Cognitive Medium Access: Constraining Interference Based on Experimental Models

Stefan Geirhofer, *Student Member, IEEE,* Lang Tong, *Fellow, IEEE,* and Brian M. Sadler, *Fellow, IEEE*

Abstract— In this paper we design a cognitive radio that can coexist with multiple parallel WLAN channels while abiding by an interference constraint. The interaction between both systems is characterized by measurement and coexistence is enhanced by predicting the WLAN's behavior based on a continuous-time Markov chain model. Cognitive Medium Access (CMA) is derived from this model by recasting the problem as one of constrained Markov decision processes. Solutions are obtained by linear programming. Furthermore, we show that optimal CMA admits structured solutions, simplifying practical implementations. Preliminary results for the partially observable case are presented. The performance of the proposed schemes is evaluated for a typical WLAN coexistence setup and shows a significant performance improvement.

Index Terms— Cognitive Radio, Dynamic Spectrum Access, Standards Coexistence, IEEE 802.11b, Resource Management.

I. INTRODUCTION

THE GROWING demand in wireless technology has resulted in a dense allocation of relevant frequency bands. So far, regulators avoid mutual interference by assigning bands that do not overlap in frequency. This traditional static approach, however, leads to inefficient usage in both spatial and temporal domains. In fact, recent measurements [1] report that generally less than 5% of the spectrum are used at any given time and location. A static frequency allocation not only leads to inefficient spectrum usage, it moreover confines many services to unlicensed bands. A sizeable portion of the wireless consumer equipment falls into this category. Consider, for instance, WLAN, Bluetooth, cordless phones, and similar applications.

As a consequence, we face today overly crowded unlicensed bands whose performance is typically limited by mutual interference, and inefficient spectrum usage in those bands statically assigned by regulators. Is inefficient usage the price we have to pay for avoiding interference?

The emerging area of dynamic spectrum access (DSA) sparked by recent advances in software-defined and cognitive

radios presents a possible solution to this problem. This paper focuses on hierarchical schemes [2], in which the secondary system, *i.e.* the cognitive radio, is designed such that no or only insignificant interference is generated towards the primary user. Orthogonality between both systems can be achieved by exploiting different degrees-of-freedom in designing the system. While most of the research has focused on limiting mutual interference by spatial separation [3], we achieve orthogonality in the time domain by reusing idle periods that remain between the bursty packet transmissions of the primary system, represented by multiple, independently evolving WLAN channels. The cognitive radio considered in this paper is based on a frequency hopping setup with a physical layer similar to Bluetooth. This allows us to draw some conceptual parallels to Bluetooth/WLAN coexistence setups. We emphasize that implementing Cognitive Medium Access (CMA) as an extension of Bluetooth requires additional sensing and processing that current hardware designs may not easily accommodate. Our results indicate the significant potential for incorporating CMA techniques in future systems.

A. Main contribution

The main contribution of this paper is the derivation of CMA, a protocol that enhances WLAN coexistence based on sensing and prediction. Our analysis is based on measurement-based models, both for predicting the WLAN's behavior and for characterizing the cognitive radio's impact.

CMA is fundamentally based on a stochastic model for the WLAN's packet transmissions. We have previously shown through theory and experiment that a semi-Markov model captures this stochastic behavior accurately [4]. We consider this model and derive practical access schemes based on a continuous-time Markov chain (CTMC) approximation. This model, together with a sense-before-transmit strategy, allows us to constrain the interference generated towards the primary user.

The cognitive radio's throughput is optimized by recasting the problem as a constrained Markov decision process (CMDP). The formulation of the CMDP depends on the sensing capabilities of the radio frontend. If all parallel channels can be observed simultaneously we say the system is fully observable. On the other hand, if only a limited number of channels can be sensed at a time, we have to address this partial observability, which significantly complicates the derivation of CMA.

It is well known that CMDPs can be solved via linear programs, and we briefly state standard solution techniques. Additionally, we show that our problem setup admits struc-

tured solutions that allow us to gain further insight into the problem and can be used to solve the optimization problem with reduced complexity.

The numerical assessment of our algorithms is based on typical WLAN deployments. We examine the performance for the CTMC model, and investigate its robustness by running the algorithms on data generated using the accurate semi-Markov model. A comparison with a blind reference scheme, which does not perform any sensing, shows a significant performance improvement.

B. Related work

Dynamically accessing spectrum in the time domain has received increasing interest [3], [5]. A taxonomy of existing architectures is introduced in [2]. Among the first to consider time domain spectrum sharing are [6], [7], where it is assumed that primary and secondary system share the same slot structure. Based on a Partially-Observed Markov Decision Process (POMDP) framework, access strategies for the secondary system are derived. The assumption of both systems having the same slot structure is, however, too restrictive in many cases.

A semi-Markov model for predicting the WLAN transmissions has been introduced in [4], [8]. In order to simplify the derivation of access schemes, a CTMC approximation has been considered in [9] and [10].

Coexistence in unlicensed bands has previously received significant attention outside the cognitive radio community [11] with Bluetooth/WLAN coexistence being a prominent example of practical concern. A possible approach to interference mitigation in this context is adaptive frequency hopping [12]. The CMA protocol is conceptually similar to such schemes but differs in its physical layer sensing. In current Bluetooth devices such sensing is not required, and consequently interference information needs to be inferred from higher layers.

The remainder of the paper is organized as follows. In Sec. II, the system setup is introduced. Sec. III presents the measurement-based interference models. In Sec. IV, optimal CMA is derived for fully and partially observed systems, constituting this paper's main contribution. Numerical performance results are presented in Sec. V.

II. SYSTEM SETUP

The physical layer setup considered in this paper consists of M parallel, independently evolving WLAN channels, as shown in Fig. 1. The cognitive radio shares the same frequency band and dynamically hops through the WLAN bands based on sensing and statistical prediction. A value of $M = 3$ bands is typical since practical WLAN setups in the ISM band at 2.4 GHz support three non-overlapping channels [13].

A. Physical layer setup

Although there are no restrictions in designing the cognitive radio, other than our ultimate goal of minimizing mutual interference, we shall focus on the frequency hopping (FH) setup depicted in Fig. 1. Let each of the M WLAN bands overlap with N narrowband hopping channels.
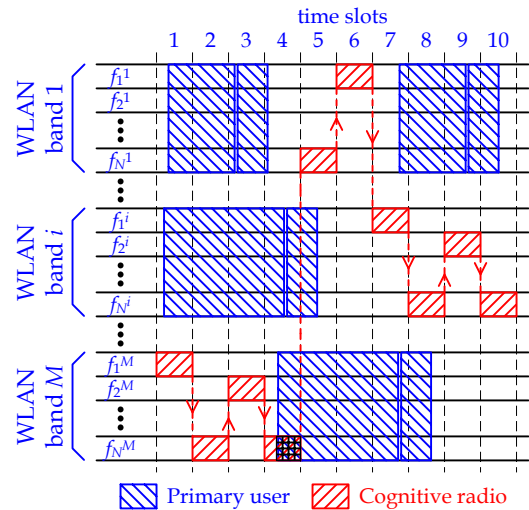

Fig. 1. System setup. The cognitive radio is a time-slotted frequency hopping system.

The choice of a FH setup is considered for two reasons. First, WLAN is an *unslotted* system that performs medium access based on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). A logical approach toward enforcing an interference constraint is thus to sense the medium periodically, and transmit in a slotted fashion. Second, mutual interference is reduced by exploiting the fact that WLAN uses spread spectrum communications and thus has some inherent robustness to narrowband interference.

The above considerations are motivated by practical experience. The FH setup considered in this work has conceptual similarities to Bluetooth. We use this similarity to find realistic deployment parameters (slot size, modulation characteristics, etc.) based on the standard [14]. CMA is therefore a potential cognitive extension to FH systems such as Bluetooth, based on sensing and statistical prediction of the WLAN.

