

Cooperative Carrier Signaling: Harmonizing Coexisting WPAN and WLAN Devices

Xinyu Zhang and Kang G. Shin, *Fellow, IEEE, ACM*

Abstract—The unlicensed ISM spectrum is getting crowded by wireless local area network (WLAN) and wireless personal area network (WPAN) users and devices. Spectrum sharing within the same network of devices can be arbitrated by existing MAC protocols, but the coexistence between WPAN and WLAN (e.g., ZigBee and WiFi) remains a challenging problem. The traditional MAC protocols are ineffective in dealing with the disparate transmit-power levels, asynchronous time-slots, and incompatible PHY layers of such heterogeneous networks. Recent measurement studies have shown moderate-to-high WiFi traffic to severely impair the performance of coexisting ZigBee. We propose a novel mechanism, called *cooperative carrier signaling* (CCS), that exploits the inherent cooperation among ZigBee nodes to harmonize their coexistence with WiFi WLANs. CCS employs a separate ZigBee node to emit a carrier signal (busy tone) concurrently with the desired ZigBee's data transmission, thereby enhancing the ZigBee's visibility to WiFi. It employs an innovative way to concurrently schedule a busy tone and a data transmission without causing interference between them. We have implemented and evaluated CCS on the TinyOS/MICAZ and GNURadio/USRP platforms. Our extensive experimental evaluation has shown that CCS reduces collision between ZigBee and WiFi by 50% for most cases, and by up to 90% in the presence of a high-level interference, all at negligible WiFi performance loss.

Index Terms—802.11, 802.15.4, coexistence, GNURadio/USRP, heterogeneous networks, software radio, TinyOS/MICAZ, WiFi, ZigBee.

I. INTRODUCTION

THE PAST decade has witnessed the proliferation of wireless MAC/PHY standards for establishing wireless local area networks (WLANs) and personal area networks (WPANs) on the ISM band. Allowing spectrum sharing among these networks will undoubtedly improve spectrum utilization. However, it also creates unprecedented challenges, especially the coexistence of incompatible MAC/PHY protocols. Two such networks, WiFi (IEEE 802.11 WLAN) and ZigBee (IEEE 802.15.4 WPAN), that operate in the 2.4 GHz license-exempt band have received considerable attention. WiFi is designed for Internet access, video streaming, etc., whereas ZigBee targets low duty-cycle monitoring and control applications such as healthcare and home/industrial automation. They are expected to run simultaneously in close proximity, e.g., inside

Manuscript received June 22, 2011; revised February 09, 2012; accepted May 11, 2012; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Papagiannaki. Date of publication June 05, 2012; date of current version April 12, 2013.

The authors are with the Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI 48109 USA (e-mail: xzyzhang@eecs.umich.edu; kgshin@eecs.umich.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2012.2200499

a residential or office or hospital building. However, recent measurement studies have shown that ZigBee's performance is severely degraded in the presence of moderate to high WiFi traffic. For example, in an enterprise WLAN and a collocated 90-node WPAN for building energy management, more than a half of the ZigBee links in the WPAN were observed to suffer connection loss during peak hours due to WiFi interference [1]. The harmful coexistence has also been observed in previous small- to medium-scale experimental deployments [2]–[6].

A straightforward way to avoid such interference is to allocate ZigBee devices to channels that are not or less used by WiFi devices [7], [8]. However, the prevailing frequency planning methods presume ZigBee to fall in the same management domain as WiFi, without resolution of the short-term collisions caused by unmanaged, bursty WiFi traffic. Moreover, there are only a limited number of orthogonal WiFi channels on the ISM band, and thus frequency separation is difficult to achieve in a densely deployed WiFi environment [1], [9], where collocated APs tend to occupy different parts of the spectrum to avoid intercell interference. Therefore, it is important to devise a complementary approach that enables ZigBee to share the same frequency band with WiFi, but multiplex the channel over time.

The CSMA-style spectrum etiquette in such networks may seem to be an effective means to achieve this. However, system heterogeneity poses a serious challenge and may severely degrade CSMA-based coexistence schemes. Through link-level measurement of coexisting WPAN and WLAN, we find that the legacy ZigBee MAC experiences a 51% collision rate even when WiFi leaves the channel unused for 67% of the time. A further microscopic study has revealed the root cause of such harmful coexistence. First, ZigBee's transmit power is 20 dB lower than WiFi's, yielding a smaller spatial footprint and, hence, its poor visibility to WiFi. The ZigBee's MAC-layer time resolution is also 16 times coarser, and it can easily be preempted by WiFi in the middle of a rx/tx transition (e.g., sensing-to-transmission or data-to-ACK transition), thus causing collision, even if they can sense each other. In addition, ZigBee allows for TDMA mode, which operates without carrier sensing, and may arbitrarily collide with an ongoing WiFi transmission. The disparate transmit power, time resolution, and scheduling mode also make it difficult to deal with other coexisting networks, such as WiFi and Bluetooth [10] and WiFi and WiMax [11]. Therefore, by resolving the coexistence between ZigBee and WiFi, one could potentially extend the solution to other heterogeneous networks.

Based on the above observations, we propose a new mechanism, called *cooperative carrier signaling* (CCS), to facilitate ZigBee's coexistence with WiFi. CCS builds atop the ZigBee MAC/PHY, but enhances it with a new coexistence management framework, making WiFi better aware of ZigBee's presence,

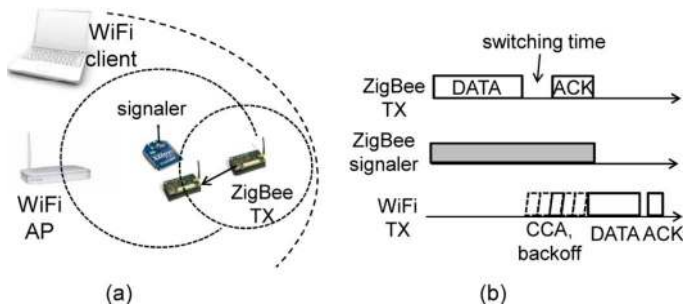


Fig. 1. Principle behind CCS. (a) Spatial domain: allowing WiFi to indirectly sense weak ZigBee signals. (b) Temporal domain: avoiding WiFi preempting in the rx/tx switching time of ZigBee.

and hence achieving better channel sharing. Unlike the traditional CSMA that relies on a data packet as an implicit carrier signal (busy tone), CCS assigns a separate ZigBee node called *signaller* as a proxy to perform the carrier signaling. The signaller may have a higher power than normal ZigBee transmitters, thus allowing the WiFi nodes to sense the ZigBee transmitter's presence indirectly by detecting the busy tone [Fig. 1(a)]. The busy tone persists throughout the data and ACK round trip, thus preventing the WiFi's preemption in the rx/tx switching gap [Fig. 1(b)].

A key challenge to CCS is that the signaller's busy tone must occur concurrently with the data transmission (without interrupting it). To overcome this difficulty, we design a *temporary channel-hopping* mechanism that separates the carrier signaling from the data transmission in frequency domain, yet ensures WiFi to sense the presence of ZigBee transmission. Furthermore, CCS must schedule the busy tone to protect both the TDMA and CSMA packets of ZigBee, without adversely affecting WiFi's performance. We extend the ZigBee's built-in beacon synchronization mechanism to be a scheduler that synchronizes the signaller and the transmitter. The scheduler preserves the carrier-sensing capability on the signaller, so that the busy tone is emitted only when the channel is unoccupied by WiFi.

The cooperative signaling mechanism is triggered when WiFi interference is present and causes severe collisions. This adaptation is facilitated by a two-dimensional carrier-sensing scheme that estimates the WiFi's interference intensity and distinguishes it from ZigBee. We further introduce a simple *signaller configuration* scheme that configures the power and location of the signaller, so that its busy tone may be sensed by multiple, randomly located WiFi interferers. With signal configuration, CCS is applicable even when the WPAN nodes and nearby WiFi interferers are mobile.

We have implemented the CCS framework in TinyOS 2.0, which runs on a ZigBee-based WPAN. We also implemented a dedicated signaller on the GNURadio [12] software radio platform. Extensive experiments on the MICAz motes [13] and dc-powered USRP2 [14] have shown that under moderate to high WiFi traffic, CCS reduces the ZigBee's packet collision by more than 50% in most cases. This translates to significant reductions of packet delay and energy consumption. More importantly, when the WPAN runs low duty-cycle traffic, CCS does not degrade the performance of WiFi, compared to the legacy ZigBee. Since CCS does not require any hardware or firmware

modification, it can be easily integrated into the ZigBee network stack.

The main contributions of this paper are threefold:

- a microscopic analysis of coexisting WiFi and ZigBee signals via software radios, which identifies key factors that account for collision between them;
- a new cooperative framework, CCS, that allows ZigBee WPAN to avoid WiFi-caused collisions. CCS overcomes the inherent limitations of traditional CSMA and can also be extended to enhance the coexistence of other heterogeneous networks;
- implementation and evaluation of the CCS framework on ZigBee motes and a USRP2-based software radio.

The remainder of this paper is organized as follows. In Section II, we review the related work on ZigBee–WiFi coexistence and introduce the differences of these two protocols. Section III analyzes the key factors causing the conflicting coexistence based on fine-grained measurements. Sections IV and V detail the design and implementation of CCS to solve the coexistence problem. Section VI presents our testbed-based evaluation of CCS, while Section VII concludes the paper.

II. BACKGROUND

A. Related Work

1) *Conflicting Coexistence of ZigBee With WiFi*: The interference between WiFi and ZigBee has been extensively studied in both the industry and the research communities. Under light WiFi traffic, ZigBee is known to suffer less from collision with WiFi and can recover loss via retransmission [15], [16]. However, under moderate to high WiFi traffic, ZigBee performance is severely degraded. In an indoor testbed with randomly deployed nodes, Gummadi *et al.* [2] reported that ZigBee experiences a median packet-loss rate of 20%, and the loss could exceed 85% due to WiFi interference. This occurs even when the carrier sensing and packet retransmission are enabled. Pollin *et al.* [3] found that WiFi may interrupt ZigBee transmissions even when they are located close to each other. Similar results have been observed in other measurement studies and real-world applications [1], [4], [5]. With the proliferation of WiFi devices (e.g., smartphones) and high-rate applications (e.g., HD-videos), the amount of WiFi traffic in a typical home or enterprise environment will keep increasing, thus severely affecting the reliability of ZigBee WPANs for monitoring and control applications.

