Cognitive Frequency Hopping Based on Interference Prediction: Theory and Experimental Results *

Stefan Geirhofer[a] *sg355@cornell.edu*
John Z. Sun[b] *johnsun@mit.edu*
Lang Tong[a] *ltong@ece.cornell.edu*
Brian M. Sadler[c] *bsadler@arl.army.mil*
[a]School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853
[b]Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139
[c]U.S. Army Research Laboratory, Adelphi, MD 20783

Wireless services in the unlicensed bands are proliferating but frequently face high interference from other devices due to a lack of coordination among heterogeneous technologies. In this paper we study how cognitive radio concepts enable systems to sense and predict interference patterns and adapt their spectrum access accordingly. This leads to a new cognitive coexistence paradigm, in which cognitive radio implicitly coordinates the spectrum access of heterogeneous systems. Within this framework, we investigate coexistence with a set of parallel WLAN bands: based on predicting WLAN activity, the cognitive radio dynamically hops between the bands to avoid collisions and reduce interference. The development of a real-time test bed is presented, and used to corroborate theoretical results and model assumptions. Numerical results show a good fit between theory and experiment and demonstrate that sensing and prediction can mitigate interference effectively.

I. Introduction

Wireless devices and services are growing rapidly and are becoming ubiquitous. The unlicensed bands have played a fundamental role in this trend as illustrated by the rapid proliferation of WiFi and Bluetooth devices. As new technologies gain momentum, however, the unlicensed bands are becoming increasingly crowded, and consequently interference management needs to be addressed.

The lack of coordination in the unlicensed bands may lead to high interference among heterogeneous technologies that operate in close proximity. At the same time, however, average spectrum utilization is typically low because of traffic burstiness and the inefficiencies associated with distributed medium access protocols. This contrast motivates a study on how better coordination may improve spectral efficiency and lead to reduced interference.

Throughout this paper we assume that there is no direct information exchange among interfering systems. Instead we study how a cognitive radio may reduce collisions by predicting the WLAN's activity and adapting its medium access accordingly. The cognitive concept of sensing and adaptation consequently serves as an implicit means for coordinating the medium access of the interfering systems. Ultimately, this leads to a new coexistence framework which we refer to as *cognitive coexistence*.

Cognitive coexistence is based on the same principles as dynamic spectrum access (DSA) but faces different challenges and tradeoffs in typical deployments (see [1] for a review of DSA). In many coexistence scenarios radios are collocated and separating transmissions in the temporal domain therefore seems a natural design choice. Compared to DSA, sensing is much easier accomplished in such setups since radios are in close proximity of each other. Furthermore, in unlicensed bands CR concepts can be used in a similar fashion to orthogonalize interfering systems dynamically. While the notion of primary and secondary users does not arise in this scenario because all systems have equal access rights to the medium, a hierarchical setup is typically a natural design approach.

I.A. Main Contribution

This paper presents an experimental study of coexistence between a cognitive radio (CR) and a set of parallel WLAN channels; see Fig. 1. Assuming that the

*Manuscript submitted October 31, 2008; accepted March 17, 2009. This work is supported in part by the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. This work has been presented in part at the Asilomar Conference on Signals, Systems and Computers, Monterey, CA, November 2007 and the IEEE Military Communications Conference, Orlando, FL, October 2007. The present manuscript extends previous work by including new results on the open- and closed-loop behavior of the test bed.

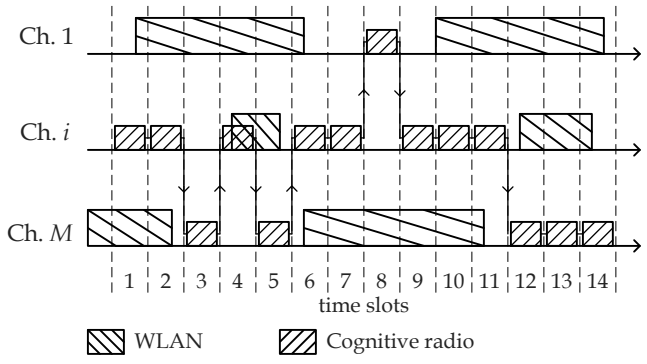

Figure 1: System setup. A frequency hopping cognitive radio coexists with a set of parallel WLAN channels.

CR operates in a frequency range that overlaps with a set of WLAN channels, we study a cognitive frequency hopping (CFH) protocol, which avoids packet collisions by periodic spectrum sensing and predicting temporal WLAN activity. By hopping to temporarily inactive channels, packet collisions are reduced and coexistence improved. Theoretically, CFH is based on an ON/OFF continuous-time Markov chain (CTMC) model, which approximates WLAN's decentralized medium access while being tractable enough to derive the optimal hopping pattern.

This paper focuses on comparing theoretical and experimental performance results. The previous work [2] is taken as a basis for CFH but the protocol is closely related to other contributions on time domain OSA [3, 4]. Theoretical assumptions are corroborated by a real-time test bed, which is used for validating model assumptions and evaluating the dynamic interaction between the CR and WLAN. For example, our theoretical analysis is based on assuming that collisions do not impact the validity of the temporal prediction model as long as the collision rate is kept below some interference constraint. In an actual system, however, does the CTMC prediction model remain approximately valid or are there significant changes in the WLAN's temporal activity? Similarly, is the WLAN's sensing-based medium access affected? The experimental approach helps quantifying the impact of retransmissions, a result that would be hard to obtain analytically. Our findings show a good match between theory and practice and demonstrate that the CR can reduce interference effectively. To the best of our knowledge this paper presents the first test bed to validate the above model assumptions experimentally.

Limitations. The experimental test bed has been tailored to the above objectives. As such, it is necessary to clarify some limitations due to hardware and complexity constraints.

First, the test bed only supports a single WLAN channel due to bandwidth limitations of the down-conversion module. Consequently, the test bed's medium access reduces to transmitting or not transmitting at the beginning of each slot. Nevertheless, even though no frequency hopping takes place, the experimental results for the single-channel case can be generalized to multi-channel setups. Specifically, we will demonstrate that by characterizing the behavior of a single WLAN channel for arbitrary traffic intensities, it is possible to extrapolate performance metrics to multi-channel scenarios as long as the behavior of parallel channels is statistically independent.

The test bed is further limited to a single cognitive transmitter; no receiver has been implemented. This removes the need to maintain slot synchronization but does not limit our ability to measure CFH's throughput and interference through monitoring channel activity. While methods such as collaborative sensing [5] or acknowledgement feedback [6] may be used for synchronization, their implementation in real-time is difficult and goes beyond the scope of this work. Similarly, multi-user aspects of the cognitive system are not addressed. However, well-studied concepts could be applied, such as Bluetooth's piconet structure [7], where a master node acts as a central controller.

