

Enabling Coexistence of Heterogeneous Wireless Systems: Case for ZigBee and WiFi

Xinyu Zhang and Kang G. Shin

Department of Electrical Engineering and Computer Science

The University of Michigan

{xyzhang, kgshin}@eecs.umich.edu

ABSTRACT

The ISM spectrum is becoming increasingly populated by emerging wireless networks. Spectrum sharing among the same network of devices can be arbitrated by MAC protocols (*e.g.*, CSMA), but the coexistence between heterogeneous networks remains a challenge. The disparate power levels, asynchronous time slots, and incompatible PHY layers of heterogeneous networks severely degrade the effectiveness of traditional MAC. In this paper, we propose a new mechanism, called the *Cooperative Busy Tone* (CBT), that enables the reliable coexistence between two such networks, ZigBee and WiFi. CBT allows a separate ZigBee node to schedule a busy tone concurrently with the desired transmission, thereby improving the visibility of ZigBee devices to WiFi. Its core components include a *frequency flip* scheme that prevents the mutual interference between cooperative ZigBee nodes, and a *busy tone scheduler* that minimizes the interference to WiFi, for both CSMA and TDMA packets. To optimize CBT, we establish an analytical framework that relates its key design parameters to performance and cost. Both the analytical and detailed simulation results demonstrate CBT's significant throughput improvement over the legacy ZigBee protocol, with negligible performance loss to WiFi. The results are validated further by implementing CBT on sensor motes and software radios.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design; Network Protocols

General Terms

Design, Performance, Reliability, Theory

Keywords

Coexistence of heterogeneous networks, coexistence of ZigBee (802.15.4) and WiFi (802.11) networks, software-radio-based protocol, modeling and analysis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiHoc'11, May 16–20, 2011, Paris, France.

Copyright 2011 ACM 978-1-4503-0722-2/11/05 ...\$10.00.

1. INTRODUCTION

Spectrum scarcity is known to be a main obstacle to the scaling of wireless network capacity. Spectrum sharing has been advocated as a key remedy for this problem, especially after the successful deployment of WLAN and WPAN devices on an unlicensed band. However, severe performance degradation has been observed when heterogeneous devices share the same frequency band (*e.g.*, WiFi & Bluetooth [13], WiFi & ZigBee [18], WiFi & WiMax [22]). Such a coexistence problem is rooted at their mutual interference due to the lack of coordination. Although most systems incorporate interference avoidance mechanisms, such as listen-before-talk, they are designed to resolve the collision between the same type of networks. These built-in mechanisms become less effective for heterogeneous MAC/PHY protocols/standards, which adopt asynchronous time slots, different scheduling modes (*e.g.*, TDMA vs. CSMA), disparate transmission/interference ranges, and incompatible communication mechanisms. The problem is likely to persist and exacerbate in future, especially within the recently opened-up TV white-space [10] for unlicensed users.

We address a key question related to this trend: *how should heterogeneous wireless MAC/PHY protocols coexist to share spectrum?* We will focus on two such protocols, WiFi (IEEE 802.11) and ZigBee (IEEE 802.15.4), that share the 2.4GHz ISM band. WiFi is typically deployed for pervasive Internet access or medium-scale WLANs, whereas ZigBee targets monitoring and control applications for home, hospital, or enterprise environments [14]. The conflicting coexistence between them has been observed in existing measurement studies [12, 18], and their underlying cause is representative of many other coexisting networks. In particular, ZigBee packets are transmitted with 20dB lower power than WiFi packets, and tend to be invisible to, and often interrupted by, WiFi transmitters. Even when it can be sensed by WiFi, a ZigBee transceiver has a 16× longer response time, and is often preempted by WiFi, when it switches from sensing to transmission, or transmission to reception mode. Besides, ZigBee allows for TDMA mode, which operates without carrier sensing, and may arbitrarily collide with an ongoing WiFi transmission. Therefore, by resolving the coexistence between ZigBee and WiFi, one could naturally extend the solution to other heterogeneous networks facing similar problems.

To meet this goal, we propose a new paradigm, called *Cooperative Busy Tone* (CBT), that enhances the mutual observability between ZigBee and WiFi, thereby improving their coexistence. CBT builds atop the legacy ZigBee MAC,

but allows the clients to cooperatively strengthen their visibility to WiFi. Unlike the traditional CSMA that relies on a data packet as an implicit busy tone, CBT designates a separate node (either a ZigBee client closer to the WiFi transmitter, or a dedicated high-power ZigBee transceiver) as a *signaler* that emits the busy tone. The busy tone harbingers the actual data transmission, and continues throughout the DATA-ACK transmission, so as to prevent WiFi preemption.

An immediate challenge to CBT is: “how to prevent the busy signal from interfering with the data packet?” We introduce an innovative *frequency flip* mechanism that temporarily re-locates the signaler to an orthogonal ZigBee band, but still ensures that the busy tone is perceived by the WiFi transmitter.

There is an additional concern: “how much performance improvement will CBT bring to ZigBee, and what is the cost to WiFi?” We develop an analytical framework that quantifies the network performance. Our analysis reveals that the legacy ZigBee MAC suffers a 11–23% collision rate even when WiFi leaves the channel unused for 80% of time, and suffers an up to 79% collision rate when WiFi becomes saturated. With CBT, the collision rate can be reduced to below 5% under medium to low WiFi interference, and to below 20% under saturated WiFi traffic. The performance can be improved further by tuning the design parameters, such as the start time and duration of the busy tone. Our analysis also shows that for typical low duty-cycle applications, CBT introduces negligible performance degradation to WiFi, as compared to the legacy ZigBee.

The above analytical results are validated via detailed simulation of CBT in ns-2. We have also prototyped CBT based on TinyOS and the GNURadio library [2]. Our experiments on the MicaZ motes and USRP2 [9] software radio platform further corroborate the feasibility and effectiveness of CBT.

The remainder of this paper is organized as follows. Sec. 2 reviews existing studies on the coexistence of heterogeneous wireless networks. Sec. 3 introduces the key components in CBT. Sec. 4 establishes a theoretical framework to analyze the performance of the ZigBee-WiFi network, with and without CBT. Sec. 5 validates CBT with ns-2 simulation and real experiments. Finally, Sec. 6 concludes the paper.

2. RELATED WORK

Coexistence has long been a problem for protocols operating on the ISM band. Industrial associations, such as the ZigBee Alliance [20], demonstrated that ZigBee can coexist well with WiFi in home networks. However, their experiments were conducted under light WiFi traffic conditions. Many empirical studies revealed severe collision when ZigBee coexists with medium to high WiFi traffic [12, 18].

The IEEE 802.15.2 [3] proposed an adaptive frequency hopping (AFH) mechanism to smooth the coexistence among incompatible MAC/PHY protocols, such as Bluetooth/ZigBee and WiFi. However, AFH is ineffective at WiFi hotspots where the entire 2.4GHz spectrum is congested by multiple WLAN cells configured to orthogonal channels. AFH also incurs substantial overhead to a ZigBee WPAN, as the network coordinator needs to scan the entire 16 channels and re-establish connections with clients. This problem becomes more pronounced in a dynamic network with mobile WiFi nodes and bursty interference.

Alternatively, coexistence can be arbitrated in space by

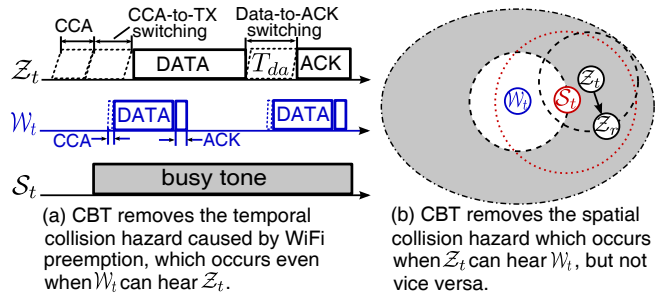


Figure 1: Principles behind CBT. Z_t , Z_r , S_t , and W_t are the ZigBee transmitter, receiver, signaler, and WiFi transmitter, respectively.

adjusting the transmit power and carrier sensing threshold. Gummadi *et al.* [12] proposed a policy framework that assigns such parameters to coexisting networks, so as to minimize mutual interference. This framework requires an arbitrator that can communicate with different network devices. It is only applicable to static networks, as any node movement would require the arbitrator to re-initiate a spectrum survey and re-allocate the parameters.

Another approach, called WISE [15], aims to enhance coexistence in the temporal domain. WISE harnesses the white spaces between WiFi transmissions, and opportunistically schedules ZigBee traffic therein. However, WISE needs to suspend ZigBee transmissions during each WiFi burst. It is unsuitable for TDMA mode, and for delay-sensitive applications.

