rate and power assignments in the network. Note that although hello messages are periodic for one node, they are asynchronous among different nodes in the network, which helps avoid hello message interference from different nodes. Many factors contribute to avoiding several simultaneous transmissions of hello messages:

- Because the hello message period in each node depends on the node's stability, hello periods are not the same across nodes.
- Clock skews contribute favorably to collision avoidance in hello messages of nodes that have

the same hello message period and that enter the network at the same time.

- The transmission time of hello messages is much shorter than typical hello message periods which further reduces the chances of collisions.

Hello messages implicitly provide nodes with ranging information about neighbors, and they explicitly advertise important local parameters to neighboring nodes. Storing recent neighbor state information locally enables a node to make decisions on rate and power assignments for new links, and to make routing decisions for multi-hop links.

First, nodes use hello messages to determine distances of neighboring nodes. Each node sends its hello message at a fixed power level know *a priori* to all nodes. Whenever a node receives a hello message from a neighbor, it can estimate the current distance of that neighbor by examining the received signal strength of the hello message and by applying the appropriate propagation model. The current distance from a neighbor enables a node to compute the path loss to that neighbor locally.

3.4.1. Format

In addition to providing ranging information, hello messages advertise local state information to neighbors. Figure 3 shows the format of hello messages in distributed mode.

Mean time between failures (MTBF) is a measure of a node’s communication reliability [9], which is a general attribute of a node. The next two fields in Figure 3 pertain to DB traffic. The PRCNT DB field in a hello message holds the parameter λ , which we introduced in Subsection 3.1. In the ‘DB links’ field of a hello message, a node indicates the number of its active DB links. This field, along with ‘PRCNT DB’ enables a neighbor with bursty traffic to choose a fair rate and power level for a DB link. The last two fields in Figure 3 are common to both traffic types. Maximum sustainable interference (MSI) information in a hello message presents a node with an upper bound of the tolerable interference at a neighbor [13]. Finally, each node also advertises the aggregate received power of all the active links in its range. This field provides neighboring nodes with recent interference

Node ID	MTBF	PRCNT DB	DB Links	MSI	Current Interfer.
---------	------	----------	----------	-----	-------------------

Fig. 3. Hello message format in distributed mode.

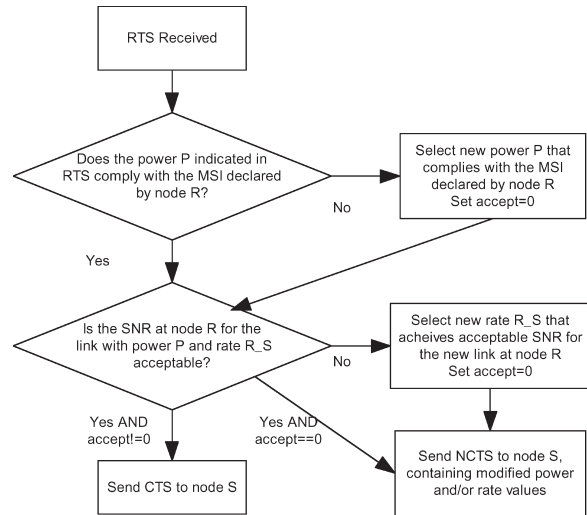


Fig. 4. Receiver R behavior.

levels, which is useful for selecting rates and power values locally.

Each node compiles information contained in incoming hello messages into a small neighbor table, and the node clears a neighbor entry in the table when it no longer detects the neighbor’s hello messages. The storage capacity for the neighbor table is not a major issue for current memory technology.

3.4.2. Period

Because state changes in nodes occur with varying frequency, the hello message period at each node should adapt to the frequency of the node’s state changes. More specifically, the period of hello messages should increase with increasing node stability so that unnecessary hello messages are avoided. Node stability combines the effects of the node’s mobility, its physical reliability, and its degree of state changes. To quantify the first two parameters, we assume that each node can estimate its velocity and its communication reliability (MTBF), which account for positional and physical stability, respectively. A node is physically reliable if its hardware and software components are robust and do not experience frequent interruptions in service. A node has positional and communication reliability if its velocity relative to its neighbors is small. Baldi *et al.* [9] provide a combined measure of these two factors: $C(\text{quality})$.

To monitor the stability of its interference state, each node can compare its current state to the state it advertised in its last hello message. Significant changes

in MSI or current interference levels trigger the node to send an early hello message and lower its hello message period. The creation of a new link could result in state changes at more than one node and trigger them to issue hello messages. Consequently, nodes that detect local state changes upon the creation of a new link must wait a random time (within a maximum wait time) before sending a hello message in order to minimize interference on the hello message TH code.

We define a new boolean cost metric $C(\text{state})$ which takes the value of 1 if either MSI or the current interference at a node vary beyond their respective thresholds, and takes the value of 0 otherwise. When $C(\text{state})$ has a value of 1, a hello message is triggered. We also define a compound cost metric:

$$C(\text{stability}) = C(\text{quality}) + C(\text{state}) \quad (5)$$

which represents the overall stability of a node. U-MAC varies the hello message period at each node linearly with $C(\text{stability})$ at that node, as indicated below:

$$T_{\text{hello}} = \begin{cases} T_{\min}, & C(\text{stability}) \geq C_{\max} \\ K + T_{\max}, & C_{\min} < C(\text{stability}) < C_{\max} \\ T_{\max}, & C(\text{stability}) \leq C_{\min} \end{cases}$$

$$K = \frac{T_{\min} - T_{\max}}{C_{\max} - C_{\min}} \times C(\text{stability})$$

where T_{\max} and T_{\min} are the maximum and minimum time between hello messages, respectively, and C_{\max} and C_{\min} represent the upper and lower bounds of $C(\text{stability})$, respectively. Finally, nodes that do not experience state changes between two consecutive hello messages increase their hello message period by 1 s as long as the period is lower than T_{\max} .