On the other hand, ZigBee seldom interferes with WiFi since it targets low duty-cycle applications with low channel occupancy (typically below 1% [17]). Moreover, WiFi has much higher transmit power, which easily forces ZigBee nodes to back off [2], [4], and can dominate the ZigBee interference. The testbed measurement in [2] has shown that WiFi experiences no packet losses when colocated with multiple ZigBee devices, except its TCP latency increases by about 5%. Controlled experiments reveal that ZigBee may also disturb the 802.11 packet decoding, but only with saturated traffic [3], or without MAC-layer sensing [10].

2) *Alternative Coexistence Mechanisms*: A straightforward way of enabling ZigBee–WiFi coexistence is frequency planning. However, this approach is ineffective when WiFi WLANs

are: 1) unmanaged and may change channels unpredictably; or 2) densely deployed, since three nonoverlapping channels (two for the 802.11n 40-MHz mode) can occupy the majority of 2.4-GHz ISM bands. Strict frequency separation may also underutilize bandwidth, a well-known problem already present in the TV spectrum. When WiFi traffic becomes intensive, ZigBee may adaptively switch to other idle channels [8]. However, this approach does not resolve bursty collisions—it responds only after collision had already occurred. Adaptive channel allocation also incurs long “blackout time” due to scanning and reassociation [18], which can be on the order of several seconds and increases with the network size.

An alternative approach, called WISE [6], changes ZigBee frame size adaptively according to the estimation of the idle interval between WiFi transmissions. WISE needs to suspend ZigBee transmissions in each WiFi burst and is unsuitable for TDMA packets or delay-sensitive applications. Liang *et al.* [1] proposed use of error-correcting code to recover partially corrupted ZigBee packets due to WiFi interference. The mechanism is corrective in nature, whereas CCS is a preventive approach that avoids collisions.

In [2], a spectrum-survey-based method is introduced to improve WPAN–WLAN coexistence by adjusting transmit power and carrier-sensing threshold. This approach is suitable for static networks and aims for long-term throughput fairness between heterogeneous networks. However, short-term performance metrics (e.g., delay and collision rate) are equally important to the monitoring and control applications typically seen on ZigBee networks.

Rahul *et al.* [19] proposed an interference-nulling approach that forces wideband devices to allocate idle spectrum for narrowband devices. This approach requires hardware modification to perform customized signal processing. Moreover, the devices must be able to sense each other so as to determine which part of the spectrum should be nulled.

In contrast to the above approaches, CCS attempts to overcome the inherent limitations of traditional CSMA and enable it to coordinate heterogeneous networks. Our key observation is that a sufficient idle channel time exists and can be exploited by ZigBee, but the WiFi’s unawareness causes severe collisions. By enhancing the visibility of ZigBee to WiFi while preserving the CSMA-based spectrum etiquette, CCS can substantially improve ZigBee’s channel utilization without compromising WiFi’s performance.

Busy-tone-based signaling mechanism was first adopted by Tobagi and Kleinrock [20] as a solution to the hidden terminal problem in CSMA protocols. Many subsequent proposals extended the mechanism to enhance the performance of CSMA for ad hoc and sensor networks [21], [22]. These protocols usually require two radios on each transceiver: one for data transmission, and the other emitting a dedicated narrowband signal as the busy tone. In contrast, CCS adopts a cooperative busy-tone mechanism, which can be realized using off-the-shelf ZigBee devices. The objective of CCS differs from previous busy-tone mechanisms in that it enables the coexistence between *heterogeneous* MAC protocols.

We also proposed a protocol, called *Cooperative Busy-Tone* (CBT) [23], that adopts a mechanism similar to CCS, i.e., a separate signaller is employed to make WiFi aware of ZigBee




Fig. 2. Superframe in ZigBee. The beacon interval (BI) and superframe duration (SD) depend on beacon order (BO) and superframe order (SO), such that $0 \leq SO \leq BO \leq 14$. The BD value equals 15.36 ms.

transmissions, thus reducing collision between ZigBee and WiFi. However, that work focused mainly on establishing a theoretical framework to analyze the performance of CBT and using simulation to validate the analysis. In this paper, we focus on detailed system design, testbed implementation, and evaluation of the CCS mechanism. We perform detailed experiments to investigate the cause and effect of the conflicting coexistence between ZigBee and WiFi and further redesign the signaller’s functionalities (e.g., CSMA scheduler, signaller configuration, and signal classifier) to facilitate practical deployment.

B. ZigBee Versus WiFi: MAC/PHY Layers

ZigBee specifies the MAC/PHY functionalities to establish low-power, low-rate WPANs. Each ZigBee WPAN assigns a unique *coordinator* to perform association control and beacon scheduling for *clients*. The coordinator schedules a mixture of TDMA and CSMA frames periodically. Each scheduling period is called a *superframe*, which starts with a beacon, followed by a number of CSMA slots (called *CAP*) and TDMA slots [called *CFP* or *Guaranteed Time Slot* (GTS)], and then an inactive period, as illustrated in Fig. 2.

The CFP slots are allocated to clients and deallocated on demand. In the CAP slots, ZigBee enforces slotted-CSMA access control, which differs from WiFi as follows. First, contention access must start from the boundaries of basic time units called *backoff slots*, each lasting $320 \mu\text{s}$. In contrast, the WiFi backoff slot is only 20 (802.11b) or $9 \mu\text{s}$ (802.11a/g/n). Second, when sensing a busy channel, ZigBee resumes its backoff and clear channel assessment (CCA) and aborts after five consecutive attempts. In contrast, WiFi persists in backoff and sensing until it finds an idle slot for transmission. Third, each backoff in ZigBee consists of two contention windows, i.e., a transmitter must ensure an idle channel for two slots ($640 \mu\text{s}$) before sending data, whereas WiFi needs only one (20 or $9 \mu\text{s}$) idle slot.

On the PHY layer, ZigBee’s bit rate is limited to 250 kb/s. Its CCA operation takes $128 \mu\text{s}$, and the rx/tx switching time can be $192 \mu\text{s}$ due to the hardware limitation. Both the long backoff time at the MAC layer and slow response at the PHY render ZigBee low-priority when WiFi transmitters attempt to access the channel. Since the WiFi CCA duration is much shorter, a WiFi transmitter can easily preempt ZigBee in the middle of a ZigBee’s rx/tx switching slot (thus causing collision), even when both can perfectly sense each other’s presence. Moreover, the maximum transmit power of ZigBee is only 0 dBm, whereas WiFi typically transmits at 15–20 dBm. Hence, ZigBee has a much shorter interference range than WiFi [10]. Therefore, ZigBee may not be effectively sensed by WiFi when colocated.




Fig. 3. Experimental testbed with WiFi and ZigBee nodes. By default, the WiFi link is $A \rightarrow C$. ZigBee location is varied in different experiments.

It should be noted that ZigBee also produces special devices on the 868/915-MHz band, but the rate (20/40 kb/s) may not be suited for all applications. Some WiFi standards (e.g., 802.11n) can operate over the 5-GHz band that is orthogonal to ZigBee. However, for compatibility and extended range, 802.11n tends to be configured to the 2.4-GHz band where the dominating 802.11b/g reside.

III. CAUSE AND EFFECT OF CONFLICTING COEXISTENCE

In this section, based on extensive testbed experiments, we anatomize the inherent deficiency of the collision avoidance mechanism in WiFi and ZigBee, which accounts for their conflicting coexistence.

Our testbed includes both WiFi and ZigBee nodes, deployed in an office environment. The floorplan and measurement points are shown in Fig. 3. The WiFi transceivers adopt Atheros 5413 NIC with the MadWiFi v0.9.4 driver, default transmit power 15 dBm. The ZigBee nodes are MICAz motes programmed with openzb [24], an open-source implementation of the ZigBee MAC/PHY. The motes use default transmit power -5 dBm. To isolate the ambient interference, the WiFi link is tuned to a channel least used by nearby WLANs, and all measurements are made in the night.

We found the ZigBee link loss rate strongly depends on WiFi's airtime usage Γ_w , defined as the fraction of time spent on transmission (approximately equal to the ratio of application-layer traffic rate to PHY-layer data rate). Among eight randomly selected transmitter/receiver pairs running ZigBee's TDMA mode, the median collision probability ranges from 9% when $\Gamma_w = 5.6\%$ and to 51% when $\Gamma_w = 33\%$. The CSMA mode performs only slightly better than TDMA. This link-level experiment implies that ZigBee's performance is severely degraded even though WiFi leaves sufficient idle time. The result is consistent with existing measurement studies [1]–[6]. More detailed results are omitted to avoid repetition. In what follows, we focus on more fine-grained experiments that reveal the root effects behind the conflicting coexistence, which are also the motivation behind the CCS principle.

A. Spatial Collision Hazards

1) *Asymmetric Interference*: Due to their disparate transmit power levels, there exists a “gray region” where ZigBee can hear WiFi, but WiFi is oblivious of ZigBee and can arbitrarily interrupt its transmission, so called *asymmetric interference*. To

TABLE I
ASYMMETRIC INTERFERENCE REGION IN THE TESTBED

txpower	Nodes under asymmetric interference
0dBm	L, I
-5dBm	I, L, P
-10dBm	I, L, P, H, G, N, O
-15dBm	I, L, P, H, G, N, O, F, C
-25dBm	I, L, P, H, G, N, O, F, C

profile this effect, we measure the region within the testbed where WiFi's signal strength exceeds ZigBee's carrier-sensing threshold, but not vice versa. We use an interference classification scheme (which will be detailed in Section IV-D) to measure the WiFi signal power received by a MICAz mote. To verify if WiFi can sense ZigBee, we calibrate the power level of the USRP2 software radio [14] to that of MICAz¹ and allow it to send a continuous stream of ZigBee packets. When WiFi can sense the USRP2 transmitter, it keeps deferring the transmission, thus increasing the packet delay dramatically.