Finally, mutual interference among collocated radio systems crucially depends on the underlying propagation conditions. In this paper, we focus on the worst-case collision model in which any overlap in frequency and time results in a packet drop. This approximates a setup in which WLAN devices and CR nodes operate in close proximity.

Organization and notation. The paper is structured as follows. Sec. II introduces CFH analytically, and Sec. III describes the test bed implementation. The measurement methodology and performance results are presented in Sec. IV and Sec. V, respectively. Performance trends and tradeoffs are discussed in Sec. V.C.

I.B. Related Work

While there is a growing body of literature on cognitive radio and dynamic spectrum access, see, *e.g.*, [1, 8], reports on experimental and implementation aspects of CR and DSA have been limited. This work is most closely related to contributions on temporal DSA. Among the first to address this problem, Zhao *et al.* investigate the problem of accessing slots that are left idle by primary users [9]. The optimal medium access is derived within a Markovian decision-theoretic

framework and a separation principle between sensing and medium access is shown [10]. Similar setups have recently received increasing interest [4, 11, 12]. An experimental test bed which heuristically accesses white space in WLAN is presented in [13]. Recent experimental studies and related work include [14–17].

Coexistence between WLAN and Bluetooth devices has conceptual similarities with our contribution because of a similar physical-layer setup. Different coexistence methods have been proposed for this scenario and range from interference cancelation at the physical layer [18] to changes in MAC layer scheduling at the WLAN stations [19] and adaptive frequency hopping (AFH) [20].

AFH techniques are most closely related to CFH: both schemes adapt their hopping sequence to reduce interference. The major difference lies in how interference is detected and modeled. AFH typically classifies channels as being either "good" or "bad" according to the empirical error rates of its own transmission attempts. However, the WLAN's temporal activity is not modeled explicitly and no spectrum sensing is performed. Bad channels are simply avoided by reducing the hopping set to good channels, if possible. Naturally, this approach is well suited to suppress interference that is static or slowly time varying with respect to the Bluetooth slot length. In many cases, however, WLAN packets are only slightly longer than the slot length, reducing the benefit of this modeling technique. In contrast, CFH's sensing and prediction framework is well suited to account for the time variant behavior of WLAN traffic.

Some aspects of CFH have been investigated in prior work by the authors of this paper. The modeling of WLAN traffic has been proposed in [21], and medium access schemes based on these models were introduced in [2]. This paper goes beyond these contributions in presenting an experimental test bed and comparing theoretical with experimental results. Our findings give rise to system-level insights on the performance of this method.

II. CFH Protocol Design

To improve coexistence between the WLAN channels and the frequency-hopping CR, CFH adapts the CR's hopping sequence such that it preferably transmits in bands that will likely remain idle for the duration of the transmission.

As depicted in Fig. 1, the CR transmission is slotted in time with duration T. At the beginning of each slot the activity (idle or busy) of the WLAN channels is detected and prediction is based on a two-state continuous-time Markov chain model. We assume perfect sensing in our theoretical work.

II.A. Modeling Spectrum Opportunities

Since the WLAN channels do not overlap in frequency it is reasonable to assume that they evolve independently in time. Hence, we can focus on a single channel, and deal with N independent models (possibly with different parameters) rather than a single model encompassing all channels.

The activity of each WLAN band is modeled by a two-state CTMC $\{X(t), t \geq 0\}$ with idle state $X(t) = 0$ and busy state $X(t) = 1$. The sojourn times are exponentially distributed with parameter λ in the idle state and μ in the busy state.

This simple model approximates the statistics of the idle periods of real WLAN systems. As shown in Fig. 2, by comparing with empirical data gathered by measurement, we see that the overall behavior of the empirical cumulative distribution function is approximated by this exponential fit. The measurement setup and assumptions will be described in more detail in Sec. III.

We have shown in previous work [2] that a semi-Markov model (SMM) is able to achieve a better fit with the empirical data by approximating idle periods based on a mixture of two generalized Pareto distributions (one corresponding to the WLAN contention behavior, the other modeling random packet arrivals). This leads to the cumulative distribution function (CDF)

$$F_m(t) = p_1 F_1(t; k_1, \sigma_1) + (1 - p_1) F_2(t; k_2, \sigma_2), \quad (1)$$

where p_1 denotes the mixture coefficient and

$$F_i(t; k_i, \sigma_i) = 1 - \left(1 + k_i \frac{t}{\sigma_i}\right)^{-1/k} \quad (2)$$

represents the CDF of the generalized Pareto distribution with shape parameter k_i (measuring the deviation from an exponential distribution) and scale parameter σ_i (quantifying the decay rate). This model shows an excellent fit with the empirical data as can be seen in Fig. 2 and has also been validated statistically through the Kolmogorov-Smirnov test [21].

While a mixture distribution leads to an accurate fit, it is less tractable analytically. Our theoretical analysis is therefore based on the more tractable CTMC model. While its deviation from the empirical data is significant in Fig. 2, our results will show that the exponential model predicts overall system performance quite accurately and is a useful theoretical tool.

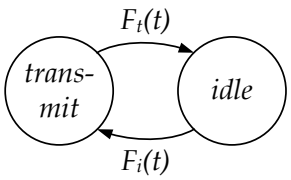

(a) State transition model

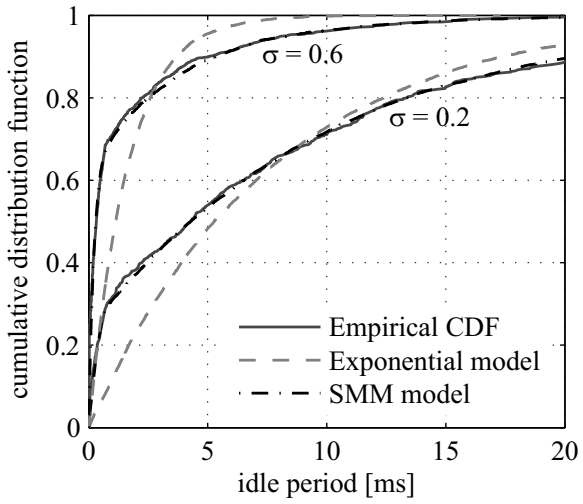

(b) Sojourn time distribution in the idle state

Figure 2: Temporary prediction model of WLAN activity. The semi-Markov model approximates the empirical sojourn time in the idle state with good accuracy (shown for two WLAN traffic intensities σ).