3.5. Rate and Power Assignment

In a centralized UWB network, the access point determines optimal rate and power assignments. Nodes that are out of range of an access point and nodes in UWB ad hoc networks must assign channel rates and transmit power levels in a distributed way. Choosing an appropriate channel rate and power is not simple, since nodes do not have a global view of the network state. Each node can use its neighbor state information to select appropriate parameters for a new link request.

Most conventional wireless networks that use multiple channels require an explicit channel separation mechanism, such as TDMA or FDMA, to accommodate multiple users [5]. In UWB networks, each pseudo-random time-hopping (TH) code constitutes a separate channel.

In U-MAC, all nodes use a known TH code as a common control channel. We also assign another fixed TH code to a dedicated channel for hello messages. Occasional hello message losses are not as critical as control message losses. We assume that nodes synchronize prior to sending and receiving hello messages and control messages. To achieve on-demand synchronization on these two channels, one node could send a short beacon prior to sending its control or hello message to allow neighbors to synchronize to its transmission, which is similar to IEEE 802.11 [16] synchronization for the distributed case. Finally, each of the remaining TH codes is a potential one-way separate data channel. Synchronization on data channels is only required between each sender and receiver pair of an active link.

3.5.1. RB traffic

The RB traffic class accommodates data streams that require a particular QoS. The two quality parameters of interest are the link transmission rate and the SNR at the receiver of the link. Providing a link rate guarantee often prevents any adjustment of the transmit power level while the link is active, in order to maintain the signal integrity at the receiver. Thus, the goal is locally assigning link rate and power values that make efficient use of the medium, achieve fairness among nodes, and ensure that the quality guarantees (transmission rate and minimum SNR) can be maintained for the lifetime of the link.

First, a node S determines the maximum allowable transmit power level for all neighbors, using the following equation [6]:

$$P_{\text{allowed}} = \min \left\{ \frac{(1 - \lambda_i) \text{MSI}_i \times \text{PL}_{si}}{T_f \sigma^2} \right\} \quad (6)$$

where MSI_i is the MSI value announced by the i th node, and PL_{si} is the path loss from node S to neighbor i . In short, node S must select a power level that does not violate the interference threshold of any active links at its neighbors. If there are no active links in the network, then P_{allowed} takes the value of P_{\max} .

Next, node S must select an appropriate rate for the new link to the receiver R , through the expression:

$$R_S = \frac{\min(P_{\text{allowed}}, P_{\text{max}})/\text{PL}_{sr}}{\text{SNR}_{\text{min}} \times \mu(\eta_r + U_r)} \quad (7)$$

where P_{max} is the maximum allowable emitted power from an UWB transmitter, U_r is the combined received power level at the intended receiver, and η_r is the thermal noise level.

R_S should also meet the QoS requirements of the higher layers at node S . Suppose that the network layer at node S requested a desired rate R_{QoS} and minimum acceptable rate R_{min} . If the value of R_S from Equation (7) is higher than R_{QoS} , then R_S is set to R_{QoS} . On the other hand, if R_S is lower than R_{min} , then node S rejects the request at the MAC layer.

After computing R_S locally, S selects a random TH code (other than the control channel and the hello channel TH codes) and initiates a sequence of control messages. If S does not detect that any of its neighbors is attempting to set up a new link, S sends a (RTS) message containing the TH code, $\min(P_{\text{allowed}}, P_{\text{max}})$, R_S , and R_{min} on the common control channel. S then listens for any replies from its neighbors on that channel. The purpose of the RTS message is to ensure that link requests are serialized and that the establishment of this link is recorded and approved by the neighbors of node S . Because S had selected R_S based on its recent local view of the network, all neighbors of S accept the transmission rate R_S with high probability.

Upon receiving an RTS, each neighbor N_i of node S must verify that the rate and additional interference of the new link are admissible. First, N_i uses the received signal strength of the RTS message to compute its current distance from S , which enables N_i to compute PL_{si} . Next, N_i calculates the received power of the new link, using the equation:

$$P_{R_i} = \frac{\min(P_{\text{allowed}}, P_{\text{max}})}{\text{PL}_{si}} \quad (8)$$

The intended receiver R must check for two additional conditions in order before admitting the new link, as shown in Figure 4: (1) the received power of the new link does not exceed the MSI that R advertised in its most recent hello message and (2) the new link has an acceptable SNR at R . If the link request satisfies both (1) and (2), then R sends a CTS message to node S immediately, otherwise R must select the appropriate rate and/or power and includes them in a NCTS message to S .

If the link request exceeds the declared MSI of node R , then R computes the allowable received power P_{Rallowed} from neighbor S . Subsequently, node R can then compute the rate at which a signal from S arriving at R with power P_{Rallowed} would have an acceptable signal quality:

$$R_r = \frac{P_{\text{Rallowed}}}{\text{SNR}_{\text{min}} \times \mu(\eta_r + U_r)} \quad (9)$$

The other case is that the link request does not violate MSI of node R but fails to achieve an acceptable SNR at node R . In that case, node R uses $P_{\text{allowed}}/\text{PL}_{si}$ instead of P_{Rallowed} in Equation (9) to get the R_r that would result in an acceptable SNR for the new link at node R . In addition, node R could check if the TH chosen by S closely correlates with one of the TH codes currently used by other links in its neighborhood [4].