Table I shows the set of ZigBee nodes vulnerable to asymmetric interference at various transmit power levels. At 0 dBm, only nodes far away from the interferer experiences asymmetric interference. As their transmit-power level decreases, the number of vulnerable nodes increases. Below -10 dBm, a majority of nodes suffer. This is because ZigBee transmitters implicitly use the data packet as a carrier signal (busy tone), expecting it can reach the WiFi interferer. However, such carrier signaling becomes less effective as transmit power decreases and when ZigBee is located farther away from the WiFi transmitter.

To combat asymmetric interference, a natural solution is to employ a proxy signaler with higher transmit power to send the busy tone. Whenever a node within the ZigBee WPAN starts transmission, the signaler would initiate the busy tone simultaneously, so as to prevent WiFi interruption. This idea constitutes a key intuition behind CCS.

2) *Hidden Terminal*: This problem has been studied extensively for 802.11 and remains a cause for spatial collision hazards in heterogeneous networks. It occurs when the WiFi and ZigBee transmitters cannot hear each other. This is the case for the location K and M , when the WiFi transmitter is at A , according to our measurement. Similar to the asymmetric interference, the hidden terminal effect can be alleviated using a proxy signaler visible to the WiFi transmitter.

B. Temporal Collision Hazards

1) *Partial Carrier Sensing*: In addition, collision hazards can occur in the time domain when a WiFi packet is partially sensed by ZigBee and is insufficient to trigger its backoff.

Intuitively, the RSSI is independent of the transmitter's duty-cycle. However, through detailed measurement, we found that the WiFi signal strength sensed by a ZigBee receiver is stable

¹Calibration is needed because the USRP2's output power level is unknown. For calibration, we place the MICAz and USRP2 transmitters near each other and allow them to send ZigBee packets alternately to the same MICAz receiver. We tune the power level of the MICA transmitter and the PHY parameters of USRP2 (including signal amplitude and transmit gain), so that the MICAz receiver reads the same RSSI values when receiving from both of them. This way, we map the parameter configuration of the USRP2 transmitter to the power level of the MICAz transmitter (in dBm scale).

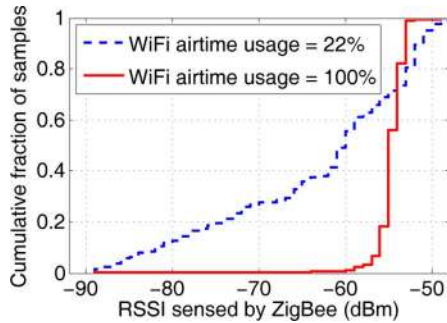


Fig. 4. Distribution of WiFi signal strength sensed by ZigBee.




Fig. 5. Collision hazards in the temporal domain. (a) Partial carrier sensing. (b) WiFi preemption.

when $\Gamma_w = 100\%$,² but it varies from -50 dBm down to the noise floor when $\Gamma_w = 22\%$ (Fig. 4). This implies that certain WiFi packets are partially sensed during the long sensing period of ZigBee, which occurs when WiFi starts transmission near the end of the ZigBee sensing duration, as illustrated in Fig. 5(a). Since the WiFi's energy level needs to be averaged over ZigBee's CCA duration ($128 \mu\text{s}$) to calculate the RSS, the resulting RSS is insufficient to trigger a ZigBee backoff, exposing it to harmful WiFi interference.

2) *Non-CSMA Transmission and WiFi Preemption*: Intuitively, the spatial collision hazards and partial carrier sensing should be less severe when WiFi and ZigBee transmitters are close to each other and RSS exceeds the CCA threshold. However, according to our measurement, even when they are 1 m apart (e.g., WiFi $A \rightarrow C$ and ZigBee $D \rightarrow E$), the collision rate can still exceed 30%. To capture the root cause, we use a software oscilloscope in GNURadio to log the sample-level traces of the channel (with 20 MHz sampling rate) and then calculate the power of the received signal. Fig. 6 illustrates a snapshot of traces when the above two links are transmitting using CSMA.

These traces reveal more deficiencies of ZigBee in avoiding WiFi interference. First, packets sent without sensing, such as GTS, ACK, and beacons, will be corrupted when encountering an ongoing WiFi session. Second, WiFi can preempt a ZigBee transmission when its carrier sensing falls in the rx/tx switching time of ZigBee transmitters ($192 \mu\text{s}$) or the time spent in waiting for the next CSMA slot boundary (up to a backoff slot, $320 \mu\text{s}$ [25]), as also illustrated in Fig. 5(b). These two cases are essentially due to the long response time of the ZigBee hardware.

²We calibrate the USRP2 transmit power to that of WiFi and use it to send a continuous stream of WiFi packets. The WiFi packets are generated using the BBN 802.11b module in GNURadio, which are compatible with legacy 802.11 packets.




Fig. 6. Temporal collision hazards captured in the traces of a GNU-Radio/USR2 oscilloscope.




Fig. 7. Architecture of the CCS framework.

Similar to the spatial conflict, temporal collision hazards can also be avoided if the carrier signaling operation is assigned to a separate ZigBee signaller. The signaller can notify WiFi before ZigBee's TDMA packets (or before CCA) by sending a busy tone and continue such signaling during the switching time of the colocated ZigBee sender, thereby preventing WiFi preemption. This establishes the key motivation and rationale behind CCS.

IV. DESIGN OF CCS

CCS harmonizes the coexistence of ZigBee WPAN with WiFi, via cooperation among the coordinator, clients, and a special ZigBee node designated as the signaller. Fig. 7 depicts the CCS architecture. CCS runs atop the ZigBee MAC/PHY layers. It incorporates a signal classifier that triggers the signaller by estimating WiFi's interference intensity and a coexistence manager that coordinates the behavior of the ZigBee nodes to prevent WiFi interruptions. The coexistence manager consists of three components:

- Frequency domain: a *temporary channel hopper* that prevents the signaller from interrupting the ongoing transmission of a ZigBee data packet;
- Temporal domain: a *signaling scheduler* that ensures the busy tone sent by the signaller protects the desired data packet, while conforming to the CSMA spectrum etiquette;
- Spatial domain: a *signaler configuration* framework that configures the location and power of the signaller, given a coarse estimation of network parameters, e.g., the maximum link distance of ZigBee.

At a high level, CCS first performs the signaller configuration when the WPAN (including one coordinator and multiple clients) is established. Using the signal classifier, the ZigBee nodes obtain an estimate of the intensity of WiFi interference and packet-loss rate. Based on this estimate, the coordinator decides whether or not to trigger the signaling mechanism. When it is activated, the signaller runs the temporary channel hopper

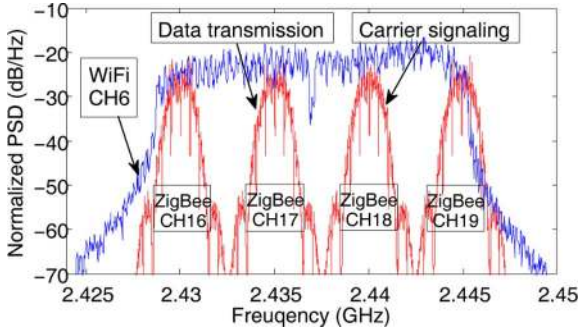


Fig. 8. CCS adopts *temporary channel hopping* to prevent mutual interference between carrier signaling and data transmission.

to prevent its transmission of busy tones from interfering with data packets. The busy tones are scheduled jointly by the coordinator, the clients, and the signaler, according to the ZigBee’s MAC mode (TDMA or CSMA).

Note that the signaler must be ZigBee-compatible, i.e., it must be able to decode ZigBee packets so as to cooperate with the coordinator/clients in the WPAN. The signaler needs special features only in the following two cases.

- When a certain link in the WPAN is long, the signaler requires higher transmit power than normal ZigBee nodes in order to ensure its visibility to potential WiFi interferers. In Section IV-C, we characterize the required signaler’s transmit power as a function of such network parameters as the maximum WPAN link distance and the WiFi’s carrier-sensing threshold.
- If the WPAN runs delay-sensitive applications, CCS needs to be active persistently to guard against bursty WiFi interference. In such cases, a dc-powered ZigBee node is necessary as a signaler.

There already exist commercial off-the-shelf ZigBee modules (e.g., the XBee Pro [26]) that meet the above requirements.

In what follows, we detail the design of CCS’s components.

A. Temporary Channel Hopping

To make the signaling effective, the signaler must emit the busy tone concurrently with the clients’ or coordinator’s data transmission. However, the signaler must not interrupt the ongoing or forthcoming transmission from the desired transmitter. A straightforward way to meet this requirement is to ensure sufficient spatial separation between the signaler and the transmitter. However, this solution is too restrictive since the transmitter needs to fall in the same WPAN as the signaler, and likely within its interference range.

CCS addresses this problem with a *temporary channel-hopper* that leverages the inherent spectrum features of ZigBee and WiFi (Fig. 8). In the 2.4-GHz spectrum, the width of each WiFi channel is 20 MHz, and the i th channel is centered at $(2.407 + 0.005i)$ GHz, $i \in [1, 11]$. Adjacent channels partially overlap with each other. For ZigBee, the k th channel is centered at $[2.405 + 0.005(k - 11)]$ GHz, $k \in [11, 26]$ [25]. Each ZigBee channel occupies 4 MHz, with 1 MHz guard band between adjacent channels. As a result, each WiFi channel overlaps with four ZigBee channels.




Fig. 9. Scheduling data transmission and carrier signaling. (a) CSMA scheduler. (b) TDMA scheduler.

When running the temporary channel-hopper, a ZigBee signaler switches to a nearby channel before its scheduled signaling and returns to the original channel immediately after the busy tone is sent. This approach ensures signaling is decoupled from ZigBee transmission in frequency domain, as adjacent ZigBee channels are orthogonal. However, the busy tone still overlaps with the WiFi spectrum and can inform WiFi of a ZigBee transmission, as long as its power exceeds the WiFi’s carrier-sensing threshold.