The busy periods are directly related to the packet length and transmission rate chosen by the scheduler of each WLAN station. In our work, we focus on traffic scenarios with constant payload packets and transmission rate, and can therefore treat the sojourn time as deterministic. For the CTMC approximation, which requires exponential sojourn times, we choose the distribution's mean to equal the deterministic sojourn time.

II.B. Designing the Optimal Control

Based on the CTMC model we discuss the optimal CFH behavior for a single WLAN channel (due to single channel operation of the test bed). The general case of M channels has been addressed in [2].

The objective of CFH is to maximize CR throughput while constraining the rate of packet collisions. The CR throughput is measured in terms of successful packet transmissions per unit time and the interference constraint is characterized by the WLAN packet error rate (that is, the number of dropped divided by the number of transmitted WLAN packets). This normalizes the interference constraint by the WLAN's traffic intensity and ensures that no more than a certain fraction of packets gets dropped.

In the single channel case, the optimal medium access reduces to finding transmission probabilities for idle and busy sensing outcomes, respectively. Whenever the channel is observed idle, we transmit with a probability p, which is designed such that the interference constraint is met with equality [2]. Whenever a busy sensing outcome is observed, it is optimal not to transmit because this would only cause higher interference while not increasing throughput.

In the multi-channel case we need to select one out of M bands for transmission. This case can be addressed through decision-theoretic analysis by formulating the problem as a constrained Markov decision process. Due to space limitations and because our experimental work focuses on the single channel case, we refer to [2] for details.

III. Test Bed and Experiment Design

Having introduced CFH, we describe the test bed implementation, compare the experimental design with our analytical setup, and discuss fundamental design objectives.

The test bed has been developed with the objective of validating some of the implicit assumptions made in our analytical work. This specifically includes dynamic effects between both systems that are difficult to characterize analytically. Our analytical work focused on designing transmission probabilities, given a stochastic model for the WLAN, such that interference constraints remained met. The idea behind this modeling approach is that as long as the interference constraints are sufficiently tight, the residual interference caused to the WLAN will not impact its temporal behavior. While analytical approaches for justifying this assumption are difficult, our experimental results enable us to confirm its validity.

A block diagram of the test bed is shown in Fig. 3. The implementation is based on a Sundance software defined radio (SDR) development kit, consisting of processing and data acquisition modules. Radio-frequency (RF) signals are down-converted using a commercial WLAN transceiver and up-converted using an Agilent vector signal generator. The baseband processing is done entirely on the SDR board. The CR's slot structure is implemented using an accurate timer, which triggers periodic interrupts with a period of $T = 625\mu$s. The analog-digital converter (ADC) is triggered at the beginning of each slot. A $1\,\mu$s block at a rate of 72 MHz is captured and passed on

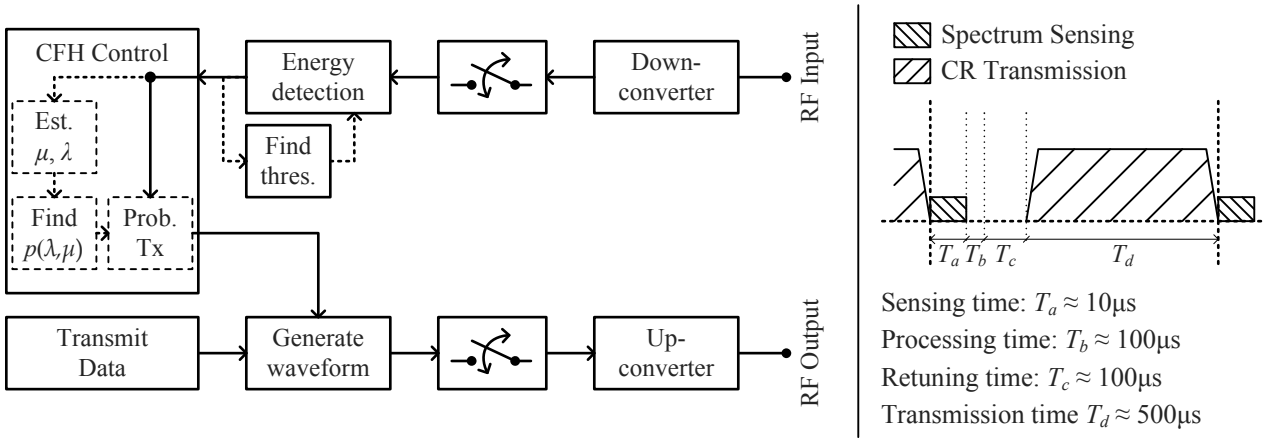

Figure 3: Block diagram of the cognitive radio (left). The slot operation and timing is shown on the right. Pictures of the cognitive radio test bed are available online at http://acsp.ece.cornell.edu/mc2r.pdf.

to the energy detector, which computes the signal's energy and compares it with a threshold. This results in a binary sensing result (idle/busy), which is used by the CFH controller to determine the CR's medium access. Depending on the outcome of the stochastic control a transmission may be initiated by using a programmable signal generator. Upon digital-to-analog (DAC) conversion, the signal is up-converted by the RF front-end.

The above operations introduce processing delays that reduce the actual transmission time per slot. Typical values for these delays are shown in Fig. 3. Spectrum sensing relies on blocks of less than $10\,\mu s$, and is almost negligible compared to T. The processing time for the sensing result and the CFH controller together amount to roughly $100\,\mu s$ in our implementation. The time it takes to re-tune the transmitter to a different channel amounts to approximately $100\,\mu s$ (this delay does not occur in our setup, however, as we only deal with a single WLAN channel). The remainder of the slot can be used for the CR's transmission.

The baseband processing can be categorized into three parts, namely (i) the spectrum sensor, (ii) the CFH controller, and (iii) the CFH transmitter. In the following each component is discussed in more detail.

III.A. Spectrum Sensor

Spectrum sensing plays a key role in CR systems and the challenges associated with reliably detecting signals at very low SNR have been the subject of much investigation. Compared to some DSA setups, however, the burden of reliable spectrum sensing is reduced in this cognitive coexistence setup. In contrast to DSA schemes that orthogonalize systems by sufficient spatial separation (and therefore require the ability to detect weak primary signals), typical SNR values faced in this cognitive coexistence setup will be substantially larger as both systems operate in close proximity of each other.