Node S waits for incoming neighbors' replies. If S receives only a CTS message, then it sends a 'Reserve' message indicating the rate, power, and duration of its link reservation, and it immediately sets up a link to node R . The 'Reserve' message also allows the receiver to synchronize to the sender's TH code. If at least one NCTS arrives at S , then S adjusts P_{allowed} and R_S in order to satisfy the updated interference state of its neighbors. If the new value of R_S is higher than R_{min} , then S sends a 'Reserve' message and sets up the link with the newly chosen rate and power. Otherwise, the link request fails.

Upon reception of the 'Reserve' message and establishment of the link, all neighbors of S update their MSI and current interference levels. If any neighbor N_i detects an appreciable variation in either of the two parameters as a result of the update, N_i issues a hello message to inform its neighbors of the state change. If any node's hello message timer expires during a link setup phase, the node postpones sending the hello message until after the link request, to avoid inconsistent views of network state during the link request.

3.5.2. DB traffic

The purpose of DB traffic in UWB is to support best-effort delivery of data without any quality requirements. More specifically, a DB link can sacrifice performance in order to keep interference levels constant at neighbors. In U-MAC, nodes accommodate a new DB request by lowering rates of their other active DB links so that the creation of the new link keeps the

interference levels constant. Naturally, ensuring that the link rates are adaptive to network state requires the symmetric mechanism of increasing rates once a DB transmission ends. To promote fairness, the receiver can also split its DB bandwidth equally among all active DB links.

A node could monitor the traffic nature of its neighborhood through hello messages and allocate a portion of its spectrum to each traffic class. As mentioned earlier, each node allocates λ of its MSI to DB traffic, and $(1 - \lambda)$ to RB traffic. Each node further divides its DB portion equally among all active DB links, and it adjusts all DB power levels and potentially the corresponding rates whenever a new DB link is established. If a node N_i has k active DB links, a new DB link would cause it to adjust the received power level of each link based on the expression:

$$P_{Ri} = \frac{\text{MSI}_i \times \lambda_i}{k + 1} \quad (10)$$

Since nodes in the network use random TH codes, the aggregation of several transmissions appears as background noise at any receiver. The addition of another DB link with a new TH code does not add to the interference at a receiver if the overall DB received power stays the same.

When S has DB data to send to node R , it checks the information compiled from recent hello messages.[§] For each neighbor N_i , S uses a modified version of Equation 6:

$$P_{\text{allowed}} = \min \left\{ \frac{\text{MSI}_i \times \text{PL}_{si} \times \lambda_i}{T_f \sigma^2 (k + 1)} \right\} \quad (11)$$

S then proceeds as in the RB case to assign a corresponding rate with an appropriate margin, to send an RTS message, and to await neighbor replies. The intended receiver R replies with CTS if it consents to the DB request, or with NCTS if the request is not appropriate. Other neighbors of S only reply in case they do not agree with the DB request.

[§]Since λ changes rarely, the MSI triggering of hello messages ensures that neighboring can make DB rate selections based on a sufficiently up-to-date local view of network. To account for DB link changes at neighbors since the last hello message, S could use a margin which is dependent upon the traffic pattern.

Once S processes all the replies, it sends a 'Reserve' message and begins sending DB data. The neighbors of S that are sources to DB links hear the 'Reserve' message and lower their DB link power and rates as needed to accommodate the new DB link from S . However, two-hop neighbors of S do not detect the 'Reserve' message. Suppose N_j is a two-hop neighbor of S , and N_j has an active DB link with a neighbor N_i of S . When N_i detects that N_j has not reduced its power in response to the 'Reserve' message, N_i signals N_j to lower the power (and potentially the rate) of its active DB transmissions.

Modifying the power of all received DB transmissions upon the creation of a new DB stream ensures that the aggregate received power from DB traffic remains constant at each node. Based on the size of DB traffic indicated in RTS and the granted rate in 'Reserve,' each two-hop neighbor can set a timer to indicate the approximate time that this DB link will be active. When the timer for a DB link expires, each neighbor releases the link resources, recomputes the updated state parameters locally, and includes these changes in the next scheduled hello message. There is no need to trigger hello messages upon a DB link expiration, since all nodes in the area set the same timer for this link, and each of them releases its resources locally. Clock skews among the nodes only lead to instantaneous differences in local node states and do not affect the protocol behavior.

3.6. MSI Margin

So far, the discussion has focused on rate and power assignment from the point of view of a sender. Each sender must know the interference state of its neighbors when it sets up a new link. The interference state information that a node advertises in its hello messages is therefore the basis for transmission power and rate assignment at the sender. First, if N_i advertises MSI_{N_i} of its weakest active link, one neighbor N_j may set up a link with N_i that causes MSI_{N_i} to drop to zero, which would block other nodes in the vicinity from setting up new links. Thus, the first challenge is to declare an MSI value that makes efficient use of the medium and ensures fairness. The other challenge for MSI reporting is that, as we mentioned earlier, minor changes in MSI or interference at the node do not trigger hello messages, so nodes may have slightly inaccurate state information about their neighbors.

