CTU: Capturing Throughput Dependencies in UWB Networks

Ioannis Broustis*, Angelos Vlavianos*, Prashant Krishnamurthy◇, Srikanth V. Krishnamurthy*
*University of California, Riverside, ◇University of Pittsburgh
{*broustis, aggelos, krish*}*@cs.ucr.edu, prashant@mail.sis.pitt.edu*

***Abstract*—The inherent channel characteristics of impulse-based UWB networks affect the MAC layer performance significantly. Previous studies on evaluating MAC protocols are based on prolonged simulations, and do not account for the multiple-access interference that arises due to multipath delay spread. In this work, we develop CTU, an analytical framework that captures the performance of MAC protocols, while taking into account the underlying PHY layer effects. The key attributes that make CTU novel are: (a) It is modular and therefore flexible; it can be easily modified to provide a basis for characterizing and evaluating a wide range of MAC protocols designed for impulse-based UWB networks. The only requirements are that the MAC protocol under study be based on time-hopping, and the modulation scheme be pulse position modulation; these are common design decisions in most impulse based UWB networks. (b) It considers the channel characteristics in addition to MAC layer effects; in particular, CTU correlates probabilistically the multipath delay profile of the channel with the packet error rate. We employ CTU to evaluate the performance of a generic medium access procedure. We compare the results with those from extensive simulations and show the high accuracy of CTU. We use CTU to assess the impact of various system parameters on the MAC layer performance; we make several interesting observations that are discussed in depth.**

***Index Terms*—Ultra-Wide Band (UWB)[1], Wireless Communications, Multipath Delay Spread, Modulation.**

I. Introduction

The medium access performance of multi-hop impulse-based UWB networksis significantly affected by the PHY layer characteristics, and in particular by the multipath delay spread [1]. Due to reflections from various objects, multiple copies (rays) of the same transmitted signal appear at the receiver; each copy has a different amplitude, phase and delay. The copies interfere with later pulse transmissions and this significantly affects the long-term network throughput. When evaluating MAC protocols designed for UWB networks, it is critical that one accounts for these effects. Previous studies on UWB, however, are limited to either (i) possibly time-intensive simulation-based MAC protocol evaluations (most of such efforts do not account for multipath delay spread effects), or (ii) PHY layer analyses to model the behavior of a single link in the presence of delay spread. A cross-layer analytical framework that quickly and accurately quantifies the impact of the UWB PHY layer attributes on MAC layer mechanisms, could be of great help throughout the protocol design process.

In this paper, we develop an analytical framework to compute the average saturation throughput and evaluate its dependencies on PHY layer effects, in multi-hop UWB networks. The saturation throughput is defined to be the throughput when the nodes in the network always have packets to send. We call our framework "**CTU**" for short (for Capturing Throughput dependencies in UWB networks). The **key attributes** of CTU are: *(a)* it is modular and therefore flexible, and *(b)* it captures both physical layer characteristics and MAC layer effects. CTU computes the probability of packet loss due to (i) MAC layer collisions and (ii) multiple access interference (*MAI*) due to delay spread effects. It then combines them to compute the saturation throughput. The immediate benefit from our analysis is that it obviates the need for *repeated* simulations. We elaborate on the attributes of CTU below:

CTU is modular and hence applicable to a wide set of MAC protocols. The part of the CTU framework that computes packet losses due to *MAI* is generic and can be used in conjunction with any MAC layer protocol. The only requirements are that the protocol be based on time-hopping (TH), a multiple access procedure, and that the underlying modulation scheme be Binary Pulse Position Modulation (BPPM); this is common to most PHY/MAC layers designed for UWB networks. The building block for CTU is a basic MAC layer procedure that is likely to be the foundation for more complex TH-based MAC layer protocols.

CTU accounts for cross layer dependencies. Multiple-access interference in UWB networks not only depends on the extent of delay spread but also on the access procedure defined by the MAC protocol. CTU provides a seamless methodology of integrating the access procedure with the effects of delay spread; in other words, it captures the cross layer dependencies between the PHY and the MAC layers. In particular, CTU incorporates the impact of delay spread on the multiple access performance based on a modified version of the Saleh-Valenzuela (SV) model that has been adopted by the IEEE 802.15.3a task group [2], (discussed in more detail in section II).

Evaluation of CTU. We apply CTU on a simple MAC protocol, and we compute the per node throughput. We validate the accuracy of our analytical results, by comparing them with results from extensive simulation experiments that we perform. We also find that while obtaining simulation results takes days, generating results with CTU requires less than one hour.

Using CTU to assess system performance. With CTU, we examine the impact of various system parameters on the saturation throughput. We discuss in detail, many interesting trends that are observed. We believe that CTU can help designers and implementers choose UWB MAC layer system parameters appropriately during network design.

The remainder of the paper is structured as follows. In section II we provide background on the UWB PHY layer and discuss related work. In section III, we describe our analytical framework in detail. In section IV, we compare the simulated behavior of the candidate MAC protocol with the behavior observed with CTU. We also describe the impact of various system parameters on network throughput under saturation conditions, report the trends that we observe and the interpretations thereof. In section V we discuss the scope of CTU in terms of its tunability and its limitations. We conclude in section VI.

[1]This work is supported in part by the NSF CAREER Grant No. 0237920 and the NSF NRT grant No. 0335302.

II. Background and Related Work

In this section, we briefly discuss the UWB PHY layer, the SV model and related previous studies [3], [4], [5].

Time-hopped Binary-PPM: FCC has allocated the frequency band between 3.1 and 10.6 GHz for UWB communications. It regulates that UWB signals span at least 500 MHz, or more than 25% of a center frequency [1]. BPPM signals are comprised of pulses that are represented by Gaussian-shape doublets. Each pulse represents an encoded bit, which can be a *'0' or a '1'*, ass determined by the pulse position during a *chip* time, T_c. A " T_c time duration" is divided into two consecutive time *slots*. If the pulse occupies the first slot, it represents a *"0"*, otherwise a *"1"*. The pulses that form a signal are transmitted as per a predetermined sequence in the time space, called the Time Hopping Sequence, or THS. Each specific THS maps onto pulse transmissions in certain chips, with exactly one pulse transmission during a fixed duration of a *sequence frame*, T_f. In other words, only one pulse is sent, per transmitter, during T_f, within one of the possible T_f/T_c chips. With this, parallel transmissions are possible on the same channel, as long as a transmitter-receiver pair uses a unique THS. When the used time hopping sequences are not orthogonal, the parallel transmissions may interfere with each other due to overlap among the sequences. We denote by $N_f = T_f/T_c$, the number of chips during a sequence frame.

The modified SV model: Due to reflections on various surfaces, attenuated copies of transmitted a pulse (called rays) arrive at the receiver at different time instances. This is due to the *multipath delay spread* of the channel [4]. The time-delayed rays are likely to interfere with subsequent pulse transmissions. When such rays are due to the same transmitter, the use of an equalizer at the receiver can help detect the original pulse. If, however, there are may interfering transmitters in the receiver's vicinity, the rays from the different transmissions add to interference noise at an *unintended* receiver.

