Exploiting Microscopic Spectrum Opportunities in Cognitive Radio Networks via Coordinated Channel Access

Tao Shu and Marwan Krunz
Department of Electrical and Computer Engineering
University of Arizona

Abstract—Under the current opportunistic spectrum access (OSA) paradigm, a common belief is that a cognitive radio (CR) can use a channel only when this channel is not being used by any neighboring primary radio (PR). Therefore, the existence of a spectrum opportunity hinges on the absence of active co-channel PRs in a *macroscopic* region. In this paper, we propose the concept of *microscopic* spectrum opportunity and show that CRs can still utilize this type of opportunities without interfering with active co-channel PRs, even when these PRs are close to them. As a result, a channel may at the same time present different levels of availability to different CRs. Channel access needs to be carefully coordinated between these CRs to avoid collisions, and more importantly, ensure efficient utilization of the spectrum opportunity from a network's standpoint. In this paper, we formulate the coordinated channel access as a joint power/rate control and channel assignment optimization problem, with the objective of maximizing the sum-rate achieved by the cognitive radio network (CRN). We develop both centralized and distributed algorithms to solve this problem. Our simulation results show that even when accounting for the implementation overhead, significant throughput gain is achieved under our designs.

Index Terms—cognitive radio networks, opportunistic spectrum access, power/rate control, spectrum leasing.

1 Introduction

1.1 Motivation

Cognitive radios (CRs) are an enabling technology for opportunistic spectrum access (OSA). These radios rely on channel sensing to identify idle frequency bands (channels) and dynamically hop between them to avoid interfering with licensed users (a.k.a., primary radios (PRs)). Because channel sensing is based on detecting the activity of a PR transmitter whereas interference takes place at the PR receiver, hidden-terminal problems can occur during the sensing process. Specifically, even if the CR does not detect any nearby PR activity, it is still possible for a PR receiver in the CR's neighborhood to be receiving signals from a PR transmitter that is outside the CR's sensing range. To alleviate this problem, a common way is to increase the CR's detection sensitivity, leading to an enlarged detection range. Because the CR can now access the channel only when all PR transmitters within this relatively large range are silent, it may miss some available spectrum opportunities.

To illustrate this situation, consider the example in Figure 1 (a), which depicts two PR transmitters, a and b, and one CR transmitter (CR1). The transmission range of a PR transmitter is denoted by $R_{tx}^{(PR)}$. This is the maximum distance for the transmitted signal to be correctly received by a PR receiver. We denote the interference range of CR1 by $R_I^{(CR)}$. This is the maximum range that CR1 can cause interference to a PR receiver

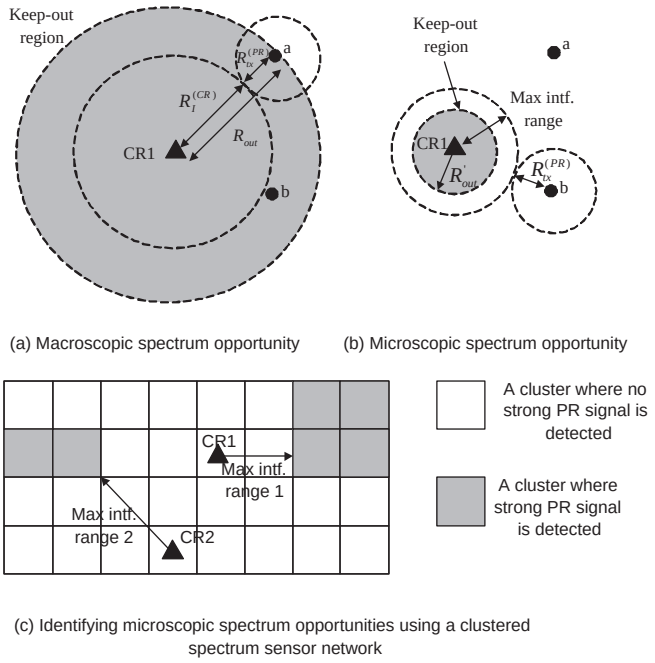

Fig. 1. Macroscopic vs. microscopic spectrum opportunities.