To the best of our knowledge, CBT is the first attempt that allows ZigBee to coexist and even contend with WiFi in frequency, spatial and temporal domains. Our key observation is that a sufficient idle channel time exists and can be exploited by ZigBee, but the WiFi’s unawareness of ZigBee causes severe collisions. By enhancing the visibility of ZigBee to WiFi while preserving the carrier-sensing-based spectrum etiquette, CBT can substantially improve channel utilization without compromising WiFi performance.

This paper is also the first that establishes a comprehensive analytical framework to quantify the performance of coexisting ZigBee and WiFi networks. Our analysis is inspired by the pioneering efforts on renewal process models for 802.11 and 802.15.4 MAC protocols [17, 21]. The key challenge lies in modeling the disparate MAC-layer operations. Using reasonable simplifications, our analysis can accurately capture different performance metrics, such as collision probability and throughput. The results are also used to balance the cost and effectiveness of CBT.

3. COOPERATIVE BUSY TONE (CBT)

In this section, we present the key principles and components of CBT. CBT is built atop the ZigBee MAC/PHY, but adopts an innovative way of signaling a busy channel to WiFi. It employs a separate ZigBee node (signaler) to emit a busy tone concurrently with the desired data transmission, thereby eliminating the following collision hazards induced by MAC/PHY heterogeneity.

Temporal collision hazards.

Due to their disparate time resolutions, ZigBee transmissions may be easily preempted by WiFi transmissions. Zig-

Bee takes $128 \mu\text{s}$ to perform CCA (clear channel assessment), and an additional $192 \mu\text{s}$ to switch from the CCA to transmission mode, and even longer from receiving a packet to sending the ACK [4]. In contrast, WiFi nodes take only $28 \mu\text{s}$ for CCA and an average of $72 \mu\text{s}$ for a backoff (with the default backoff window size in 802.11a/g/n) [5]. Therefore, a WiFi node may finish the entire backoff process and start transmission within the switching time of ZigBee, thus causing collision (Fig. 1(a)). CBT reduces such temporal collision hazards by allowing the signaler to emit a busy tone, which is long enough to cover the data packet, the switching time of ZigBee, thus causing collision (Fig. 1(a)). CBT reduces such temporal collision hazards by allowing the signaler to emit a busy tone, which is long enough to cover the data packet, the switching time of ZigBee, thus causing collision (Fig. 1(a)). CBT reduces such temporal collision hazards by allowing the signaler to emit a busy tone, which is long enough to cover the data packet, the switching time of ZigBee, thus causing collision (Fig. 1(a)).

Spatial collision hazards.

Due to their disparate power levels (-25 to 0dBm for ZigBee vs. 15 to 20dBm for WiFi), ZigBee signals may not be effectively sensed by WiFi. As illustrated in Fig. 1(b), there exists a “gray region” where ZigBee can hear WiFi, but WiFi is oblivious of ZigBee and may thus interrupt it arbitrarily. To combat such spatial collision hazards, CBT allows the ZigBee node close to WiFi interferers (or a dedicated high-power ZigBee node such as XBee [8]) to work as the signaler, by transmitting a busy tone synchronously, thus notifying WiFi to suspend its transmission.

An immediate challenge to the above principles is: how to prevent the signaler from interfering with the transmitter, and how to synchronize the busy-tone and data transmission? We resolve these challenges using a *frequency flip* scheme and a *busy tone scheduler*.

3.1 Frequency Flip

The frequency flip exploits the inherent spectrum heterogeneity between ZigBee and WiFi. On the 2.4GHz ISM band, each WiFi channel occupies 22MHz, and overlaps with 4 orthogonal ZigBee channels. When running the frequency flip, the signaler hops to an adjacent channel before starting the busy tone, and hops back to the original channel immediately after the busy tone is transmitted. This way, CBT ensures the busy tone is orthogonal to the data packet, but still overlaps with the WiFi channel and can cause it to defer transmission.

Frequency flip incurs overhead to the signaler due to channel switching. However, the switching time is limited to $192 \mu\text{s}$ in ZigBee [4], and can be overlapped with the CCA-to-TX switching time (Fig. 1(a)). CBT assumes WiFi will defer when the ambient signal level is above its CCA threshold. This is a mandatory operation for 802.11a/g/n [5, Sec. 17.3.10.5]. However, CBT may become ineffective when it coexists with 802.11b, which can be configured to defer only for valid WiFi signals [5, Sec. 18.4.8.4].

3.2 Busy Tone Scheduler

In a ZigBee WPAN, a unique *coordinator* schedules a mixture of TDMA and CSMA slots periodically. Each scheduling period (called a *superframe*) starts with a beacon, followed by a number of CSMA slots and TDMA slots and then an inactive period.

CBT maintains the legacy scheduling protocol, but requires the signaler to dispatch the busy tone at an appropriate time, such that: *i*) it reduces the WiFi preemptions of ongoing or forthcoming ZigBee transmissions and *ii*) it min-

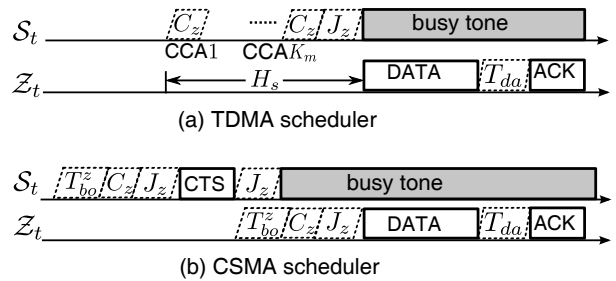


Figure 2: CBT scheduler. T_{bo}^z denotes the backoff time; $C_z(128\mu\text{s})$ is the CCA duration; $J_z(192\mu\text{s})$ the CCA-to-tx switching time (or channel switching time); T_{da} the data-to-ACK switching time.

imizes the potential influence on WiFi performance. The busy tone scheduler is designed to address this tradeoff. It allows both the TDMA and CSMA mode of ZigBee to coexist with WiFi.

3.2.1 TDMA scheduler

Fig. 2(a) illustrates the procedure to send a TDMA packet in CBT. CBT maintains the original TDMA slot allocation mechanism in ZigBee, but ensures the start time of each slot is conveyed to the signaler as well as the target clients, through persistent transmissions from the coordinator. Before each slot, the signaler performs CCA (for at most K_m times) in order to avoid interfering with WiFi. At the first idle CCA, it runs the frequency flip, switches to the adjacent channel, and starts the busy tone immediately. The busy tone lasts from the first idle CCA to the end of the TDMA slot. In this way, both the data and ACK packets can be protected from WiFi interruption.

A key parameter in the TDMA scheduler is the *harbinger time* H_s , defined according to how early the signaler starts the first CCA ($H_s = K_m C_z + J_z$, following Fig. 2(a)). If H_s is too long, the busy tone may occupy an unduly amount of channel time, thus reducing the channel utilization. If H_s is too short, the signaler may not be able to identify an idle slot before the scheduled transmission. Thus, it often has to abort the busy tone, degrading the effectiveness of CBT. In Sec. 4.2, we balance this tradeoff using a model-driven approach that relates H_s to network performance.

3.2.2 CSMA scheduler

Fig. 2(b) shows CBT’s operations in CSMA mode. Each CSMA transmission is initiated by the signaler, which performs CCA and backoff just as a normal client. Upon detection of an idle CCA, the signaler broadcasts a notification message (referred to as *CTS*) to the clients, switches to the adjacent channel and starts emitting the busy tone, as specified by the frequency flip. After receiving the CTS, the clients will contend for the channel access, following the same specification as in legacy ZigBee. However, if a client fails to acquire the channel, it needs to wait for the next CTS.

In designing the CSMA scheduler, we assume the signaler is able to obtain a rough estimate of the CSMA traffic demand in each superframe. This can be achieved by allowing the clients to report to the coordinator the number of pending packets in the current superframe. The coordinator then conveys the aggregated amount of unsatisfied traffic demand to the signaler, who then adjusts the number of CTS

attempts in the next superframe.

Note that the CTS preceding the actual CSMA data induces extra overhead. However, the CTS has equal length with an ACK packet, which contains only 5 bytes payload (11 bytes in total if the PHY layer preamble is included), much smaller than a typical data packet size. To further reduce the overhead, we adopt a *busy tone aggregation* scheme that allows $G > 1$ packets to be sent following the CTS, *i.e.*, each client can participate in G channel access contention upon receiving a CTS.