To address these challenges, each node can declare a fraction of its MSI, in order to avoid starvation of some nodes and to account for unreported minor changes. The portion of MSI that a node declares should depend on how busy a receiver it is. For example, if the average number of active links at N_i in the recent past is low, then N_i can declare a larger portion of its MSI, since it does not expect to receive many more requests. On the other hand, if N_i has many active links on average, then it advertises a smaller MSI. We use the following expression to compute the declared MSI for hello messages at node N_i :

$$\text{MSI} = \frac{\text{MSI}_{\text{total}} \times \delta}{\text{active}} \quad (12)$$

where $\text{MSI}_{\text{total}}$ is the full MSI of the weakest link at N_i , active is the number of active links at N_i , and δ is a topology dependent adjustable margin that trades off fairness for throughput. In Section 5, we explore the effect of varying the values of δ on fairness and throughput.

The margin δ enforces power control, which can contribute to fairness among near and far nodes. Consider the case of Figure 5, where both nodes A and C wish to send data to node B . We expect P_{allowed} at C to be higher than at A for an equal contribution to the interference at node B . However, the value of P_{allowed} at C is constrained by two other factors: the absolute upper limit on transmit power (set by regulatory entities) and the maximum allowable interference at the neighbors of C .

The upper limit on UWB emissions affects network behavior when P_{allowed} is higher than P_{max} (see Equation (7)). In that case, both nodes A and C may set up a link with B at P_{max} , so the rate of a link from A to B can be up to d_2^2/d_1^2 times larger than the rate of a link from C to B .

If the interference state at node B causes the values of P_{allowed} from nodes A and C to be lower than P_{max} ,




Fig. 5. Two senders at different distances from a receiver.

P_{allowed} at C is constrained by the maximum allowable interference at C neighbors. Because C is further away from B than A , C has a higher value for P_{allowed} for the same MSI declared by B . It is therefore more likely that a new link between C and B at P_{allowed} violates the MSI of one of C neighbors. As a result, C selects a power level lower than P_{allowed} for a link with B , which results in a lower link rate. This situation is less likely to occur at node A since P_{allowed} at A is relatively low.

Thus, the relative distances of nodes in the network are a dominant factor for allocating rates in the absence of power control. In Section 4, we investigate the impact of distance and the MSI margin δ on throughput. To this end, nodes must choose rate and power values for new data links based on their view on past and current network state, and based on their projection of future network conditions.

4. Simulation and Results

We used OPNET modeler [17] to implement our protocol model, and to examine the protocol performance in two network settings with RB traffic. First, we consider a case where all the nodes communicate with one central receiver. This case is representative of personal area network settings, such as a home network in which multiple multimedia devices send high quality video or audio to a central screen or computer [18]. It is also applicable to monitoring sensor networks, which typically have a single data sink. Furthermore, this scenario illustrates the performance of our protocol for a highly loaded receiver, and analyzes the degree of favorability for nodes at varying distances from the receiver.

We also consider the case of a network with symmetric traffic that applies to typical wireless local area networks. For the symmetric traffic case, we compare the performance of U-MAC to that of the reactive approach for different traffic loads.

4.1. Simulation Parameters

The upper limit on R_{QoS} for our simulations is 10 Mbps, and the minimum rate of a link R_{min} is based on a uniform distribution with a maximum of 1 Mbps. The minimum acceptable SNR for any link is 14.7 dB [4]. The maximum and minimum hello periods, T_{max} and T_{min} , are 10 and 1 s, respectively. C_{max} and C_{min} are 2 and 0, respectively. The ratio of P_{max} at any node to the thermal noise level is 10^{20} [4]. The size in bits of

hello messages, RTS, CTS, NCTS, and Reserve messages are 64, 40, 16, 32, and 88, respectively. We set the MSI threshold to 10%, and the interference threshold to 50%.[¶]

We assumed a free space path loss model for our simulations. The simulation results provide the upper bound of performance improvement for U-MAC in a line-of-sight (LOS) environment and minimal channel variation. In non-LOS cases or cases where the channel conditions vary frequently, nodes running U-MAC have to provide a higher margin for transmit power to account for potential ranging errors. Note that the SNR quality margin and the MSI margin already offset ranging errors by providing a safety margin above the minimum transmit power values. In our simulations, the SNR quality margin μ is set to 2.

The MAC layer at each node receives requests from the network layer according to a poisson process, and selects the receiver at random in the symmetric case (there is only one receiver in the loaded receiver case). If the network layer at a node S requests a new link while some other node N has a link request in progress, the new request is buffered at S until N completes its current link request. The serialization of link requests achieved by the RTS/CTS exchange ensures that the MSI and interference levels at a node remain the same during the handshaking process. Once the MAC layer fetches a link request at the head of the request queue, it attempts to send RTS and wait for replies. If RTS times out, the node sends RTS again. If there is no reply after 3 RTS messages, the link request fails. Also, more than one node may have buffered link requests, so if all of these nodes attempt to send RTS at the same time, then collisions will occur. Thus, each node waits for a random time within 0.2 s before servicing a queued link request to reduce the probability of RTS collisions. A node that has just completed a link setup cycle must choose a random time within 0.3 s before it services any buffered requests. This mechanism helps promote fairness, since it gives the node with the most recent link a lower chance of immediately starting a new link request. Finally, each node may have multiple active links at the same time, by using a separate TH code for each link.

[¶]We set the interference threshold for triggering hello messages higher than that for MSI because the former changes more frequently. Since changes within the threshold do not trigger hello messages, nodes use margins in their rate and power assignments (see Subsection 3.6).

4.2. Results

4.2.1. Loaded receiver

The topology for the loaded receiver case has 25 nodes, where 24 nodes are located at distances varying from 5 to 27 m from the common receiver. We observe the impact of distance from the receiver on transmission rate using our MAC protocol, and we demonstrate how power control can be used to adjust the radius of favorable senders around the receiver.