Thanks to the fairly high SNR conditions, energy detection can be used efficiently with very little complexity. Energy detection is mathematically formulated as a binary hypothesis testing problem on a set of N samples that either represent just noise, or a signal in noise, respectively. This leads to

$$\begin{aligned} \mathcal{H}_0 &: \quad Y_i = V_i,\ i = 1, \ldots, N \\ \mathcal{H}_1 &: \quad Y_i = S_i + V_i,\ i = 1, \ldots, N, \end{aligned} \tag{3}$$

where Y_i denotes the complex baseband samples, V_i are noise samples, $V_i \sim \mathcal{CN}(0, \sigma_0^2)$, and S_i denotes the signal samples drawn from a complex Gaussian, $S_i \sim \mathcal{CN}(0, \sigma_1^2)$. This hypothesis testing problem is standard [22]. The optimal Neyman-Pearson detector is given by

$$T(\mathbf{y}) = \sum_{i=1}^{N} |y_i|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma, \tag{4}$$

where the threshold γ needs to be chosen such that the probability of false alarm (*i.e.*, erroneously declaring a busy channel) is no larger than a specific value. The Neyman-Pearson detector then yields the optimal probability of detection.

In theory, when (3) holds exactly, the threshold γ can be determined analytically by finding closed-form expressions for the probability of false-alarm and detection. In the experimental domain, however, numerous other factors need to be taken into account. A fundamental problem in this work is the fact that hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$ cannot be observed isolated from

each other because the test bed and the WLAN packet transmission are not synchronized. Therefore it is not possible to sample exclusively during idle periods or exclusively during busy periods. This results in a mixture of the distributions under either hypothesis where the mixture parameter depends on the long term probability of observing idle and busy slots. The hypotheses can however be separated by statistical analysis.

Empirical observations of the sufficient statistic (4) are plotted in Fig. 4. We expect to observe a mixture of chi-square distributions because both $T(\mathbf{y}|\mathcal{H}_0)$ and $T(\mathbf{y}|\mathcal{H}_1)$ are chi-square distributed. More than two mixture components may be necessary, however, due to slightly different power levels of the WLAN terminals. Indeed, the empirical CDF can be well approximated by a mixture of three chi-square distributions,

$$f(x) = \sum_{i=1}^{3} p_i f_i(x; \alpha_i, \beta_i), \tag{5}$$

where $p_i \geq 0$ for all i, $\sum_{i=1}^{3} p_i = 1$, and

$$f_i(x; \alpha_i, \beta_i) = x^{\alpha_i - 1} \frac{\beta_i^{\alpha_i} e^{-\beta_i x}}{\Gamma(\alpha_i)} \tag{6}$$

represents the PDF of a Gamma distribution with shape parameter α_i and rate parameter β_i. An Expectation-Maximization algorithm [23] is used to find the model parameters of (5) and the fitting result is shown in Fig. 4. A good match with the empirical data is observed. We also found that busy and idle mixture components have very different rate parameters. This illustrates the very different energy levels present in idle and busy sensing slots.

Numerical performance analysis demonstrates that the test bed's spectrum sensor works reliably. By choosing the decision threshold appropriately a detection probability of 98.5% can be realized with less than 1% false alarms. The good performance of energy detection is, of course, due to the moderate to high SNR conditions, which make spectrum sensing similar to the carrier-sensing employed in systems such as IEEE 802.11.

III.B. Cognitive Controller

The role of the cognitive controller is to initiate a transmission probabilistically, whenever an idle sensing result is observed. Transmissions are never initiated after a busy sensing result because this would lead to a collision with high probability. A full implementation of CFH encompasses the tracking of traffic

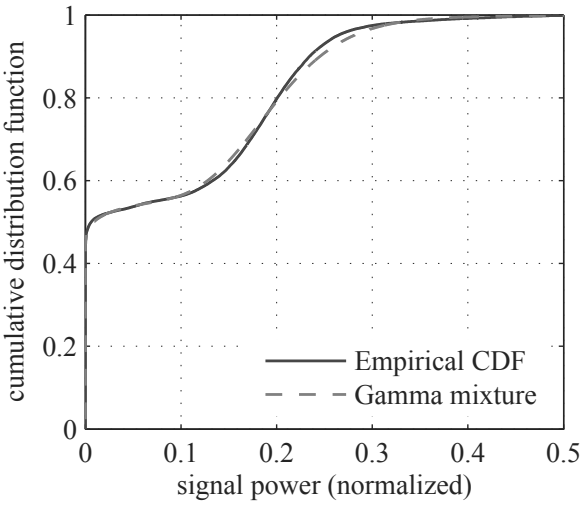

Figure 4: Decision-statistic for the energy detector. A mixture of three gamma distributed components is observed.

variations and estimation of the CTMC model parameters λ and μ, the computation of the optimal transmission probabilities $p(\lambda, \mu)$, and the randomized selection of the control action based on these probabilities. Due to the complexity associated with estimating λ and μ, the experimental test bed uses a fixed transmission probability p. Despite this limitation, it is possible to infer the throughput and collision rate for the case in which the transmission probability is adjusted with the traffic intensity. Since transmissions are never initiated in busy slots, any choice of p corresponds to a certain throughput and collision rate. Furthermore, increasing p leads to an increased throughput but also to a larger collision rate. For the single-channel case in which the medium access is described by a single parameter it is therefore possible to measure throughput and collision rate experimentally and then normalize accordingly such that the interference constraint is met. Thus, the performance of CFH can be evaluated without much loss of generality.

III.C. Cognitive Transmitter

If a CR transmission is initiated, it lasts for the remainder of the slot duration. For the CFH operation it is not relevant what specific transmission scheme is used. For optimal usage of the white space between consecutive packets the CR could use the same frequency bands as the WLAN. For simplicity and motivated by the conceptual similarity with Bluetooth/WLAN coexistence, however, we designed the transmitter to resemble that of Bluetooth. It therefore transmits in narrowband channels of 1 MHz and similar modula-

tion parameters [7]. Gaussian Frequency Shift Keying (GFSK) with a time-bandwidth product of 0.3 was used at a symbol rate of 1 MSps. We should note that the bandwidth of the sensing frontend, however, amounted to 36 MHz as discussed before.

The test bed's transmitter is implemented based on a programmable signal generator which is integrated into the acquisition module, and can be triggered in software. Data contained in an internal buffer is then transferred to the DAC and played back in an infinite loop (for further implementation details we refer to [24] due to space limitations).

IV. Measurement Methodology

The previous section described the test bed's implementation. As CFH relies on sensing and predicting packet collisions, its performance naturally depends on the specifics of the coexistence setup, such as propagation conditions, traffic intensity, system parameters, etc. This section is devoted to describing the measurement methodology that underlies the performance assessment.