Apart from the FH setup, we also considered a direct-sequence spread-spectrum physical layer for the cognitive radio, spanning the same M frequency bands as the WLAN. The spread-spectrum setup is also slotted and designed to dynamically hop across the WLAN bands. For the spread-spectrum setup, we found by experiment [15] that as long as the spreading code is different from the WLAN's, interference properties are similar to the FH setup. As a consequence, CMA extends to this setup as well, as it is only concerned with finding a hopping sequence across bands (but not within bands for the FH setup).

B. Operational block diagram

The operation of the cognitive radio is illustrated in Fig. 2. An RF frontend is used for up- and down-conversion of the signals, and sampling is performed using an acquisition board. At the beginning of every slot a spectrum sensor determines whether the medium is busy, either based on energy detection or by exploiting features of the WLAN standard. The sensing result is processed within the CMA controller, which determines whether it is 'safe' to transmit, and if yes, in which
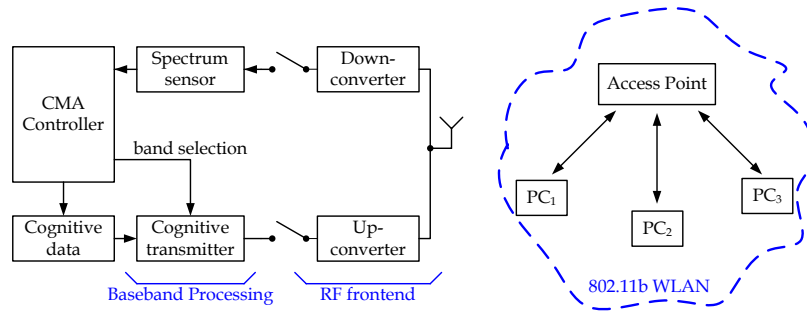
Fig. 2. Block diagram of the cognitive radio's operation. The cognitive radio is shown on the left and coexists with the WLAN on the right.

band. The secondary transmitter is tuned accordingly and a transmission may be initiated.

According to the above, on a slot level, the following operations are performed in sequence. At the beginning of every slot the medium is sensed, the CMA controller decides on which band, if any, to transmit on, the transmitter is retuned accordingly, and a transmission takes place for the remainder of the slot period.

The sensing time, the run time of the controller, and the time it takes to retune the transmitter contribute to the overhead of the system. We have previously analyzed the performance of energy detection on WLAN signals [8]. With a detection error probability of 10^{-5} and a signal-to-noise ratio (SNR) of 5 dB we require a sensing time of less than 5 µs. Thus the sensing overhead at the beginning of every slot is fairly small. The CMA controller implements a randomized control policy (to be described in detail later), which can be implemented by biased coin flips. The delay associated with it is small as well. Finally, current technology yields a frequency retuning time on the order of 100 µs [16], and thus dominates the processing overhead time. In our numerical evaluations we choose a slot size of $T_s = 625$ µs, which is the same as in Bluetooth. We believe this is a practical choice given the similarities in the physical layer setup.

We stress that since every WLAN channel overlaps with N hopping channels, the hopping within each band decouples from selecting one of the M WLAN bands. This paper solely deals with the optimal selection of one of the M bands. Hopping within each band can, for instance, be performed pseudo randomly.

In this paper, we assume that the secondary system is synchronized. In a practical system, maintaining identical hop sequences within the cognitive system may be challenging, as sensing results obtained at different nodes could potentially differ. While it is beyond the scope of this paper to address this issue in detail, we believe that collaborative sensing techniques could be used to provide hop sequence coordination. By exchanging sensing metrics across subsequent slots, master and slave could perform sensing jointly and thus arrive at identical results. Another potential scheme for maintaining synchronization is to employ acknowledgement feedback [10].

III. MEASUREMENT-BASED INTERFERENCE MODEL

CMA is a protocol that dynamically hops across multiple parallel WLAN bands in an optimal fashion. It is fundamentally based on a measurement-based prediction model, which

will be described in detail in Sec. IV. In addition we present physical-layer coexistence models, that describe the interaction between both systems, should a collision occur. This experimental coexistence model is described in the following.

In particular, we first evaluate whether the cognitive radio affects the WLAN's carrier sensing. Second, we obtain empirical results for the probability that a collision between both systems leads to a WLAN packet error. Due to space limitations, we focus on a qualitative assessment; detailed quantitative results can be found in a separate technical report, available online [15].

A. Impact on WLAN carrier sensing

The design of the cognitive radio needs to ensure that its transmissions remain transparent to the WLAN. This implies that the WLAN's carrier sensing must not be altered. Otherwise not only our paradigm of hierarchical DSA would be undermined, but also the cognitive radio's dynamic effect on the WLAN would render our prediction model useless, unless it incorporated the WLAN's retransmission behavior.

We evaluated the cognitive radio's impact by measuring the probability that the WLAN detects the cognitive radio's transmission. The experimental setup is shown in Fig. 3(a). It consists of an 802.11b router and an RF signal source, generating the WLAN and the FH signal, respectively. More precisely, we consider a static (non-hopping) FH signal with Bluetooth's modulation parameters [14]. As the signal remains static in one of the channels it is possible to examine the mutual interference resulting from a specific channel.

The WLAN adapter and the signal source are connected via circulators, which couple generator and router while providing isolation in the reverse direction. A WLAN adapter is used to capture the received signal, and a vector signal analyzer is used to verify the correct operation of the setup.

The impact of the FH signal is assessed by measuring its effect on the WLAN packet rate. The WLAN router continuously transmits packets, and a WLAN adapter card is used to measure its average rate (by capturing packets over long periods of time). In the presence of the cognitive system, if the WLAN detected the interference, its rate would decrease as back-off periods would need to be accommodated.

The impact on the WLAN largely depends on the interferer's channel index. Our results suggest that the WLAN adapter performs energy detection in narrow bands spaced about 10 MHz apart. No impact was observed outside these

(a) Impact on WLAN's carrier sensing
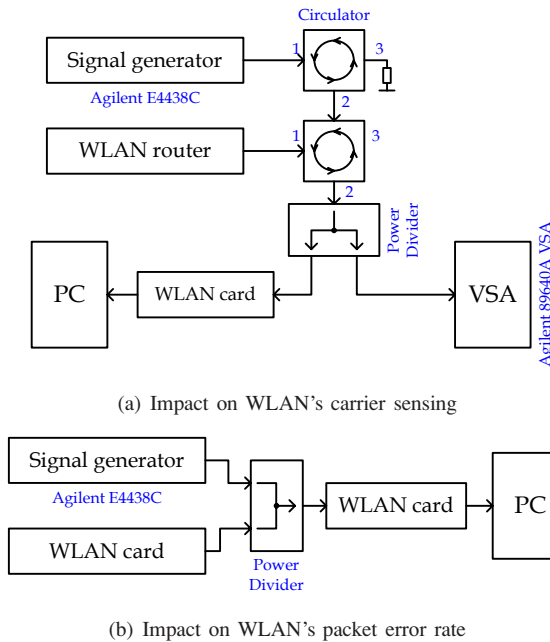

(b) Impact on WLAN's packet error rate

Fig. 3. Experimental setup for evaluating the cognitive radio's impact on the WLAN's carrier sensing and packet error rate.

channels, even for fairly high interference powers. Due to space limitations quantitative results are not included in this paper, but given in a technical report, available online [15].

Based on our quantitative results [15] and given typical setups [17] and path loss models [18], we conclude that the cognitive radio does not alter the WLAN's medium access. This solidifies our hierarchical approach and renders the stochastic prediction model applicable.

Furthermore, we have analyzed whether measuring rate changes is an appropriate method to determine the cognitive radio's impact. We validated our measurement approach by performing the same analysis for a different type of interferer, namely a WLAN-type signal using the *same* spreading code as standardized for 802.11b [19]. We observed that the adapter card is significantly more sensitive to this type of signal. The power level above which an impact occurs was determined to be -77 dBm (see [15] for details). This is in accordance with the 802.11b standard [19, p.58] which mandates the sensitivity to be -76 dBm or better.