A preliminary version of this paper was presented at the IEEE INFOCOM 2009 Mini-conference, Rio de Janeiro, Brazil, Apr. 2009. Part of this work was conducted while M. Krunz was a visiting researcher at the University of Carlos III, Madrid, and IMDEA Networks, Spain. This research was supported in part by NSF (under grants CNS-0721935, CNS-0904681, IIP-0832238), Raytheon, and the Connection One center. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

when CR1 is transmitting at full power $P_{\max}$. Under this setup, it is easy to show that, to avoid colliding with a potential PR reception, CR1 should be able to detect the activity of every PR transmitter in the range $R_{tx}^{(PR)} + R_I^{(CR)}$. In other words, the energy detection threshold at CR1 should be sensitive enough to maintain a "keep-out" distance $R_{out} = R_{tx}^{(PR)} + R_I^{(CR)}$ from any active PR transmitter[1]. If a PR transmitter within this range, e.g., node b, is active, then CR1 will not use the channel. The spectrum opportunity identified under this setting implies the absence of active PR transmitters in a relatively large keep-out region. For this reason, we refer to it as *macroscopic* spectrum opportunity. This is the common type of spectrum opportunities studied in the literature (see [24] for a good survey).

From Figure 1 (a), we observe that even when node b is active, CR1 can still transmit over the same channel without causing interference to any receiver of b, provided that CR1's transmission power is controlled such that its effective interference range is smaller than $d_{1b} - R_{tx}^{(PR)}$, where d_{1b} is the distance between CR1 and node b (see Figure 1 (b)). As long as the distance between CR1 and node b is greater than $R_{tx}^{(PR)}$, CR1 can always use the channel with an appropriately selected transmission power. Thus, the keep-out distance for CR1 becomes $R'_{out} = R_{tx}^{(PR)} < R_{out}$, implying more chances for CR1 to use the channel (without interfering with PRs). Because the availability of such a spectrum opportunity requires a smaller keep-out region, we refer to it as *microscopic*. We also refer to the maximum allowable transmission power a CR can use over a channel without interfering with co-channel PRs as the *power mask* of the CR over that channel.

In practice, the identification of microscopic spectrum opportunities can be realized using cluster-based collaborative spectrum sensing techniques (e.g., see [12][7][4]), as illustrated in Figure 1 (c). Specifically, spectrum sensors are grouped into clusters. Each cluster covers a geographic grid. Sensors in the same cluster collaborate to detect whether the received PR signal, if any, is strong enough for valid reception by a PR receiver. This detection capability essentially represents the reception sensitivity of a PR receiver, and thus is lower than that used in the macroscopic case. The grids in which a strong PR signal is detected are marked, indicating the potential presence of active PR receivers. The maximum allowable interference range for each CR is simply the distance from the CR to its nearest marked grid. The power masks can be computed accordingly. In contrast to the binary-type (0 or $P_{\max}$) power mask used with macroscopic spectrum opportunities, the microscopic case has an intrinsic multi-level structure, ranging between 0 and $P_{\max}$. This is because the CR's power mask changes with the location of the closest marked grid, as shown in Figure 1 (c). It is also easy to see that the macroscopic spectrum opportunity is a special case of the microscopic opportunity (when all PR transmitters within distance $R_{tx}^{(PR)} + R_I^{(CR)}$ are silent). Because the multi-level structure is unique to microscopic spectrum opportunities, we will use these two terms interchangeably in the subsequent discussion. Similarly, we will also use the terms, "binary" and "macroscopic" spectrum opportunities interchangeably.

In this paper, we are interested in studying the CR network (CRN) throughput achieved under the microscopic spectrum opportunity setting. The main challenge here stems from the fact that the same channel may simultaneously present different levels of availability to different CRs. Therefore, channel access needs to be carefully coordinated between these CRs to avoid collisions, and more importantly, ensure efficient utilization of the spectrum opportunity from a network-wide standpoint.

We study the coordinated channel access problem by formulating it as a joint power/rate control and channel assignment optimization problem. Given the available channels at different CRs, we need to specify for each CR which channels to use and at what powers and rates. We are interested in both centralized (for better performance) and distributed (for better implementability) solutions. In contrast to previous works that aim at maximizing the information-theoretic capacity of the system, our objective is to maximize the sum-rate achieved by each CR. Unlike the information-theoretic (Shannon) capacity, the achievable rate in our setup depends on the PHY-layer implementation. In other words, our problem has a wider scope and can be applied to any arbitrarily given rate-SINR function.