IV.A. Hardware Setup

The experimental coexistence setup is depicted in Fig. 5 and consists of the WLAN system (composed of an access point and three workstations), the CR, as well as a vector signal analyzer which was used to monitor the operation of the system.

Two fundamentally different configurations were considered. In the *open-loop* setup the CR's output is not fed back to the WLAN system but only used to determine the packet error rate. While this does not reflect what would occur in practice, this setup enables us to draw a direct comparison with our analytical results. In the *closed-loop* setup, on the other hand, interference impacts the WLAN and leads to frequent retransmissions. We analyze the CR's impact and relate the results to the open-loop setup.

WLAN configuration. The WLAN consists of commercial off-the-shelf IEEE 802.11 devices and includes an access point and three workstations with adapter cards. The access point is connected to a fourth workstation using a wired ethernet connection. All wireless devices are configured identically to operate in channel 6 (corresponding to a center frequency of 2.437 GHz) and use a transmission rate of 5.5 Mbps.

The wireless devices' RF outputs are all connected to a resistive power divider using coaxial cables. This isolates the transmissions from the environment and reduces interference that could otherwise result from unrelated transmissions in the unlicensed bands (measurements were taken in an office building with a number of unrelated WLAN access points). In addition, this configuration removed any propagation effects and allowed for repeatable results. While the propagation conditions encountered in practice will deviate from this idealized setup, our results correspond to a worst-case scenario in which packet overlaps inevitably result in collisions.

Open-loop setup. The open-loop configuration is shown in Fig. 5(a). While the combined WLAN signals serve as the input to the CR, its output is not fed back to the WLAN system. Instead it is combined with the WLAN signal and detected by a separate workstation computer with a WLAN adapter card. The two circulators isolate the output of the test bed and prevent it from impacting the WLAN.

The workstation capturing the combined signal of the WLAN and the test bed uses an adapter card in promiscuous mode together with commercial WLAN analysis software to detect WLAN packets. The output power level of the CR is set large enough such that a collision between both systems will prevent the adapter card from successfully receiving the packet (either a packet error will be displayed or the packet will be missed altogether, depending on whether the synchronization preamble or only the payload is affected). The rate of successful packet reception is thus measured and, by comparison with settings of the traffic generators, the rate of packet losses can be inferred.

Closed-loop setup. The closed-loop setup is depicted in Fig. 5(b). The test bed again receives the combined WLAN signal, but its output is now connected to the same power-divider as the WLAN devices. Therefore, packet collisions lead to packet drops at the WLAN devices themselves (and not at a reference station as in the open-loop setup). Packet loss will therefore initiate WLAN retransmission, which may in turn affect the test bed's performance. The closed-loop setup therefore allows for dynamic interaction between both devices.

In the closed-loop setup, PC_1, which is connected to the WLAN access point by ethernet, runs traffic analysis software. The traffic generators are configured such that PC_1 is the intended receiver, and hence the successful portion of the WLAN traffic (including any retransmissions that may occur) is measured. By comparing with the case of no interference, the packet loss rate is inferred.

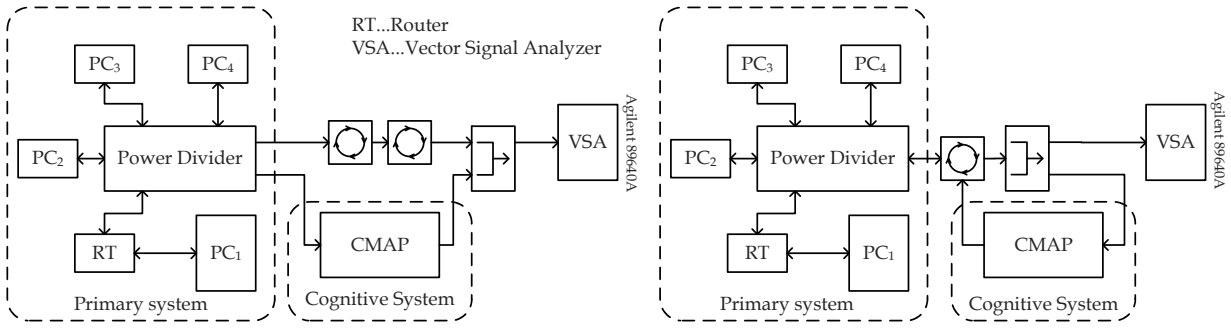

(a) Open-loop setup (b) Closed-loop setup

Figure 5: Measurement methodology. In the open-loop setup (left) the CR's transmission are not fed back to the WLAN. The closed-loop setup (right) is used to quantify the impact of retransmissions.

IV.B. Traffic Characteristics

WLAN transmissions are generated by using a traffic generator on each workstation. The Distributed Internet Traffic Generator (D-ITG) [25] is used and allows to specify packet lengths, the distribution of inter-arrival times, and transmission rates. The traffic properties form an important component of the measurement methodology.

This work focuses on UDP traffic with constant payload of 1024 bytes and plots performance with respect to traffic intensity. Specifically, we define the WLAN traffic load σ, which is normalized such that $\sigma = 0$ corresponds to an inactive WLAN and $\sigma = 1$ represents a WLAN operating at full capacity.

Experimentally, the WLAN traffic load σ is measured as follows. First, the rates of all traffic generators are set to a value that is too large to be supported by the WLAN (even in the case of no interference) and by measuring the actual rate, the WLAN capacity is found. Then, the settings of the traffic generator are normalized by this value, leading to values $0 \leq \sigma \leq 1$. For example, given the traffic and propagation settings in our setup, a maximum of approximately 450 packets per second could be supported by the WLAN. Configuring each of the traffic generators to transmit at a rate of 50 packets per second therefore corresponds to $\sigma = 1/3$.

The measurements focus on the average throughput and interference for stationary traffic scenarios with different intensities σ. Measurement results are compared with simulations using the model parameters in Tab. 1. These parameters were obtained based on statistical analysis, as discussed in Sec. II. The results can be extended to non-stationary traffic setups, provided that parameters of the traffic model are tracked over time [21].

IV.C. Measurement Process

The measurements are performed in the following manner. First, with the CR portion of the test bed turned off, the WLAN traffic generators are adjusted such that a specific traffic load σ is realized. The CR is then enabled and the successfully received WLAN packets are counted. By comparing with the nominal packet rate in the interference-free case, the packet error rate can be obtained. At the same time, the CR keeps statistics of the number of initiated and successful slot transmissions. Lacking a CFH receiver, a successful CFH slot transmission is defined by initiating slot transmission *and* observing an idle channel at the beginning of the next slot. Due to the reliable spectrum sensing performance and the fact that, in our setup, WLAN packets are always longer than the slot duration T, this is a valid metric.