B. Effect on packet error rate

The second component of our interference model focuses on the cognitive radio's impact on WLAN's packet error rate. Specifically, we measure the probability that a collision between both systems leads to a WLAN packet error. The measurement setup is shown in Fig. 3(b). It consists of a WLAN adapter card and a signal source generating the WLAN signal and the FH interferer, respectively. The signals are combined and captured via another WLAN card and commercial packet capturing software. A vector signal analyzer is used to verify the operation of the setup.

The packet error probability is measured in the following way. A continuous stream of packets is generated and captured at the receiver to determine the rate of packets with the interferer turned off. Subsequently, in the presence of interference
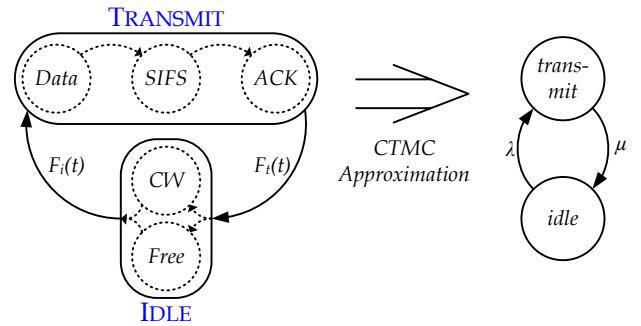

Fig. 4. Semi-Markov model (SMM) and its continuous-time Markov chain (CTMC) approximation. The SMM is shown on the left. By lumping states and approximating the holding times by exponential distributions we arrive at the CTMC approximation on the right.

the rate will decrease since some packets will be too distorted to be captured by the adapter. Other packets will be captured but will show an invalid redundancy check. By comparing the number of successfully received packets with the interference-free case, we can determine the probability of a packet error.

The impact of the FH interferer depends on the channel index. Close to the center frequency, a significant impact is observed for signal to interference ratios (SIR) of less than 0 dB. For instance, for an offset of 3 MHz from center frequency and an SIR of -3 dB, we observe a packet error rate of 85%. If the SIR drops below -5 dB virtually every packet is lost. The impact of the FH interferer decreases as we move away from the center frequency. This is not surprising and has previously been reported [20]. It is due to the downconversion and filtering performed within the WLAN receiver. More quantitative results for this measurement scenario are given in [15].

IV. COGNITIVE MEDIUM ACCESS

The derivation of CMA is based on the empirical interference model. In short, we have seen that while the WLAN's carrier sensing remains unaltered, a packet collision is likely to cause a packet error. For ease of analysis, and because such an assumption has frequently been made in other papers [11], we assume that every collision inevitably results in a packet error. This is a worst case assumption given our measurement results.

A. Empirical WLAN Model and CTMC Approximation

Since collisions cause packet errors we need to constrain the rate of collisions between both systems. The derivation of CMA is based on a previously established WLAN prediction model [4], [8]. The key components of this model are briefly reviewed, specifically the semi-Markov model and its CTMC approximation. We stress that both models have been based on and validated through experimental data gathered via a sensing testbed [8]. The model is based on empirical data for the idle and busy durations of the medium. In short, the stochastic model enables us to predict whitespace and direct our cognitive radio such that it preferably hops to bands not currently used by the WLAN.

1) Semi-Markov model: Analyzing the idle and busy durations obtained by measurement we found that almost always the transmission of a data packet is followed by a short interframe space (SIFS) and an acknowledgement from the receiver. This does not come as a surprise, given that such acknowledgements are mandated by the standard and merely correspond to a successful packet exchange [13].

We thus arrive at the transition diagram depicted in Fig. 4, which consists of alternating packet transmissions (including their mandatory acknowledgements) and idle periods, respectively. The transitions between DATA, SIFS, and ACK state are deterministic. For CMA it is crucial to predict how long the system remains in either state. In statistical terminology this corresponds to finding distributions $F_t(t)$ and $F_i(t)$ for the holding (or sojourn) times in the TRANSMIT and IDLE state, respectively.

Clearly, the holding time in the TRANSMIT state corresponds to the length of the packets, which is largely influenced by the type of traffic and the specific properties of the scheduler. Our measurements suggest that over short time periods, $F_t(t)$ consists of several (less than five) discrete values. In the traffic scenario to be analyzed in this paper the packet lengths are, in fact, almost deterministic [21].

Predicting the idle durations is more involved. As a matter of fact, an idle channel can either be due to inactivity of the medium or be a residue of the WLAN's medium access which has to guarantee each station an equal transmission opportunity. The latter is implemented via a backoff procedure requiring stations to defer medium access for uniformly-distributed random time periods. In essence, the station with the smallest backoff gets to access the medium first. As a consequence of the above we observe uniformly distributed idle durations resulting from the above operation.

If the medium is inactive, *i.e.* if none of the stations has any data to transmit, we have shown by experiment that the idle durations are well approximated by a generalized Pareto distribution [8].

The existence of both effects leads to a mixture distribution to model the holding time in the IDLE state. We arrive at

$$F_i(t) = p_{cw}F_u(t) + (1 - p_{cw})F_{gp}(t), \qquad (1)$$

where p_{cw} denotes the probability that an idle duration is due to the contention window, $F_u(t)$ is a uniform distribution within $[0, 0.7\,\mathrm{ms}]$ [8], and

$$F_{gp}(t; k, \sigma) = 1 - \left(1 + k\frac{t}{\sigma}\right)^{-1/k} \qquad (2)$$

is the CDF of a generalized Pareto distribution with shape parameter k and scale parameter σ. By fitting the mixture distribution (1) to the empirical data we can estimate the above parameters. The accuracy has been validated by statistical measures of fit in [8].

2) CTMC approximation: The semi-Markov model presented above provides for an excellent fit with empirical data. However, in deriving CMA the semi-Markov model is difficult to analyze since it does not possess the continuous-time Markov property. In order to simplify analysis we consider a CTMC approximation, which corresponds to fitting exponential distributions to the idle and busy periods. The

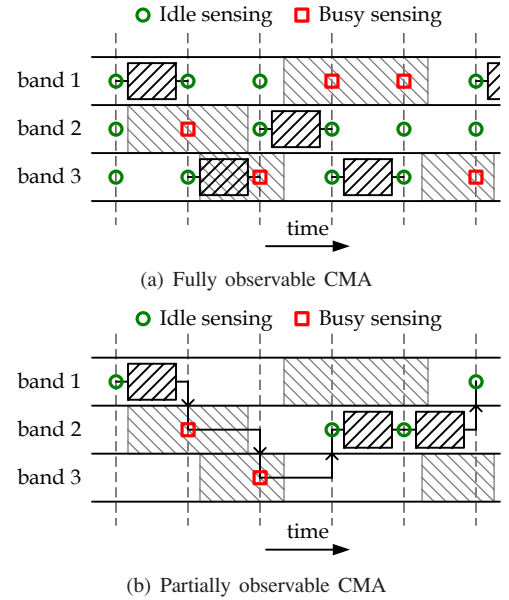(a) Fully observable CMA

(b) Partially observable CMA

Fig. 5. Illustration of fully and partially observable CMA. Circles and squares denote idle and busy sensing results, respectively. In the fully observable scenario, all bands are sensed simultaneously and depending on this result, a transmission may be initiated in one of the channels. In the partially observed case, only one band can be observed at a time. Consequently, actions are limited to either staying in the current band, or hopping to another.

exponential fit provides a good approximation although it is not strongly validated by statistical measures of fit. More details can be found in [15]. The parameters of the CTMC model are thus λ and μ, leading to $F_t(t) = 1 - \exp(-\mu t)$ and $F_i(t) = 1 - \exp(-\lambda t)$, respectively.