The SV channel model characterizes the arrival of rays at the receiver [4]. As per this model, rays arrive at the receiver in clusters. Within each cluster, rays appear in time as per the Poisson(λ) arrival process:

$$p(\tau_{\nu,l}|\tau_{(\nu-1),l}) = \lambda e^{-\lambda(\tau_{\nu,l}-\tau_{(\nu-1),l})}, \tag{1}$$

where $\tau_{\nu,l}$ is the delay of the ν_{th} ray, relative to the l^{th} cluster arrival time, T_l and the previous, $(\nu-1)^{st}$ ray. Similarly, clusters also appear as per a second Poisson(Λ) process:

$$p(T_l|T_{l-1}) = \Lambda e^{-\Lambda(T_l-T_{l-1})}. \tag{2}$$

In each cluster, the power in the component rays, decays as per an exponential-decay profile. In addition, the model determines the discrete time impulse response of the signal to be:

$$h(t) = \sum_{l=0}^{L}\sum_{\nu=0}^{V} \alpha_{\nu,l}\delta(t - T_l - \tau_{\nu,l}), \tag{3}$$

with $\alpha_{\nu,l}$ being the multipath gain coefficient. The primary version of the SV model employs a Rayleigh distribution for the multipath gain magnitude. Experimental data, using Kolmogorov-Smirnov testing with a 1% significance level suggested that the lognormal amplitude distribution seems to fit the profile better than the Rayleigh distribution [5]. Thus, in this paper we adopt the lognormal distribution for the channel coefficients. In particular, $\alpha_{\nu,l} = p_{\nu,l}\beta_{\nu,l}$, where $p_{\nu,l}$ is equiprobably ± 1 and $10\log_{10}|\beta_{\nu,l}|^2 \propto Normal(\mu_{\nu,l}, \sigma^2)$. Moreover, $\mu_{\nu,l}$ is given by:

$$\mu_{\nu,l} = \frac{10ln(P_m) - 10T_l/\Gamma - 10\tau_{\nu,l}/\gamma}{ln10} - \frac{\sigma^2 ln10}{20} \tag{4}$$

where P_m is the mean power of the first ray of the first cluster, and Γ, γ are the cluster and ray decay-factors, respectively. We will revisit this model later, in section III.

Related studies: Our work is the first effort to propose an analytical framework for MAC protocol evaluation in pulse-based UWB. However, many studies on designing and evaluating MAC protocols through simulations exist. Some of these protocols take the delay spread effects as well as overall background noise, into account. In [1] we propose a multi-band MAC protocol to alleviate the impact of both THS overlaps and multipath delay spread. There are numerous other studies on proposing new MAC protocols [6].

Studies on Impact of Delay Spread: In [7], the authors compute the conditional bit error rate for TH-BPSK (Time-Hopped Binary Phase Shift Keying) UWB, in a multi-user, multipath-channel environment. However, unlike with our work, the authors assume that an interferer collides with at most one pulse; inter-symbol interference effects are also ignored. In addition, they do not study the cross layer dependencies between the MAC and PHY layers. Gubner and Hao, in [8], derive an analytical formula for the average bit error probability at the PHY layer, considering the IEEE 802.15.3a model. They observe that, if one were to consider an observation window that accounted for 99% of the interference energy due to pulse copies, the performance over a link was identical to a case where one considers the delay spread to be infinite. In other words, they show that there is a maximum duration (maximum value of the delay spread) over which pulse copies will need to be considered. The authors however, do not consider multi-access interference or MAC level characteristics. In [9] we perform a simulation study of the impact of delay spread on the throughput in UWB networks with a simple model; we do not consider any ray fading model (as with the SV model).

Channel Models: There exists a lot of work on modeling the multipath delay spread profile. Foerster [5] proposes a variation of the SV model that adopts the Nakagami fading distribution. In [10] Cassioli et al. consider only one cluster with an exponential decay of the ray powers. Ghassemzadeh et al. [11] modify the single-cluster model of [10] to include a noise-like variation with lognormal statistics. Greenstein et al., in [12], present an interesting comparison of these models.

III. Our Analytical Framework

In this section we present CTU, a framework for estimating the MAC layer performance, in pulse-based UWB networks. In order to make the description easy to understand, we first explain the operations of a simple MAC procedure and then, progressively build CTU on top of it. Note that there could exist more intelligent MAC protocols [6] than the one considered here. CTU *can be easily tuned to characterize different impulse-based protocols* (as explained later in section V).

The Basic MAC Layer Procedure: The single-band MAC layer procedure used for modeling purposes is a simplified

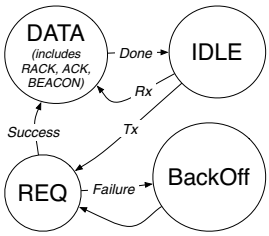

Fig. 1. MAC state diagram. The DATA state also includes the transmission of RACK, ACK and BEACON messages.

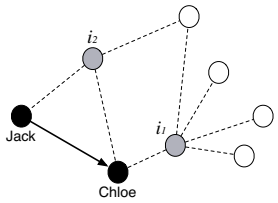

Fig. 2. Receiver's neighborhood: Jack wants to send a packet to Chloe, while there are i other neighbor transmitters (here i_1 with l_1=5, and i_2 with l_2=3).

version[2] of the one in [14]. In the simplified version, nodes utilize randomly overlapping THSs. The protocol operations are depicted in Fig. 1.

Let us assume that Jack wants to send a data packet to his neighbor, Chloe (Fig. 2). Jack is in the *REQ* state, while Chloe is in the *IDLE* state. The following events take place (Fig. 1).

***REQ* state:** Jack transmits a request (REQ) to Chloe (a randomly selected neighbor), as per Chloe's THS. He then switches to a THS that is unique to Jack and Chloe. Nodes in the REQ state are only the transmitters that send REQ messages. Their potential receivers may be in any of the other three states; however, the REQs can be received successfully only when the receiver is in the *IDLE* state.

***DATA* state:** After a successful reception of the REQ, both Jack and Chloe are in the *DATA* state, wherein Chloe responds to Jack's request with a RACK message. Note that in the *DATA* state, a THS unique to Jack and Chloe, is used. Jack now knows that Chloe has accepted his request, and he proceeds with his DATA packet transmission. Chloe acknowledges the DATA reception with an ACK message to Jack. At this point, the MAC layer session between Jack and Chloe has ended. Both Jack and Chloe further transmit a short BEACON message, on their own THSs, to inform their neighbors that they are now idle. Then, both of them transit to the *IDLE* state.

***IDLE* state:** A node resides in the *IDLE* state for a random duration, waiting for potential REQ transmissions. If no REQ arrives within this time, the node will go to the REQ state, to initiate a new session. Note that if Chloe is in the *IDLE* state, and the random duration expires before Jack completes a REQ transmission to her, Chloe will abort this REQ reception, and will transition to the *REQ* state. Again, the model for this simple procedure can be easily replaced with other possibilities.

***BackOff* State:** If Bill is transmitting a REQ to Chloe at the same time as Jack, the two REQ messages will collide since both Jack and Bill are using Chloe's THS for their REQ transmissions. Jack and Bill will realize that a collision has taken place, since Chloe will not send a RACK back. Jack and Bill will both go to the BackOff state, and will initiate their random backoff timers. When a timer expires, the corresponding node will re-transmit the REQ to Chloe. With this backoff algorithm, the contention window increases exponentially from 0 to J for the first ξ REQ collisions, and remains at J for up to ζ failures. If the REQ does not go through in ζ consecutive attempts, Jack will abort his transmission, and will send a new REQ to another, randomly selected neighbor. Moreover, if Chloe is already in the DATA state with Bill, Jack will go to the *BackOff* state, until he receives a BEACON message from Chloe, notifying him that she is now idle.