Figure 6 plots the average node throughput at various distances from the data sink. Let the distance of the closest node from the receiver equal d_{\min} . We define the radius of fair access (ROF) as the maximum distance from the receiver within which nodes get similar throughput as nodes at d_{\min} from the receiver. A node at a distance $d > \text{ROF}$ from the receiver achieves a throughput proportional to $1/(d - \text{ROF})^2$. As we lower δ , we find that ROF expands according to the following expression:

$$\delta = \frac{1}{2^n}, \quad n \geq 2; \quad (13)$$

$$\text{ROF} \simeq 2^{n-1} d_{\min}$$

Equation (13) states that cutting δ by half doubles the ROF for δ values of 0.25 or lower. Figure 6 also shows that lower values of δ improve the performance of nodes further away from the receiver, even if these nodes remain outside the ROF. To have strictly fair access to the receiver among all nodes, ROF must equal the radius of the network centered around the common receiver.

Figure 7 plots the local throughput of nodes at different distances from the receiver against δ . The




Fig. 6. Node throughput versus distance.




Fig. 7. Node throughput versus MSI margin.

maximum value of each distance curve in Figure 7 represents the highest achievable throughput for a node at that distance in this topology. For nodes that are relatively far from the receiver, the highest achievable throughput using U-MAC occurs for low values of δ . Thus, Figures 6 and 7 provide the basis for adjusting δ to favor nodes at certain ranges depending on the spatial distribution and throughput requirements of nodes in a network.

Finally, Figure 8 shows how network throughput varies with δ . Decreasing δ from 1 to 0.25 improves network throughput. This peak in throughput can be understood by examining Figure 6, which shows that an MSI Margin δ of 0.25 widens the ROF to 10 m and raises the throughput of nodes further away from the receiver with only minor decreases in the throughput of nodes closer to the receiver. Lowering δ beyond 0.25 causes significant decreases in throughput of




Fig. 8. Throughput versus MSI margin.

nodes that are close to the receiver, and thus yields lower overall throughput.

4.2.2. Symmetric traffic

We considered three different topologies with 25 nodes each to investigate the symmetric traffic case. The first topology is a grid topology with a constant node separation distance of 5 m. The second topology is a random distribution of nodes in a 100 m \times 100 m area, with an average and minimum inter-node distance of 10 and 5 m, respectively. The third topology is a random distribution of nodes in a 50 m \times 50 m area, where the average and minimum inter-node distances are 5 and 1 m, respectively. The results presented here are the average of three topologies.

We vary the arrival rate of new link requests to observe the behavior of the protocol for different traffic loads. The δ value used for this scenario is 1.

We first consider the link setup latency benefits of using U-MAC. In the reactive approach, a node that sends an RTS must wait for replies from all of its known neighbors. Each of the neighbors uses a probabilistic back-off scheme for sending its response in order to avoid collisions of replies on the control channel. In U-MAC, a node that has sent RTS only waits for replies from the receiver and any neighbor with conflicts, so there is an inherent latency improvement. Figure 9 compares the link setup latency in U-MAC to the reactive approach. The average latency in U-MAC increases steadily from 13 ms at low arrival rates to 93 ms at an arrival rate of 0.66. At low link request arrival rates, the improvement in average latency of U-MAC over the reactive approach remains between 130 and 155 ms. The gap starts to widen at a request arrival rate of 0.25 and reaches a maximum of about 36 s at very high arrival rates. The exponential

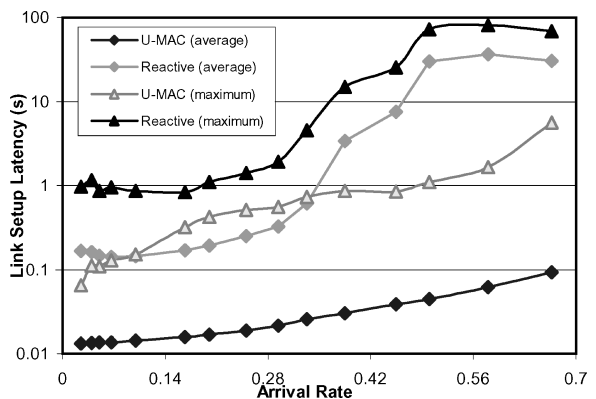


Fig. 9. Average link setup latency versus offered load.

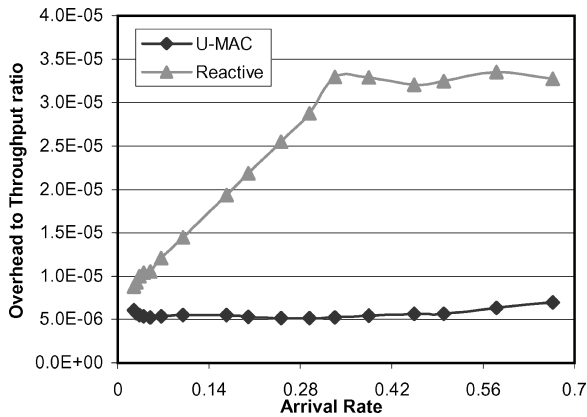


Fig. 10. Control overhead versus offered load.

increase in latency for the reactive case is attributed to the requirement that all neighbors must send their replies upon a link request. An increased frequency of link requests causes a sharp rise in link setup latency. At arrival rates of 0.5 and higher, both the average and maximum latency for the reactive protocol stabilize. The nodes reach their limit in the rate of requests they can handle, and although the arrival rate varies, the same number of link requests are serviced while the other link requests are discarded locally.