Since a data packet must follow the CTS from the signaler, there does not exist a harbinger time as in TDMA mode. However, the signaler needs to determine the busy tone duration, so that it covers the data and ACK with high probability, even after the random channel access among the contenders. Undoubtedly, setting the busy tone duration to the maximum backoff window plus the data and ACK duration would ensure full coverage. However, it may also waste channel time since the winning contender's backoff duration is random, and likely to be smaller than the maximum. In Sec. 4.3, we derive the busy tone duration that probabilistically makes this tradeoff.

It should be noted that CBT cannot completely remove the temporal collision hazards, because the signaler has the same time-resolution as a normal ZigBee node, and may be preempted by WiFi before sending the busy-tone. However, CBT can significantly reduce the collision hazards by augmenting the signaler's CCA capability in the ZigBee's TDMA mode, and by combining the signaler's CCA with the normal CCA/backoff in the ZigBee's CSMA mode. Its potential benefits will become clear as shown in our analysis and experimentation below.

4. PERFORMANCE ANALYSIS AND OPTIMIZATION

In this section, we establish a theoretical framework to analyze the performance of CBT in comparison with the legacy ZigBee. Our analysis pinpoints the key design parameters that affect the effectiveness of CBT in improving channel utilization while causing minimal interference to WiFi.

4.1 Network Model

We consider a ZigBee WPAN co-located with a WiFi WLAN, both sharing the same spectrum, and adopting the energy sensing based CCA as a spectrum etiquette. We mainly focus on the case with unsaturated WiFi and ZigBee links. As we will clarify, the saturated WiFi traffic results in almost zero throughput for the legacy ZigBee, and is less relevant for coexistence analysis. We assume the packet arrival follows a Poisson distribution. With unsaturated traffic, the aggregated traffic pattern is still approximately Poisson [7]. Hence, it is reasonable to deem the aggregated traffic as coexisting transmissions between one ZigBee and WiFi link. The Poisson assumption here is used for analytical tractability. The rationale behind parameter optimization (*e.g.*, the busy tone duration) does not depend on the traffic pattern.

We introduce the following notations beside those in Fig. 1 and Fig. 2. τ_w and τ_{wa} denote W_t 's data and ACK packet duration, respectively. T_w and λ_w denote the data packets' mean inter-arrival time and arrival rate ($T_w = \lambda_w^{-1}$). T_{bo}^w is the backoff duration, uniformly distributed between 0 and the backoff window (which may grow from CW_{min}^w to CW_{max}^w). After a backoff, WiFi must ensure the chan-

nel is idle for $DIFS(28\mu s)$ before transmission. β_w denotes the duration from backoff until an ACK when channel is idle. $\tau_z, \tau_{za}, T_z, \lambda_z, T_{bo}^z$ are the corresponding parameters for Z_t . Further, we denote γ_z as ZigBee's data/ACK duration (including the switching time between them, *i.e.*, T_{da}), and thus $\gamma_z = \tau_z + T_{da} + \tau_{za}$. U_z is ZigBee's slot duration ($U_z = 320\mu s$ [4]) and R_z the retransmission limit of a packet (default to 3 [4]).

We use the \subset notation to denote the observability between transmitters. We assume S_t is a high-power, ZigBee compatible node (*e.g.*, [8]), and $S_t \subset W_t$, *i.e.*, the S_t 's busy-tone can be sensed by W_t . Moreover, we assume the CTS packet from S_t will capture WiFi's packet even when collision occurs. Since Z_t has around 20dBm lower power than W_t , we assume the common case where collision affects Z_t 's packets, but not W_t 's. These assumptions will be further justified in Sec. 5.1.3.

Our analysis incorporates both the TDMA and CSMA mode, for both the legacy ZigBee and CBT, considering both $Z_t \not\subset W_t$ and $Z_t \subset W_t$. The primary method is to derive the collision probability in each case, and then relate it to the network's performance metric. We first analyze the *temporal collision probability* (Sec. 4.2 and Sec. 4.3), *i.e.*, probability that packets from both networks overlap with each other, thus causing collision. Later in Sec. 4.4 we analyze the *spatial collision probability*, probability that overlapped packets (from *randomly* located transmitters) fail to be decoded, taking into account the capture effect. We focus on each network's normalized throughput as the performance metric, denoted as Γ_z (Γ_w), which is essentially the ratio between the data packet duration and the average packet service time (including CCA, backoff, ACK, and retransmissions).

4.2 ZigBee's TDMA Coexistence with WiFi

4.2.1 Collision probability

As Z_t usually runs in low duty-cycle mode ($T_z \gg T_w$), we can tag an arbitrary packet from Z_t , and observe the collision with W_t . For simplicity, we introduce the concept of *vulnerable period*. A W_t packet arrival within the vulnerable period will overlap with the tagged packet from Z_t , resulting in collision. Let v be the duration of vulnerable period, then the collision probability becomes:

$$1 - \mathbb{P}[\text{no WiFi packet arrival in } v] = 1 - e^{-v\lambda_w}. \quad (1)$$

The following analysis derives the collision probability by analyzing v , depending on whether W_t can sense Z_t and CBT is adopted or not.

Case 1: Legacy ZigBee, $Z_t \not\subset W_t$. Collision occurs whenever W_t starts or ends its transmission within Z_t 's packet duration. Hence, the vulnerable period is $\beta_w + \tau_z$ for Z_t 's data packet, and $\beta_w + \tau_{za}$ for ACK packet. β_w involves a random variable T_{bo}^w , but for tractability, we approximate it with the mean \bar{T}_{bo}^w . When unsaturated, WiFi's backoff window tends to stay in CW_{min}^w . Hence, $\bar{T}_{bo}^w \approx \frac{1}{2}CW_{min}^w$. The data/ACK collision probability (denoted as P_1^d and P_1^a , respectively) is readily obtained by following Eq. (1):

$$P_1^d = 1 - e^{-\lambda_w(\beta_w + \tau_z)} \quad (2)$$

$$P_1^a = 1 - e^{-\lambda_w(\beta_w + \tau_{za})} \quad (3)$$

Case 2: Legacy ZigBee, $\mathcal{Z}_t \subset \mathcal{W}_t$. In this case, \mathcal{W}_t defers transmission if it senses \mathcal{Z}_t 's packets. Hence, collision occurs only if \mathcal{Z}_t 's packet starts while \mathcal{W}_t is transmitting. Equivalently, at least one packet of \mathcal{W}_t arrives within a vulnerable period β_w before \mathcal{Z}_t starts.

In addition, \mathcal{Z}_t 's ACK packet is corrupted under two conditions: *i*) \mathcal{W}_t starts its backoff and CCA within the data-to-ACK switching time T_{da} of \mathcal{Z}_t , corresponding to a vulnerable period $T_{da} - (\bar{T}_{bo}^w + DIFS)$. In slotted CSMA mode, T_{da} ranges from J_z to $U_z + J_z$, depending on the data packet size [4]. $(\bar{T}_{bo}^w + DIFS)$ is \mathcal{W}_t 's mean backoff plus defer time preceding each data packet. *ii*) The tail part of \mathcal{W}_t 's packet overlaps \mathcal{Z}_t 's ACK, corresponding to a vulnerable period β_w . Overall, the vulnerable period for the ACK packet is the minimum for condition *i* and *ii*. Again, following Eq. (1), the data/ACK collision probability for case 2 is:

$$P_2^d = 1 - e^{-\lambda_w \beta_w} \quad (4)$$

$$P_2^a = 1 - e^{-\lambda_w \cdot \min\{T_{da} - (\bar{T}_{bo}^w + DIFS), \beta_w\}} \quad (5)$$

It is easy to see that Case 2 has a shorter vulnerable period for both data and ACK packets, and thus $P_2^d < P_1^d$ and $P_2^a < P_1^a$, which agrees with the intuition that collision is reduced when \mathcal{W}_t can sense \mathcal{Z}_t .

Case 3: CBT is adopted. When running CBT, a key problem is how the signaler's harbinger time H_s (Sec. 3.2.1) affects CBT's performance. The follow analysis derives the relation between H_s and the failure rate of busy tone.