[2]We have used this protocol version in the past [1], [13], [9]

Constructing CTU: Our primary task is to capture the impact of the multipath delay spread on the MAC layer performance, taking into account *(a)* the inherent properties of BPPM-enabled links, and *(b)* the MAC protocol scheduling aspects. In particular, we utilize the SV model, to approximate the level of interference within a chip time, due to ray arrivals. This enables us to compute the probability of successful decoding of a bit, and progressively a MAC-layer packet. The packet success, however, also depends on the set of actions dictated by the MAC layer. Using the basic MAC protocol described above, CTU combines these interdependent factors, in order to compute the probability of a MAC-layer packet success.

For analytical tractability, we make the following assumptions; we elaborate on some of them in section V.

1) Nodes always have packets for transmission in their queues. In other words, we consider conditions of saturation.

2) A packet reception can only be interfered with, by transmissions from nodes that are within the transmission range R, of the receiver. Furthermore, the transmission range is much smaller than the deployment region.

3) Each bit is represented by PB repetitions of a pulse transmission (repetition coding). A bit is successfully decoded, if *at least one* of the PB pulses is correctly deciphered.

4) The duration of the multipath delay spread is fixed and equal to Δ; this implies that we are interested in rays for which the received power is above a certain threshold, a reasonable approach as shown in [8]. We assume that Δ corresponds to an integer number of chips: $\Delta = N_c \cdot T_c$.

5) Receivers are equipped with an equalizer. Therefore, rays do not interfere with subsequent pulses from the *same* transmitter.

6) Regarding the path-loss, we assume an attenuation of 20 dB for a transmission rate of 110 Mbps, over a distance of 10m [5], and for both the given transmitter and interferers.

7) Based on the SV model, we assume that there is no overlap between the rays induced by a single transmitter in any slot and that the received pulses/rays are not distorted. The rays are assumed to arrive within slot boundaries.

8) We assume that the number of active transmitters that can cause interference remains a constant during a packet's transmission. This number can actually increase in some cases while decrease in others; the former effect is detrimental to packet receptions while the latter aids the process. Thus, on average, one might expect the behavior to be similar to that of the assumed system.

Let us now consider a network deployment consisting of N nodes, randomly and uniformly deployed across a square region A. CTU estimates the MAC-layer network throughput, by computing the packet failure probability, $P_{failure}$, for the MAC protocol under study, as follows.

Step 1: Identifying the cases for packet failure: A packet transmission may fail due to any of the following two reasons:

1) The packet collides because more than one senders transmit simultaneously towards the same receiver. Here each transmission fails because the same THS is being used. For the protocol described earlier, only REQ messages are likely to suffer these collisions; with REQ messages, Chloe's THS may be used by both Jack and Bill at the same time; for RACK, DATA, ACK

and BEACON transmissions, unique THSs are used. We denote the probability of such a collision as P_{MAC}.

2) The packet does not collide in the above case; however, it is partially corrupted, due to high multi-access interference levels (*MAI*). This partial corruption may occur because: *(a)* individual pulse transmissions may collide due to partial THS overlaps, and *(b)* there is a high interference energy due to multipath components. We argue that the impact due to (a) is negligible (later in the section) and denote the probability of such a partial corruption due to (b) by P_{MAI}. In particular:

$$P_{failure} = P_{MAC} + \ (1 - P_{MAC}) \cdot P_{MAI} \tag{5}$$

In what follows, we derive P_{MAC} and P_{MAI}.

Step 2: Computing the collision probability P_{MAC}: Recall that with our simplistic MAC protocol, P_{MAC} can be non-zero for REQ transmissions only. Let us denote by P_{REQ}, $P_{BackOff}$, P_{IDLE} and P_{DATA}, the probabilities of a node being in the *REQ, BackOff, IDLE* and *DATA* states, respectively. The transition from one state to the other is a Markov process [15]; thus the probability P_{REQ} that a node resides in the REQ state depends on the expected time E_t in each state, i.e.,:

$$P_{REQ}=\frac{E_t[REQ]}{E_t[REQ]+E_t[BackOff]+E_t[DATA]+E_t[IDLE]}. \tag{6}$$

Let us again assume that Jack wants to send a REQ to Chloe (Fig. 2). If at least one more transmitter (besides Jack) sends a REQ to Chloe in parallel with Jack, both REQs will collide. In this case, P_{MAC} is equal to the probability that at least one other of Chloe's neighbors transmits a REQ to Chloe, at the same time as Jack. Clearly, this is heavily dependent on the *transmitter* density around Chloe. The probability that there are K *nodes* within Chloe's sensing vicinity is:

$$P_{range}(K) = \binom{N}{K} \cdot (\frac{\pi R^2}{A})^K \cdot (1 - \frac{\pi R^2}{A})^{N-K}, \tag{7}$$

assuming a radial coverage of range R. Let Chloe have M active neighbor transmitters, including Jack. The probability P_{Chloe}, that none of the other $M-1$ transmitters picks Chloe, equals:

$$P_{Chloe} = \sum_{i=0}^{M-1} \binom{M-1}{i} (P_{REQ})^i (1 - P_{REQ})^{M-1-i} P_{notChloe}(i),$$

where $P_{notChloe}(i)$ is the probability that the specific i transmitters *do not* pick Chloe to send a REQ at this point in time. To compute $P_{notChloe}(i)$, we consider the probability P_ϑ that the ϑ_{th} neighbor transmitter ($\vartheta \in i$) will not pick Chloe, given that ϑ has l_ϑ neighbor potential receivers (Fig. 2). With this,

$$P_{notChloe}(i) = \prod_{\vartheta=0}^{i} \left[\sum_{l_\vartheta=1}^{N-1} \left((1 - \frac{1}{l_\vartheta}) \cdot P_{range}(l_\vartheta). \right) \right] \tag{8}$$

where $P_{range}(l_\vartheta)$ is given by Eq. 7. We now define $P_{tx} = P_{REQ} + P_{DATA}$, where P_{DATA} is the probability that a node is in the *DATA* state, and is computed in a manner similar to the calculation of P_{REQ} in Eq. 6, with $E_t[DATA]$ in the numerator, instead of $E_t[REQ]$. The probability that the REQ from Jack to Chloe does not collide with another REQ towards Chloe, is equal to P_{Chloe} times $P_{TX}(M-1)$, the probability that Chloe actually has M-1 neighbor transmitters:

$$P_{TX}(M-1)=\sum_{n=M-1}^{N} \binom{n}{M-1} P_{tx}^{M-1} (1 - P_{tx})^{n-(M-1)} P_{range}(n). \tag{9}$$

As a consequence, for all possible values of M:

$$P_{MAC} = 1 - \sum_{M=2}^{N-1} P_{Chloe} \cdot P_{TX}(M-1). \tag{10}$$

We now have an expression for P_{MAC}, which depends on the expected time of residing in each state, as per Eq. 6. We still need to calculate the constituent terms of Eq. 6. $E_t[DATA]$ is easy to compute; it corresponds to the expected number of chips needed for the DATA, RACK, ACK and BEACON transmissions. This is equivalent to:

$$E_t[DATA] = N_f \cdot PB \cdot (DATA + RACK + ACK + BEACON) \tag{11}$$

in number of chips[3].