B. Fully observable CMA (FO-CMA)

The sensing based access schemes presented in this paper can be categorized according to the sensing capabilities of the cognitive radio. If the spectrum sensor supports a high enough bandwidth, the state of all M channels can be observed simultaneously at the beginning of every slot. This is illustrated in Fig. 5(a). Based on the sensing result the cognitive controller decides in which, if any, channel to transmit.

If the spectrum sensor has limited bandwidth, we assume that only one of the M bands can be sensed at a time; the state of the other bands remains hidden. Furthermore, we assume that for practical reasons, a transmission can only be initiated in the channel that has just been sensed. This assumption makes the system *partially observable* and significantly complicates the analysis. This is illustrated in Fig. 5(b), where only one channel is observed at the beginning of any slot.

We start with the fully observed case and first recast the problem mathematically as a constrained Markov decision process. The standard solution technique [22] is briefly reviewed. We then show that our problem setup admits a structured solution. The partially observed case is addressed separately in Sec. IV-E.

In order to find optimal access strategies we need to formulate a constrained optimization problem [9]. Let each of the WLAN bands $i = 1, \ldots, M$ evolve as a CTMC $\{X_i(t), t \geq 0\}$ with states '0' (IDLE) and '1' (TRANSMIT). The holding

times are exponentially distributed with parameters λ_i and μ_i, respectively, as shown in Fig. 4.

The generator matrix $\mathbf{G}_i$ for channel i is hence given by

$$\mathbf{G}_i = \begin{bmatrix} -\lambda_i & \lambda_i \\ \mu_i & -\mu_i \end{bmatrix}, \tag{3}$$

which leads to the stationary distribution

$$\eta_0^{(i)} = \frac{\mu_i}{\lambda_i + \mu_i}, \quad \eta_1^{(i)} = \frac{\lambda_i}{\lambda_i + \mu_i} \tag{4}$$

and transition matrix [23, p.391]

$$\mathbf{P}^{(i)} = \frac{1}{\lambda_i + \mu_i} \begin{bmatrix} \mu_i + \lambda_i e^{-(\lambda_i + \mu_i)t} & \lambda_i - \lambda_i e^{-(\lambda_i + \mu_i)t} \\ \mu_i - \mu_i e^{-(\lambda_i + \mu_i)t} & \lambda_i + \mu_i e^{-(\lambda_i + \mu_i)t} \end{bmatrix}. \tag{5}$$

The secondary system senses the state of the entire system at the beginning of every slot, inducing a discrete-time chain $\{Y_i[k], k \geq 0\}$ of sensing results for each channel i. For notational convenience let us define the vector valued random process $\{\mathbf{Y}[k], k \geq 0\}$ that contains the latest sensing result for all channels,

$$\mathbf{Y}[k] = \begin{bmatrix} Y_1[k], \dots, Y_M[k] \end{bmatrix}^T. \tag{6}$$

It is straightforward to verify that $\mathbf{Y}[k]$ is a discrete-time Markov chain with state space $\mathbb{X} = \{0,1\}^M$. The transition matrix becomes, due to the independence of the WLAN bands,

$$\mathbf{P_{xy}} = \prod_{i=1}^{M} \mathbf{P}_{x_i y_i}^{(i)}, \quad \mathbf{x}, \mathbf{y} \in \mathbb{X} \tag{7}$$

and we arrive at the following expression for the stationary distribution

$$\eta_{\mathbf{x}} = \prod_{i=1}^{M} \eta_{x_i}^{(i)}. \tag{8}$$

Given the sensing results in each slot, the CMA controller decides whether to transmit and if yes, in which channel. The action set is thus $\mathbb{A} = \{0, 1, \dots, M\}$, where $a = 0$ denotes that no transmission takes place, and $a \geq 1$ means that a transmission is scheduled in channel a.

Transmitting across channel a accrues a unit reward provided no collision occurs. The expected immediate reward of choosing action a in state $\mathbf{y}$ thus becomes

$$r(\mathbf{y}, a) = \begin{cases} 1_{[y_a = 0]} e^{-\lambda_a T_s}, & a \geq 1 \\ 0, & a = 0 \end{cases}, \tag{9}$$

where $1_{[\cdot]}$ represents the indicator function and T_s denotes the slot duration.

The interference constraint can be formulated in several ways. We shall call slot k in band i busy, and denote this as $\mathcal{C}_i[k] = 1$, if

$$\exists t \in [kT_s, (k+1)T_s) \text{ s.t. } X_i(t) \neq 0. \tag{10}$$

The interference constraint can then be formulated as

$$D_{\text{CIC}} = \lim_{N \to \infty} \frac{\sum_{k=1}^{N} 1_{[A_k = i, \mathcal{C}_i[k] = 1]}}{N}, \tag{11}$$

where A_k refers to the action taken in slot k. It is capitalized to stress that A_k is random; it depends on the current sensing result and the action (randomly) chosen by the CMA controller. The above equation corresponds to the long run

fraction of slot collisions per unit time. We will refer to (11) as the *cumulative interference constraint* (CIC).

While the CIC seems an intuitive measure, it quantifies the interference from the secondary user's perspective. The density of the WLAN traffic is not taken into account. The formulation can be better tailored to the primary user by imposing a *packet error rate constraint* (PERC) for all bands

$$D_{\text{PERC}}^{(i)} = \lim_{N \to \infty} \frac{\sum_{k=1}^{N} 1_{[A_k = i, \mathcal{C}_i[k] = 1]}}{N_i(NT_s)}, \quad 1 \leq i \leq M, \tag{12}$$

where $N_i(t)$ counts the number of transmitted WLAN packets in band i up to time t,

$$N_i(t) = \sum_{n=0}^{\infty} 1_{[S_n^{(i)} \leq t]}. \tag{13}$$

In the above equation, $S_n^{(i)}$ denotes the arrival times of WLAN packets in band i. Therefore (12) is the long-run fraction of collisions per transmitted WLAN packets. In short, this represents the fraction of WLAN packets that get dropped due to the cognitive radio's interference.

Based on the above definitions, we define the expected immediate costs for the CIC,

$$d_{\text{CIC}}(\mathbf{y}, a) = \begin{cases} 1 - e^{-\lambda_a T_s} & \text{if } y_a = 0, \, a \geq 1 \\ 1 & \text{if } y_a = 1, \, a \geq 1 \\ 0 & \text{if } a = 0 \end{cases}, \tag{14}$$

and for the PERC,

$$d_{\text{PERC}}(\mathbf{y}, a) = \begin{cases} \frac{(\lambda_a + \mu_a)(1 - e^{-\lambda_a T_s})}{\mu_a \lambda_a T_s} & \text{if } y_a = 0, \, a \geq 1 \\ 1 & \text{if } y_a = 1, \, a \geq 1 \\ 0 & \text{if } a = 0 \end{cases}. \tag{15}$$

Having introduced rewards and costs, the CMDP can now be defined. CMA maximizes

$$J(\beta, \pi) = \lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} \mathbb{E}_{\beta}^{\pi} r(\mathbf{Y}_t, A_t) \tag{16}$$

with respect to policy π, subject to a CIC

$$D_{\text{CIC}}(\beta, \pi) = \lim_{N \to \infty} \frac{1}{N} \sum_{t=1}^{N} \mathbb{E}_{\beta}^{\pi} d_{\text{CIC}}(\mathbf{Y}_t, A_t) \leq \alpha, \tag{17}$$

or subject to PERCs, (see Eq. 18).

In the above formulas β denotes the initial distribution of the system, and π the policy we maximize for. The expectation operator is thus taken with respect to the probability distribution induced by π given initial distribution β.