Proposition 1 *When the harbinger time $H_s = K_m C_z + J_z$, the probability that CBT fails to send the busy tone is: $P_{b0} = P_b(1 - P_{i|b})^{K_m - 1}$, where P_b and $P_{i|b}$ are defined in Eq. (6) and (10).*

PROOF. The probability that the first CCA returns busy, denoted as P_b , can be approximated by the fraction of time that the channel is occupied by \mathcal{W}_t , *i.e.*,

$$P_b = \gamma_w T_w^{-1} = \lambda_w \gamma_w, \quad (6)$$

where $\gamma_w \triangleq \beta_w - \bar{T}_{bo}^w - DIFS$ is equivalent to the total airtime of \mathcal{W}_t 's packet. Subsequent CCA events are complicated because they are correlated throughout \mathcal{W}_t 's airtime. Let $P_{i|b}$ be the probability that a subsequent CCA returns idle conditioned on the fact that the previous one is busy. This event occurs if the previous CCA falls in the tail of one WiFi packet, and the second CCA falls in the gap to the next WiFi packets. The gap, denoted as I_w , is a random variable that depends on \mathcal{W}_t 's backoff and inter-arrival time, which needs to be examined first.

Since \mathcal{Z}_t runs in low duty-cycle, it is reasonable to assume the busy/idle state created by \mathcal{W}_t is independent of \mathcal{Z}_t . If additional WiFi packets arrive during the service time of one packet (with probability P_q), then consecutive transmissions are separated solely by the backoff and defer times. It follows immediately that $I_w \sim U(DIFS, D_m)$, where D_m is the maximum backoff plus defer time preceding a transmission ($D_m = CW_{min}^w + DIFS$). Otherwise, I_w equals the inter-arrival time minus the duration of the previous transmission. In summary, we have:

$$\mathbb{P}[I_w \leq t] = \begin{cases} 1 - e^{-\lambda_w(\tau_w + t)}, & t > D_m \\ \frac{P_q t}{D_m - DIFS}, & DIFS < t \leq D_m \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

We proceed to derive P_q , the backlog probability. The evolution of \mathcal{W}_t 's state can be modelled as an $M/D/1$ queue, with arrival rate λ_w and service time β_w . Let P_{qi} be the probability that i packets are held by \mathcal{W}_t (one in service and others queued), then from well-established results in queuing theory [6], we have: $P_{q0} = 1 - \lambda_w \beta_w$, $P_{q1} = (e^{\lambda_w \beta_w} - 1)P_{q0}$, and $P_q = 1 - P_{q0} - P_{q1}$.

Back to the conditional probability $P_{i|b}$, we have:

$$P_{i|b} = \int_0^{C_z} \beta_w^{-1} (1 - \mathbb{P}[I_w \leq t]) dt \quad (\text{Note: } C_z < D_m) \quad (8)$$

$$= \int_0^{C_z} \beta_w^{-1} (1 - P_q(C_z - t)(D_m - DIFS)^{-1}) dt \quad (9)$$

$$= C_z \beta_w^{-1} - 0.5 P_q C_z^2 (D_m - DIFS)^{-1} \beta_w^{-1} \quad (10)$$

$$\approx C_z \beta_w^{-1} - 0.5 P_q C_z^2 D_m^{-1} \beta_w^{-1} \quad (11)$$

Consequently, the probability of aborting busy tone after K_m CCA attempts is: $P_{b0} = P_b(1 - P_{i|b})^{K_m - 1}$. \square

Conditioned on the fact that CCA succeeds and the busy tone is sent, \mathcal{W}_t may still preempt in the switching time of the signaler (Fig. 2(a)), with probability P_x :

$$P_x = \int_0^{C_z} C_z^{-1} \mathbb{P}[t \leq I_w \leq t + J_z] dt \quad (12)$$

$$= C_z^{-1} \int_0^{C_z} P_q (D_m - t) D_m^{-1} dt \approx \frac{1}{2} P_q C_z^{-1} D_m \quad (13)$$

Note that WiFi preemption results in collision only if the preemption time is after the K_s -th CCA, where $K_s = K_m - \lceil (\beta_w - J_z) C_z^{-1} \rceil$ is the time beyond which the end of \mathcal{W}_t 's packet always overlaps with \mathcal{Z}_t 's.

Summarizing the above analysis, the collision probability for data packets when using CBT is:

$$P_3^d = P_{b0} \mathbb{P}[\text{collision} | \text{CCA fail}] + \mathbb{P}[\text{CCA succeeds at } k, k > K_s] P_x \quad (14)$$

$$= P_b P_{b|b}^{K_m - 1} P_i^d + (1 - P_b P_{b|b}^{K_s}) P_x \quad (15)$$

where $P_{b|b} = 1 - P_{i|b}$. The ACK packet can only be corrupted if CCA fails, exposing the ACK to WiFi collision, just as the previous two cases. Hence,

$$P_3^a = P_{b0} P_i^a = P_b P_{b|b}^{K_m - 1} P_i^a, i \in \{1, 2\} \quad (16)$$

4.2.2 Network performance

Based on the above analysis, we can derive the throughput of \mathcal{Z}_t for all 3 cases using a renewal model. We model the transmission attempts by \mathcal{Z}_t as a renewal reward process. Each *renewal interval* is the service time of a packet, which starts with a transmission attempt, and ends with a successful ACK, or with a transmission failure if the retry limit R_z is exceeded. The *reward* is the amount of time in transmitting data without collision. Hence, the mean reward equals $[1 - (1 - P_{si})^{R_z}] \tau_z$, where P_{si} is the probability that both data and ACK are successfully delivered in case i (thus $P_{si} = (1 - P_i^d)(1 - P_i^a)$). The resulting throughput of \mathcal{Z}_t equals the reward rate Γ_{zi} . Let T_{zi}^{sv} be the service time of a packet in case i , then:

$$\Gamma_{zi} = [1 - (1 - P_{si})^{R_z}] \tau_z (\bar{T}_{zi}^{sv})^{-1}, i \in \{1, 2, 3\} \quad (17)$$

The mean service time \bar{T}_{zi}^{sv} depends on P_{si} , and the duration of each transmission attempt, denoted as T_i . For legacy ZigBee, we have $T_1 = T_2 = \gamma_z$. For CBT, from Fig. 2(a), we obtain $T_3 = U_z K_m + J_z + \gamma_z$. Further, note in all the

cases, the first transmission attempt occurs with probability 1, and the second occurs only if the first fails. Following this reasoning, the expected service time:

$$\bar{T}_{zi}^{sv} = \mathbb{E}[T_{zi}^{sv}] = \sum_{k=0}^{R_z-1} T_i (1 - P_{si})^k \quad (18)$$

Next, consider the performance of WiFi. In case 1 ($\mathcal{Z}_t \not\subset \mathcal{W}_t$), \mathcal{W}_t is unaffected by \mathcal{Z}_t , thus its mean service time $\bar{T}_{w1}^{sv} = \beta_w$. In case 2 ($\mathcal{Z}_t \subset \mathcal{W}_t$), \mathcal{W}_t 's arrival overlaps \mathcal{Z}_t 's data or ACK packet, with probability $1 - e^{-\lambda_z(\tau_z + \tau_{za})} \approx \lambda_z(\tau_z + \tau_{za})$. Since \mathcal{W}_t freezes its backoff upon a busy CCA, the corresponding service time equals β_w plus the extra time that it has to wait until \mathcal{Z}_t finishes transmission. Otherwise, its service time remains β_w . Therefore, the mean service time is:

$$\begin{aligned} \bar{T}_{w2}^{sv} &= \beta_w [1 - \lambda_z(\tau_z + \tau_{za})] \\ &+ \lambda_z \tau_z \left(\frac{\tau_z}{2} + \beta_w \right) + \lambda_z \tau_{za} \left(\frac{\tau_{za}}{2} + \beta_w \right) \end{aligned} \quad (19)$$

When CBT is used (case 3), its CCA succeeds at the i -th attempt with probability $(1 - P_b)$ for $i = 1$, and $P_b P_{b|i}^{i-2} P_{b|i}$ otherwise. Since each CCA attempt takes one slot, the busy tone duration is $\tau_b = U_z(K_m - i) + J_z + \tau_z + T_{da} + \tau_{za}$. The mean duration of a transmission attempt in CBT is thus:

$$\begin{aligned} \Phi_b &= (1 - P_b)(C_z(K_m - 1) + \tau_b) \\ &+ \sum_{i=2}^{K_m} P_b P_{b|i}^{i-2} P_{b|i} (C_z(K_m - i) + \tau_b) \\ &+ P_b P_{b|i}^{K_m-1} \left(\frac{\tau_z}{\tau_z + \tau_{za}} \tau_z + \frac{\tau_{za}}{\tau_z + \tau_{za}} \tau_{za} \right) \end{aligned} \quad (20)$$

and the WiFi service time can be approximated as:

$$\bar{T}_{w3}^{sv} = (1 - \lambda_z \Phi_b) \tau_w + \lambda_z \Phi_b \left(\frac{\Phi_b}{2} + \tau_w \right) \quad (21)$$

Following a similar renewal model for \mathcal{Z}_t , the throughput of \mathcal{W}_t for the 3 cases is:

$$\Gamma_{wi} = \tau_w (\bar{T}_{wi}^{sv})^{-1}, i \in \{1, 2, 3\} \quad (22)$$

4.3 ZigBee's CSMA Coexistence with WiFi

4.3.1 ZigBee's legacy CSMA mode (Case 4)

We model ZigBee's CSMA mode using a Markov chain shown in Fig. 3, where BS_k denotes the k -th backoff and CCA attempt. Similar to the TDMA analysis, we make a key simplification that decouples the channel status from \mathcal{Z}_t 's transmission, *i.e.*, assuming the busy/idle state of WiFi is unaffected by \mathcal{Z}_t . This assumption will be removed when analyzing the effects of \mathcal{Z}_t on \mathcal{W}_t .