Next we consider the control overhead of U-MAC and compare it to the reactive case. For each approach, we obtain the ratio of the bit rate used by control messages to the overall bit rate in the network, which we refer to as overhead to throughput ratio. Figure 10 reveals that this ratio for U-MAC remains constant and only starts to increase slightly at arrival rates above 0.56. For the reactive case, the overhead to throughput ratio increases at a constant rate with increasing link requests because the increase in control messaging exceeds the throughput increase. At arrival rates above 0.33, the ratio in the reactive approach stabilizes as both the control overhead and the network throughput remain almost unchanged.

Figure 11 plots the admitted rate as a function of the offered load for both U-MAC and the reactive approach. In an ideal scenario, the network would admit all of the requested transmission rate, which corresponds to the line $y = x$. In U-MAC, the admitted rate is the same as the ideal case for loads up to 20 Mbps. As the link request rate grows, state changes occur more frequently, and as a result nodes have less accurate information about their neighbors' states. Consequently, the admitted rate starts falling short of the requested rate, but the behavior remains close to

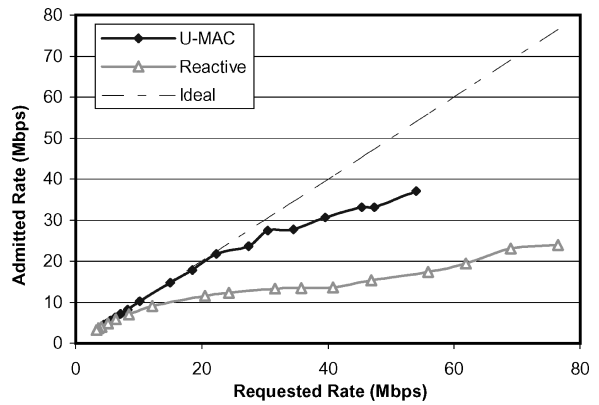


Fig. 11. Admitted load versus offered load.

the ideal case at offered loads above 20 Mbps. In the reactive approach, the requested rate is fully admitted only for loads below 10 Mbps. As the network load increases, the admitted rate in the reactive approach is increasingly lower than the requested rate. Figure 11 also reveals that nodes in the reactive approach request more bandwidth than in U-MAC because of their lack of information on network state. In U-MAC, nodes request only as much bandwidth as can be supported by the network according to their local view of network state.

Figure 12 compares the overall network throughput in both the reactive approach and U-MAC. When the link request arrival rate is low, the throughput of both cases is similar because few links are active simultaneously, so protocol mechanisms have minimal effect. As the link request arrival rate increases, nodes in the reactive case grab bandwidth greedily and limit the potential number of coexisting links. In U-MAC,




Fig. 12. Throughput versus offered load.

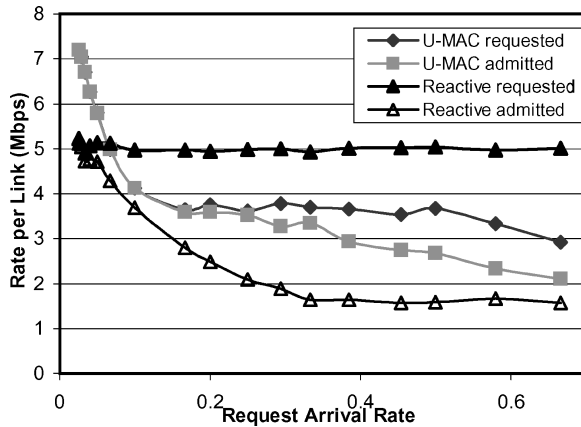


Fig. 13. Rate per link versus request arrival rate.

nodes stay updated about the network state which allows for more efficient use of the medium. Thus, there is a growing gap in the throughput as arrival rates increase above 0.1. At a request arrival rate of 0.66, the throughput for U-MAC is about double the throughput in the reactive case.

Figure 13 plots both the admitted link rate and the requested link rate as functions of the link request arrival rate for U-MAC and the reactive approach. In the reactive approach, the requested link rate is independent of network state, so it does not vary with network load. As a result, the gap between admitted and requested link rates grows with network load, and stabilizes for link arrival rates above 0.3. In U-MAC, nodes adapt their requested link rate to the interference and number of active links in the network for new link requests. The requested link rate is generally admitted for link request arrival rates up to 0.25. For link request arrival rates between 0.25 and 0.5, there is a growing gap between the requested and admitted link rates. This indicates that nodes make less accurate local rate and power assignments due to a higher rate of change in network state. However, the gap between the requested and admitted rate stabilizes for arrival rates between 0.5 and 0.66, which indicates that the portion of inaccurate rate and power assignments remains the same for those loads.