In many existing power/rate control and channel assignment problems, a convex formulation is obtained through the capacity approximation $\log(1 + \text{SINR}) \approx \text{SINR}$ if $\text{SINR} \ll 1$ (low-SINR regime, e.g., see [16]) or $\log(1 + \text{SINR}) \approx \log(\text{SINR})$ if $\text{SINR} \gg 1$ (high-SINR regime, e.g., see [20]). However, due to the multi-level structure of its power mask, a CR is expected to operate over a wide range of SINR values. Even at a given time instance, different CRs may be operating in different SINR regimes. Therefore, approximation techniques that are adopted separately for low- and high-SINR regimes are no longer appropriate here. A new SINR-independent treatment is needed for the problem. We attempt to accommodate this more general form in our optimization.

1.2 Contributions and Paper Organization

The contributions of this work are as follows. We first show that the joint power/rate control and channel assignment problem can be formulated as a mixed integer nonlinear programming (MINLP) problem that is NP-hard. By exploiting the discrete nature of a CR's multi-rate capability, we transform this MINLP into a

1. Note that the keep-out region in this work is defined from the CR transmitter's standpoint (i.e., a circle centered at the CR transmitter). This is in contrast with the conventional PR-transmitter-based definition, in which the keep-out region is centered at the PR transmitter.

binary linear program (BLP) that contains only binary variables and linear objective function and constraints. This transformation applies to any arbitrary rate-SINR relationship. We then develop two polynomial-time approximate algorithms for the BLP. The first one is the centralized *LPSF* algorithm, and is based on iteratively solving a series of linear programming problems and sequentially fixing the variables to either 1 or 0 in each iteration. The second is the distributed *EF-based* algorithm. It involves iterative and on-line adjustment of the powers and rates of each CR over each channel based on some *economic factor* that accounts for the efficiency of expending power over a given channel. We show that this distributed algorithm is provably efficient, i.e., it can achieve a provable fraction of the optimal performance. As a byproduct, the centralized algorithm gives an upper bound on the optimal solution. This allows us to explicitly evaluate the performance gap between the approximate solutions and the optimal one. Simulation results are used to verify the accuracy of the LPSF and EF algorithms. They indicate that the observed performance gap between the approximation and the exact (optimal) solution is always less than 10%.

To evaluate the benefit brought by the microscopic spectrum opportunity, we subsequently apply our algorithms to a spectrum-leasing system, whereby a CRN shares the spectrum with an infrastructure PR network (PRN). We illustrate how the multi-level spectrum opportunity can be calculated under the assistance of a broadcast-subscription mechanism. The interesting question is how much gain the multi-level spectrum opportunity scheme can attain over the conventional binary scheme, with the protocol overhead being accounted for. We evaluate this gain for various levels of overhead. Our simulation results show that significant gain (e.g., over 100% at best) can be achieved by the multi-level scheme.

The rest of this paper is organized as follows. We review the related work in Section 2. We describe the models and formulate the optimization problem in Section 3. The transformation to a BLP, the LPSF, and the EF algorithms are presented in Section 4. We apply our algorithms to the infrastructured PRN/CRN system in Section 5. Simulation evaluation and discussion are provided in Section 6, and we conclude the work in Section 7.

2 Related Work

Prior work was mostly focused on exploiting macroscopic spectrum opportunities. Early works provide collision-free channel assignment for CR nodes given a set of available channels at each node. This problem can be described as an interference-graph vertex-coloring problem [18], [27]. To obtain a fast solution, various distributed approximations were proposed, based on observing local interference patterns [26], local bargaining [1], or coordinations between CR nodes that aim at maximizing some system utility [2][22]. Because of the graph-theoretic nature of these algorithms, they take transmission power as input, and thus are not applicable to power/rate control problems.

Another body of work considers the optimal sensing/channel access decision-making process from a single CR's viewpoint. This is also termed as *MAC-layer sensing*. Existing works include the POMDP model [25], the constrained Markov decision processes (CMDPs) model [23], and the optimal stopping-rule models [3] [10]. Assuming a semi-Markov process for the PR traffic, Kim and Shin [11] proposed a sensing-period adaptation algorithm that maximizes the discovery of spectrum opportunities and minimizes the delay in finding an available channel. Based on a similar PR traffic model, the authors in [9] studied a dynamic access scheme subject to a constraint on the CR-to-PR violation rate, but only for a system of one PRN and one CR link. In contrast to these works, ours aim at optimizing the spectrum utilization for the entire CRN, rather than for a single CR.