We derive $E_t[REQ]$, as follows. If Jack's REQ transmission fails, for each subsequent REQ retransmission for that particular request, Jack will reside in the *REQ* state. The probability of having exactly r successive REQ failures is:

$$P_{r\ failures}(REQ) = (P_{failure}(REQ))^r (1 - P_{failure}(REQ)).$$

We wish to point out here that $P_{failure}(REQ)$ will ultimately be computed recursively by Eq. 5. From the above:

$$E_t[REQ] = \sum_{r=0}^{\zeta} (r+1) N_f \cdot PB \cdot REQ \cdot P_{r\ failures}(REQ). \tag{12}$$

With regards to computing $E_t[BackOff]$, the backoff contention window for the first ξ attempts follows the discrete uniform distribution in the value field $(0, 2^x - 1)$, with $x = 3, 4, ..., 8$. The average value for the contention window is $\frac{(2^x-1)+1}{2} = \frac{2^x}{2}$. Hence:

$$E_t[r\ backoffs] = \begin{cases} \sum_{b=1}^{r} \frac{2^{b+2}}{2}, & r \leq \xi \\ \sum_{b=1}^{6} \frac{2^{b+2}}{2} + \sum_{b=7}^{r} \frac{2^8}{2}, & \xi < r \leq \zeta \end{cases}$$

$$\text{and } E_t[BackOff] = \sum_{b=1}^{r} \left(E_t[r\ backoffs] \cdot P_{failure}(REQ) \right).$$

Finally to compute $E_t[IDLE]$, recall that the receivers for potential REQ messages are in the *IDLE* state. Thus, nodes should stay in the *IDLE* state for at least $E_t[REQ]$. Moreover, if Jack's REQ towards Chloe collides, she should wait long enough for Jack to come back from the *BackOff* to the *REQ* state, at least to try to transmit the request again. Therefore, we consider $E_t[IDLE]$ to be equal to $2 \cdot E_t[REQ] + E_t[BackOff]$, plus a uniformly distributed random interval with values between 0 and $E_t[REQ]$. In particular:

$$E_t[IDLE] = 2 \cdot E_t[REQ] + E_t[BackOff] + \frac{E_t[REQ]+1}{2}$$

The decision on the duration of this random interval depends entirely on the protocol designer; here we employ the specific strategy for completeness (other strategies could be easily incorporated).

At this stage, we have derived the expected time for being in each state, we may recursively compute P_{REQ} from Eq. 6,

[3]We make an implicit assumption that the various packets are of fixed size. If they vary, then an average value will have to be considered.

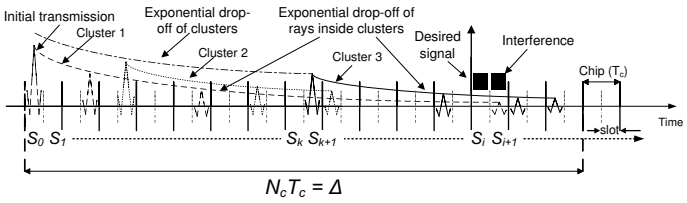

Fig. 3. The multipath delay profile, due to a pulse transmission in slot S_0. The pulse reception is expected in slot S_i. Rays appearing in S_{i+1} may cause interference to the reception of the pulse in the neighboring slot S_i.

and P_{MAC} from Eq. 10.

Step 3: Computing P_{MAI}: Unlike P_{MAC} which can be only non-zero for REQ transmissions, P_{MAI} can be non-zero for any transmission. Let us consider the scenario in Fig. 3. Using TH-BPPM, we assume that the desired signal is transmitted in slot S_i. Note here that slots S_i and S_{i+1} belong to the same chip of interest. In this case, the interference is due to the presence of energy in slot S_{i+1}. A correlator function at the receiver will compare the energies (plus noise) in the two slots and will decide on whether a *"0"* or a *"1"* was transmitted. For mathematical tractability we consider a simplified analysis of a successful reception of a bit, ignoring the receiver functionality and additive Gaussian noise. We assume that the bit is correctly detected if the power in slot S_i is greater than the power in slot S_{i+1} by g dB. As shown in Fig. 3, the desired power (that in slot S_i) will be equal to the received power due to the desired pulse plus the sum of powers of all rays appearing within slot S_i. Moreover, the interference power will be equal to the sum of powers of all rays that arrive during S_{i+1}.

The l_{th} cluster starts at T_l (assuming the first cluster starts at time $t = 0$ – slot S_0 in our scenario). The ν_{th} ray in this cluster is $\tau_{\nu,l}$ away from T_l. Given the short range of UWB signals, we assume that pathloss is the same for all the received rays irrespective of the transmitter; the pulse energy $p^2(t - \tau_{\nu,l})$ is fixed. With this, one can ignore the effects of these factors on the detection probability. In addition, the power (dB) in the ν_{th} ray in the l_{th} cluster is normally distributed with mean $\mu_{\nu,l}$ (as per Eq. 4) and standard deviation σ_{ch} (a constant for the channel). The channel coefficients themselves have a lognormal distribution. When we sum the power from different pulses, we have to add lognormal random variables (as described later in this section).

Given a pulse transmission in slot S_0, there is a cluster that starts in slot S_l with probability $P_{cluster}(S_l|S_0)$. This probability depends on the probability that a previous cluster started in any of the slots prior to slot S_l, and is recursively given by:

$$P_{cluster}(S_l|S_0) = \sum_{n=0}^{l-1} \Lambda e^{-\Lambda(l-n)\frac{T_c}{2}} P_{cluster}(S_n|S_0) \qquad (13)$$

Given that a cluster starts in slot S_l, a ray appears in slot S_i with probability $P_{ray}(S_i|S_l)$. In order to find the probability that a ray appears in slot S_i, we need to find the probability that the previous ray (caused by the same pulse transmission in slot S_0 resulting in a cluster starting in S_l) appeared in any slot in the past, between slots S_l and S_{i-1}. The first ray (corresponding to the clusterhead) appears in slot S_l. Hence, the probability that a ray appears in slot S_i given that the clusterhead occurred in slot S_l is:

$$P_{ray}(S_i|S_l) = \sum_{j=l}^{i-1} \lambda e^{-\lambda(i-j-1)\frac{T_c}{2}} P_{ray}(S_j|S_l) \qquad (14)$$

The power in such a ray is lognormally distributed with the associated normal distribution (in dB) having a mean $\mu_{i-l,l}$ (as per Eq. 4) and standard deviation σ_{ch}. Similarly, the probability that a ray occurs in slot S_{i+1} (i.e., $P_{ray}(S_{i+1}|S_l)$) can be computed.

The probability that there is a ray in S_i, but not in S_{i+1}, given that a cluster started in S_l is:

$$P_A = (1 - \lambda e^{-\lambda\frac{T_c}{2}}) P_{ray}(S_i|S_l)$$

In this case, the (constructive) interference is lognormal with associated normal parameters $\mu_{i-l,l}$ and σ_{ch}. The probability that there is no ray in S_i, but a ray appears in $S_i + 1$ is:

$$P_B = \sum_{j=l}^{i-1} \lambda e^{-\lambda(i-j)\frac{T_c}{2}} P_{ray}(S_j|S_l)\,(1 - P_{ray}(S_i|S_l))$$

In this case, the (destructive) interference is again lognormal with associated normal parameters $\mu_{i-l+1,l}$ and σ_{ch}. Furthermore, the probability that there are rays in both slots S_i and S_{i+1} will be:

$$P_C = \lambda e^{-\lambda\frac{T_c}{2}} P_{ray}(S_i|S_l)$$

The interference in this case is the difference of two lognormal random variables U (with associated normal parameters $\mu_{i-l,l}$ and σ_{ch}) and V (with associated normal parameters $\mu_{i-l+1,l}$ and σ_{ch}) i.e., the interference is $W = U - V$. We discuss the distribution of W below. It is also possible that there are no rays in either of the slots S_i and S_{i+1}, however that situation does not create any interference.