Note that in practice, a ZigBee channel under interference may reside at the edge of a WiFi band, so CCS must decide on hopping to the left- or right-side adjacent channel. In Fig. 8, for example, when ZigBee runs on channel 19 and is interfered with by WiFi channel 6, the signaling would be ineffective if it hops to the right-side channel 20. To solve this problem, the ZigBee signaler first sends busy tones on the left-side channel 18. If the packet loss is not alleviated, it switches to the right-side channel 20 instead for signaling. If packet loss persists, CCS recognizes that the current channel 19 is interfered with by two partially overlapping WiFi bands (channels 6 and 9, spanning 2.426–2.448 and 2.441–2.463 GHz, respectively). Such a problem occurs only when both bands carry intensive traffic, and can be resolved using two signalers that hop to channel 18 and channel 20, respectively.

The channel-hopper incurs channel-switching overhead to the signaler. However, the switching time is limited to 192 μ s in ZigBee [25]. This overhead is equivalent to only 4 B of airtime and is the same as the rx/tx switching time.

B. Signaling Scheduler

CCS maintains the legacy scheduling protocol in ZigBee, but requires the signaler to dispatch the busy tone at an appropriate time, such that: 1) it reduces the WiFi preemption over ongoing or forthcoming ZigBee transmissions; and 2) it minimizes the potential influence on WiFi performance. The signaling scheduler is designed to address this tradeoff. It allows both the CSMA and TDMA modes of ZigBee to coexist with WiFi.

1) *CSMA Scheduler*: The CSMA scheduler is akin to the *indirect transmission* mode in ZigBee [25]. Specifically, a sender polls the receiver before delivering data, and the receiver returns a 5-B confirmation packet when it is ready to receive (we refer to the polling and confirmation packet as *RTS*, *CTS*, respectively). Upon overhearing the CTS confirmation, the signaler starts the temporary channel-hopper and emits the busy tone immediately. This handshake process is illustrated in Fig. 9(a).

The busy-tone duration equals the data packet length plus the ACK wait duration and a guard period. The data packet length is a 1-B field piggybacked in the CTS. The ACK wait duration

includes the airtime of ACK packet (352 μ s), the rx/tx switching time (192 μ s), plus a backoff slot (320 μ s) that is needed to ensure slot-boundary alignment [25].

To avoid unnecessary interruptions to WiFi, the signaler also senses the channel after switching to the new channel. It starts signaling only if the channel is idle throughout one CCA slot, and aborts the signaling if detecting a busy channel over five consecutive CCA attempts. The other ZigBee transmitters are oblivious of the signaler's behavior and need to perform CCA independently. Since the signaler does not perform backoff and only needs one CCA slot to assess a clear channel, it tends to send the busy tone before the data transmission. To compensate for this time offset, the busy tone is extended by a guard period, which equals a CCA slot plus half of the maximum backoff window (four backoff slots).

To reduce the overhead due to excessive CCA and backoff, the RTS/CTS packets are sent without carrier sensing. As a result of this, RTS/CTS may be lost due to collision with WiFi packets. However, since the RTS/CTS packets are much shorter than data packets, the collision probability is lower. To further reduce such losses, each RTS packet is retransmitted for RETX times, and the receiver replies with a CTS whenever it receives an RTS. The RETX value is piggybacked as a 1-B tag in the CTS packet. The signaler schedules its signaling time according to the first tagged CTS packet it receives.

2) *TDMA Scheduler*: In TDMA mode, CCS exploits the RTS mechanism in ZigBee to allocate fixed slots to clients, thus eliminating the need for per-packet handshake. The slot allocation information is carried in the coordinator's beacon message. Both the client and the signaler send a confirmation packet via CSMA after receiving such a beacon. The beacon is retransmitted if the confirmation is missing. Following a successful slot allocation, the signaler will send the busy tone whenever a scheduled TDMA slot fires, as shown in Fig. 9(b). To reduce unnecessary interference to an ongoing WiFi transmission, the signaler starts CCA δ units of time earlier than the scheduled ZigBee transmission (δ is called *presignaling time*). It starts signaling on the first idle CCA and cancels the signaling if the channel remains busy before the TDMA transmission. The presignaling time δ is a design knob serving two purposes: 1) it tolerates imperfect synchronization due to the clock skew between the signaler and transmitter; 2) it can be used to raise the priority level of ZigBee when running delay-sensitive applications (e.g., real-time monitoring). A larger value of δ allows the signaler to find an idle slot with a higher probability, but at the cost of extra channel time. In our current design, δ is set to the duration of five CCA attempts.

In addition to the data packets, the TDMA scheduler can protect beacon messages. Beacon protection is critical for reducing the packet delay since ZigBee allows for TDMA packet transmission only after successful reception of a beacon for the corresponding superframe. Otherwise, the packet must be postponed to future superframes, implying a typical delay of 1 s. In CCS, the coordinator transmits a CTS packet immediately before the beacon, which will be used by the signaler as a sync message for starting a busy tone to protect the beacon. Since beacons are short (11 B by default), sent infrequently, and tend to have high priority, the signaler is forced to send the busy tone at the due time of each beacon, assuming that WiFi can promptly recover

itself via backoffs and retransmissions, even if collision occurs and corrupts its data transmission.

C. Signaler Configuration

When using a dedicated signaler, one must carefully configure its location and transmit power so that the busy tone may be sensed by the potential WiFi interferers randomly located near the ZigBee WPAN or moving around. Consider a WPAN with one coordinator and multiple clients, with either uplink or downlink traffic. An optimal configuration method may place one signaler near each receiver. However, this will significantly increase the deployment cost. It might also be possible to adaptively place the signaler for the receiver that needs protection. However, it is difficult to realize this due to the random network topology, dynamic traffic pattern, and nodes' mobility.

We adopt a more practical solution in CCS: A single signaler is placed near the coordinator so that all nodes in the WPAN may be equally protected. We verify this intuition based on the insights gained from an empirical propagation model. We show that with an appropriate level of signaling power, the signaler's busy tone can be visible to any potential WiFi interferers with high probability and can protect packets addressed not just to the coordinator (uplink traffic), but also to all clients within the WPAN (downlink traffic).

Let D_m be the maximum separation between the coordinator and any client. Let Λ_w , Λ_z , and Λ_s be the transmit power of WiFi, ZigBee, and the signaler, respectively. The WiFi transmitter can cause packet loss (hence becoming a potential interferer) only if its signal power exceeds the desired ZigBee packet power by a capture threshold C_a . To model the signal attenuation, we use an empirical propagation model recommended by the IEEE 802.15 for 2.4 GHz indoor environment [17]. At distance d , the signal's path loss (in dB) is

$$L_{\text{dB}}(d) = \begin{cases} 40.2 + 20 \log_{10}(d), & d \leq 8m \\ 58.5 + 33 \log_{10}\left(\frac{d}{8}\right), & d > 8m. \end{cases}$$

Consider uplink first (i.e., the coordinator is the receiver). Given D_m , there exists a maximum distance M_{wr} between the coordinator and potential interferer that can cause packet loss, which satisfies

$$\Lambda_w [L(M_{\text{wr}})]^{-1} \left(\frac{4}{20}\right) = \Lambda_z [L(D_m)]^{-1} C_a^{-1} \quad (1)$$

where $L(d) = 10^{0.1L_{\text{dB}}(d)}$ is the path loss in normal scale. The term $\frac{4}{20}$ represents the effective power received by WiFi since ZigBee and WiFi occupy 4- and 20-MHz bandwidth, respectively. Note that M_{wr} is also the maximum separation between the signaler and any potential interferer. To make the busy tone sensible by WiFi, the signaler's transmit power Λ_s must satisfy

$$\Lambda_s [L(M_{\text{wr}})]^{-1} \geq W_{\text{cs}} \quad (2)$$

where W_{cs} is the carrier sensing threshold of WiFi. From (1) and (2), we obtain the minimum signaler power for protecting the entire WPAN under arbitrarily located interferers

$$\Lambda_s \geq \frac{\Lambda_w}{5\Lambda_z} W_{\text{cs}} L(D_m) C_a. \quad (3)$$




Fig. 10. Required signaling power to ensure the WPAN is sensible by WiFi: (a) uplink; (b) downlink.

For downlink traffic, consider WiFi interferers around a ZigBee client, again with the maximum separation M_{wr} . In the worst case, the distance between the interferer and the signaler is $M_{wr} + D_m$. Therefore, the required signaller power for protecting ZigBee in the worst case is

$$\Lambda_s \geq \frac{\Lambda_w}{5\Lambda_z} W_{cs} L(D_m + M_{wr}) C_a. \quad (4)$$

Fig. 10(a) shows the required power that makes the signaller visible to any potential interferers colocated with the WPAN. We configure the parameters to their typical values (converted to dB scale): $\Lambda_w = 15$ dBm, $\Lambda_z = 0$ dBm, $C_a = 10$ dB. We consider two carrier-sensing thresholds: $W_{cs} = -81$ dBm, which is the default for many commercial WiFi cards [27], and $W_{cs} = -62$ dBm, which is the maximum allowable energy sensing threshold for 802.11a/g/n [28, Sec. 17.3.10.5]. From these results, we observe that the downlink requires higher signaling power since the signaller is located farther away from the clients than from the coordinator. However, the difference is insignificant since the signal strength drops sharply with distance. The required signaling power increases with the ZigBee link distance D_m and with the WiFi carrier sensing threshold W_{cs} . However, even when $W_{cs} = -62$ dBm, a normal ZigBee signaller with 0 dBm power can be effective for short-range WPANs (e.g., $D_m < 2$ m, as in body-area networks). For longer ZigBee links, we need a dedicated signaller that has power comparable to WiFi's (above 15 dBm). According to FCC part 15 rules [29], the maximum transmit power of devices on the 2.4-GHz ISM band is 30 dBm (1 W). Fig. 10 shows that a signaller with 30 dBm transmit power can protect a WPAN even when the link distance is up to 20 m (a typical ZigBee link distance is only around 6 m [10]) and the WiFi's carrier sensing threshold is as low as -81 dBm.