Finally, Figure 14 compares the sustained throughput of a node central to the topology and a node on the periphery versus the link request arrival rate. At low arrival rates, the gap between the throughput values of the two nodes is narrow regardless of MAC protocol mechanisms. As the arrival rate increases both nodes exhibit higher throughput, but the gap increases indicating that nodes central to the topology grab more

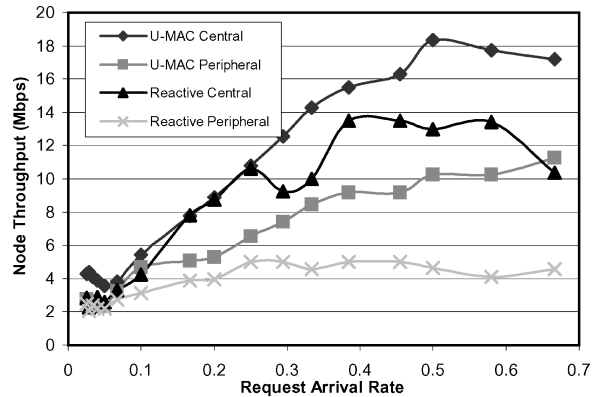


Fig. 14. Throughput of central and peripheral nodes.

bandwidth as the traffic load increases. The gap stabilizes for a link request arrival rate of about 0.5, where network throughput starts to saturate.

The gap in throughput between the two nodes is due to the difference in their average distances from other nodes in the network. Decreasing δ to reduce this gap is not always an attractive option in symmetric traffic networks, since it affects overall throughput more severely than in the single receiver case. Both the reactive approach and U-MAC exhibit this behavior; however, the gap for the reactive case is larger due to the greedy approach in bandwidth assignment.

5. Discussion and Conclusion

Previous work has addressed the joint rate assignment problem under different assumptions and conditions. Lal and Sousa [13] proposed a reactive protocol that addresses the problem for direct sequence code division multiple access (DS-SS) networks. Their protocol involves a set of handshaking messages to negotiate resource allocation and leverages the concept of MSI along with several techniques for resource allocation based on minimizing power, maximizing rate, or maximizing SNR.

The CDMA model of Lal and Sousa was adapted by Cuomo *et al.* [6] for UWB networks. Their work presented a reactive approach [6] to address the joint rate power assignment problem. The approach specifies that nodes request neighbor information on-demand for setting up a new link. More specifically, a sender first polls its neighbors for their MSI measurements. Each neighbor must send an MSI measurement to the sender, and neighbor replies may overlap

in time. Once the sender gets replies from all its neighbors, it selects the appropriate rate and power for the link and initiates another handshake to confirm the selected parameters. This protocol requires the sender to receive and differentiate between replies from all its neighbors at the same time. As the number of neighbors grows, so does the number of simultaneous replies that must be processed by the sender. This presents a technical challenge since all replies use the same control channel code. The work in Reference [6] disregards this challenge and assumes all control messages are successfully received without taking up any radio resources. Ensuring that all neighbor replies are received successfully requires some probabilistic back-off scheme at each neighbor, which delays links setup. Going through two handshakes further contributes to link setup latency.

The comparison of U-MAC and the reactive approach that is similar to Reference [6] in a realistic scenario has shown that U-MAC decreases control overhead and link setup latency considerably while making more efficient use of the medium. The decrease in control overhead and improvement in efficiency are attributed to the availability of neighbor state information locally at each node, and to the fact that only some neighbors reply to each link request. The latency decrease also benefits from the selective neighbor replies, as well as the elimination of one control message in the handshaking sequence.

The work in References [7,8] independently proposed a proactive protocol that is related to U-MAC in that it uses periodic broadcast messages. The authors presented techniques for setting and adjusting MSI margins and simulation results for access probabilities under different medium conditions. U-MAC further expands these ideas in two directions. First, U-MAC dynamically sets the hello message period to adapt the degree of state changes at nodes in order to avoid unnecessary state advertisements. Second, the protocol in References [7,8] specifies that a node setting up a link requires replies from all neighbors. In U-MAC, only the subset of the neighbors that do not agree with the link parameters send replies, which is similar to the selective reject (SREJ) concept in Reference [13]. As a result, U-MAC reduces control overhead and link setup latency.

In addition, our study on U-MAC provides the first comparative study of reactive and proactive approaches to the joint rate and power assignment problem in UWB. The study reveals that the proactive approach doubles network throughput under high traffic conditions. Our other simulation study addresses fair access

between nodes at varying distances from a common receiver, and explores the tradeoff between achieving maximum network throughput and promoting fair access to individual nodes. The results reveal that there is a radius of fair access for each receiver, within which all nodes achieve comparable throughput. The MSI margin δ in U-MAC controls the radius of fair access and determines the balance between fair access to the receiver and overall network throughput.

One direction for future research is to integrate multi-hop links into our protocol. Using information provided by hello messages, the nodes can make local decisions on least-cost paths to a particular destination [9].

Another issue for future investigation is the effect of mobility on the protocol. The protocol framework provides measures of positional reliability, but our simulations only considered stationary nodes. It would be interesting to explore techniques for using UWB radar capability to keep track of mobile nodes, and to study the impact of mobility on the hello message period and the resulting protocol behavior.

Finally, coupling UWB positioning capability with directional antennas [19] can reduce power consumption since the receiver captures most of the transmitted power. Once a sender knows the receiver's location, it can direct the antenna beam towards the receiver. Smart antennas [20], which have been considered for UWB transmissions [21,22], are directional antennas that can physically steer themselves towards the receiver. We can modify U-MAC to operate with smart antennas by considering interference in each sector around the receiver independently. We expect that the modified protocol would enable more simultaneous links provided the links are evenly distributed in all sectors.

In sum, we have presented a new proactive and adaptive MAC protocol for UWB networks called U-MAC. It provides well-defined parameters to control the fairness/throughput tradeoff and that reduces control overhead and connection latency, while increasing network throughput. We have also presented our simulations to quantify the tradeoffs involved and the benefits of our proactive approach over the reactive approach.

Acknowledgments

The authors thank Dr Murad Jurdak for his valuable input on improving the organization of the paper. The authors also thank the anonymous reviewers.