The distribution of W can be approximated as lognormal with associated normal parameters μ_W and σ_W. There is no closed form solution for the distribution of W. However, several approximations have been suggested for sums of lognormal random variables. We use the *Wilkinson's* approach, due to its simplicity [16]. Even though methods with greater accuracy exist, the Wilkinson's approach has been found to provide reasonable cumulative distribution functions for the sum of lognormal random variables [17]. The approach is summarized as follows. Let $I = \sum_{n=1}^{F} I_n$ be the sum of F lognormal random variables I_n. Then, I is also lognormal with the associated normal distribution having a mean μ and standard deviation σ. The μ and σ can be approximated in the following way:

- Let $X_i = 10\log_{10} I_i$ be normal with mean μ_{X_i} and standard deviation σ_{X_i}. Let $Y_i = cX_i$ have mean μ_{Y_i} and standard deviation σ_{Y_i} where $c = \ln(10)/10$.
- Let $X = 10\log_{10} I$ be the normal distribution associated with the sum I of the F lognormal random variables. Then the mean and standard deviation of I are:

  $$\mu_X = \frac{1}{c}\left(2\ln u_1 - 0.5\ln u_2\right);\ \ \sigma_X = \frac{1}{c}\sqrt{\ln u_2 - 2\ln u_1}$$

  where $u_1 = \sum_{n=1}^{F} e^{\mu_{Y_n} + \sigma^2_{Y_n}/2}$, and

  $$u_2 = \sum_{n=1}^{F} e^{2\mu_{Y_n} + 2\sigma^2_{Y_n}} + 2\sum_{i=1}^{F-1}\sum_{j=i+1}^{F} e^{\mu_{Y_i} + \mu_{Y_j}} \cdot e^{0.5(\sigma^2_{Y_i} + \sigma^2_{Y_j})}$$

  assuming that the random variables Y_n are uncorrelated.
- We denote the Wilkinson's sum of the F variables as:

  $$(\mu_X, \sigma_X) = \bigoplus_{i=1}^{F} (\mu_{X_i}, \sigma_{X_i})$$

These equations can be modified to account for the difference between two lognormal random variables.

The "average" parameters for the lognormal interference in slot S_i, conditioned on the clusterhead being in S_l can now be expressed as

$$\mu_{av|l} = P_A \cdot \mu_{i-l,l} - P_B \cdot \mu_{i-l+1,l} + P_C \cdot \mu_W$$

$$\sigma_{av|l} = P_A \cdot \sigma_{ch} - P_B \cdot \sigma_{ch} + P_C \cdot \sigma_W$$

which is what we use to compute the successful reception of a pulse, described next. Note that these average parameters are a function of the slot S_l.

As shown in Fig. 4, let $\Delta = KT_f + LT_c + MT_c$, so that for any *one* interferer, there are K or $K-1$ frames contained completely within Δ (depending on the position of slot S_i, L chips of a partial frame before and M chips of a partial frame after the block of full frames. In addition, $(L+M)T_c < T_f$. The frames are numbered 0 through $K+1$, and S_i, S_{i+1} belong to frame 0, as shown in Fig. 4. Consider that an active transmitter has transmissions in each of the full frames and (possibly) the two partial frames. If we assume that a transmission in a chip is equally likely (due to the randomness of the overlapping time hopping sequences), the probability of transmission that could cause interference in slot S_i in frame 0 is M/N_f and that in frame $K+1$ is L/N_f. In all full frames, this probability is 1.

Case for full frames: Consider the transmission of a pulse in some frame u, where $u \in [1, K]$. The pulse transmission occurs in slot S_0 within frame u, such that the time span between S_0 and S_i is given by: $t_u = MT_c + (u-1)T_f + xT_c/2$, where $x \in [1, 2N_f)$. Thus, there are $\delta = 2M + 2(u-1)N_f + x$ slots between S_0 and S_i including S_0. Given that a pulse transmission occurred in S_0, a clusterhead may appear in any slot $l \in [0, \delta]$, with probability $P_{cluster}(S_l|S_0)$ as per Eq. 13. Since the cluster can begin in any slot between S_i and S_0, the average values of the parameters of the lognormal distribution of the interference over all possible clusterhead locations are:

$$\mu(S_0) = \sum_{l=0}^{\delta-1} \left(P_{cluster}(l|S_0)\mu_{av|l}\right),\ \sigma(S_0) = \sum_{l=0}^{\delta-1} \left(P_{cluster}(l|S_0)\sigma_{av|l}\right). \tag{15}$$

Notice that $\mu(S_0)$ and $\sigma(S_0)$ are functions of δ and hence, functions of x. Thus, the average parameters of the lognormally distributed ray power over all possible slots in frame u are:

$$\mu(u) = \frac{1}{2N_f} \sum_{x=1}^{2N_f} \mu(S_0),\ \sigma(u) = \frac{1}{2N_f} \sum_{x=1}^{2N_f} \sigma(S_0). \tag{16}$$

There is a similar contribution from transmissions in each of the full frames within Δ. The total interference from these transmissions will have a power that is lognormally distributed with average parameters $\mu_{av_{full}}$ and $\sigma_{av_{full}}$. These are computed using the Wilkinson's approach described earlier:

$$(\mu_{av_{full}}, \sigma_{av_{full}}) = \bigoplus_{\text{Full Frames}} (\mu(u), \sigma(u)).$$

Case for partial frames: We initially assume that the values for M and L are known. We consider the case for M first. The pulse transmission occurs in some slot S_0, in frame 0, such that the time span between slots S_0 and S_i is given by: $t_u = xT_c/2$, where $x \in [i, 2M]$. Thus, there are $\delta = x$ slots between S_0 and S_i, including S_0. Here, $\mu(S_0)$ and $\sigma(S_0)$ are computed to be (like in Eq. 16):

$$\mu_M(u) = \frac{1}{2M} \sum_{x=1}^{2M} \mu(S_0),\ \text{and } \sigma_M^2(u) = \frac{1}{2M} \sum_{x=1}^{2M} \sigma(S_0)$$

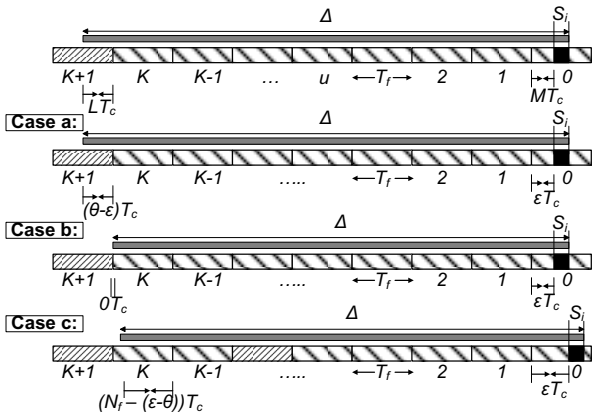

Fig. 4. Handling the cases for which Δ covers partially sequence frames.

For the case of L, we have: $t_u = MT_c + KT_f + xT_c/2$, where $x \in [i, 2L]$ and $\delta = 2M + 2K \cdot N_f + x$. Moreover:

$$\mu_L(u) = \frac{1}{2L} \sum_{x=1}^{2L} \mu(S_0),\ \text{and } \sigma_L^2(u) = \frac{1}{2L} \sum_{x=1}^{2L} \sigma(S_0).$$

We set $\theta = L + M$. Thus $\Delta = KT_f + \theta T_c$. We assume there are ε chips between the start of the "0"th frame and S_i. Depending on ε, there are three possible cases.