V. Performance Results

This section presents experimental performance results and compares them to simulation results obtained using the SMM and CTMC models discussed in Sec. II. For the open-loop setup, where theoretical and experimental results should coincide, we observe an excellent match. In the closed-loop setup, where mandatory re-transmissions affect the experimental performance results, we introduce a simple heuristic formulation that allows us to approximate the WLAN's behavior based on open-loop results.

The performance assessment in this section focuses on the CR's throughput and the interference it causes (in terms of WLAN packet errors induced by the interference). Clearly, both performance metrics are interrelated: by increasing the CR's transmission probability p we can increase throughput at the expense of a larger number of collisions and vice versa. By measuring both metrics, we can get a better insight on how

aggressive the CR transmission policy can be without significantly affecting WLAN performance.

We compare the performance of CFH with a blind reference scheme that neither detects nor predicts WLAN activity but obliviously initiates transmissions with probability p_r regardless of the state of the medium. Our results demonstrate that CFH introduces a significant performance gain and increases CR throughput while reducing WLAN interference.

V.A. Open-Loop Measurement Result

The open-loop measurement result is shown in Fig. 6. The left panel shows the CR throughput in terms of the expected number of successful slot transmissions per unit time and the right panel shows the packet error rate of the WLAN. We stress that the experimental curves have been obtained by counting the rate of successfully received WLAN packets and the packet error rate is computed by comparing with the average number of packets that should have been received during that time period. The results are compared with simulations based on the SMM and CTMC model. The SMM-based simulation also incorporates the fact that the CR does not use the entire slot for transmission and should therefore approximate the measurement results with good accuracy.

The SMM-based simulation results indeed show an excellent match with the experimental results. The throughput curves shown in the left panel of Fig. 6 almost coincide (the largest aberration amounts to less than two percentage points) and for the packet error rate (shown on the right) we observe a maximum abberation of two percentage points. By comparison, the CTMC model is less accurate. While the throughput results match fairly well, the predicted packet error rate deviates significantly. From a modeling standpoint, this is not surprising because the exponential approximation does not capture the WLAN's contention behavior.

Similar performance trends are observed for the reference scheme without spectrum sensing. As shown in Fig. 6, measurement and SMM-based simulation again match very well while the CTMC approximation shows noticeable deviation in terms of predicted packet error rate.

V.B. Closed-Loop Measurement Result

The results for the closed-loop setup are shown in Fig. 7. Here, the WLAN terminals are strongly interfered with by the CR and therefore initiate retransmissions whenever a packet is dropped. If a retransmission attempt is ultimately successful, the packet is counted as successfully received (the retransmission packets themselves are not counted toward the WLAN traffic load).

Compared to the open-loop setup, the throughput of the cognitive system stays approximately the same, but the interference to the WLAN changes drastically due to the retransmission behavior. Up to approximately $\sigma = 0.8$, no packet loss is observed. At $\sigma = 1.0$, it increases to roughly 5%.

The non-cognitive reference scheme exhibits a similar behavior. No packet errors occur for $\sigma = 0.5$ or below, and an approximately linear increase is observed for higher values of σ. The packet error rate reaches a maximum of roughly 35% at $\sigma = 1.0$.

The reason for the reduced packet errors lies in the WLAN's retransmission behavior. If a packet transmission fails, the standard mandates that a retransmission be initiated. At low WLAN rates, it is likely that these retransmissions will be successful because the medium is predominantly idle. At high rates, however, it may no longer be possible to accommodate such retransmissions, leading to an increasing packet error rate.

In order to compare experimental and theoretical results, we consider a simple approximation that incorporates the retransmission behavior. For simplicity, we assume that packets are retransmitted indefinitely until they are received successfully. A traffic load σ encountering a collision probability q leads to the cumulative traffic load

$$\sigma' = \sigma(1 + q + q^2 + \cdots) = \frac{\sigma}{1-q}. \qquad (7)$$

We assume that as long as the traffic load σ' is below the capacity of the WLAN channel $\bar{\sigma}$, no packet loss occurs because retransmissions can be accommodated. However, for $\sigma' > \bar{\sigma}$ this is no longer possible and there will be a packet error rate of approximately

$$\frac{\sigma' - \bar{\sigma}}{\sigma'}. \qquad (8)$$

This simple heuristic approximation is used to compare closed and open-loop results. Clearly, in order to compute (8) only open-loop performance result (obtained by measurement or simulation) are needed. Based on these it is possible to approximate the packet error rate in the closed-loop setup.

The performance curves gathered in this way are shown in Fig. 7 and again show CR throughput (left) and packet error rate (right). The measurement curves are obtained directly by measurement from the closed-loop setup. The simulation curves are obtained based

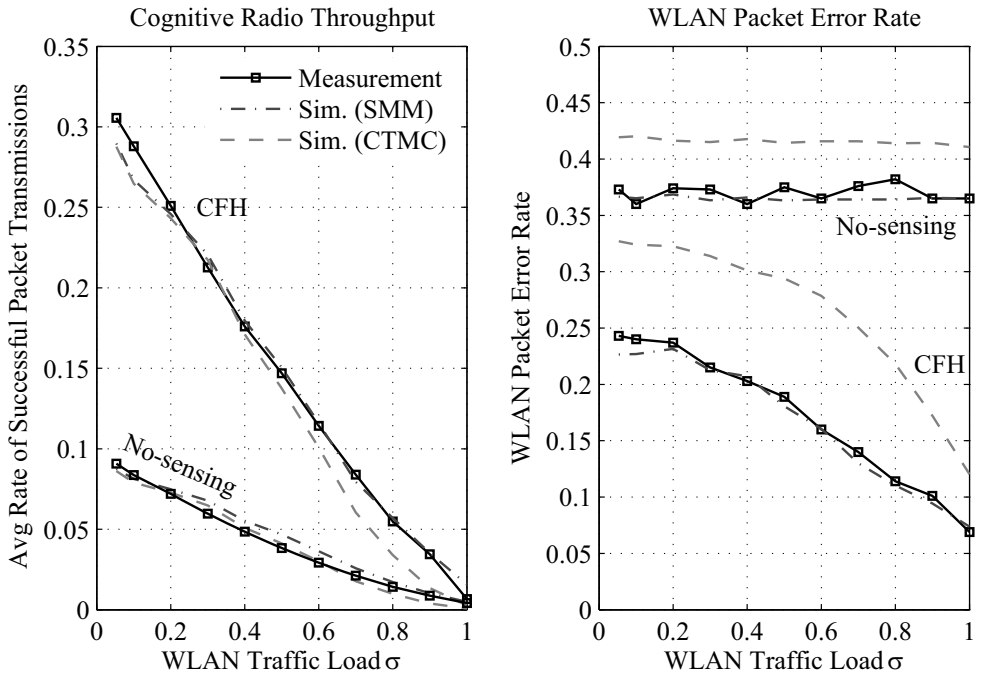

Figure 6: Open-loop performance result. The CR throughput (left) and WLAN packet error rate (right) are shown. The measurement results match closely with SMM-based simulations. The results obtained under the CTMC model show a larger deviation.

on open-loop simulations with the SMM and CTMC models and applying (8). We observe a good match with the measurement results for the closed-loop case suggesting the simple heuristic predicts the closed-loop behavior quite accurately.