In CSMA mode, \mathcal{Z}_t must perform backoff and ensure two consecutive slots (CCA1 and CCA2) are idle before transmission, and abort the transmission if sensing a busy channel even after K attempts. Straightforwardly, the steady state probability that \mathcal{Z}_t senses an idle channel in CCA1 is $P_i = 1 - P_b$ (see Eq. (6) for P_b). Conditioned on the event that CCA1 is idle, CCA2 returns idle if no packets for \mathcal{W}_t arrive between them, which has a probability: $P_{i|i} = e^{-\lambda_w U_z}$. Let P_{tx} denote the transmission probability after a backoff and CCA attempt, then:

$$P_{tx} = P_i P_{i|i} = (1 - P_b) e^{-\lambda_w U_z} \quad (23)$$

The duration of the k -th backoff attempt is uniformly distributed between 0 and $[4]$:

$$B_k = \min\{2^{\min BE + k - 1}, 2^{\max BE}\} \cdot U_z, k \in [1, K] \quad (24)$$




Figure 3: Markov chain model for analyzing ZigBee's throughput when it coexists with WiFi.




Figure 4: Analyzing spatial collision probability under random link locations.

where $\min BE$ and $\max BE$ are the minimum and maximum backoff exponent, default to 3 and 5, respectively. Note that CCA1 must be performed, but CCA2 is needed only if CCA1 returns idle. The mean time spent in a CCA attempt is thus: $U_z + P_i U_z = (1 + P_i) U_z$.

Given the above components, we can derive \mathcal{Z}_t 's service time. From the Markov chain model, the mean first passage time to data transmission is:

$$F_{TX} = \frac{b_1 + (1 - P_{tx})b_2 + \dots + (1 - P_{tx}^{K-1})b_K}{1 - (1 - P_{tx})^K} \quad (25)$$

where $b_k = 0.5B_k + (1 + P_i)U_z$ is the expected duration of the k -th backoff plus CCA attempt (BS_k). Intuitively, the numerator of Eq. (25) is the mean duration of all the attempts until data transmission or backoff failure. The number of attempts follows a geometrical distribution with success probability $1 - (1 - P_{tx})^K$. Consider $\mathcal{Z}_t \subset \mathcal{W}_t$, and let P_{s4} be the probability that both data and ACK are transmitted without collision, then the expected number of trials until transmission succeeds or retry failure is:

$$1 + (1 - P_{s4}) + \dots + (1 - P_{s4})^{R_z-1} = \sum_{k=0}^{R_z-1} (1 - P_{s4})^k$$

Each transmission takes F_{TX} plus the data and ACK (or timeout) duration γ_z . Thus the mean service time:

$$\bar{T}_{z4}^{sv} = (F_{TX} + \gamma_z) \sum_{k=0}^{R_z-1} (1 - P_{s4})^k \quad (26)$$

Note that WiFi may preempt the CCA-to-TX and data-to-ACK switching time, with probability P_d and P_a , respectively (Fig. 3). Following similar analysis of the vulnerable period in the TDMA case 2, we obtain $P_d = 1 - e^{-\lambda_w J_z}$, and $P_a = P_d^a$, from which P_{s4} can be derived. In case when $\mathcal{Z}_t \not\subset \mathcal{W}_t$, the success probability is equivalent to P_{s1} , and mean service time can be obtained in the same way as Eq. (26).

4.3.2 CBT's CSMA mode (Case 5)

The key problem in this case is to determine the busy tone duration T_b , which should cover the entire data transmission without wasting extra channel time. The following analytical results state that this can be realized with probabilistic guarantee.

Proposition 2 For CBT with N clients and a single busy tone, with duration $T_b = \frac{kB_1}{N} + \gamma_z$, the busy tone covers the data and ACK packets with probability of at least $1 - e^{-k}$. For CBT with G aggregated busy tones, with duration $T_b' = \frac{kGB_1}{N+1} + \gamma_z$, the busy tone covers all the packets with probability of at least $1 - \frac{1}{k}$.

PROOF. Let t_i be the backoff counter set by node i , $t_i \sim U(0, B_1)$. The probability that the busy tone covers the data/ACK duration of the winning contender is:

$$\mathbb{P}\left[\min_i t_i + \gamma_z \leq T_b\right] = \mathbb{P}\left[\min_i t_i \leq T_b - \gamma_z\right] \quad (27)$$

$$= 1 - \left(1 - \frac{T_b - \gamma_z}{B_1}\right)^N \quad (\text{Let } T_b = \frac{kB_1}{N} + \gamma_z) \quad (28)$$

$$= 1 - \left(1 - \frac{k}{N}\right)^N \geq 1 - e^{-k} \quad (29)$$

For the case with busy tone aggregation, we first derive the expectation of the minimum backoff counter among all the clients, denoted as M .

$$\mathbb{E}[M] = \int_0^{B_1} [1 - F_M(x)] dx \quad (\text{cf. [11, Sec. 1.3.5]}) \quad (30)$$

$$= \int_0^{B_1} \left(1 - \frac{x}{B_1}\right)^N dx = \frac{B_1}{N+1} \quad (31)$$

Further, note that:

$$\mathbb{P}\left[\gamma_z + \sum_{i=1}^G b_i \geq T'_b\right] \quad (32)$$

$$\leq \frac{\mathbb{E}\left[\sum_{i=1}^G b_i\right]}{T'_b - \gamma_z} \quad (\text{Markov inequality}) \quad (33)$$

$$= \frac{G \cdot B_1}{N+1} \cdot \frac{1}{T'_b - \gamma_z} \quad (\text{Let } T'_b = \frac{kG \cdot B_1}{N+1} + \gamma_z) \quad (34)$$

$$= k^{-1} \quad (35)$$

Hence $\mathbb{P}[\sum_{i=1}^G b_i \leq T'_b] \geq 1 - \frac{1}{k}$. \square

Note that the above result is independent of the WiFi arrival time distribution. It implies that the busy tone needs extra number of slots (denoted as K_b) to compensate for the random backoff, *i.e.*, $T_b = K_b U_z + \gamma_z$. If the busy tone covers the backoff duration B_1 plus the data/ACK duration: $T_b^* \triangleq G(K_b^* U_z + \gamma_z)$, $K_b^* = 2^{\min^{BE}}$, then collision rate can be reduced to 0.

When $G = 1$ and $T_b < T_b^*$, the end of the ACK packet may still be exposed to WiFi collision. Suppose $\mathcal{Z}_t \subset \mathcal{W}_t$, then this occurs when the data-to-ACK switching time is beyond the protection of the busy tone. Let $u = \tau_z + T_{da} - \bar{T}_{bo}^w$, then the ACK collision probability can be derived similarly to the TDMA mode:

$$P_5^a = \mathbb{P}\left[\min_i b_i + u \geq T_b\right] (1 - e^{-\lambda_w(\min_i b_i + u)}) \quad (36)$$

$$\approx (1 - B_1^{-1}(T_b - u))^N (1 - e^{-\lambda_w(\frac{B_1}{N+1} + u)}) \quad (37)$$

Given T_b , we can further derive the service time, denoted as T_{z5}^{sv} . Note that the signaler's backoff and CCA operations are similar to the legacy case (Fig. 3), except that its CCA attempt always takes one slot U_z , and the transmission probability $P_{tx} = P_i$ and success probability of each transmission attempt $P_{s5} = 1 - P_5^d$. Therefore, the service time starts from the first backoff stage, and ends after the busy tone or the ACK packet, whichever lasts longer. Denote T_{m5} as the expected duration of a transmission attempt after backoff and CCA succeed, then $T_{m5} = \max\{\frac{B_1}{N+1} + \gamma_z, T_b\}$. Similar to Eq. (25), the mean service time of a packet:

$$\bar{T}_{z5}^{sv} = (F_{TX} + T_{m5}) \sum_{k=0}^{R_z-1} (1 - P_{s5})^k \quad (38)$$

Given the mean service time and collision probability in CSMA, the throughput readily follows from Eq. (17).