Case a: From Fig. 4 it is clear that $M = \varepsilon$ and $L = \theta - \varepsilon$, where $\varepsilon \in [1, \theta - 1]$ and there are K full frames in Δ. Considering a fixed value of ε, the expected parameters for the lognormally distributed power in slot S_i can be computed using the Wilkinson's approach:

$$(\mu_{av_1}, \sigma_{av_1}) = \bigoplus (\mu_{av_{full}}, \mu_M(u), \mu_L(u), \sigma_{av_{full}}, \sigma_M(u), \sigma_L(u)).$$

The above notation indicates that the Wilkinson's approach is used to consider the sum of K lognormal random variables from the full frames and the two lognormal random variables from the partial frames, to compute the associated normal parameters. These parameters depend upon the value of ε. Removing the condition on ε:

$$(\mu_{av_1}, \sigma_{av_1}) = \frac{1}{\theta - 1} \sum_{\varepsilon=1}^{\theta-1} \bigoplus (\mu_{av_{full}}, \mu_M(u), \mu_L(u), \sigma_{av_{full}}, \sigma_M(u), \sigma_L(u)).$$

Case b: As shown in Fig. 4, in this case $M = \varepsilon = \theta$ and $L = 0$ (no partial frames before the block of K full frames). In this case:

$$(\mu_{av_2}, \sigma_{av_2}) = \frac{\theta}{N_f} \bigoplus (\mu_{av_{full}}, \mu_M(u), \sigma_{av_{full}}, \sigma_M(u))$$

Here, the K full frames and the partial frame containing S_i are considered in the Wilkinson's sum.

Case c: In this case the slot S_i is more than θ slots away from the start of frame "0" (i.e., $\varepsilon > \theta$). We will have $K-1$ full frames and two partial frames where $M = \varepsilon$ and $L = N_f - (\varepsilon - \theta)$ as shown in Fig. 4. In this case:

$$(\mu_{av_3}, \sigma_{av_3}) = \frac{1}{N_f - \theta} \sum_{\varepsilon=\theta+1}^{N_f} \bigoplus (\mu_{av_{full}}, \mu_M(u), \mu_L(u), \sigma_{av_{full}}, \sigma_M(u), \sigma_L(u)),$$

where the Wilkinson's sum is computed over $K-1$ full frames and the partial frames with L and M chips respectively

(averaged over the appropriate slots as before). Considering all three cases now, the total average of the normal parameters of the lognormally distributed power at S_i are:

$$\mu_{avg} = \mu_{av_1} \cdot \frac{\theta - 1}{N_f} + \mu_{av_2} \cdot \frac{1}{N_f} + \mu_{av_3} \cdot \frac{N_f - \theta}{N_f}$$

$$\text{and } \sigma^2_{avg} = \sigma^2_{av_1} \cdot \frac{\theta - 1}{N_f} + \sigma^2_{av_2} \cdot \frac{1}{N_f} + \sigma^2_{av_3} \cdot \frac{N_f - \theta}{N_f}.$$

Recall that the above calculation involves only one interferer in Chloe's neighborhood. Considering all possible cases of multiple active neighbor transmitters we have:

$$(\mu_{avg_{total}}, \sigma_{avg_{total}}) = \bigoplus_{n=1}^{N-1} \left(\mu_{avg} \cdot P_{TX}(n), \sigma^2_{avg} \cdot P_{TX}(n)\right). \quad (17)$$

Again the Wilkinson's sum of the interfering powers is taken; $P_{TX}(n)$ (given by Eq. 9) is the probability that n transmitters contribute to the interference.

Impact of Partial Overlaps: Due to partial overlaps of THSs, there may be an interfering transmission in the chip of interest. This interfering transmission could be in either slot S_i or in slot S_{i+1}, with equal probability. First, the two cases where this interfering transmission occurs in one of the slots and there is no ray in the other slot, tend to compensate for each other, *on average*, in terms of the interference powers. In the last case, the transmission occurs in slot S_i and a ray due to this transmission appears in slot S_{i+1}. In this case, due to the proximity of the two signals in time, one may expect the powers to be similar and thus, they together contribute little to the interference power (these effects were implicitly validated via simulations described later).

Deriving the probability of a pulse failure: In order for a pulse to be successfully deciphered, the desired power plus the (constructive) interference must be greater than g dB. This constructive interference is the relative interference power in slot S_i and was computed as the difference in interference power in slots S_i and S_{i+1}. The analysis thus far, assumed that the intended transmission was in slot S_i; if on the other hand, the intended transmission was in slot S_{i+1}, the computed interference acts as *destructive* interference (since it is the relative power in slot S_i). Let μ_z and σ_z be the mean and standard deviation of the normal distribution associated with the lognormal decision variable Z, if the intended transmission is in slot S_i. These values can be obtained by first considering the Wilkinson's sum of the power in the desired pulse (P_r) and the lognormal interference. If the intended transmission is in slot S_{i+1}, the mean and the standard deviation μ'_z and σ'_z are computed by considering the difference in the desired pulse (P_r) and the lognormal interference computed earlier, using the Wilkinson's approach. Once the corresponding normal parameters in dB are obtained, the value g can be subtracted from the resulting mean. We are interested in the probability that $Z > 0$. Accounting for the fact that the intended transmission could be either in slot S_i or in slot S_{i+1} with equal probability, this desired probability is computed as:

$$P_{pulse} = P(Z > 0) = \frac{1}{4} erfc(-\frac{\mu_z}{\sigma_z\sqrt{2}}) + \frac{1}{4} erfc(-\frac{\mu'_z}{\sigma'_z\sqrt{2}}) \quad (18)$$

The above expression is the probability that one pulse is received successfully in the presence of *MAI*. Since PB pulses make up one bit, the probability that a bit is correctly decoded is given by:

$$P(bit\ success) = 1 - (1 - P_{pulse})^{PB}. \quad (19)$$

Finally the probability of packet corruption due to MAI is:

$$P_{MAI} = 1 - [P(bit\ success)]^{Bits/packet} \quad (20)$$

Having computed both P_{MAC} and P_{MAI}, we are able to calculate the probability of packet failure, as per Eq. 5. However, a packet may be lost due to the corruption of either of the REQ, RACK, DATA or ACK transmissions. Specifically:

$$P_{failure}(RACK) = (1 - P_{failure}(REQ))P_{MAI}(RACK) \text{ and,}$$

$$P_{failure}(DATA) = \{(1 - P_{failure}(REQ)) \cdot (1 - P_{failure}(RACK))P_{MAI}(DATA)\} \text{ and,}$$

$$P_{failure}(ACK) = \{(1 - P_{failure}(REQ)) \cdot (1 - P_{failure}(RACK))(1 - P_{failure}(DATA))P_{MAI}(ACK)\}.$$

$$\text{Finally: } P_{failure} = P_{failure}(REQ) + P_{failure}(RACK) + P_{failure}(DATA) + P_{failure}(ACK).$$

Next, we compute the per node saturation throughput. This throughput is defined to be the product of (i) the fraction of time for which a node is able to send useful pulses and (ii) the raw data rate; the throughput is expressed in units of bits/sec. The duration during which useful data is transferred is simply $(1 - P_{failure}(DATA)).\{E_t[DATA].(DATA/[DATA + RACK + ACK + BEACON])\}$. However, the transmitter is able to send a pulse, only once every sequence frame; therefore, this time has to be scaled down by a factor N_f. If the raw data rate of the channel is $Rate$, the achievable rate is scaled down by a factor PB due to the use of repetition codes. Thus, using Eq. 11,

$$Tput = \frac{DATA \cdot (1 - P_{failure}) \cdot Rate}{E_t[REQ] + E_t[DATA] + E_t[BackOff] + E_t[IDLE]}. \quad (21)$$

IV. Evaluating our Framework

In this section, we evaluate the accuracy of CTU and discuss how CTU can guide MAC protocol design decisions.