Lastly, the results show that even at high traffic load, there exists a residual throughput of the CR system. This appears to be the result of WLAN's retransmission behavior, which due to frequent collisions enlarges some contention windows to accommodate other stations. As can be seen in the left panel of Fig. 7, this results in a residual throughput of approximately 0.03 even when the WLAN is fully loaded.

V.C. Performance Trends

Based on the performance results presented in the previous section, we discuss some tradeoffs and challenges which may arise in a complete implementation of such a system.

The choice of system parameters in this paper is motivated by facilitating a comparison between theory and practice. Consequently, we focus on a worst-case propagation scenario in which any time overlap between transmissions results in packet errors. This approximates the case in which the devices are located in close proximity.

An important design parameter in trading off throughput versus interference is the transmission probability p. Clearly, CR throughput increases with p, but so does the interference that is inflicted upon the WLAN. A natural design approach is to choose p based on the traffic parameters such that a specific constraint on the WLAN packet error rate is met with equality. In the multi-channel case the optimal vector of transmission probabilities can be found through decision-theoretic analysis, which leads to a linear programming solution with fairly low implementation complexity. Further, transmission probabilities only need to be recomputed whenever the prediction parameters change significantly; on a slot level, the random access can be implemented by storing the vector of transmission probability in a lookup table.

A tradeoff arises in selecting the slot length. If there was no overhead associated with sensing and re-tuning the CR, reducing the slot duration would enable us to more efficiently "fill up" the idle periods between packets. Due to this overhead, however, choosing a small slot duration leads to small CR payload and consequently reduced performance. Because of conceptual similarity we have chosen a slot length of 625 μs, which equals the value used in Bluetooth. Further, we have found that around this value, the throughput of the CR changes very little.

Throughout this paper we have focused on stationary traffic scenarios with varying WLAN traffic inten-

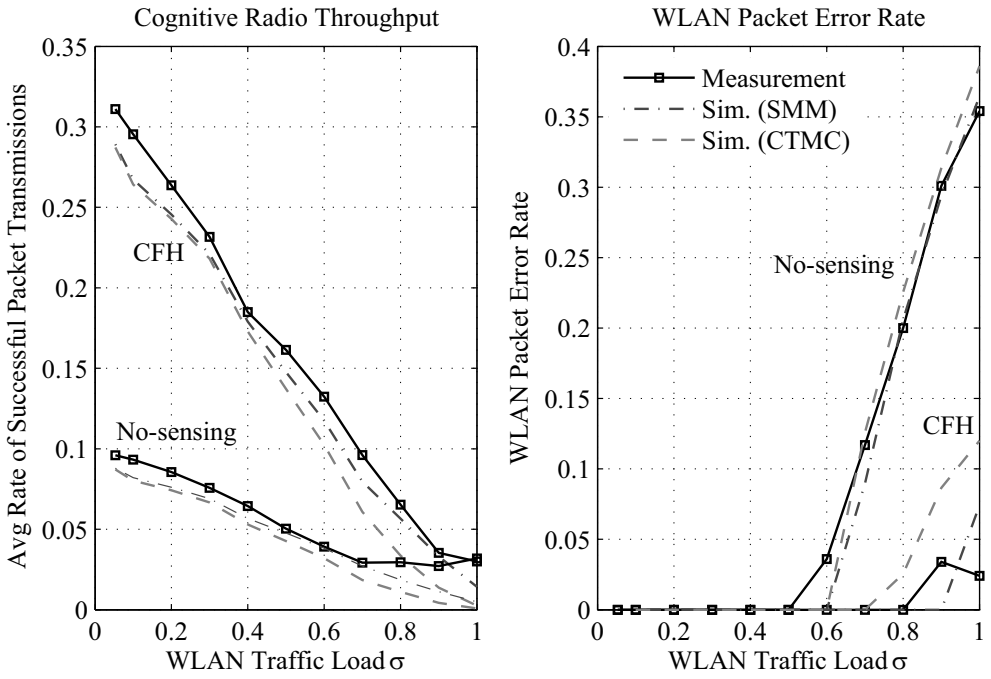

Figure 7: Closed-loop performance result. CR throughput (left) and WLAN packet error rate (right) are shown. The packet error rate includes retransmissions and is therefore smaller than in the open-loop case.

sity. This enabled us to measure throughput and interference by recording long packet traces and computing time averages. Nevertheless, we believe that our results will extend to the non-stationary traffic scenarios observed in practice. In fact, the time scale of our prediction model is in the order of tens of milliseconds, which is much smaller than the typical time scale associated with changes in usage patterns. Tracking non-stationarities by adapting model parameters is therefore a viable approach. We have verified that this leads to an accurate fit in previous work [21].

While our comparison between theory and experiment shed some light on implementation aspects, our work represents the first steps toward developing a fully functional prototype. Synchronization is one aspect that goes beyond the scope of this paper. Due to the dependence on local sensing results, synchronization is more difficult to achieve than in related AFH setups. Collaborative sensing concepts are a possible solution approach. By exchanging sensing metrics, the detection process could be coordinated, and the stochastic control actions can be synchronized by using identical random seeds. Methods such as acknowledgement feedback are also applicable [6].

VI. Conclusion

In conclusion, this paper presented an experimental cognitive radio test bed which uses sensing and prediction to exploit temporal white space between primary WLAN transmissions. By comparing measurement results with simulations based on an idealized mathematical model, we were able to validate model assumptions that have frequently been used in other works.

The experimental aspect of our work resulted in some test bed limitations which were due to hardware and complexity constraints. However, while the test bed is limited to single-channel operation, it captures fundamental deviations from the idealized mathematical model and incorporates carrier sensing and retransmission effects. We have further shown that generalizations to the multi-channel case are possible by extrapolating performance measurements for the single-channel case.