From the above analysis, we conclude that *by placing a single signaller near the coordinator, CCS can prevent the entire WPAN from being interfered with by randomly located (including mobile) WiFi transmitters*. This approach makes CCS deployable in mobile WPANs, as the signaller can always be colocated and moving together with the coordinator. In practice, due to the additional variations caused by small-scale fading (e.g., multipath and Doppler effects), the empirical propagation equation cannot model all cases accurately. However, it allows us to flexibly explore the effects of all possible power-location configurations and verify the effectiveness of CCS in the average cases. The propagation (1) can also be replaced by other empirical models in order to obtain more accurate estimation for the required signaling power.

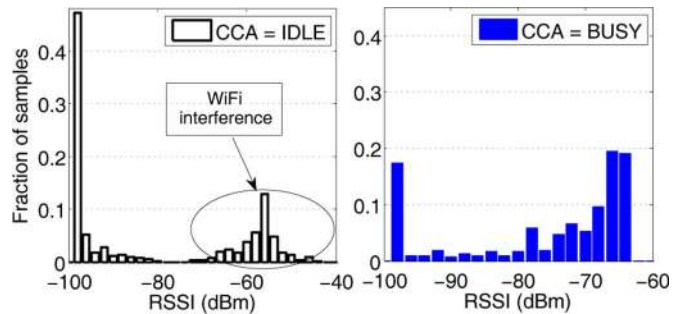


Fig. 11. Distribution of RSSI for different CCA (mode 2) decisions. CCS singles out WiFi interference based on the fact that WiFi signals' RSSI values may exceed the energy-sensing threshold, but CCA still returns "idle."

D. Signal Classifier

Recall that the signaller needs a signal classifier to differentiate WiFi interference and estimate its intensity. CCS exploits the hardware CCA capability of ZigBee devices to achieve this. ZigBee devices can perform three different modes of CCA: energy sensing (mode 1), feature detection (mode 2), and a mixed mode (mode 3). Mode 1 returns busy if the RSSI exceeds a carrier sensing threshold. Mode 2 returns busy only when a valid 802.15.4 signal with the specific spreading and modulation features has been detected [25]. Mode 3 performs a logical OR over the above two modes.

To classify the interference, CCS uses a two-dimensional carrier-sensing scheme that combines the RSSI measurement and CCA feature detection mode (mode 2). In particular, *signals that have high RSSI (i.e., exceeding the energy sensing threshold) but cannot be detected via CCA mode 2 are classified as WiFi interference*. In typical ZigBee-compatible hardware such as the MICAz mote, although only a single CCA mode is allowed at any time, the RSSI register can be read simultaneously with the CCA, thus making the signal classifier feasible.

To verify this simple scheme, we use a MICAz mote E (Fig. 3) to collect the RSSI values and CCA decisions (as required in CCS), when a ZigBee D and WiFi node A are transmitting. We allow the MICAz sensor to read the RSSI register every $128 \mu\text{s}$, the minimum time needed for a valid RSSI value [30].

Fig. 11 shows a histogram of RSSI values sampled over 60 s. Although partial sensing persists and RSSI varies, WiFi interference can still be singled out by combining the two CCA modes. The four different combinations between two logical decisions (if RSSI is larger than the sensing threshold, and if CCA returns idle) are clustered around different RSSI values and are separated by more than 20 dB from each other. Note that WiFi signals below ZigBee's CCA threshold may be misclassified, but this does not fundamentally affect CCS since such signals are less likely to interfere with ZigBee transmissions.

After classifying WiFi interference, the clients periodically report the mean RSSI reading to the coordinator as an indication of the interference intensity. The coordinator triggers the carrier signaling if the interference level exceeds the RSS of any link by the capture threshold C_a . Carrier signaling is also triggered if any link experiences a packet-loss rate that exceeds an application-defined threshold.

E. Application to Other Heterogeneous Networks

CCS can be applied to the coexistence of different protocols that have heterogeneous PHY characteristics and incompatible MAC layers. Specifically, CCS can be used when two coexisting networks have heterogeneous transmit power, scheduling mode, or time resolution. To enable the temporary channel hopping, these networks must be able to tune to different bands. Moreover, at least one of the them must be able to defer its transmission when sensing a busy tone. We briefly discuss one potential extension of CCS to heterogeneous networks consisting of Bluetooth and WiFi. A detailed study of this is left as our future work.

Bluetooth has been a popular means of establishing WPANs. It has a similar level of transmit power as ZigBee, but runs frequency-hopping based on a TDMA schedule. Existing measurement studies have revealed severe collision problems when a Bluetooth WPAN is colocated with WiFi WLANs (e.g., [2] and [10]). To prevent interference between Bluetooth and WiFi on the same hardware platform (e.g., laptop), existing solutions mostly allow them to access the medium alternately [11]. For nodes on different platforms, the IEEE 802.15.2 standard [31] proposed adaptive channel hopping that allows Bluetooth to hop within the spectrum unused by WiFi. The limitation of this approach has been discussed in Section II-A.

CCS can be used as a complementary approach to allow Bluetooth and WiFi to share the same spectrum via a dedicated Bluetooth-compatible signaler. Since the hopping sequence is known to all nodes in the WPAN, the signaler can hop to the next frequency band before the clients, and perform CCA and signaling, similar to the TDMA mode for ZigBee (Section IV-B). Bluetooth has only 1-MHz bandwidth, and consecutive channels are orthogonal to each other, hence the signaling does not interrupt the data transmission. Note that Bluetooth radios can measure RSS and perform CCA directly, but enabling the deferral in the signaler requires firmware modification.

V. IMPLEMENTATION

A. Implementation on TinyOS/MICAz

We have implemented CCS on the TinyOS/MICAz platform using the nesC language. Our implementation builds atop `openzb` [24], an open-source implementation of the ZigBee MAC/PHY in TinyOS2.0. We have built the major components of CCS, including the CSMA/TDMA scheduler, temporary channel-hopper, and interference classification algorithms.

We use the 32 768-Hz alarm clock in MICAz as an internal timer for the CCS scheduler. The resolution of this timer is $30.5 \mu\text{s}$, and the best approximation to the $320 \mu\text{s}$ slot resolution in ZigBee is 10 ticks ($305 \mu\text{s}$). Currently, we have not incorporated any sophisticated clock-calibration mechanism into the CCS scheduler. Using the USRP2 software radio as a sniffer, we found the clock jitter of the ZigBee timer (without calibration) is on the order of several milliseconds over one superframe (986 ms), which is comparable to the airtime of a data packet. Therefore, it is impractical to rely on the built-in timer to follow a fixed schedule to protect the beacon and GTS packets. We thus allow the coordinator to send two CTS packets before beaconing, and before each CFP slot, as synchronization pilots. Due to CTS

losses, the actual performance of our implementation would be lower than one with a well-calibrated timer.

B. Implementation on GNURadio/USRP2 and ns-2

We have also implemented a high-power *dedicated ZigBee signaler* based on the GNURadio/USRP2 software radio platform [12], [14]. One advantage in using such a dedicated signaler is that it does not restrict packet size, unlike the 127-B limitation of typical ZigBee devices. A single USRP2 can emit a sufficiently long carrier signal that covers the data-ack round-trip time of ZigBee packets. In addition, USRP2 has a maximum power level of 20 dBm that is comparable to the WiFi transmitter. Hence, it can protect ZigBee WPANs with long link distance (Section IV-C).

The key challenge in realizing CCS on GNURadio/USRP2 is that it does not yet support delay-sensitive MAC operations because of its inefficient user-mode signal processing modules. For instance, our first implementation of the beacon protection mechanism was based on the GNURadio 802.15.4 library [32] (we modified the library so that it can support USRP2 and can transmit/receive packets compatible with the MICAz motes). This direct implementation results in a several-milliseconds response time (from receiving a CTS packet to hopping channel and sending the busy-tone packet), whereas the response time of a MICAz mote is around $192 \mu\text{s}$. Therefore, we performed the following optimization to reduce the USRP2 signaler's response time to a level comparable with a MICAz signaler.

We migrate relevant python components (used to glue the signal processing modules) to C++ level and implement a radio controller that interfaces directly with the USRP2. We incorporate the transmit path and the receive path into a single module. These two paths are separated in GNURadio, and the context switching between them is a major source of latency. Furthermore, since the signaling effectiveness only depends on the power level rather than the actual content of the signaling packet, we preprocess the signaling packet and log the digital samples corresponding to its modulated signals, which will be emitted by the USRP2 signaler as a busy tone.

In CSMA mode, whenever the signaler receives a CTS packet, it first switches to the signaling channel, blocks waiting for the scheduled signaling time, and then feeds the prefetched signaling packet directly into the USRP2 transmit buffer via the radio controller. With this measure, we are able to reduce the mean response time of signaling to around $200 \mu\text{s}$ and jitter to $100 \mu\text{s}$, which is even shorter than the slot resolution of ZigBee ($320 \mu\text{s}$). An early version of our implementation also attempted to use a passband filter to separate the signaling channel from the data channel. However, the filter involves intensive floating point multiplication, and the processing delay (4–5 ms) is unacceptably long for our purpose.

For TDMA mode, the PC host's timer is unable to achieve a microsecond scheduling accuracy, especially when running concurrently with the signal processing blocks. In our actual implementation, the MICAz transmitter needs to send a CTS message so as to sync with the USRP2 signaler on a per-packet basis. Since its switching delay between sensing and transmission is still on the order of milliseconds, the USRP2 signaler is unable to perform carrier sensing before emitting the busy tone. Hence, it will affect WiFi performance more than a full-fledged

implementation on carrier-sensing-based ZigBee nodes. Further evaluation of this approach will be discussed in Section VI-E. Note that both the CTS overhead and the lack of carrier sensing result from the limitation of the software radio platform. Implementation of CCS on a ZigBee-compatible high-power transceiver [26] would not suffer from these problems and is left as our future work.

Besides the testbed implementation, we have also implemented CCS in ns-2 (version 2.33) to perform trace-based simulation. Following the 802.15.4 specification, we developed a GTS management module in addition to the existing CSMA module in ns-2. We also implemented the core components of CCS, including the scheduler and channel-hopper schemes, in order to flexibly evaluate the energy consumption of CCS.