For the WiFi transmitter \mathcal{W}_t , the service time depends on the \mathcal{Z}_t 's traffic load. Consider the service time of each packet from \mathcal{Z}_t as a renewal interval with length L_{zi} . Note that \mathcal{W}_t tends to have higher priority than \mathcal{Z}_t , and thus its load determines the length of the L_{zi} : $L_{zi} = \min\{T_z, T_{zi}^{sv}\}$, $i \in \{4, 5\}$. In case 4, the mean number of transmission attempts by \mathcal{Z}_t in its renewal interval is

$$N_{d4} = \sum_{r=1}^{R_z} [(1 - (1 - P_{tx})^K)(1 - P_{si})]^{r-1} \quad (39)$$

which is also the mean number of times that \mathcal{W}_t defers its transmission. For case 5, since one busy tone for \mathcal{Z}_t is sent in each renewal interval, $N_{d5} = 1$. For both cases, \mathcal{W}_t 's packets are disrupted with probability $P_{ri} = \frac{N_{di} T_w}{L_{zi}}$, $i \in \{4, 5\}$, resulting in mean service delay:

$$\bar{T}_{w4}^{sv} = (1 - P_{r4})\beta_w + P_{r4}(\tau_z + \beta_w) \quad (40)$$

$$\bar{T}_{w5}^{sv} = (1 - P_{r5})\beta_w + P_{r5}(T_{m5} - T_w + \beta_w) \quad (41)$$

The throughput can be derived similarly to the TDMA case (Eq. (22)).

4.4 Spatial Collision Probability

The above analysis focused on the collision probability between co-located WiFi and ZigBee. In practice, collision does not necessarily cause packet loss. Due to the capture effect, the desired packet can still be decoded if its power is sufficiently higher than the interfering packet. Such opportunities depend on the relative location of WiFi and ZigBee links.

Consider a randomly located ZigBee link and WiFi interferer, as shown in Fig. 4. For simplicity, we denote $d_t = d(\mathcal{W}_t, \mathcal{Z}_t)$, $d_z = d(\mathcal{Z}_t, \mathcal{Z}_r)$, $d_r = d(\mathcal{W}_t, \mathcal{Z}_r)$, where the function $d(\cdot, \cdot)$ represents the distance between two nodes. We assume $\theta \sim U(0, 2\pi)$, and $d_z \sim U(0, \rho_z)$, where ρ_z is the maximum distance between \mathcal{Z}_t and \mathcal{Z}_r . Since ZigBee and WiFi have similar receiver sensitivity (around -86dBm [4, 5]), we assume \mathcal{W}_t and \mathcal{Z}_t have the same carrier sensing threshold. We use a generic pathloss model, where the ratio between received power and transmit power is $cd_z^{-\alpha}$. The constant c depends on the propagation characteristics (*e.g.*, free-space or two-ray ground model), and α is the environment dependent pathloss exponent. Denote C_a as the capture threshold, Λ_z and Λ_w as the transmit power of \mathcal{Z}_t and \mathcal{W}_t , respectively, then \mathcal{Z}_r fails to decode the packet if:

$$\frac{\Lambda_z cd_z^{-\alpha}}{\Lambda_w cd_r^{-\alpha}} \leq C_a, \text{ or equivalently, } \frac{d_r^2}{d_z^2} \leq \left(C_a \frac{\Lambda_w}{\Lambda_z}\right)^{\frac{2}{\alpha}} \quad (42)$$

The key to our analysis is to derive the probability that the two link distances satisfy the above collision condition, which we name as *spatial collision probability*. Following the cosine rule $d_r^2 = d_t^2 + d_z^2 - 2 \cos \theta d_t d_z$, the spatial collision probability becomes:

$$\frac{d_r^2}{d_z^2} = \frac{d_t^2}{d_z^2} - 2 \cos \theta \frac{d_t}{d_z} \leq \left(C_a \frac{\Lambda_w}{\Lambda_z}\right)^{\frac{2}{\alpha}} - 1 \triangleq c_1 \quad (43)$$

which can be transformed into:

$$\left(d_z + \frac{d_t \cos \theta}{c_1}\right)^2 \geq \frac{d_t^2}{c_1} + \frac{d_t^2 \cos^2 \theta}{c_1^2} = \frac{d_t^2}{c_1^2} (c_1 + \cos^2 \theta) \approx \frac{d_t^2}{c_1}$$

The geometrical meaning of the above equation, combined with $d_z \sim U(0, \rho_z)$ (*i.e.*, $d_z^2 \leq \rho_z^2$), is the exclusion region of two circular areas. To further simplify, observe that $d_z \gg \frac{d_t \cos \theta}{c_1}$, and the random variable $\cos \theta$ has mean 0, hence the

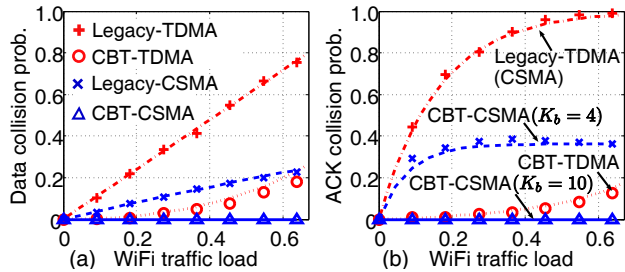


Figure 5: Collision probability of data and ACK packets. Markers and lines represent simulation and analytical results, respectively.

above two inequalities can be approximated as: $\rho_z^2 \geq d_z^2 \geq \frac{d_t^2}{c_1^2}$, which reduces the original circular intersection problem into a 1-dimensional, line-segment intersection problem (for a given d_t). Consequently, the exclusion region:

$$I_e = \max\{0, 1 - d_t c_1^{-\frac{1}{2}} \rho_z^{-1}\} \quad (44)$$

which is exactly the spatial collision probability. The above approximation will be verified in Sec. 5.1.3.

5. EXPERIMENTAL RESULTS

This section compares the qualitative prediction of the analytical model with detailed simulation in ns-2. We find them matching well across a broad range of experiments. We further explore the design parameters of CBT using the analytical model, and demonstrate its performance in a real experimental testbed.

5.1 Simulation experiments

We have implemented the CBT protocol based on the ZigBee (CSMA) module in ns-2 (version 2.33). Following the IEEE 802.15.4 standard, we also developed a TDMA module for ZigBee. The main components of CBT, *i.e.*, the frequency flip and busy tone scheduler, are implemented on top of the TDMA/CSMA modules. The PHY-layer parameters are set to their default values. The WiFi module is configured consistently with the 802.11g standard. Our experiments first focus on the case where the two networks are in close proximity and can sense each other (Sec. 5.1.1, Sec. 5.1.2), and later explore the effects of link locations (Sec. 5.1.3, Sec. 5.1.4).

5.1.1 Temporal collision probability

We vary the *traffic load* of one network and observe its impacts on the other, where:

$$\text{Traffic load} = \frac{\text{packet size} \times \text{packet arrival rate}}{\text{PHY layer bit-rate}} \quad (45)$$

Throughout the experiments, ZigBee uses a fixed bit-rate of 250Kbps [4], packet size 63 bytes, and arrival rate 8 packets/second. WiFi bit-rate is set to 18Mbps, and packet size to 1KB. The corresponding saturation traffic load is around 0.67. Fig. 5(a) shows probability that ZigBee’s data packets collides with WiFi packets, under varying WiFi traffic load. We see a close match between analysis and simulation, with a deviation of less than 1% for most experiments. In TDMA mode, coexisting WiFi traffic is detrimental to legacy ZigBee — the collision probability grows from 0 to above 0.79 as WiFi load increases from 0 to 0.67 (corresponding to the saturation load). With $K_m = 8$ (*i.e.*, harbinger

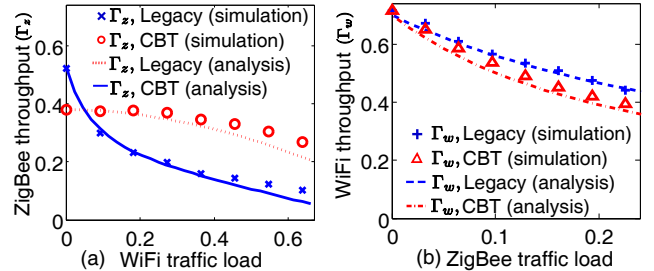


Figure 6: Throughput performance in TDMA mode.