References

[1] Q. Zhao and B. M. Sadler, "Dynamic Spectrum Access: Signal Processing, Networking, and Regulatory Policy," *IEEE Signal Process. Mag.*, vol. 55, no. 5, pp. 2294–2309, May 2007.

[2] S. Geirhofer, L. Tong, and B. M. Sadler, "Cognitive Medium Access: Constraining Interference Based on Experimental Models," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 1, pp. 95–105, Jan. 2008.

Table 1: Model parameters for SMM and CTMC formulations. The parameters were obtained by experiment (see [21] for details) and used to obtain the simulation results presented in the previous section.

		WLAN Traffic Load										
		0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
CTMC	λ [ms]	23.3	11.6	7.89	5.42	3.32	2.34	1.63	1.01	0.68	0.43	0.24
	μ [ms]	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0
SMM	p_1	0.11	0.33	0.27	0.34	0.41	0.54	0.62	0.76	0.82	0.88	0.95
	k_1	0.75	-0.39	-0.33	-0.40	-0.52	-0.3	-0.09	0.02	0.12	0.14	0.07
	σ_1 [ms]	29.6	18.6	10.7	8.14	5.48	4.81	3.43	3.63	2.59	2.42	2.17
	k_2	0.75	-0.39	-0.33	-0.40	-0.52	-0.3	-0.09	0.02	0.12	0.14	0.07
	σ_2 [ms]	0.07	0.34	0.32	0.34	0.41	0.29	0.22	0.18	0.15	0.12	0.11

[3] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized Cognitive MAC for Opportunistic Spectrum Access in Ad Hoc Networks: A POMDP Framework," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 3, pp. 589–600, Apr. 2007.

[4] S. Huang, X. Liu, and Z. Ding, "Opportunistic spectrum access in cognitive radio networks," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 1427–1435.

[5] A. Sahai, N. Hoven, S. Mishra, and R. Tandra, "Fundamental tradeoffs in robust spectrum sensing for opportunistic frequency reuse," University of California, Berkeley, Tech. Rep., 2006.

[6] Q. Zhao, S. Geirhofer, L. Tong, and B. M. Sadler, "Opportunistic Spectrum Access via Periodic Channel Sensing," *IEEE Trans. Signal Process.*, vol. 56, no. 2, pp. 785–796, Feb. 2008.

[7] Bluetooth Special Interest Group, "Specification of the Bluetooth System," Nov. 2004.

[8] I. Akyildiz, W. Lee, M. Vuran, and S. Mohanty, "NeXt generation/dynamic spectrum access/cognitive radio wireless networks: A survey," *Computer Networks*, vol. 50, no. 13, pp. 2127–2159, Sep. 2006.

[9] Q. Zhao, L. Tong, and A. Swami, "Decentralized Cognitive MAC for Dynamic Spectrum Access," in *Proc. First IEEE Internatial Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005, pp. 224–232.

[10] Y. Chen, Q. Zhao, and A. Swami, "Joint Design and Separation Principle for Opportunistic Spectrum Access in the Presence of Sensing Errors," *IEEE Trans. Inf. Theory*, vol. 54, no. 5, pp. 2053–2071, May 2008.

[11] A. N. Mody, S. R. Blatt, D. G. Mills, T. P. McElwain, McElwain, N. B. Thammakhoune, J. D. Niedzwiecki, M. J. Sherman, C. S. Myers, and P. D. Fiore, "Recent Advances in Cognitive Communications," *IEEE Commun. Mag.*, vol. 45, no. 10, pp. 54–61, Oct. 2007.

[12] A. Motamedi and A. Bahai, "MAC Protocol Design for Spectrum-agile Wireless Networks: Stochastic Control Approach," in *Proc. IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Apr. 2007, pp. 448–451.

[13] S. D. Jones, N. Merheb, and I.-J. Wang, "An experiment for sensing-based opportunistic spectrum access in CSMA/CA networks," in *First IEEE Internatial Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005, pp. 593–596.

[14] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: Slotted Seeded Channel Hopping for Capacity Improvement in IEEE 802.11 Ad-Hoc Wireless Networks," in *Proc. ACM MobiCom*, 2004.

[15] K. G. Shin, H. Kim, C. Cordeiro, and K. Challapali, "An Experimental Approach to Spectrum Sensing in Cognitive Radio Networks with Off-the-Shelf IEEE 802.11 Devices," in *Proc. IEEE Consumer Communications and Networking Conference (CCNC)*, Jan. 2007, pp. 1154–1158.

[16] H. Kim and K. G. Shin, "In-band spectrum sensing in cognitive radio networks: energy detection or feature detection?" in *Proc. ACM MobiCom*, 2008, pp. 14–25.

[17] F. Perich, R. Foster, P. Tenhula, and M. McHenry, "Experimental Field Test Results on Feasibility of Declarative Spectrum

Management," in *Proc. IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, Oct. 2008, pp. 1–10.

[18] M. Ghosh and V. Gaddam, "Bluetooth interference cancellation for 802.11g WLAN receivers," in *Proc. IEEE International Communications Conference (ICC)*, May 2003, pp. 1169–1173.

[19] N. Golmie, "Bluetooth Dynamic Scheduling and Interference Mitigation," *Mobile Networks and Applications*, vol. 9, pp. 21–31, 2004.

[20] N. Golmie, N. Chevrollier, and O. Rebala, "Bluetooth and WLAN Coexistence: Challenges and Solutions," *IEEE Trans. Wireless Commun.*, vol. 10, no. 6, pp. 22–29, Dec. 2003.

[21] S. Geirhofer, L. Tong, and B. M. Sadler, "Dynamic Spectrum Access in the Time Domain: Modeling and Exploiting Whitespace," *IEEE Commun. Mag.*, vol. 45, no. 5, pp. 66–72, May 2007.

[22] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Springer-Verlag, 1994.

[23] Z. Liu, J. Almhana, V. Choulakian, and R. McGorman, "Traffic Modeling with Gamma Mixtures and Dynamical Bandwidth Provisioning," in *Proc. IEEE Communication Networks and Services Research Conference (CNSR)*, May 2006.

[24] S. Geirhofer, J. Z. Sun, L. Tong, and B. M. Sadler, "CMAP: A Real-Time Prototype for Cognitive Medium Access," in *Proc. IEEE Military Communications Conference (MILCOM)*, Oct. 2007.

[25] S. Avallone, A. Botta, D. Emma, S. Guadagno, and A. Pescape, "D-ITG V.2.4 Manual," University of Napoli "Federio II", Tech. Rep., Dec. 2004.