VI. EVALUATION

In this section, we evaluate CCS's performance via testbed and simulation experiments. Our testbed configuration is the same as the measurement setup in Section III. The main findings from our evaluation are summarized as follows.

- CCS reduces the WiFi-caused collision by 42%–90%, depending on the WiFi airtime usage, relative location, scheduling mode, etc. Compared to the retransmission mechanism, it reduces not only packet loss, but also the average packet delay by up to 63%.
- CCS can improve the performance of an entire WPAN with multiple ZigBee nodes, which coexist with randomly located WiFi transmitters. Compared to an ideal adaptive channel allocation protocol, it achieves a similar long-term throughput, but 59.4% less disruption time.
- CCS may consume 10+% more energy when running the CSMA scheduler and a ZigBee signaller. In TDMA mode, however, it saves energy by 73.2%–83.3% with a dedicated signaller.
- WiFi performance is degraded when running CCS or ZigBee in case of high airtime usage (60%), but unaffected under low duty-cycle traffic (below 10%).

A. Link-Level Performance Improvement

We now quantify the effectiveness of CCS in alleviating collisions caused by spatial and temporal collision hazards, thereby improving ZigBee's link-level performance. In this set of experiments, we use the MICAz mote directly as the signaller. Since the maximum packet size of MICAz is only 127 B, a busy tone sent by MICAz may be insufficient to cover the period for presignaling, data, data/ACK switching, and ACK packet. Thus, we employ two MICAz signalers to send overlapping (to a certain degree) busy tones consecutively, thus extending the busy-tone duration. The signalers are configured to the maximum transmit power level in order to protect the target link. The beacon and superframe orders (Fig. 2) are set to 4 and 6, respectively, resulting in a 986-ms beacon interval.

1) *Reducing Spatial Collision Hazards:* We select link $G \rightarrow H$ and set its transmit power to -10 dBm, such that it experiences packet losses due to asymmetric interference (Table I). To focus on collision rate and isolate the temporal collision hazards (which will be evaluated separately), the transmitter operates in CSMA mode and sends unidirectional data without ACK.

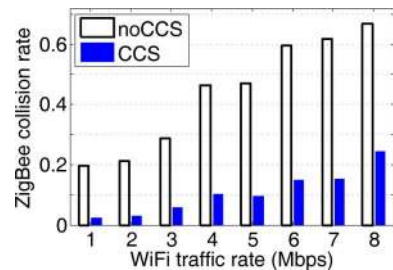


Fig. 12. Packet-loss rate under asymmetric interference.




Fig. 13. Percentage of backoff failures in ZigBee CSMA mode.

The packet size is 64 B, and two packets are sent in each superframe, corresponding to airtime usage 0.48%. WiFi packet size is 1024 B. Its bit rate is fixed at 18 Mb/s, and the traffic rate is varied from 1 to 8 Mb/s, corresponding to airtime usage (Γ_w) from 5.6% to 44.4%, which is consistent with existing measurement studies [1], [9].

Fig. 12 illustrates the ZigBee packet-loss rate due to WiFi-caused collisions. As Γ_w increases, the collision rate increases drastically. CCS significantly reduces the collision rate—the reduction ranges from 64% to more than 90%. Ideally, CCS should be able to fully protect the data packets. However, its packet-loss rate also increases under high WiFi interference due to the increased CTS losses.

In CSMA mode, ZigBee may lose packets due to backoff failures since its MAC protocol aborts the transmission if the sense-and-backoff fails after four retries. Fig. 13 illustrates the fraction of backoff failures among all transmission attempts. Using the signaling mechanism in the CSMA scheduler, CCS can reduce the failure rate by 50% in typical cases. Under high interference, the signaller has less opportunity to find idle slots, so the failure rate increases. In this sense, CCS provides best-effort signaling in CSMA mode in order to avoid unnecessary interference to WiFi.

2) *Reducing Temporal Collision Hazards:* To evaluate how CCS alleviates the temporal collision hazards that corrupt beacon packets and TDMA-scheduled packets, we focus on the nodes $F \rightarrow G$ close to the WiFi transmitter A and set their transmit power to the maximum value to isolate the effects of asymmetric interference.

Fig. 14 shows the loss rate of beacon and TDMA data packets. Without carrier sensing, these packets suffer from a high loss rate, especially as WiFi airtime usage increases. Data packets suffer a higher loss rate due to their larger size (while a beacon is only 11 B long). By notifying WiFi via the signaller's busy tone, CCS reduces the packet-loss rate by 42%–83%, without affecting the normal TDMA schedule.

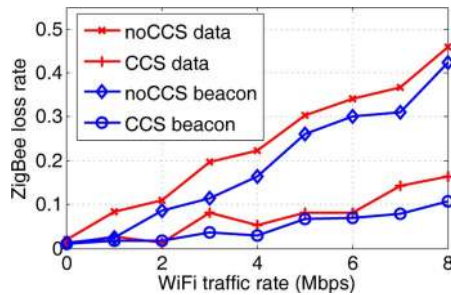


Fig. 14. ZigBee packet-loss rate in TDMA mode.

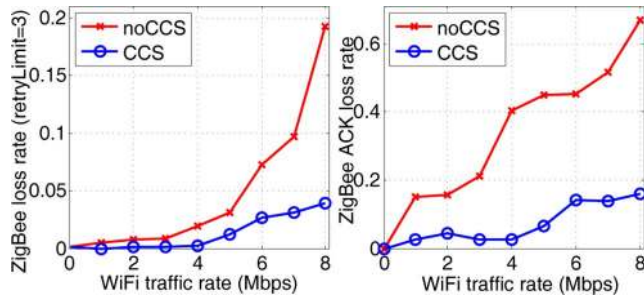


Fig. 15. Loss rate due to WiFi preemption.

Temporal collision hazards are also caused by WiFi preemption, which occurs when WiFi starts transmission within the tx/rx switching time between ZigBee data and ACK. We evaluate such an effect by enabling both data and ACK, log the packet sequences, and calculate the percentage of ACK losses due to WiFi preemption and overall packet-loss rate after retransmission. The results (Fig. 15) indicate that ZigBee ACK loss rate can be up to 63% when $\Gamma_w = 44.4\%$. By filling in the switching time with a busy tone, CCS reduces the ACK loss rate by more than 50% in most cases. Its packet-loss rate is consistently more than 55% lower than the non-CCS (legacy ZigBee). Although retransmission alleviates the non-CCS's packet loss and keeps the loss rate below 19%, it should be emphasized that the experiment is focused on link $F \rightarrow G$, which can be sensed by WiFi (as verified in Section III-A) and suffers from relatively less collision.

3) *Reducing Packet Delay*: When retransmission is enabled, WiFi interference can cause excessive retransmissions, increasing the delivery delay of each packet. In this experiment, we run a full featured ZigBee MAC that retransmits data until an ACK is received or retransmission attempts exceed the retryLimit (default set to 3). Fig. 16 plots the average delay between the generation of a CSMA packet and the corresponding ACK (for the same link $F \rightarrow G$ as in the above experiment). Due to the additional cost of RTS/CTS packets, CCS extends the delay when the WiFi interference level is low, compared to the built-in retransmission scheme in ZigBee. However, this delay cost is offset by the retransmission time under high WiFi interference. The packets that are lost even after retransmissions are not counted in.

In TDMA mode, packets are sent without the RTS/CTS handshake. Without such overhead, CCS reduces packet delay by up to 63% when the WiFi airtime usage $\Gamma_w = 44.4\%$, as shown in Fig. 16. However, it should be noted that the additional delay due to beacon losses is not counted here. Beacon-loss rate for




Fig. 16. Packet delivery delay (retryLimit = 3).

non-CCS can be up to 42% (Fig. 14). This means that with probability 0.42, each TDMA packet sent from the application layer will be delayed by one additional superframe duration (0.986 s). By preventing the TDMA packets from being delayed to the next superframe due to a corrupted beacon, CCS can further reduce the actual packet delay perceived by applications.

B. Operational Validation for Coexisting WPAN and WLAN

The above microbenchmarks showed the effectiveness of CCS in removing the collision hazards for a single link. We proceed to evaluate its performance in a WPAN running multiple ZigBee nodes and with the USRP2 signaller.

The WPAN coordinator is located in E , and clients are F, G, H, N, Q . Each client requests one GTS slot and transmits one 64-B data packet to the coordinator within each superframe, using transmission power 0 dBm. The WiFi link's airtime usage is fixed at $\Gamma_w = 30\%$. Its AP location varies between A, C, K, and P, and client is fixed at B. Under different scenarios, collision can occur due to both spatial and temporal hazards. In all the tests, the USRP2 signaller is located close to the WPAN coordinator, consistent with the signaller configuration (Section IV-C). Its transmit gain is set to half of the maximum gain, corresponding to an approximate output power of 10 dBm.

Fig. 17(a) and (b) plots the collision rate (one-way packet-loss ratio) for the data and beacon packets, respectively. The collision rate varies with the relative location between the WiFi transmitter and the ZigBee transmitter/receiver. The collision rate is also asymmetric. Data packet (uplink) collision is more severe when the WiFi interferer (locations A and C) is closest to the ZigBee coordinator. The beacon (downlink) collision is more severe when the interferer (location P) is closer to the ZigBee clients. When the interferer (location A) is located in between the ZigBee nodes, the collision rate is lower since WiFi senses the ZigBee nodes more effectively. In all of the location settings, CCS reduces the collision rate for both data packets and beacons, with mean reduction 72.6% and 61.1%, demonstrating its effectiveness in removing the collision hazards.

Fig. 17(c) shows the packet-loss rate when retransmission and ACK are enabled. Even with retryLimit 3, the original ZigBee protocol (labeled as *noCCS*) experiences severe losses. By protecting the data and ACK packets, CCS reduces the loss rate by 37.7%–87.5%, depending on the AP locations. By reducing retransmissions, CCS reduces the packet delay by 22.9%–28.6% [Fig. 17(d)]. These experiments justify that even though the signaller's power and location are fixed, it can facilitate the WPAN's coexistence with randomly located WLANs.