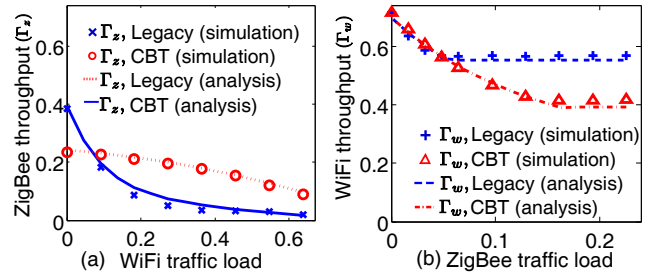


Figure 7: Throughput performance in CSMA mode.

time $H_s = 8C_z + J_z$), CBT reduces the collision probability to below 0.05 for medium to low WiFi traffic, and below 0.2 even when WiFi is saturated.

Measurement of WiFi hotspots observed traffic load of around 0.6 at peak hours [19]. Even under this level of interference, the data and ACK collision probability in CBT is 0.20 and 0.16, in contrast with 0.71 and 0.97 in legacy ZigBee, more than 72% reduction for both types of packet. With a low collision probability, packet loss can be easily recovered via retransmissions.

In CSMA mode, the CCA and backoff mechanism in legacy ZigBee alleviates the collision of data packets, but becomes ineffective when WiFi is heavily-loaded due to the preemption problem (Sec. 3). By preventing WiFi’s preemption in ZigBee’s CCA-to-TX switching period, CBT reduces the data collision probability to 0.

In addition, for legacy ZigBee, the ACK collision probability grows from 0 to 0.97 for both TDMA and CSMA mode (Fig. 5(b), only TDMA mode is plotted for clarity), implying a substantial number of redundant retransmissions due to ACK losses. For CBT, the probability is consistently below 0.2 in TDMA mode. In CSMA mode, when $K_b = 10$, the ACK collision probability is reduced to 0. As K_b is reduced to 4, the collision probability increases up to 0.4.

5.1.2 Throughput performance

Fig. 6 shows the throughput (Γ_z) of ZigBee in TDMA mode. Γ_z decreases as WiFi increases its load, because severe collision can cost extra service time, and lower the efficiency of channel utilization. When WiFi becomes saturated, ZigBee throughput is reduced to 0.09, 82% lower than the case without coexisting traffic. With CBT, throughput can be boosted by 47% to 112% when WiFi load is above 0.18. Notably, CBT may have lower throughput than ZigBee under mild interference, due to its CCA overhead. This implies CBT’s benefit outweighs its overhead beyond a “sweet spot”, which can be obtained by numerically solving for λ_w from the equation $\Gamma_{z2} = \Gamma_{z3}$ (see Eq. (17)). To optimize performance, the signaler can monitor the busy/idle status of WiFi and trigger CBT only when λ_w is beyond this point.




Figure 8: Spatial collision probability ($\Lambda_w = 15\text{dBm}$).

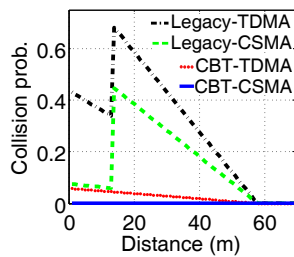


Figure 9: Spatial-temporal collision probability

We leave such optimization as future work.

Intuitively, the extra channel time taken by CBT will degrade WiFi's performance. However, the figure shows such effect is negligible when its traffic load is below 5%, and is comparable to legacy ZigBee even under heavy traffic. This is because CBT saves the retransmissions caused by collision, thereby counter-balancing the extra overhead. In practical ZigBee applications, the duty-cycle is typically below 1%, and up to 10% at its maximum [16]. Hence, in practice, CBT does not cause any additional performance degradation to WiFi performance.

Fig. 7 plots the throughput performance in CSMA mode. CBT (with $K_b = 10$ and without busy tone aggregation) consistently achieves a throughput of above 0.15, whereas the legacy ZigBee's throughput quickly drops below 0.03 as WiFi becomes saturated. CBT's advantage comes with the overhead from CTS and busy tone. Similar to the TDMA case, a "sweet spot" of WiFi load can be derived, beyond which the overhead is quickly outweighed by the saving in retransmission time, resulting in several folds performance improvement over the legacy ZigBee. The overhead can be further reduced by reducing K_b and enabling busy tone aggregation.

In addition, although both CBT and legacy reduce the throughput of WiFi, the reduction quickly converges to a certain level corresponding to the saturation throughput of ZigBee under the current WiFi traffic load. This again justifies the modeling assumption that WiFi has higher priority than ZigBee when both are running CSMA. The achievable throughput of ZigBee is determined by the load of coexisting WiFi traffic, but not vice versa. In addition, note that heavily loaded CBT degrades WiFi performance, but under practical low duty-cycle applications (load < 5%), the degradation is negligible compared with the legacy ZigBee.

5.1.3 Spatial collision probability

We use Monte Carlo simulation to verify the analysis in Sec. 4.4. We first fix the locations of \mathcal{W}_t and \mathcal{Z}_t , and then randomly uniformly generate the location of \mathcal{Z}_r within the maximum range ρ_z (set to 6m) from \mathcal{Z}_t . Fig. 8(a) shows the spatial collision probability I_e for different d_t , each obtained from 10^4 samples. Consistent with the analytical results (solid lines), I_e decreases almost linearly with d_t . Beyond a certain threshold (as indicated in Eq. (44)), the ratio between the received power from \mathcal{Z}_t and that from \mathcal{W}_t is larger than the capture threshold (set to 10dB) with probability 1, and thus I_e becomes 0.

In addition, I_e increases as ZigBee's transmit power decreases from 0dBm to -20dBm, implying that a low-power mode suffers more from collisions, and may not save energy due to the potential retransmission overhead. Note that the

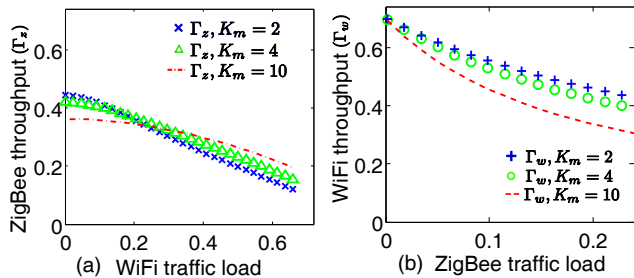


Figure 10: Impact of harbinger time in CBT TDMA mode.




Figure 11: Impact of busy-tone duration in CBT CSMA mode.

15dBm transmit power is only practical for DC-powered ZigBee node (*e.g.*, the XBee module [8]), which can be used as the signaler. As this power level is comparable to WiFi, spatial collision occurs only when the two networks are extremely close. Thus, it is reasonable to assume the CTS packet from the signaler is unaffected by WiFi in common cases. In addition, note that a larger ρ_z results in lower mean SINR for \mathcal{Z}_r and even larger I_e . We omit the detailed experiments as the effect is similar to the decrease in transmit power.

5.1.4 Joint spatial-temporal effects

Combining the analysis for temporal and spatial factors, we can analyze how ZigBee is affected by a WiFi transmitter at an arbitrary location. We fix the WiFi traffic load to a medium value 0.36, the ZigBee's transmit power to -10dBm, and then calculate the probability of packet loss, *i.e.*, both spatial and temporal collision occurs. As can be seen from the results (Fig. 9), the joint collision probability is bisected for legacy ZigBee, according to whether \mathcal{W}_t and \mathcal{Z}_t can sense each other. When $\mathcal{Z}_t \subset \mathcal{W}_t$, the joint collision probability is relatively low (below 41% and 17% for TDMA/CSMA), implying a low packet loss rate. However, when $\mathcal{Z}_t \not\subset \mathcal{W}_t$, the collision rate increases dramatically (to above 43%) for both TDMA and CSMA. Using the signaler, CBT extends the range where \mathcal{Z}_t is visible to \mathcal{W}_t and prevents WiFi pre-emption in the time domain. These two advantages together bring the collision probability to below 7%.

5.1.5 Exploring the design space

Having validated the accuracy of the analysis, we now employ it to flexibly navigate the impact of design parameters on the performance vs. cost tradeoff in CBT. Fig. 10 shows the throughput in TDMA mode for different harbinger time (determined by K_m). A small K_m induces less overhead, and gains higher throughput for ZigBee when the WiFi load is small (below 0.2), but becomes less effective under medium to high WiFi load. In addition, under low duty-cycle ZigBee traffic (below 0.05), WiFi throughput is virtually unaffected




Figure 12: Location of ZigBee and WiFi links in the testbed. A→B is the WiFi link. All others are ZigBee nodes.