The third type of works simplifies the problem by restricting the treatment to CR nodes only. So the CR-to-PR and PR-to-CR interferences are not accounted for. As a result, the power mask for every CR on each channel is implicitly assumed to be $P_{\max}$. Within this category, Hou et al. [8] considered the joint optimization of spectrum, scheduling, and routing in a multi-hop software-defined-radio (SDR) network. Yi and Hou [14][15] studied the joint optimization of power control, scheduling, and routing for a multi-hop SDR network, assuming a logarithmic rate-SINR relationship. Yuan et al. [21] introduced the concept of time-spectrum blocks to study spectrum allocation in CRNs. Based on a continuous-time Markov model, Xing et al. [19] proposed a random access protocol that achieves airtime fairness among CRs. The work in [20] considers spectrum access for CRs under an interference temperature constraint. However, because this constraint is defined only at a single location, compliance to it does not necessarily prevent interference to PR nodes.

3 System Model and Problem Formulation

We consider a distributed (ad hoc) CRN that coexists with M legacy (fixed spectrum) PRNs over a finite area. PRN m, $m = 1, \ldots, M$, is licensed to operate over its own frequency channel of bandwidth B_m. In reality, a PRN may occupy several channels. Such a network can be easily captured in our model by using multiple (virtual) PRNs that operate over different channels.

Let the number of CR links be N. For CR link i, we denote its sender and receiver by $S(i)$ and $D(i)$, respectively. A CR link can simultaneously transmit over multiple non-contiguous channels. Let the transmission power on channel m be $P_i^{(m)}$. To avoid unacceptable CR-to-PR interference, this transmission power must be constrained below a certain power mask $\hat{P}_i^{(m)}$. The value

of $\hat{P}_i^{(m)}$ is related to the status of neighboring PRs and thus changes over time. For now, we assume that $\hat{P}_i^{(m)}$, $i = 1, \ldots, N$ and $m = 1, \ldots, M$, is given in each snapshot as an input parameter to the joint power/rate control and channel assignment problem. We consider the calculation of $\hat{P}_i^{(m)}$ in Section 5.

We say that CR links i and j are interfering links on channel m if $\hat{P}_i^{(m)} h_{S(i)D(j)} > P_{I,CR}$ or $\hat{P}_j^{(m)} h_{S(j)D(i)} > P_{I,CR}$, where $h_{S(i)D(j)}$ and $h_{S(j)D(i)}$ are the cross-link channel gains of the two links and $P_{I,CR}$ is the sensitivity of the CR receiver (fixed). Any received power below $P_{I,CR}$ is deemed ignorable. We assume that an *exclusive channel occupancy* policy is used to resolve collisions between CRs: For any two *interfering* CR links on channel m, only one of them can access the channel at any given time.

Treating interference as noise, the rate of CR link i on channel m is given by

$$R_i^{(m)} = B_m f\left(\frac{P_i^{(m)} h_i^{(m)}}{q_{D(i)}^{(m)} + N_0}\right) \tag{1}$$

where f is any arbitrary rate-SINR function, decided by the PHY layer, $h_i^{(m)}$ is the channel gain of link i on channel m, $q_{D(i)}^{(m)}$ is the received interference over channel m at $D(i)$, and N_0 is the AWGN. Because of the exclusive channel occupancy policy, the interference $q_{D(i)}^{(m)}$ comes only from active co-channel PRs. It can be measured by the CR receiver $D(i)$.

For $i = 1, \ldots, N$ and $m = 1, \ldots, M$, let

$$x_i^{(m)} \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if channel } m \text{ is used by CR link } i, \\ & \text{i.e., } R_i^{(m)} > 0 \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

Our objective is to maximize the sum-rate of all CR links over all channels, i.e.,

$$\text{maximize} \sum_{i=1}^{N} \sum_{m=1}^{M} x_i^{(m)} R_i^{(m)} \tag{3}$$

where the maximization is to be carried out with respect to $x_i^{(m)}$'s and $R_i^{(m)}$'s.