C. Linear programming solution

It is well known that a CMDP's optimal policy is in the space of Markovian randomized policies [24]. The optimal policy π^* is hence a function that maps state-action pairs $(\mathbf{y}, a)$ to the probability of choosing action a in state $\mathbf{y}$, $\pi^* : \mathbb{X} \times \mathbb{A} \to [0, 1]$. Since, in CMDPs both the reward and the constraints can be expressed using the frequency of state-action pairs $(\mathbf{y}, a)$ the optimal policy can be found by linear programming [22].

$$D_{\text{PERC}}^{(i)}(\beta,\pi) = \lim_{N\to\infty} \frac{1}{N}\sum_{t=1}^{N}\mathbb{E}_\beta^\pi 1_{[A_t=i]}d_{\text{PERC}}(\mathbf{Y}_t,A_t) \le \alpha_i,\ 1\le i\le M. \tag{18}$$

Theorem 1 [22, p.38]: The linear program

$$\max_{\rho(\mathbf{y},a)} \sum_{\mathbf{y}\in\mathbb{X}}\sum_{a\in\mathbb{A}(\mathbf{y})} \rho(\mathbf{y},a)r(\mathbf{y},a) \tag{19}$$

subject to

$$\sum_{\mathbf{y}\in\mathbb{X}}\sum_{a\in\mathbb{A}(\mathbf{y})} \rho(\mathbf{y},a)d_i(\mathbf{y},a) \le \alpha_i,\ 1\le i\le M, \tag{20}$$

where $\rho(\mathbf{y},a)\in\mathbb{Q}(\beta)$ and

$$\mathbb{Q}(\beta) = \left\{ \begin{array}{l} \rho(\mathbf{y},a), \mathbf{y}\in\mathbb{X}, a\in\mathbb{A}(\mathbf{y}): \\ \sum_{\mathbf{y}\in\mathbb{X}}\sum_{a\in\mathbb{A}(\mathbf{y})} \rho(\mathbf{y},a)(\delta_\mathbf{y}(\mathbf{x})-P_{\mathbf{x}a\mathbf{y}})=0 \\ \sum_{\mathbf{y}\in\mathbb{X}}\sum_{a\in\mathbb{A}(\mathbf{y})} \rho(\mathbf{y},a)=1,\ \rho(\mathbf{y},a)\ge 0 \end{array} \right\} \tag{21}$$

is equivalent to the CMDP formulations (16)-(18). □

The CMDP's optimal policy is completely determined by the state-action frequencies. After obtaining $\rho(\mathbf{y},a)$ via the linear program the probability $w_\mathbf{y}(a)$ of choosing action a in state $\mathbf{y}$ simply becomes,

$$w_\mathbf{y}(a) = \frac{\rho(\mathbf{y},a)}{\sum_{a\in\mathbb{A}(\mathbf{y})}\rho(\mathbf{y},a)},\ \mathbf{y}\in\mathbb{X},\ a\in\mathbb{A}(\mathbf{y}), \tag{22}$$

provided the denominator is non-zero (arbitrary otherwise).

D. Optimal FO-CMA structure

In this section we show that the above linear programs, under some conditions, admit structured solutions that help to gain insight into the problem. Moreover, using the structured results simplifies the implementation of CMA.

1) Cumulative interference constraint: First, consider the case of the cumulative interference constraint (20). Without loss of generality assume that $\lambda_1 \le \cdots \le \lambda_M$. We show that it is optimal to use only channels with small λ. How many channels should be used can be derived from a threshold model with respect to the constraint α. The solution structure is depicted in Fig. 6.

Algorithm 1 (Threshold solution)

i. Define the maximum interference level for channel i as

$$\xi_i = \sum_{\mathbf{y}\in\mathbb{X}} 1_{[y_i=0]}1_{[y_j=1,\forall j<i]}\eta_\mathbf{y}d_{\text{CIC}}(\mathbf{y},i) \tag{23}$$

ii. Based on the $\{\xi_i\}$, find the smallest k such that

$$\xi_1 + \cdots + \xi_k > \alpha \tag{24}$$

iii. Adopt the following randomized policy. With probability w_i, transmit in the idle channel with the lowest λ_i, *i.e.*, in state $\mathbf{y}$ transmit in channel i if and only if

$$y_i = 0 \text{ and } y_j = 1\ \forall j < i, \tag{25}$$

and choose not to transmit otherwise. The probabilities w_i are given with k as defined in step (ii),

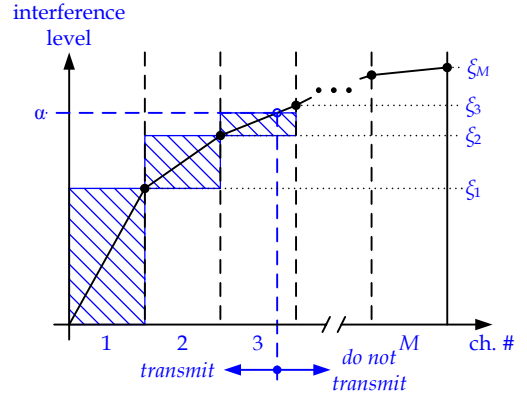$$w_j = 1,\ 1\le j < k,\ w_k = \frac{\alpha-\xi_{k-1}}{\xi_k-\xi_{k-1}},\ w_j=0,\ j>k \tag{26}$$

Fig. 6. Solution structure for the cumulative interference case.

Theorem 2: The policy induced by Algorithm 1 is a solution to the LP (19)-(20) and hence equivalent to CMA. □

Proof: see appendix.

2) Packet error rate constraints: In the case of PERCs, separate constraints for each channel need to be considered simultaneously. Intuitively, the maximum reward would be achieved if all constraints could be made tight (otherwise transmission opportunities would be wasted). It may not be feasible, however, to tighten all constraints, given that a transmission can be initiated in, at most, one channel per slot. Whether this is possible, in fact, depends on how loose the PERCs are chosen.

If the M constraints can be made tight, we show that the problem decouples, and the optimal policy can be found by considering each channel individually. This is the case if the following condition is met for all channels,

$$\xi_a = \sum_{\mathbf{x}\in\mathbb{X}} \frac{1_{[x_a=0]}\eta_\mathbf{x}}{\sum_{l=1}^{M}1_{[x_l=0]}} \ge \frac{\alpha_a}{d_a},\ \forall a\in\mathbb{A}, \tag{27}$$

where

$$d_a = \frac{(\lambda_a+\mu_a)(1-e^{-\lambda_a T_s})}{\mu_a\lambda_a T_s}. \tag{28}$$

is the expected average cost associated with a collision.

The intuition behind (27) is to find a condition under which the interference constraint can be made tight by considering channels separately. In fact, a tradeoff on which channel to transmit in need only be struck if $\mathbf{x}$ contains multiple zeros. The denominator accordingly normalizes by the number of transmission opportunities in state $\mathbf{x}$ and thus ensures that, although the M bands are considered separately, the probability of transmission for a given state will never exceed one.

We can thus adopt the following algorithm.

Algorithm 2:

i. The maximum activity level ξ_i in channel i is given by (27). Disregarding other transmission opportunities, we transmit

in band i with probability

$$w_i = \frac{\alpha_i}{d_i \xi_i}, \quad \forall i \in \mathbb{A} \tag{29}$$

ii. Since (27) is met, the constraints can be made tight, and we obtain

$$w_{\mathbf{y}}(a) = \frac{w_a 1_{[y_a=0]}}{\sum_{l=1}^{M} 1_{[y_l=0]}}. \tag{30}$$

Theorem 3: The policy induced by Algorithm 2 is a solution to the LP (19)-(21) and hence equivalent to CMA. □

Proof: see appendix.

E. Partially Observable CMA (PO-CMA)

In the last section we assumed that the state of all M channels can be observed simultaneously. In this section we alleviate this constraint and assume that only a single channel can be observed at a time, as shown in Fig. 5(b). Furthermore, a transmission can only be initiated in the channel that has just been sensed. The action set thus reduces to $\mathbb{A} = \{0, 1\}$ denoting whether or not a transmission is initiated.