Figure 13: Collision rate for different ZigBee link locations.

by K_m .

In CSMA mode, the key parameter is the busy tone duration (determined by K_b). As shown in Fig. 11, a large K_b ($K_b = 10$) gains $2\times$ higher throughput for ZigBee than a small K_b ($K_b = 2$), only at the cost of minor throughput degradation for WiFi (less than 6%), for a broad range of traffic load values. Therefore, a large busy tone duration is always preferred when CBT runs in CSMA mode.

5.2 Testbed experiments

We have developed a preliminary version of CBT (TDMA mode) on TinyOS 2.0 and GNURadio 3.2.2. We implement a TDMA scheduling module based on openzb [1], a TinyOS branch for IEEE 802.15.4. The ZigBee hardware that we use, the MicaZ mote, has a maximum power of 0dBm and packet size limit of 127 bytes, precluding a direct implementation of the signaller’s functionalities. Therefore, we implement the signaller based on the 802.15.4 PHY in GNURadio, and run it on the USRP2 software radio [9], which does not have such limitations and can communicate seamlessly with MicaZ. The coordinator and clients are running on micaZ motes.

Due to the inefficient user-space signal processing, the GNURadio/USRP2 platform cannot perform carrier sensing in real time, and cannot be synchronized to follow the TDMA schedule set by the coordinator. Therefore, we disable the carrier sensing, fix the harbinger time to $H_s = J_z$, and allow the coordinator to send a 5-byte notification message to the USRP2 immediately before the harbinger time arrives. The message is sent without carrier sensing and may still be lost due to collision with WiFi (though with a lower probability compared with a larger data packet). Thus, the performance is expected to be lower than a full-fledged implementation on ZigBee-compatible hardware, such as the XBee module [8]. To alleviate the loss, two back-to-back notification messages are used in the actual implementation. Such a patch costs additional channel time, and should affect WiFi more than a legitimate CBT.

We deploy the ZigBee coordinator, clients and signaller in




Figure 14: Impact of ZigBee on WiFi.

an office environment, co-located with a pair of WiFi nodes. Fig. 12 shows a map of the node locations. The relative distances between ZigBee and WiFi links satisfy: $1m < d_z < 3m$, $0.5m < d_t < 7m$. The distance between WiFi and ZigBee transmitters (d_t) is limited to $7m$ because ZigBee link distances (d_z) are typically short, and ZigBee signal tends to capture the WiFi interference when d_t is large. Each ZigBee link consists of one coordinator and one client (the transmitter is randomly selected between them), and the signaller is placed near the coordinator. The ZigBee link sends TDMA packets at a duty cycle of 8 packets/second. The WiFi link is running constant-bit-rate UDP traffic, with traffic load 0.22, packet size and bit rate settings following Sec. 5.1.1.

Fig. 13 plots the collision rate between ZigBee and WiFi, which is measured by the one-way packet loss-rate of ZigBee. The combined effects of spatial and temporal collision result in diverse collision rate for different links. In general, when the ZigBee transmitter is close to the WiFi transmitter, the collision rate is lower due to ZigBee’s better visibility to WiFi (e.g., for ZigBee link C→D). When the ZigBee transmitter moves far away from the WiFi, it can no longer trigger the deferring and backoff, but may still be exposed to the interference from WiFi, thus causing severe collision (e.g., for ZigBee link L→K). Overall, the collision rate is above 12% for half of the legacy ZigBee links. In contrast, when running CBT, collision rate for all the links is reduced to below 8%, and the reduction is above 60% for all links. In summary, the topological effects on collision rate and the effectiveness of CBT is consistent with the trend predicted by the analysis in Fig. 9. The actual value of collision rate does not match because of the simplified propagation model used in analysis (Sec. 4.4).

To examine the impact of ZigBee on WiFi, we create a *worst case* scenario for WiFi where the USRP2 sends legacy ZigBee packets and busy tones without carrier sensing. Fig. 14 plots the WiFi packet delay as a function of ZigBee load. The delay performance remains virtually unaffected for lightly loaded ZigBee traffic, e.g., below 2%. More importantly, the extra overhead in CBT does not degrade the WiFi performance significantly compared to the legacy ZigBee, even though the benefit of CBT in saving retransmissions is unaccounted for here. Since ZigBee targets low duty-cycle applications, the experiment essentially shows that neither the legacy ZigBee nor CBT affect WiFi in the common cases.

6. CONCLUSION

In this paper, we proposed CBT, a cooperative paradigm enabling low-power, low-rate ZigBee WPANs to coexist with a high-end WiFi WLAN. CBT maintains the spectrum etiquette based on energy detection, but overcomes its limitations in heterogeneous networks by separating the busy tone

signaling from data transmission. We establish an analytical framework that quantitatively compares CBT with the legacy ZigBee, and allows flexible exploration of its design parameters. The analysis, combined with detailed simulation and prototype implementation, demonstrates several-fold performance improvement (in terms of collision probability and throughput), with negligible cost to the WiFi side. In future, we plan to extend CBT to a general framework allowing coexistence of heterogeneous MAC/PHY protocols in the recently-opened TV band white-spaces.

Acknowledgement

The work reported in this paper was supported in part by the NSF under Grant CNS-0721529.

7. REFERENCES

- [1] open-ZB. <http://www.open-zb.net>.
- [2] The GNU Software Radio. <http://gnuradio.org/trac/wiki>.
- [3] Coexistence of Wireless Personal Area Networks With Other Wireless Devices Operating in Unlicensed Frequency Bands. *IEEE Std 802.15.2*, 2003.
- [4] Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs). *IEEE Std. 802.15.4*, 2003.
- [5] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. *IEEE Std. 802.11*, 2007.
- [6] O. Brun and J.-M. Garcia. Analytical Solution of Finite Capacity M/D/1 Queues. *Journal of Applied Probability*, 37(4), 2000.
- [7] F. Daneshgaran, M. Laddomada, F. Mesiti, and M. Mondin. On the Linear Behaviour of the Throughput of IEEE 802.11 DCF in Non-Saturated Conditions. *IEEE Communications Letters*, 11(11), 2007.
- [8] Digi International Inc. XBee-PRO 802.15.4 OEM RF Modules. <http://www.digi.com/>.
- [9] Ettus Research LLC. Universal Software Radio Peripheral (USRP). <http://www.ettus.com/>.
- [10] FCC. Second Memorandum Opinion and Order, Sep. 2010.
- [11] R. Gallager. *Discrete Stochastic Processes (2nd Ed.) Draft*. 2009.
- [12] R. Gummadi, H. Balakrishnan, and S. Seshan. Metronome: Coordinating Spectrum Sharing in Heterogeneous Wireless Networks. In *First International Workshop on Communication Systems and Networks (COMSNETS)*, 2009.
- [13] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan. Understanding and Mitigating the Impact of RF Interference on 802.11 Networks. In *Proc. of ACM SIGCOMM*, 2007.
- [14] J. Hou, B. Chang, D.-K. Cho, and M. Gerla. Minimizing 802.11 Interference on ZigBee Medical Sensors. In *Proc. of the International Conference on Body Area Networks*, 2009.
- [15] J. Huang, G. Xing, G. Zhou, and R. Zhou. Beyond Co-existence: Exploiting WiFi White Space for ZigBee Performance Assurance. In *Proc. of IEEE ICNP*, 2010.
- [16] IEEE 802.15 Working Group. Coexistence Analysis of IEEE Std 802.15.4 With Other IEEE Standards and Proposed Standards, 2010.
- [17] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New Insights From a Fixed-Point Analysis of Single Cell IEEE 802.11 WLANs. *IEEE/ACM Trans. on Networking*, 15(3), 2007.
- [18] S. Pollin, I. Tan, B. Hodge, C. Chun, and A. Bahai. Harmful Coexistence Between 802.15.4 and 802.11: A Measurement-based Study. In *Proc. of CrownCom*, 2008.
- [19] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan. Measurement-based Characterization of 802.11 in a Hotspot Setting. In *Proc. of SIGCOMM E-WIND*, 2005.
- [20] Schneider Electric. ZigBee WiFi Coexistence. <http://www.zigbee.org/LearnMore/WhitePapers.aspx>, 2008.
- [21] C.-K. Singh, A. Kumar, and P. M. Ameer. Performance Evaluation of an IEEE 802.15.4 Sensor Network With a Star Topology. *Wireless Networks*, 14(4), 2008.
- [22] J. Zhu, A. Waltho, X. Yang, and X. Guo. Multi-Radio Coexistence: Challenges and Opportunities. In *Proc. of IEEE ICCCN*, 2007.