At the same time, CR link i should satisfy the following constraints:

C1: CR-to-PR constraint: The transmission power of link i on channel m should not exceed $\hat{P}_i^{(m)}$. From (1), this constraint can be written in terms of $R_i^{(m)}$ as

$$\frac{1}{h_i^{(m)}}(q_{D(i)}^{(m)} + N_0) f^{-1}(r_i^{(m)}) \le \hat{P}_i^{(m)}, \quad m = 1, \ldots, M \tag{4}$$

where f^{-1} is the inverse of f and $r_i^{(m)} = \frac{R_i^{(m)}}{B_m}$ is the spectrum efficiency of link i on channel m.

C2: Power supply constraint: The sum of the transmission powers over all channels should not exceed the maximum power provided by the battery, i.e.,

$$\sum_{m=1}^{M} \frac{1}{h_i^{(m)}}(q_{D(i)}^{(m)} + N_0) f^{-1}(r_i^{(m)}) \le P_{\max,i}. \tag{5}$$

C3: CR-to-CR collision constraint: If channel m is being used by CR link i, then it cannot be used by another CR link that interferes with link i on channel m, and vice versa:

$$x_i^{(m)} + x_j^{(m)} \le 1, \quad \forall j \in I_i^{(m)} \tag{6}$$

where $I_i^{(m)} = \left\{j : j \ne i, \hat{P}_i^{(m)} h_{S(i)D(j)}^{(m)} > P_{I,CR}\right\} \cup \left\{j : j \ne i, \hat{P}_j^{(m)} h_{S(j)D(i)}^{(m)} > P_{I,CR}\right\}$ is the set of CR links that interfere with link i on channel m.

4 SOLUTIONS

4.1 Transformation to BLP

An observation of (3) and the constraints C1-C3 shows that this formulation constitutes a MINLP problem. In general, the solution to such a problem is NP-hard. To make this formulation more amenable to further processing, we exploit the fact that practical communication systems support only a finite set of transmission rates. Denote this set by $\mathbf{U} = \{0, u_1, u_2, \ldots, u_K\}$ (in bits/sec/Hz), where $0 < u_1 < \ldots < u_K$. Let $\gamma_k \stackrel{\text{def}}{=} f^{-1}(u_k)$ for $k = 1, \ldots, K$; γ_k is the received symbol-energy-to-interference-plus-noise density ratio (E_S/I_0) required to support the kth rate under the power-rate relationship defined by (1). Let $C_i^{(m)} \stackrel{\text{def}}{=} \frac{1}{h_i^{(m)}}\left(q_{D(i)}^{(m)} + N_0\right)$ for $i = 1, \ldots, N$ and $m = 1, \ldots, M$. $C_i^{(m)}$ is a known quantity for each CR link on each channel. We further define the new variable $y_{k,i}^{(m)}$ for all $k = 1, \ldots, K$, $i = 1, \ldots, N$, and $m = 1, \ldots, M$:

$$y_{k,i}^{(m)} \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if link } i \text{ is transmitting on channel } m \text{ using rate } u_k \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

In addition, we impose the following constraint on $y_{k,i}^{(m)}$:

$$\sum_{k=1}^{K} y_{k,i}^{(m)} \le 1 \tag{8}$$

which says that a link can use at most one rate on a given channel at a given time. It is easy to show that:

$$x_i^{(m)} = \sum_{k=1}^{K} y_{k,i}^{(m)}. \tag{9}$$

Similarly, we can rewrite the spectrum efficiency $r_i^{(m)}$ in terms of $y_{k,i}^{(m)}$ and u_k:

$$r_i^{(m)} = \sum_{k=1}^{K} u_k y_{k,i}^{(m)}. \tag{10}$$

Substituting (9) and (10) into (3) through (6), we get the following equivalent formulation to the original MINLP

problem:

$$\begin{array}{lll} & \text{maximize} & \sum_{i=1}^{N}\sum_{m=1}^{M}\sum_{k=1}^{K} B_m u_k y_{k,i}^{(m)} \\ & \text{such that} & \\ \tilde{C}1: & & C_i^{(m)}\sum_{k=1}^{K}\gamma_k y_{k,i}^{(m)} \le \hat{P}_i^{(m)} \\ \tilde{C}2: & & \sum_{m=1}^{m} C_i^{(m)}\sum_{k=1}^{K}\gamma_k y_{k,i}^{(m)} \le P_{\max,i} \\ \tilde{C}3: & & \sum_{k=1}^{K} y_{k,i}^{(m)} + \sum_{k=1}^{K} y_{k,j}^{(m)} \le 1, \quad \forall j \in I_i^{(m)} \end{array} \tag{11}$$

where the maximization is w.r.t. the $y_{k,i}^{(m)}$'s.