Validating the accuracy of CTU: We compare the analytical results produced by CTU with the simulated behavior of the network. The considered input parameters are summarized in Table I. In more detail: **(1)** We perform an extensive set of simulations of the network, using a C++ simulation platform [1]. We have incorporated the lognormal fading distribution in the SV model for the various pulse transmissions, as per [5]. Our modifications to the simulator consist of: *(a)* generating the Poisson-based arrival of clusters and rays, for a period Δ (N_c chips) after each pulse transmission (Eq. 1, 2), and *(b)* the impulse response of transmitted pulses (as per Eq. 3). **(2)** We implement CTU in both C++ and Matlab, and we use it to compute the probability of packet transmission failure. We then estimate the average MAC layer throughput per node, as per Eq. 21. The Stirling's approximation [18] was used to compute factorials and this considerably speeds up the computations. The set of all our analytical results are obtained in less than one hour, while the simulations take days. We compare the results produced by CTU and the simulator, in terms of per

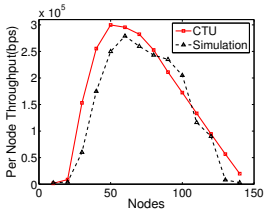

Fig. 5. Average throughput per node, for $N_f = 10$ and $PB = 7$.

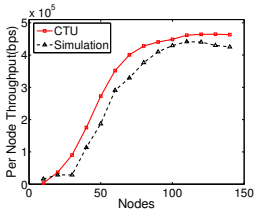

Fig. 6. Average throughput per node, for $N_f = 20$ and $PB = 7$.

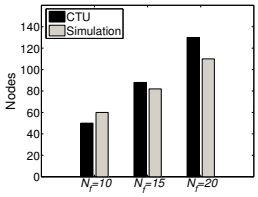

Fig. 7. CTU accurately estimates the supported number of users.

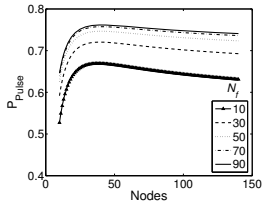

Fig. 8. The probability of pulse success increases with N_f.

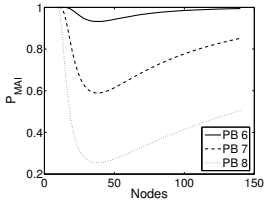

Fig. 9. Larger values of PB reduce interference.

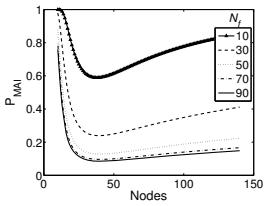

Fig. 10. The value of P_{MAI} reduces with N_f.

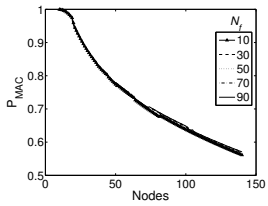

Fig. 11. The value of P_{MAC} is not significantly affected by N_f.

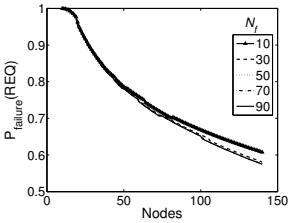

Fig. 12. $P_{failure}(REQ)$ depends on MAC operations.

Parameter	Value
Tc	1 nsec
A	80×80 m^2
R	11 m
N_f	10
PB	7
ξ	6
J	256
Bandwidth	7.5 GHz
DATA	2048 bits
ACK	120 bits
REQ	120 bits
RACK	120 bits

Parameter	Value
Γ	14.93
γ	7.03
Λ	0.0667
λ	3
P_m	5.59 nW
ζ	15
Rate	110 Mbps
σ	4.8 dB
g	10 dB
N_c	120
BEACON	32 bits
Sim. duration	10^8 slots

TABLE I
SIMULATION PARAMETERS, WHEN NOT VARIABLE

node throughput, on average, for different values of N_f. Figs. 5 and 6 depict this comparison, for the cases with $N_f = 10$ and 20, respectively, and for $PB = 7$. We observe that the results from CTU are very close to those from the simulator; the throughputs corresponding to CTU are shifted only by a few nodes with regards to the throughputs computed by the simulator (Fig. 7). The discrepancy in the results are due to: (a) the approximations from the Wilkinson's approach and (b) while a constant pathloss of 20 dB is assumed in the analysis, the Friis path loss model [19] is used in the simulations.

Interpreting the results: Next, we use CTU to examine the impact of system and environmental factors on the performance.

On the impact of system parameters on MAI: We first examine the impact of system parameters such as N (the number of nodes in the network), N_f and PB on the MAI. Fig. 8 depicts the probability of a pulse success versus N. Interestingly, as N initially increases, the probability of pulse success also increases. This contradicts the common intuition that fewer pulses are likely to be successful as N increases. This is a direct consequence of the impact of the number of interfering users in the receiver's (Chloe's) vicinity and in particular Eq. 17. The *MAI*, as seen from Eq. 17, is dependent on the number of active interfering transmitters. When N is small, the number of nodes in Chloe's range is even smaller. As the number of neighbors increases initially, the transmitters will compete among each other for the few available receivers and will often end up in the *BackOff* state. Thus, the number of active transmitters actually reduces and as a result, the *MAI* reduces as well. Note that this happens with few nodes ($N < 20$). However, beyond a certain density, the nodes are able to find receivers and thus, are able to stay as "active" transmitters. From this point on, the *MAI* begins to increase again. The point at which the transition occurs depends on the area of deployment (A), the transmission range, and N. As one might expect, this point is somewhat independent of the value of N_f; the decrease is due to the inability of transmitters to find receivers and not because of collisions due to overlaps between time-hopping sequences. Note here that this results due to the MAC protocol design choice of selecting receivers at random. It is possible that a cleverer receiver selection strategy, adopted by another potential protocol, will not present such a behavior.

Notice that the value of N_f clearly determines the extent of interference. The larger the value of N_f, the smaller is the amount of destructive interference experienced by receivers. In addition to N_f, the chosen PR value impacts the achieved performance. In Fig. 9, we plot the reduction in P_{MAI}, as the value of PR increases. As one might expect, by increasing the redundancy in the number of transmitted pulses per bit, the negative impact of interference is decreased and hence the probability of packet failure due to interference is reduced. Note here that extremely high N_f values impact the network throughput, due to significant delays between pulse transmissions [9]. Finally, as one might expect P_{MAI} varies conversely with P_{pulse}, as depicted in Fig. 10.

On the behavior of the MAC layer: Next, we examine the MAC layer success probability P_{MAC} with the simple MAC procedure considered (Fig. 11). With small N, transmitters have a very limited set of potential receivers to choose from. This causes a significant number of REQ collisions; therefore P_{MAC} initially has a very high value. As N increases, more and more potential receivers become available; this reduces the probability that a receiver is selected by more than one transmitter. The level of interference affects the value of P_{MAC} (since REQ messages may suffer interference as seen in Eq. 12). From Figs. 10 and 11, we observe that as N increases, P_{MAI} increases as well, and *drives* P_{MAC} higher. Thus, P_{MAC} does not drop sharply (but gradually) with the increased N. Since the control messages (REQ, RACK, ACK, BEACON) are of short length, the impact

of *MAI* on them is less severe. This is shown in Fig. 12. The probability of REQ transmission failure is iteratively computed by Eq. 5 and 12 and hence is dependent on P_{MAI} for the REQ message. Notice that for the different values of N_f, the values of P_{MAC} and $P_{failure}(REQ)$ change.