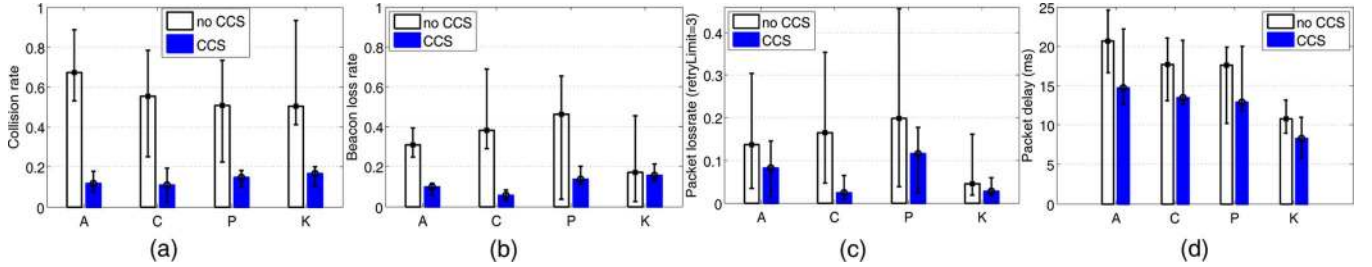


Fig. 17. TDMA performance of a ZigBee WPAN when it coexists with a WiFi WLAN. The x -axis denotes four different WiFi AP locations. Bars denote the performance averaged over all ZigBee links. Error bars denote the maximum and minimum values. Performance metrics include: (a) collision rate, (b) beacon-loss rate, (c) packet-loss rate, and (d) packet delay.




Fig. 18. Performance of CCS versus ideal adaptive channel allocation (*adaptch*). (a) Normalized throughput (ratio of the number of uniquely received packets to those sent). (b) Blackout time (the number of consecutive superframes with beacon losses) when WiFi interrupts 10 times in 10 min. Error bars show the maximum and minimum values.

C. Resilience to Bursty Interference

To evaluate CCS under changing network dynamics, we generate bursty WiFi sessions, each running a UDP file transfer (airtime usage $\Gamma_w = 0.22$) for 20 s. As a benchmark comparison, we have also implemented an adaptive frequency-allocation protocol following [33] (referred to as *adaptch*). A coordinator in *adaptch* detects interference by observing packet losses during one beacon frame. It then scans other channels and notifies all clients to switch to a new channel when available.

In the experiments, the ZigBee WPAN is running in TDMA mode. WiFi session starts periodically and interferes with the current ZigBee channel. CCS maintains a dedicated USRP2 signaler, whereas *adaptch* hops to another channel. We assume an *ideal* scenario for *adaptch*, where it can always find an unoccupied channel. Fig. 18(a) shows the long-term throughput performance. As the number of WiFi interruptions (within 10 min) increases, the legacy ZigBee protocol suffers severe throughput degradation. Both CCS and *adaptch* can reduce packet losses and maintain more than 95% throughput. Their difference lies in the response time when WiFi interference occurs. As shown in Fig. 18(b), *adaptch* has larger average and maximum blackout time due to the need for channel scanning, notification, and re-association. In comparison, CCS reduces the mean and maximum blackout time by 59.4% and 64%, so it is better suited for real-time applications that are sensitive to network outage.

D. Energy Consumption

CCS's performance improvement comes at the cost of additional energy consumption at the signalers. Using the ns-2 simulator, we evaluate CCS's energy consumption averaged over the amount of data that is successfully delivered. We adopt




Fig. 19. Energy consumption in (a) CSMA and (b) TDMA mode.

the device-specific parameters for ZigBee from the TI CC2420 datasheet [30]. Specifically, we set transmit power to 0 dBm, transmit current to 17.4 mA, receive current to 19.7 mA, idle current to 0.426 mA, and supply voltage to 2.4 V. Other parameters, including carrier-sensing threshold, packet size, and WiFi data rate, are configured consistently with the link-level experiments. The ns-2 802.15.4 module simulates packet-level energy consumption by transmission, idle listening, and reception. However, it does not model detailed PHY-layer interference and channel. Therefore, we use the measured packet collision probability over the links in the WPAN as the packet-loss rate model.

Fig. 19 illustrates the energy consumed per byte ($\mu\text{J}/\text{B}$) of the entire WPAN. In CSMA mode, CCS consumes more energy due to the additional overhead in synchronizing the signaling and data transmission. However, in TDMA mode, especially under intensive interference, the additional energy cost is offset by the savings in retransmissions, and the per-packet energy consumption is much lower than simple retransmission. In addition, when using dedicated dc-powered signalers (CCS-DC), CCS reduces the network energy consumption consistently by 73.2%–83.3%, without incurring any overhead to the low-power motes.

E. Impact on WiFi

Our experimental results have shown that CCS improves ZigBee's performance under moderate to high WiFi traffic. We now evaluate its influence on WiFi's normal operation. We focus on the ZigBee link $D \rightarrow E$, which is close to the WiFi link $A \rightarrow C$. The ZigBee runs in TDMA mode, which is more aggressive in channel acquisition. We fix the WiFi airtime usage at $\Gamma_w = 0.22$ and vary the ZigBee airtime usage Γ_z . Fig. 20 shows that WiFi suffers from severe throughput degradation

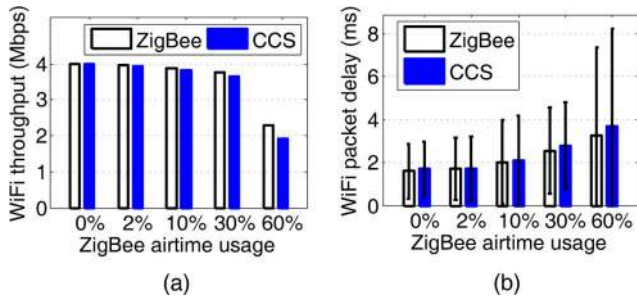


Fig. 20. Impact of ZigBee (TDMA mode) and CCS on WiFi performance in terms of (a) throughput and (b) delay.

and longer latency when Γ_z approaches a high value, say 60%. However, this extreme case rarely occurs since ZigBee targets low-rate applications with typical airtime usage lower than 1% and up to 10% at its maximum [17]. In such common cases, WiFi performance is virtually unaffected by ZigBee, as shown in the figure. More importantly, CCS incurs negligible additional delay or throughput loss compared to the legacy ZigBee, as it may reduce the ZigBee's channel occupancy by reducing retransmissions.

F. Discussion and Future Work

From the above experiments, CCS is found best applicable for ZigBee WPANs under moderate to high WiFi interference, especially when WiFi traffic is bursty. CCS makes greater performance improvements in TDMA mode due to low signaling overhead and energy cost. In addition, a dc-powered ZigBee-compatible transceiver (e.g., the XBee node [26]) is preferred as a dedicated signaler. Such a signaler can have comparable transmit power as WiFi and can persistently protect the entire WPAN under bursty WiFi interference. The dc power for the signaler is a reasonable condition in practice because CCS places the signaler near the WPAN coordinator, which usually serves as a gateway and has access to dc-source.

Throughout our experiments, the WiFi transmitter is running energy-detection-based CCA with the default threshold (around -81 dBm [27]). A lower threshold improves its sensitivity to ZigBee, but the threshold cannot be arbitrarily tuned (minimum -90 dBm for typical devices [34]), and too low a threshold reduces the spatial reuse between WLAN cells. Also, note that the 802.11b allows for either energy detection or preamble detection mode [28, Sec. 18.4.8.4]. Since the CCS signaler cannot emit 802.11b preambles, it can only help coexistence with the former mode. However, for 802.11a/g/n (which uses OFDM modulation), both modes are mandatory [28, Sec. 17.3.10.5], albeit with different default thresholds (-6 and -82 dBm, respectively).

Our current evaluation of CCS has left several issues as future work. For instance, how should the priority of ZigBee for delay-sensitive applications be tuned? CCS can use the presignaling time as a key control knob. If the busy-tone signal is emitted much earlier before the data transmission, the probability of sensing failure due to WiFi preemption would be even lower. However, early presignaling increases CCS's airtime usage and may adversely affect WiFi performance, especially when WiFi is running bandwidth-intensive applications. Our current implementation of CCS does not allow for the evaluation of this delicate tradeoff since the ZigBee timers are not calibrated and not

synchronized at the microseconds level. The USRP2 incurs several milliseconds delay between CCA and data transmission—a much coarser resolution than the ZigBee nodes. We plan to explore the presignaling issue on more capable platforms such as the XBee module [26].

The current deployment of CCS is restricted to a single WPAN with star topology. However, it can be extended to the multihop cluster-tree topology proposed in IEEE 802.15.4 [25]. By augmenting a signaler for each cluster-head, CCS can ensure the entire ZigBee network can coexist with nearby WLANs. Exploring this feasibility is part of our future work.

It should also be noted that we have focused on a single WPAN coexisting with randomly located WiFi WLANs. When multiple colocated WPANs are running CCS, they can simply use the 16 orthogonal ZigBee channels on the 2.4-GHz ISM band. Since each WPAN needs two channels (one for the data transmitter and the other for the signaler), CCS can support a maximum of eight colocated WPANs, which is well above the density of typical WPANs. Moreover, ZigBee WPANs usually target low duty-cycle applications, and hence their throughput is given a much lower priority than their packet loss and delay—our primary evaluation metrics.

VII. CONCLUSION


Prior studies have revealed severe ZigBee performance degradation in the presence of WiFi traffic, even if WiFi leaves a sufficient idle time (e.g., more than 67% airtime unused). In this paper, we presented fine-grained experimental evaluation to identify the causes—the spatial and temporal effects that fail the traditional CSMA mechanisms. We introduced CCS, a network-level framework to enhance CSMA and enable ZigBee-WiFi coexistence. CCS adopts a separate ZigBee signaler to emit carrier signals, which improves WiFi's awareness of, and prevents unnecessary interruptions to, ZigBee. CCS incorporates a scheduler to synchronize the signaling with ZigBee data transmission and a temporary channel-hopping mechanism to avoid interference between the signaler and the transmitter. We have implemented CCS on the TinyOS and GNURadio platforms. Our experimental results have shown a more than 50% reduction of packet collisions when ZigBee operates in the presence of moderate to high WiFi traffic. Both CCS and the original ZigBee have negligible impact on WiFi's throughput and delay when running low duty-cycle applications.