An examination of (11) shows that the former MINLP problem has been transformed into a binary linear program (BLP) that contains only binary variables and linear objective function and constraints. A nice property of (11) is that the rate levels u_k, $k = 1, \ldots, K$, and the corresponding γ_k's are inputted into the BLP formulation as tuples (u_k, γ_k). In other words, the BLP formulation does not rely on the specific functional relationship between u_k and γ_k, and thus can accommodate any arbitrary rate-power relation (e.g., a staircase-like function that characterizes practical multi-rate systems).

4.2 LPSF Centralized Algorithm

A BLP is a combinatorial problem. Its solution is, in general, NP-hard. A typical approximation approach to such a problem is provided by the so-called *branch-and-bound* algorithm, whose worst-case time complexity is exponential.

Instead of employing a branch-and-bound algorithm, we exploit the special structure of the problem to develop polynomial-time approximate algorithms. An observation of (11) indicates that if we relax $y_{k,i}^{(m)}$'s from their binary values and allow them to take real values between 0 and 1, then the formulation becomes a linear program (LP) that is solvable in polynomial time. In addition, the constraint $\tilde{C}3$ dictates that if for some m, k, and i, $y_{k,i}^{(m)} = 1$, then $y_{h,i}^{(m)} = 0$ for all $h \neq k$ and $y_{l,j}^{(m)} = 0$ for all $j \in I_i^{(m)}$ and $1 \le l \le K$. In other words, a strong dependence exists between the $y_{k,i}^{(m)}$'s that belong to the same set of interfering links. The main idea behind our approximate solution is to sequentially fix the values of $y_{k,i}^{(m)}$'s by solving a series of relaxed LP problems, where in each iteration the binary value of at least one $y_{k,i}^{(m)}$ is finalized.

An overview of our approximation algorithm, called LP with sequential fixing (LPSF), is given in Table 1. In the first iteration, we append the constraint $0 \le y_{k,i}^{(m)} \le 1$ to (11) and relax all $y_{k,i}^{(m)}$'s to real values between 0 and 1. We refer to the resulting formulation as LP$^{(1)}$, which must have a feasible solution according to Lemma 1 (introduced later). The solution to LP$^{(1)}$ is an upper bound on the optimal solution to (11), because the feasibility region of the BLP is a subset of that of LP$^{(1)}$. However, the solution of LP$^{(1)}$ is, in general, not a feasible solution to the original BLP problem, because the $y_{k,i}^{(m)}$'s can now take any real values between 0 and 1. Among all $y_{k,i}^{(m)}$'s, we pick the largest one, and we denote this $y_{k,i}^{(m)}$ by $Y_{k,i}^{(m)}$ for ease of identification. We set $Y_{k,i}^{(m)} = 1$. Accordingly, all $y_{h,i}^{(m)}$'s for $h \neq k$ and all $y_{l,j}^{(m)}$'s for $j \in I_i^{(m)}$ and $1 \le l \le K$ must now be set to 0. Replacing these $y_{k,i}^{(m)}$'s by their fixed values in LP$^{(1)}$, we get a new LP, called LP$^{(2)}$, whose variables do not include those that have been fixed after the execution of LP$^{(1)}$. A feasibility check is then conducted on LP$^{(2)}$. An empty feasible region for LP$^{(2)}$ means the first fixing in this iteration, i.e., $Y_{k,i}^{(m)} = 1$, is not correct. So we reset $Y_{k,i}^{(m)}$ to 0. This change means all those variables that belong to the same interfering CR link set as $Y_{k,i}^{(m)}$ and whose values have been fixed to 0 in this iteration must become variables again. The revised fix, i.e. $Y_{k,i}^{(m)} = 0$, is then substituted into LP$^{(1)}$, giving rise to LP$^{(3)}$. LP$^{(3)}$ must be feasible (see Lemma 2). In a nutshell, at this point we either have a feasible LP$^{(2)}$ or a feasible LP$^{(3)}$. Whichever formulation is feasible is renamed as LP$^{(1)}$, and a new iteration starts following the same process above. The process is repeated until all $y_{k,i}^{(m)}$'s are set to either 0 or 1. The final rate allocation for each link on each channel is calculated according to (10).