The combined behavior of P_{MAC} and P_{MAI} determines the average achievable MAC layer throughput, as per Eq. 21. As shown in Figs. 5 and 6, for very small N, the average per node throughput is very low, since P_{MAC} and P_{MAI} are high. For moderate populations (40 to 70 nodes) the average throughput is high; however as N increases further, the throughput decreases due to increased P_{MAI}.

V. Scope of CTU

In this section, we discuss the generality of CTU and deliberate on its limitations.

CTU is a generic tunable framework: Our description of CTU in section III and its evaluation was based on a simplified MAC protocol and a specific (although popularly considered) PHY layer. However, we argue here that the modular nature of CTU allows one to tune it to evaluate a wide class of different TH-BPPM based MAC protocols and PHY models that are similar to the one considered. In particular:

1) The state diagram in Fig. 1 can be modified to reflect the states of nodes in the new protocol. The expected times in each state can be tuned as appropriate to the specific MAC protocol under consideration, in order to derive P_{MAC} as per Eq. 10. Consider two specific examples. In the first, let there be no request handshake i.e., nodes simply transmit data as per the receiver's THS. It is easy to see that now, the transition probabilities into the REQ and BACKOFF states are zero. A packet collision is computed similar to the REQ collision. As a second example, consider a case where the packet sizes vary and the back-off procedure is different from the one considered. These features can again be easily incorporated into CTU by appropriately modifying the specific portions of the framework.

2) While developing CTU, we assumed that nodes have a radial range R. However, note that this assumption is primarily used to (a) fix a pathloss value and (b) determine the number of possible interferers. One can easily incorporate a different range distribution in the CTU framework. The only change required will be in Eq. 7.

3) One may incorporate a different (than the SV) channel model, as long as the model describes the ray inter-arrival times and gains. For example, adopting the Nakagami fading distribution instead of the lognormal distributions [5] is easy, due to the similarities in Eq. 4. A more accurate pathloss model can also be integrated when computing the probability of successful reception of a pulse in Eq. 18; however, one would have to compute the average interference, conditional on the channel attenuations to the different users.

4) CTU has been currently designed for computing the throughput under saturation conditions. If the user traffic characteristics are known, it may be possible to extend the analysis to include the throughput under more general conditions. As an example, if the users generate traffic as per a Poisson process, it is easy to compute the service time for a packet given its probability of success; each user may then be modeled as an M/G/1 queue.

Limitations of CTU: The CTU framework however, has the following limitations. **(a)** CTU does not account for the use of convolutional or block codes. Incorporating the effects of coding will increase the complexity of CTU. **(b)** We have based our analytical framework on TH-BPPM based communications. Thus, other multiple access strategies (such as M-PPM, CDMA or OFDMA) are not applicable. As an alternative transmission technique, the Wimedia Alliance supports an OFDM (Orthogonal Frequency Division Multiplexing) based specification [20] for UWB.

VI. Conclusion

We design and evaluate CTU, an analytical framework for MAC layer throughput estimation in impulse-based multihop UWB networks. The modularity of CTU, enables one to easily tune it to model a wide range of MAC protocols; the only requirements are that the underlying modulation scheme be BPPM and time-hopping be used as the multi-access procedure. CTU is the first analytical framework that captures the cross layer interactions between the PHY and the MAC layers. We validate the accuracy of CTU with extensive simulations. We also use CTU to analyze a simple MAC procedure; the analysis yields insights on several interesting properties of UWB networks.

References

[1] I. Broustis, S. Krishnamurthy, M. Faloutsos, M. Molle, and J. Foerster. Multiband Media Access Control in Impulse-Based UWB Ad hoc Networks. In *IEEE Trans. on Mobile Comp., Vol. 6, No. 4*, Apr. 2007.

[2] A. Molisch, J. R. Foerster, and M. Pendergrass. Channel Models for Ultra-wideband Personal Area Networks. In *IEEE Wireless Comm. Mag., Vol. 10, No. 6*, Dec. 2003.

[3] M. Z. Win and R. A. Scholtz. Impulse Radio: How It Works. In *IEEE Comm. Letters, Vol.2, No.1*, Jan. 1998.

[4] A. Saleh and R. Valenzuela. A statistical model for indoor multipath propagation. In *IEEE JSAC, Vol. SAC-5, No. 2*, Feb. 1987.

[5] J. R. Foerster. Path Loss Proposed Text and S-V Model Information. In *Intel R&D document, IEEE P802.15-02/xxxr0-SG3a*, Sept. 2002.

[6] A. Gupta and P. Mohapatra. A Survey on Ultra Wide Band Medium Access Control Schemes. In *Tech. Report CSE-2006-4, UC Davis*, 2006.

[7] R. Merz and J. Y. Le Boudec. Conditional Bit Error Rate for an Impulse Radio UWB Channel with Interfering Users. In *IEEE ICU*, 2005.

[8] J. A. Gubner and K. Hao. A Computable Formula for the Average Bit Error Probability as a Function of Window Size for the IEEE 802.1.3a UWB Channel Model. In *IEEE Trans. Micr. Theory, Vol. 54, No. 4*, 2006.

[9] I. Broustis, A. Vlavianos, and S. V. Krishnamurthy. On the MAC Layer Performance of Time-Hopped UWB Ad Hoc Networks. In *ICCCN*, 2006.

[10] D. Cassioli, M. Z. Win, and A. Molisch. The Ultra-wide Bandwidth Indoor Channel: From Statistical Model to Simulations. In *IEEE JSAC, Vol. 20, No. 6*, Aug. 2002.

[11] S. S. Ghassemzadeh et al. UWB Indoor Delay Profile Model for Residential and Commercial Environments. In *IEEE VTC*, 2003.

[12] L. J. Greenstein et al. Comparison Study of UWB Indoor Channel Models. In *IEEE Trans. of Wireless Comm., Vol. 6, Issue 1*, Jan. 2007.

[13] I. Broustis, M. Molle, S. Krishnamurthy, M. Faloutsos, and J. Foerster. A New Binary Conflict Resolution Based MAC protocol for Impulse-Based UWB Ad hoc Networks. In *Wireless Comm. and Mobile Comp., Vol. 6, Issue 7*, Nov. 2006.

[14] 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. In *Broadnets*, 2004.

[15] S. M. Ross. Applied Probability Models With Optimization Applications. Dover Publications, 1992.

[16] P. Cardieri and T. S. Rappaport. Statistics of the Sum of Lognormal Variables in Wireless Communications. In *IEEE VTC*, 2000.

[17] N. Beaulieu, A. A. Abu-Dayya, and P. J. McLane. Estimating the Distribution of a Sum of Independent Lognormal Random Variables. In *IEEE Trans. Comm., Vol. 43, No. 12*, Dec. 1995.

[18] M. Abramowitz and I. Stegun. Handbook of Mathematical Functions. In *http://www.math.sfu.ca/∼cbm/aands*, 2006.

[19] S. Roy, J. R. Foerster, V. S. Somayazulu, and D. G. Leeper. Ultrawideband Radio Design: The Promise of High-Speed, Short-Range Wireless Connectivity. In *Proc. of the IEEE*, Feb. 2004.

[20] WiMedia Alliance. http://www.wimedia.org.