The partial observability severely complicates the problem. It is now necessary to trade off the exploration of the system (by frequently sensing different bands) with the exploitation of transmission opportunities. In order to illustrate how such a tradeoff can be struck, we present some preliminary results for the special case of $M = 2$ channels.

A fundamental consideration in designing PO-CMA is which of the past observations and actions are useful for making optimal decisions. Given that all bands are modeled as CTMCs, the continuous-time Markov property leads to the conclusion that the latest sensor reading of every channel is sufficient for predicting its behavior.

This is illustrated in Fig. 7 for the special case of $M = 2$ bands. The states are labeled according to whether the current sensing result is busy or idle. A busy channel is simply denoted 'b', whereas for an idle channel we also keep track of how many consecutive slots the channel has been sensed idle. Hence the states labeled '0',…,'N' all correspond to an idle sensing result.

The sensing history is incorporated in brackets, where the first number reflects the currently active channel and the second index denotes the latest sensing result in the other channel, respectively. Note that the time since the other channel has last been sensed is given, in this special case of $M = 2$, by the number of slots spent in the current channel.

According to the above, the setup is fully described by the triple (y, i, x) where $y \in \mathbb{Y} = \{b, 0, \ldots, N\}$ denotes the current channel's sensing history, i the currently active channel index, and $x \in \{0, 1\}$ the last sensing result in the other band. The transition behavior can be understood as follows. Under action 1, *i.e.* keep transmitting, the following transitions are possible

$$(y, i, x) \rightarrow (y+1, i, x), \quad 0 \le y < N \tag{31}$$
$$(y, i, x) \rightarrow (0, \bar{i}, 1) \tag{32}$$
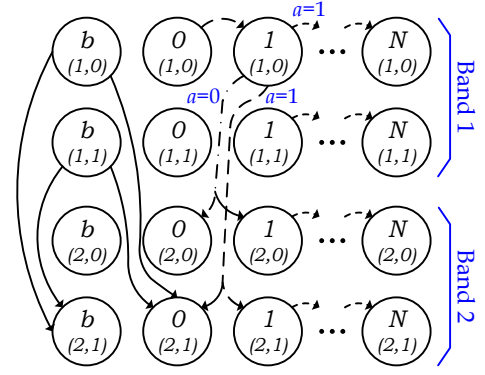$$(y, i, x) \rightarrow (b, \bar{i}, 1), \tag{33}$$

Fig. 7. Markov chain model for the partially observable case. Only some example transitions are shown. The indices denote the sensing result, the active channel number, and the last sensing result, respectively.

where the first line denotes the channel staying idle. The other two possible transitions correspond to the channel becoming busy and the cognitive radio thus relocating to the other band. This other band can in turn be either idle or busy and thus two different transitions can occur. The notation $\bar{i}$ represents the other 'band,' that is $\bar{i} = (i \mod 2) + 1$.

Under action 0, *i.e.* relocate to other band, the following transitions may occur

$$(y, i, x) \rightarrow (0, \bar{i}, 0) \tag{34}$$
$$(y, i, x) \rightarrow (b, \bar{i}, 0), \tag{35}$$

denoting a relocation to the other band and finding it either idle or busy, respectively. The transition behavior is shown in Fig. 7. In order to keep a finite state space, we cannot stay with any channel longer than for N slots. The optimal policy for the partially observable case can be found by linear programming, as shown in Sec. IV-C.

V. NUMERICAL RESULTS

In this section we present numerical results for the CMA schemes and evaluate their performance gain compared to a blind reference scheme that does not perform sensing. The schemes are evaluated in terms of throughput and interference for *varying* WLAN traffic load. Such analysis is standard in coexistence literature [25]. The results are based on simulations using the CTMC approximation. Furthermore, we examine the robustness of this approximation by running algorithms derived from the CTMC model on data generated via the semi-Markov model. We will see that the results match closely, justifying the approximation.

A. Simulation parameters

The numerical results reflect the throughput and interference behavior for varying WLAN traffic load. While performance is assessed by simulation, the model parameters (for both SMM and CTMC approximation) are extracted from real trace data. The experimental setup consisted of a wireless router and three workstations with adapter cards [8] to gather packet traces. Using the Distributed Internet Traffic Generator [26] on each workstation we generated constant-payload UDP traffic. The interdeparture time between packets was chosen

TABLE I

Measurement parameters for semi-Markov model and its continuous-time Markov chain approximation. This table was obtained by estimating model parameters from empirical data obtained from real packet traces.

Parameter	WLAN traffic load $\bar{\sigma} = \sigma/\sigma_{max}$						
	0.05	0.1	0.2	0.3	0.4	0.5	1.0
	CTMC approximation						
λ^{-1} [ms]	15.9	9.10	4.48	2.90	1.98	1.39	0.21
μ^{-1} [ms]	1.11	1.08	1.05	1.03	1.02	1.03	1.03
	semi-Markov model						
σ [ms]	18.7	11.3	5.46	3.95	3.09	2.35	0.04
$k/10^{-2}$	-2.11	-2.50	2.47	1.51	2.61	1.69	50.1
p_c [%]	13.2	18.3	21.2	30.1	40.0	47.7	98.8
T_μ [ms]	1.11	1.08	1.05	1.03	1.02	1.03	1.03

to be exponentially distributed with varying parameter σ. By initially choosing σ very large we first determined the maximum traffic load supported by the setup. Each setting of the traffic load was then normalized with respect to this maximum value $\bar{\sigma} = \sigma/\sigma_{max}$.

For each setting of the WLAN traffic load $\bar{\sigma}$, we captured the raw complex baseband data of the transmissions using a vector signal analyzer and processed the data to find the busy and idle durations of the WLAN. Based on these results we obtained the parameters of the semi-Markov model and its CTMC approximation, shown in Tab. I; see [15] for details. Remaining parameters were chosen such that they reflect a typical WLAN/Bluetooth coexistence setup. The slot duration was chosen as $T_s = 625\,\mu\text{s}$.

B. Performance of FO-CMA and PO-CMA

First, we evaluate FO-CMA's throughput for a CIC of $\alpha = 5\%$, and PERCs of $\alpha_i = 10\%$ for each channel. The throughput for both scenarios is shown in Fig. 8 for $M = 1, 2$, and 3 parallel bands. We see that for the CIC, we obtain the same throughput regardless of how many parallel bands exist. Indeed, this is to be expected, since we deal with a *cumulative* constraint. Having multiple parallel bands available does not loosen this limitation. In contrast, the throughput for a PERC in each channel is shown in the right plot of Fig. 8. In this case, the throughput does increase with the number of channels, since every constraint is only applicable within one channel and adding channels provides for additional transmission opportunities.

Second, we show the performance of PO-CMA for $M = 2$ and a CIC as well as a PERC, respectively. We note that the performance of PO-CMA is close to the corresponding FO-CMA scheme. While the performance gap may be more pronounced for larger M, this result suggests that performance degrades gracefully with the number of bands. This is an encouraging result, given that for PO-CMA it suffices to have a simple transceiver instead of a more sophisticated radio frontend.