STEP 0: Formulate LP$^{(1)}$ by appending $0 \le y_{k,i}^{(m)} \le 1$ to (11) and relaxing all binary variables to real values.
STEP 1: Solve LP$^{(1)}$.
STEP 2: Set $Y_{k,i}^{(m)} \leftarrow \max\left\{y_{l,j}^{n}, l \in (1,\ldots,K), j \in (1,\ldots,N), n \in (1,\ldots,M)\right\}$.
STEP 3: Formulate LP$^{(2)}$ by substituting $Y_{k,i}^{(m)} = 1$, $y_{h,i}^{(m)} = 0$ for $h \neq k$ and $y_{l,j}^{(m)} = 0$ $\forall j \in I_i$ and $1 \le l \le K$ into LP$^{(1)}$.
STEP 4: If LP$^{(2)}$ is feasible
LP$^{(1)}$ $\leftarrow$ LP$^{(2)}$
else
Formulate LP$^{(3)}$ by substituting $Y_{k,i}^{(m)} = 0$ into LP$^{(1)}$.
LP$^{(1)}$ $\leftarrow$ LP$^{(3)}$
End-if
STEP 5: If all variables are fixed, then **Terminate**; otherwise go to **STEP 1**.

TABLE 1
Overview of the LPSF algorithm.

Note that a similar algorithm was suggested in [8] to solve a different problem. From a methodology standpoint, the main difference between our algorithm and the one in [8] is that in [8] there is no guarantee that a feasible solution can be found at the termination of the algorithm. Our algorithm improves upon [8] by adding a revised fixing component when any intermediate fixing leads to infeasibility, so that a feasible solution can always be found.

Theorem 1: The LPSF algorithm can correctly determine the binary values of all $y_{k,i}^{(m)}$'s in no more than NMK iterations.

The proof of Theorem 1 is based on the following lemmas.

Lemma 1: In the first iteration, LP$^{(1)}$ has a nonempty fea-

sible solution set, and because its variables are bounded, it also has an optimal solution.

Proof: It is easy to show that $y_{ki}^{(m)} = 0$ for all $k = 1, \ldots, K$, $i = 1, \ldots, N$, and $m = 1, \ldots, M$, is one feasible solution to the original BLP. Thus, it is also a feasible solution to LP$^{(1)}$. Note that all variables are bounded between 0 and 1; therefore, Lemma 1 holds. ∎

Lemma 2: In the first iteration, LP$^{(3)}$ has a nonempty feasible solution set, and because its variables are bounded, it also has an optimal solution.

Proof: According to Lemma 1, LP$^{(1)}$ in the first iteration must have an optimal solution. Therefore, $Y_{ki}^{(m)} \geq 0$ before the fix. When $Y_{ki}^{(m)}$ is fixed to 0 to get LP$^{(3)}$, its value is changed from no less than 0 to 0, leading to a non-increase in the required transmission power. So none of the RHS of C1′ through C3′ could be violated by this non-increasing action on the LHS of C1′ through C3′, respectively. Therefore, LP$^{(3)}$ must have at least one feasible solution. Noting that all variables are bounded, Lemma 2 holds. ∎

Lemma 3: In all iterations, LP$^{(1)}$ and LP$^{(3)}$ have nonempty feasible solution sets, and because their variables are bounded, they also have optimal solutions.

Proof: The situation in the first iteration is proved by Lemma 1 and Lemma 2. In the second iteration, LP$^{(1)}$ is obtained either from a feasible LP$^{(2)}$ or a feasible LP$^{(3)}$ of the first iteration. So LP$^{(1)}$ must be feasible in the second iteration. Given that LP$^{(1)}$ is feasible in the second iteration, the rationale used in proving Lemma 2 also applies here to prove the feasibility of LP$^{(3)}$ in the second iteration. This induction can be repeated in all iterations. Noting that all variables are bounded, Lemma 3 holds. ∎

The proof of Theorem 1 is straightforward: Iteratively applying Lemmas 1 to 3, it is guaranteed that in each iteration at least one $y_{ki}^{(m)}$ is fixed to either 0 or 1 and a new feasible LP$^{(1)}$ is generated for the next iteration. For the last iteration, if fixing $y_{ki}^{(m)}$ to 1 does not lead to a feasible BLP solution, then changing its value to 0 must lead to a feasible BLP solution (for the same reason as in the proof of Lemma 2). ∎

Based on Theorem 1, it is easy to show that the time complexity of the LPSF algorithm is bounded by the complexity of the LP solver times NMK. An LP solver is of polynomial complexity, so the complexity of the LPSF is also polynomial. In addition, the performance gap between the approximate LPSF solution and the exact (optimal) one can be assessed by comparing the former with an upper bound on the optimal solution, given by the the solution to LP$^{(1)}$ in the first iteration. Lemma 1 guarantees the existence of this upper bound. We will later show by simulation that this gap is small (below 10%), and in most cases it is zero.