References

1. IEEE 802.15.3a. WPAN High Rate Alternative PHY Task Group 3a. Draft Standard. Available: <http://www.ieee802.org/15/pub/TG3a.html>. [Accessed: 25 January 2005].
2. IEEE 802.15.4a. WPAN Low Rate Alternative PHY Task Group 4a. Draft Standard; <http://www.ieee802.org/15/pub/TG4a.html>. [Accessed: 25 January 2005].
3. Win MZ, Scholtz R. Impulse radio: how it works. *IEEE Communications Letters* 1998; **2**(1): 36–38.
4. Win MZ, Scholtz RA. Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications. *IEEE Transactions on Communications* 2000; **48**(4): 679–691.
5. Jurdak R, Lopes CV, Baldi P. A survey, classification, and comparative analysis of wireless medium access control protocols for ad hoc networks. *IEEE Communications Surveys and Tutorials* 2004; **6**(1): 2–16.
6. Cuomo F, Martello C, Baiocchi A, et al. Radio resource sharing for ad-hoc networking with UWB. *IEEE Journal on Selected Areas in Communications* 2002; **20**(9): 1722–1732.
7. Cuomo F, Martello C, Caputo S. An interference-controlled admission control scheme for QoS support in distributed UWB networks. In *Proceedings of 1ST Mobile and Wireless Communications Summit*, 2003.
8. Cuomo F, Martello C. Improving wireless access control schemes via adaptive power regulation. In *Proceedings of PWC 2003*, 2003; pp. 114–127.
9. Baldi P, De Nardis L, Di Benedetto MG. Modelling and optimization of UWB communication networks through a flexible cost function. *IEEE Journal on Selected Areas in Communications* 2002; **20**(9): 1733–1744.
10. Buehrer RM, Ha D, Reed J, McKinstry D, Kumar N. Range extension for UWB communications. Virginia Tech, 2002. Available: <http://filebox.vt.edu/users/nikumat/>; [Accessed: 20 January 2005].
11. Cuomo F, Martello C. MAC principles for an ultra wide band wireless access. In *Proceedings of Global Telecommunications Conference*, 2001.
12. Whyless.com: the Open Mobile Access Network. Available: <http://www.whyless.org/>; [Accessed: 20 January 2005].
13. Lal S, Sousa ES. Distributed resource allocation for DS-SS-based multimedia ad-hoc wireless LAN's. *IEEE Journal on Selected Areas in Communications* 1999; **17**(5): 947–967.
14. Revision of Part 15 of the Commission's Rules Regarding Ultra-Wideband Transmission Systems. Federal Communications Commission, ET Docket 98–153, 2002.
15. Toh C-K. A novel distributed routing protocol to support ad-hoc mobile computing. In *Proceedings of IEEE 15th Annual International Phoenix Conference on Computer Communications* 1996; pp. 480–486.
16. IEEE 802.11 WG. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 1999. Standard.
17. OPNET Modeler software. Available: <http://www.opnet.com/products/modeler/home.html>; [Accessed: 25 January 2005].
18. Stroh S. Ultra-wide band: multimedia unplugged. *IEEE Spectrum* 2003; **40**(9): 23–27.
19. Smith CE. Theory and design of directional antenna systems. Cleveland Institute of Radio Electronics, 1951.
20. Barrett M, Arnott R. Adaptive antennas for mobile communications. *Electronics and Communication Engineering Journal* 1994; **6**(4):203–214.
21. Schantz HG. Introduction to ultra-wideband antennas. In *Proceedings of IEEE Conference on Ultra-Wideband Systems and Technologies*, 2003.
22. Hussain MGM. Principles of spacetime array processing for ultrawide-band impulse radar and radio communications. *IEEE Transactions on Vehicular Technology* 2002; **51**(3): 393–403.

Authors' Biographies



Raja Jurdak received his B.E. degree in Computer and Communications Engineering from the American University of Beirut in 2000, and his M.S. degree in Computer Networks and Distributed Computing from the University of California, Irvine in 2001. Currently, he is a Ph.D. candidate at the School of Information and Computer Science at UCI. His research interests include energy-efficient modeling of ad hoc networks and sensor networks, and emerging communication technologies like ultra wide band radio and acoustics. He is a student member of the IEEE.



Pierre Baldi is a professor in the School of Information and Computer Science and Application Layer Leader of the California Institute for Telecommunications and Information Technology [CAL(IT)²] at the University of California, Irvine. He received his Ph.D. in Mathematics from the California Institute of Technology in 1986. From 1986 to 1988, he was a postdoctoral fellow at the University of California, San Diego. From 1988 to 1995, he held faculty and was a member of the technical staff positions at the California Institute of Technology and at the Jet Propulsion Laboratory. He was the CEO of a startup company from 1995 to 1999 and joined UCI in 1999. He is the recipient of a 1993 Lew Allen Award at JPL and a Laurel Wilkening Faculty Innovation Award at UCI. Dr Baldi is the author of over 100 research articles and four books, including *Modeling the Internet and the Web—Probabilistic Methods and Algorithms* published by Wiley in 2003. His research focuses on probabilistic modeling and statistical inference, machine learning, bioinformatics, information theory, data mining, and communication networks.



Cristina Videira Lopes is an assistant Professor in the School of Information and Computer Science, University of California, Irvine. She conducts research in ubiquitous computing and in programming languages. She has a Ph.D. in Computer Science from Northeastern University. She is a member of the IEEE and the ACM.