From Fig. 8, we can also infer that, for small WLAN traffic load $\bar{\sigma}$, the CIC is less restrictive than the PERC. On the other hand, for large $\bar{\sigma}$, this situation is reversed. This behavior is inherently linked to the definitions of both constraints. The PERC implicitly conditions on the WLAN being active and is thus more restrictive if the number of packets is small.
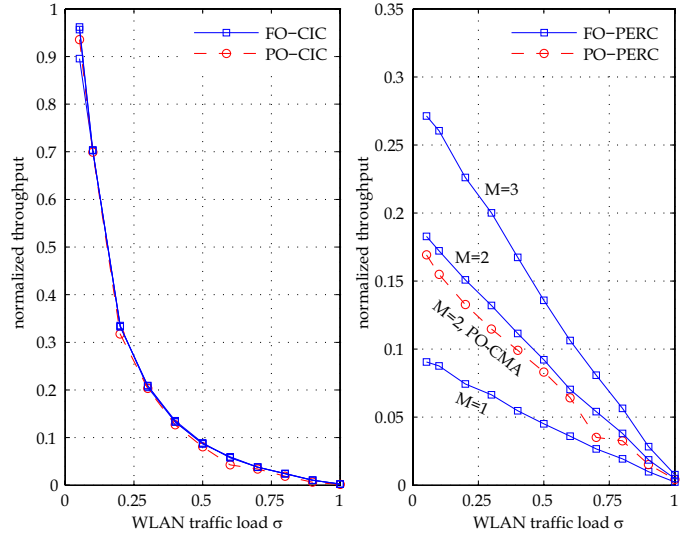

Fig. 8. Performance with CIC (left side) and PERC (right side) for increasing number of bands M. The normalized throughput represents the time-averaged fraction of successful transmissions out of the total number of slots.

C. Comparison with a blind hopping scheme

In Fig. 9 we compare the throughput *and* interference of our CMA schemes with a blind reference scheme that performs oblivious hopping throughout the entire band. We evaluate the performance for $M = 3$ since the ISM band at $2.4\,\text{GHz}$ supports exactly three parallel WLAN bands [13]. The traffic load was chosen such that the blind reference transmits in every, every other, or every third slot. The hopping pattern was pseudo-random across the entire band.

We can see that while the blind hopper's throughput can be higher than CMA's, its interference may become prohibitively large. Even if the blind reference is transmitting only in every third slot, the packet error rate can be 20% or higher. For a fully loaded system the packet error rate can be larger than 60%. We stress that for interference this high, the WLAN's throughput will inevitably decrease violating our paradigm of a hierarchical scheme. With CMA, on the other hand, a small interference level is guaranteed and as a consequence we assume that the cognitive radio's impact on the WLAN will be small.

D. Robustness to CTMC approximation

Lastly we evaluate the robustness to the CTMC approximation by heuristically running the CMA schemes derived from the CTMC model on data generated via the semi-Markov model. The simulation parameters for both models are shown in Tab. I. The performance and interference for a PERC is shown in Fig. 10. We can see that the deviation between the semi-Markov model and its CTMC approximation is small. While, the throughput curves almost coincide, the interference constraint is slightly violated for high WLAN traffic load. However, the aberration is still small enough to justify the use of the CTMC model.

VI. CONCLUSIONS

In conclusion we proposed Cognitive Medium Access (CMA) schemes for sharing spectrum with a set of parallel
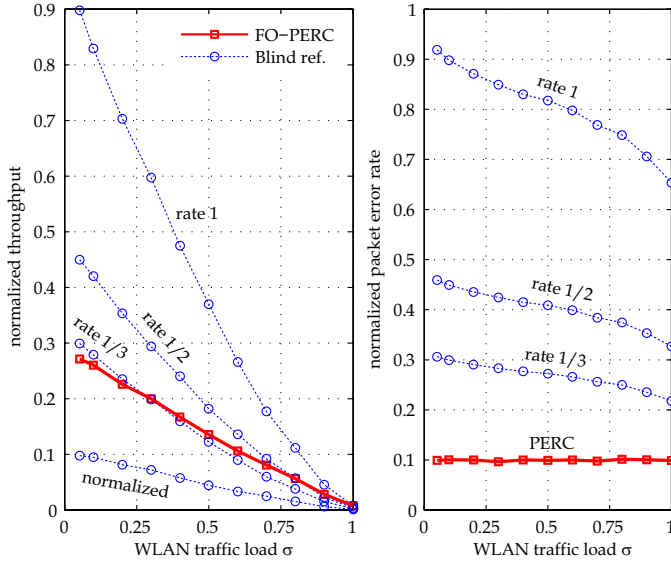
Fig. 9. Comparison with a blind reference scheme. The blind hopper operates with constant rate and is completely oblivious of the WLAN.

Fig. 10. Evaluating the robustness of the CTMC approximation.

WLAN bands. Its derivation was based on a measurement-based interference study as well as a stochastic model which captures the WLAN's medium access. The CMA schemes proposed in this paper can be classified according to fully observable and partially observable systems. In both cases optimal policies can be obtained via linear programming, but in the case of FO-CMA we furthermore specified structured solutions. The paper evaluated the performance of these schemes for various scenarios. We found that PO-CMA achieves almost the same performance as FO-CMA. Both schemes significantly outperform blind hopping schemes.

APPENDIX

Proof of Theorem 2

We need to show that Algorithm 1 is a solution to the LP (19)-(20). First, rewrite the LP in terms of its stationary distribution. This is possible since the transition behavior does not depend on the actions a. We obtain

$$\max_{w_{\mathbf{y}}(a)} \sum_{(\mathbf{y},a)\in\mathbb{X}\times\mathbb{A}} \eta_{\mathbf{y}} w_{\mathbf{y}}(a) r(\mathbf{y},a) \qquad (36)$$

subject to

$$\sum_{(\mathbf{y},a)\in\mathbb{X}\times\mathbb{A}} \eta_{\mathbf{y}} w_{\mathbf{y}}(a) d_{\mathrm{CIC}}(\mathbf{y},a). \qquad (37)$$

Note that for the CIC, $d_{\mathrm{CIC}}(\mathbf{y},a) = 1 - r(\mathbf{y},a)$. Furthermore, define $\mathbb{X}_1 = \{\mathbf{x} \in \mathbb{X} : x_i = 0\}$ and

$$\mathbb{X}_i = \{\mathbf{x} \in \mathbb{X} : x_i = 0, x_j = 1, \forall_j < i\}, \ 1 \le i \le M. \quad (38)$$

Then, states for which

$$\eta_{\mathbf{x}} w_{\mathbf{x}}(i) = \eta_{\mathbf{y}} w_{\mathbf{y}}(i), \quad \mathbf{x}, \mathbf{y} \in \mathbb{X}_i, \qquad (39)$$

clearly offer the same reward at the same cost. Due to this and our assumption that $\lambda_1 \le \cdots \le \lambda_M$, channel i offers higher (or equal) reward at lower (or equal) cost than channel j for $i < j$. Hence it is optimal to increase the transmission probability in channel 1 until $w_{\mathbf{y}}(1) = 1, \forall \mathbf{y} \in \mathbb{X}_1$. We
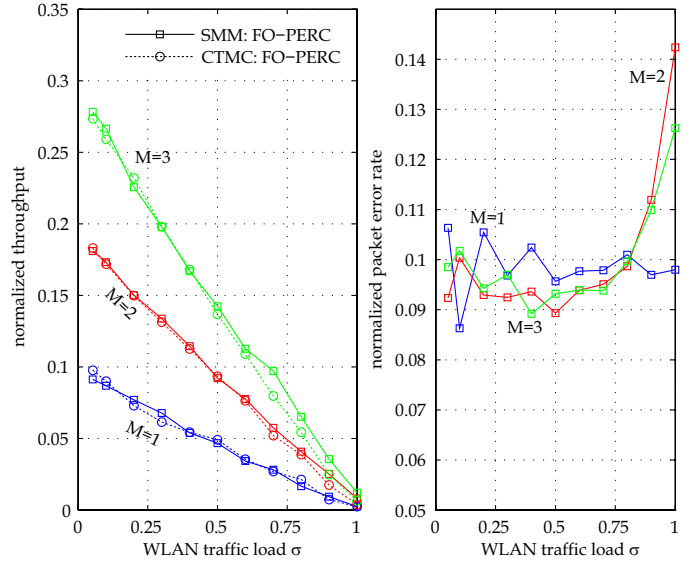
continue with channels $i = 2,\ldots,k$ until choosing $w_{\mathbf{y}}(k) = 1, \forall \mathbf{y} \in \mathbb{X}_k$ exceeds the constraint. We then choose $w_{\mathbf{y}}(k)$ such that the constraint is satisfied with equality. We thus have the optimal policy structure outlined in (26). See Fig. 6 for a graphical illustration. $\square$