There are two important implications of the CCS framework. First, CCS can be deployed in a real environment that needs both WiFi coverage for Internet access and a ZigBee network for environmental monitoring and control. It is particularly useful for real-time applications, such as medical sensing [35], which adopt ZigBee-based body area networks for patients' health monitoring in hospitals that are usually equipped with WiFi networks for data communications and Internet connections. Second, it represents a simple network-level rationale to assist the coexistence of different protocols that adopt a CSMA-style spectrum etiquette but have heterogeneous PHY characteristics and incompatible MAC layers. We have briefly discussed how CCS can enhance the CSMA mechanisms in IEEE 802.15.2 (for Bluetooth–WiFi coexistence), which cannot prevent the harmful interference caused by unmanaged WiFi devices. The design philosophy behind CCS, i.e., decoupling

carrier signaling from data transmission, can be used to facilitate the coexistence of other heterogeneous wireless networks.

REFERENCES

- [1] C.-J. M. Liang, N. B. Priyantha, J. Liu, and A. Terzis, "Surviving Wi-Fi interference in low power ZigBee networks," in *Proc. ACM SenSys*, 2010, pp. 309–322.
- [2] R. Gummadi, H. Balakrishnan, and S. Seshan, "Metronome: Coordinating spectrum sharing in heterogeneous wireless networks," in *Proc. Ist COMSNETS*, 2009, pp. 157–166.
- [3] 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. CrownCom*, 2008, pp. 1–6.
- [4] J.-H. Hauer, V. Handziski, and A. Wolisz, "Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks," in *Proc. EWSN*, 2009, pp. 17–32.
- [5] M. Petrova, L. Wu, P. Mahonen, and J. Riihijarvi, "Interference measurements on performance degradation between colocated IEEE 802.11g/n and IEEE 802.15.4 networks," in *Proc. ICN*, 2007, p. 93.
- [6] J. Huang, G. Xing, G. Zhou, and R. Zhou, "Beyond co-existence: Exploiting WiFi white space for ZigBee performance assurance," in *Proc. IEEE ICNP*, 2010, pp. 305–314.
- [7] S. Pollin, M. Ergen, M. Timmers, L. Van Der Perre, F. Catthoor, I. Moerman, and A. Bahai, "Distributed cognitive coexistence of 802.15.4 with 802.11," in *Proc. CrownCom*, 2006, pp. 1–5.
- [8] C. Won, J.-H. Youn, H. Ali, H. Sharif, and J. Deogun, "Adaptive radio channel allocation for supporting coexistence of 802.15.4 and 802.11b," in *Proc. IEEE VTC*, 2005, vol. 4, pp. 2522–2526.
- [9] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *Proc. SIGCOMM E-WIND*, 2005, pp. 5–10.
- [10] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of RF interference on 802.11 networks," in *Proc. ACM SIGCOMM*, 2007, pp. 385–396.
- [11] J. Zhu, A. Waltho, X. Yang, and X. Guo, "Multi-Radio coexistence: Challenges and opportunities," in *Proc. IEEE ICCCN*, 2007, pp. 358–364.
- [12] GNU Radio, "The GNU software radio," 2012 [Online]. Available: <http://gnuradio.org/>
- [13] Crossbow Technology, Inc., Milpitas, CA, "MICAz datasheet," 2007.
- [14] Ettus Research, Mountain View, CA, "Universal software radio peripheral (USRP)," 2012 [Online]. Available: <http://www.ettus.com/>
- [15] Crossbow Technology, Inc., Milpitas, CA, "Avoiding RF interference between WiFi and ZigBee," Tech. Rep., 2004.
- [16] Schneider Electric, Grenoble, France, "ZigBee WiFi coexistence," 2008 [Online]. Available: <http://www.zigbee.org/Learn-More/WhitePapers.aspx>
- [17] IEEE 802.15 Working Group, "Coexistence analysis of IEEE Std 802.15.4 with other IEEE standards and proposed standards," 2010.
- [18] F. Zhang, F. Wang, B. Dai, and Y. Li, "Performance evaluation of IEEE 802.15.4 beacon-enabled association process," in *Proc. AINAW*, 2008, pp. 541–546.
- [19] H. Rahul, N. Kushman, D. Katabi, C. Sodin, and F. Edalat, "Learning to share: Narrowband-friendly wideband networks," 2008.
- [20] F. Tobagi and L. Kleinrock, "Packet switching in radio channels: Part II—The hidden terminal problem in carrier sense multiple-access and the busy-tone solution," *IEEE Trans. Commun.*, vol. COM-23, no. 12, pp. 1417–1433, Dec. 1975.
- [21] Z. J. Haas and J. Deng, "Dual busy tone multiple access (DBTMA)—A multiple access control scheme for ad hoc networks," *IEEE Trans. Commun.*, vol. 50, no. 6, pp. 975–985, Jun. 2002.
- [22] I. Demirkol, C. Ersoy, and F. Alagoz, "MAC protocols for wireless sensor networks: A survey," *IEEE Commun. Mag.*, vol. 44, no. 4, pp. 115–121, Apr. 2006.
- [23] X. Zhang and K. G. Shin, "Enabling coexistence of heterogeneous wireless systems: Case for ZigBee and WiFi," in *Proc. ACM MobiHoc*, 2011, Article no. 6.
- [24] A. Koubaa, A. Cunha, and M. Alves, "An IEEE 802.15.4 protocol implementation (in nesC/TinyOS): Reference guide v1.2 Tech. Rep., HURRAY-TR-061106, 2005.
- [25] IEEE-SA Standard Board, "IEEE Standard Part 15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (LR-WPANs)," 2003.
- [26] Digi International, Inc., Minnetonka, MN, "XBee-PRO 802.15.4 OEM RF modules," 2012 [Online]. Available: <http://www.digi.com/>
- [27] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *Proc. ACM SIGCOMM*, 2006, pp. 51–62.
- [28] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std. 802.11, 2007.
- [29] FCC Part 15, Washington, DC, "FCC rules for unlicensed wireless equipment operating in the ISM bands," 2012 [Online]. Available: <http://www.afar.net/tutorials/fcc-rules/>
- [30] Texas Instruments, Dallas, TX, "CC2420 datasheet," 2004.
- [31] *Coexistence of Wireless Personal Area Networks With Other Wireless Devices Operating in Unlicensed Frequency Bands*, IEEE Std 802.15.2, 2003.
- [32] T. Schmid, "GNU radio 802.15.4 en- and decoding UCLA NESL, Los Angeles, CA, Tech. Rep., 2006.
- [33] R. C. Shah and L. Nachman, "Interference detection and mitigation in IEEE 802.15.4 networks," in *Proc. ACM/IEEE IPSN*, 2008, pp. 553–554.
- [34] E. Perahia and R. Stacey, *Next Generation Wireless LANs: Throughput, Robustness, and Reliability in 802.11n*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [35] J. Hou, B. Chang, D.-K. Cho, and M. Gerla, "Minimizing 802.11 interference on ZigBee medical sensors," in *Proc. BodyNets*, 2009, Article no. 5.



Xinyu Zhang received the B.E. degree in electronic and communications engineering from Harbin Institute of Technology, Harbin, China, in 2005, and the M.S. degree in computer engineering from the University of Toronto, Toronto, ON, Canada, in 2007, and is currently pursuing the Ph.D. degree in electrical engineering and computer science at the University of Michigan, Ann Arbor.

His research interests are in the MAC/PHY co-design of wireless networks, with applications in WLANs, WPANs, mesh networks, and white-space networks. He is a recipient of the Best Paper Award at ACM MobiCom 2011.



Kang G. Shin (S'75–M'78–SM'83–F'92) received the B.S. degree in electronics engineering from Seoul National University, Seoul, Korea, in 1970, and the M.S. and Ph.D. degrees in electrical engineering from Cornell University, Ithaca, NY, in 1976 and 1978, respectively.

He is the Kevin & Nancy O'Connor Professor of Computer Science with the Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor. He has supervised the completion of 71 Ph.D.'s and authored/coauthored more than 770 technical articles (more than 270 of these are in archival journals), one textbook, and more than 20 patents or invention disclosures. He has also cofounded a couple of startups. His current research focuses on computing systems and networks as well as on embedded real-time and cyber-physical systems, all with emphasis on timeliness, security, and dependability.

Prof. Shin is a Fellow of the Association for Computing Machinery (ACM). He served on Editorial Boards, including those for the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS and the *ACM Transactions on Embedded Systems*. He has also served or is serving on numerous government committees, such as the US NSF Cyber-Physical Systems Executive Committee and the Korean Government R&D Strategy Advisory Committee. He has chaired several major conferences, including ACM MobiCom 2009, IEEE SECON 2008, ACM/USENIX MobiSys 2005, IEEE RTAS 2000, and IEEE RTSS 1987. He has received numerous best paper awards, including the Best Paper Awards from ACM MobiCom 2011, the 2011 IEEE International Conference on Autonomic Computing, the 2010 and 2000 USENIX Annual Technical Conferences, as well as the 2003 IEEE Communications Society William R. Bennett Prize Paper Award and the 1987 Outstanding IEEE TRANSACTIONS ON AUTOMATIC CONTROL Paper Award. He has also received several institutional awards, including the Research Excellence Award in 1989, Outstanding Achievement Award in 1999, Distinguished Faculty Achievement Award in 2001, and Stephen Attwood Award in 2004 from The University of Michigan (the highest honor bestowed to Michigan Engineering faculty); a Distinguished Alumni Award of the College of Engineering, Seoul National University in 2002; 2003 IEEE RTC Technical Achievement Award; and 2006 Ho-Am Prize in Engineering (the highest honor bestowed to Korean-origin engineers).