Proof of Theorem 3

We need to show that Algorithm 2 is a solution to the LP (19) and (21). We again rewrite the LP in terms of the chain's stationary distribution

$$\max_{w_{\mathbf{y}}(a)} \sum_{(\mathbf{y},a)\in\mathbb{X}\times\mathbb{A}} \eta_{\mathbf{y}} w_{\mathbf{y}}(a) r(\mathbf{y},a) \qquad (40)$$

subject to

$$\sum_{(\mathbf{y},a)\in\mathbb{X}\times\mathbb{A}} \eta_{\mathbf{y}} w_{\mathbf{y}}(a) d_{\mathrm{PERC}}(\mathbf{y},a) \le \alpha_a, \ 1 \le a \le M. \quad (41)$$

Clearly, since $r(\mathbf{y},a) \ge 0$, an optimal solution is found if equality is achieved for all the constraints (41). In general, this need not be possible, however, since we can only transmit in one of the parallel channels at a time, even if multiple of them are idle (not possible to use all transmission opportunities). However, a sufficient condition for a structured solution is (27). In this case, $\sum_{i=1}^{M} w_{\mathbf{y}}(i) \le 1$, and the constraint can be made tight for channel i independently of all other channels. This leads to the transmission probabilities specified in (30). $\square$

REFERENCES

[1] R. Tandra and A. Sahai, "Fundamental limits on detection in low SNR under noise uncertainty," in *Proc. International Conference on Wireless Networks, Communications and Mobile Computing*, 2005.

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

[3] *Proceedings of the First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005.

[4] 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.

[5] 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, Sept. 2006.

[6] 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.

[7] 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. Select. Areas Commun.*, vol. 25, no. 3, pp. 589–600, Apr. 2007.

[8] S. Geirhofer, L. Tong, and B. M. Sadler, "Dynamic Spectrum Access in WLAN Channels: Empirical Model and Its Stochastic Analysis," in *Proc. First International Workshop on Technology and Policy for Accessing Spectrum*, 2006.

[9] Q. Zhao, S. Geirhofer, L. Tong, and B. M. Sadler, "Optimal Dynamic Spectrum Access via Periodic Channel Sensing," Proc. IEEE Wirless Communications and Networking Conference (WCNC), Mar. 2007.

[10] ——, "Opportunistic Spectrum Access via Periodic Channel Sensing," *to appear in IEEE Trans. Sig. Proc.*, 2007.

[11] 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.

[12] IEEE Computer Society, "IEEE Std 802.15.2-2003, Coexistence of Wireless Personal Area Networks with Other Wireless Devices Operating in Unlicensed Frequency Bands," Aug. 2003.

[13] ANSI/IEEE Standard 802.11, 1999 Edition (R2003), "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," IEEE/SA Standards Board, Tech. Rep., 1999.

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

[15] S. Geirhofer, L. Tong, and B. M. Sadler, "Measurement-Based Models for Cognitive Medium Access in WLAN Bands," Cornell University, Adaptive Communications and Signal Processing Group (ACSP), Technical Report ACSP-TR-02-07-02, Feb. 2007. [Online]. Available: http://acsp.ece.cornell.edu/papers/ACSP-TR-02-07-02.pdf

[16] Maxim MAX2810 Datasheet, http://datasheets.maxim-ic.com/en/ds/MAX2820-MAX2821A.pdf.

[17] N. Golmie, R. E. Van Dyck, A. Soltanian, A. Tonnerre, and O. Rébala, "Interference Evaluation of Bluetooth and IEEE 802.11b Systems," *Wireless Networks*, vol. 9, no. 3, pp. 201–211, May 2003.

[18] A. Kamerman, "Coexistence between Bluetooth and IEEE 802.11 CCK. Solutions to avoid mutual interference." Lucent Technologies, IEEE P802.11 Working Group Contribution IEEE P802.11-00/162, July 2000.

[19] ANSI/IEEE Standard 802.11b-1999 (R2003), "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Higher-Speed Physical Layer Extension in the 2.4GHz band," IEEE SA Standards Board, Tech. Rep., 1999.

[20] J. Zyren, "Extension of Bluetooth and 802.11 Direct Sequence Interference Model," Harris Semiconductor, Study Group Contribution IEEE 802.11-98/378, Nov. 1998.

[21] S. Geirhofer, L. Tong, and B. M. Sadler, "A Measurement-Based Model for Dynamic Spectrum Access," in *Proc. IEEE Conference on Military Communications (MILCOM)*, 2006.

[22] E. Altman, *Constrained Markov Decision Processes*. Chapman & Hall/CRC, 1999.

[23] S. I. Resnick, *Adventures in Stochastic Processes*. Birkhäuser, 1992.

[24] M. L. Puterman, *Markov Decision Processes. Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.

[25] C.-F. Chiasserini and R. R. Rao, "Coexistence Mechanisms for Interference Mitigation in the 2.4-GHz ISM Band," *IEEE Trans. Wireless Commun.*, vol. 2, no. 5, pp. 964–975, Sept. 2003.

[26] 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.

Stefan Geirhofer (S05) received the Dipl.-Ing. degree in Electrical Engineering from the Vienna University of Technology, Austria, in 2005. Since then he has been working toward the Ph.D. in the School of Electrical and Computer Engineering at Cornell University, Ithaca, New York. He has been a member of the Adaptive Communications and Signal Processing Group (ACSP) since May 2005. His research interests focus on signal processing and rapid prototyping in wireless communications, including cognitive radio, dynamic spectrum access, and MIMO systems.

Lang Tong (S'87, M'91, SM'01, F'05) is the Irwin and Joan Jacobs Professor in Engineering at Cornell University Ithaca, New York. Prior to joining Cornell University, he was on faculty at the West Virginia University and the University of Connecticut. He was also the 2001 Cor Wit Visiting Professor at the Delft University of Technology. Lang Tong received the B.E. degree from Tsinghua University, Beijing, China, in 1985, and M.S. and Ph.D. degrees in electrical engineering in 1987 and 1991, respectively, from the University of Notre Dame, Notre Dame, Indiana. He was a Postdoctoral Research Affiliate at the Information Systems Laboratory, Stanford University in 1991. Lang Tong is a Fellow of IEEE. He received the 1993 Outstanding Young Author Award from the IEEE Circuits and Systems Society, the 2004 best paper award (with Min Dong) from IEEE Signal Processing Society, and the 2004 Leonard G. Abraham Prize Paper Award from the IEEE Communications Society (with Parvathinathan Venkitasubramaniam and Srihari Adireddy). He is also a coauthor of five student paper awards. He received Young Investigator Award from the Office of Naval Research. Lang Tong's research is in the general area of statistical signal processing, wireless communications and networking, and information theory. He has served as an Associate Editor for the IEEE Transactions on Signal Processing, the IEEE Transactions on Information Theory, and IEEE Signal Processing Letters.

Brian M. Sadler (Fellow, IEEE) received the B.S. and M.S. degrees from the University of Maryland, College Park, and the PhD degree from the University of Virginia, Charlottesville, all in electrical engineering. He is a senior research scientist at the Army Research Laboratory (ARL) in Adelphi, MD. He was a lecturer at the University of Maryland, and has been lecturing at Johns Hopkins University since 1994 on statistical signal processing and communications. Dr. Sadler is currently an associate editor for the IEEE Signal Processing Letters and the IEEE Transactions on Signal Processing, and has been a guest editor for several journals including IEEE JSTSP, IEEE JSAC, and the IEEE SP Magazine. He is a member of the IEEE Signal Processing Society Sensor Array and Multi-channel Technical Committee, and received a Best Paper Award (with R. Kozick) from the Signal Processing Society in 2006. His research interests include signal processing for mobile wireless and ultra-wideband systems, sensor signal processing and networking, and associated security issues.