4.3 Distributed Algorithm

In this section, we develop an efficient distributed algorithm for the BLP problem in (11), which can achieve a provable fraction of the optimal performance. The intuition behind such an algorithm comes from understanding the conflicts between CRs when utilizing spectrum opportunities. There are two aspects to such conflicts. First, neighboring CRs may observe a similar level of spectrum availability over a given channel, and thus may attempt to transmit simultaneously over the same channel, causing collisions. Second, transmissions by the same CR over different channels may also conflict with each other, in the sense that the maximum transmission power provided by the battery may not be sufficient to support parallel transmissions over all these channels. In a nutshell, conflicts between transmissions occur due to their competition for both spectrum and power resources. A good design philosophy is to give priority to a transmission that can contribute a higher rate at a lower power. Following this philosophy, the proposed distributed algorithm defines an *economic factor* (EF) for each channel at each CR link. Let the current transmission rate of link i on channel m be $r_i^{(m)} = u_k$, $k \in \{0, \ldots, K-1\}$. Then, the EF of this link on channel m is defined as

$$\eta_i^{(m)} \stackrel{\text{def}}{=} \frac{\Delta P_i^{(m)}}{B_m \Delta r_i^{(m)}} = \frac{C_i^{(m)}(\gamma_{k+1} - \gamma_k)}{B_m(u_{k+1} - u_k)}. \tag{12}$$

We let $\eta_i^{(m)} \stackrel{\text{def}}{=} +\infty$ when $r_i^{(m)} = u_K$.

The basic idea of our EF-based distributed algorithm is to iteratively ramp up the rate level over each channel of every neighboring link until the power mask and maximum-power constraints are violated. In each iteration, the link-channel pair that has the smallest EF value among its interfering-link set is raised to the next higher rate. This is achieved by sequentially executing the following three procedures in each iteration (note that this algorithm is asynchronously executed by various CR transmitters). The first procedure is an *internal candidate selection* process, in which a link, say i, selects a channel m^* that has the smallest EF among all channels in a *candidate channel set* $\mathcal{C}$. The set $\mathcal{C}$ is initialized to contain all M channels. The selected channel m^* is tested for the feasibility of a rate increase. This is done by calculating the power increment $\Delta P_i^{(m^*)} = C_i^{(m^*)}(\gamma_{k+1} - \gamma_k)$. If this increment violates the power mask or the battery power constraint, then a rate increase on channel m^* is infeasible for that CR. So m^* is deleted from $\mathcal{C}$ and the above selection process is repeated. Eventually, either a feasible m^* is selected or $\mathcal{C}$ becomes empty. When $\mathcal{C}$ becomes empty, the iterative process at the CR transmitter terminates. If a feasible m^* is found, the algorithm enters the *inter-link selection* phase.

In the inter-link selection phase, neighboring CR transmitters exchange the results of their internal selection to elect the link-channel pair that has the smallest EF in the neighborhood. The internal selection result of a link i is broadcasted as a triple (link id i, channel id m^*, $\eta_i^{(m^*)}$). The link-channel pair that has the smallest EF in its neighborhood raises the corresponding transmission rate

/* CR link i */
Initialization: $r_i^{(m)} \leftarrow 0$, for $m = 1, \ldots, M$ and $\mathcal{C} \leftarrow \{1, \ldots, M\}$

while ($\mathcal{C} \neq \emptyset$)
  /* Internal candidate selection */
  $violation\_flag \leftarrow 1$
  while ($violation\_flag == 1$)
    $m^* \leftarrow \arg\min \left\{ \eta_i^{(m)} | m \in \mathcal{C} \right\}$
    calculate $\Delta P_i^{(m^*)}$
    if $((\Delta P_i^{(m^*)} + P_i^{(m^*)} \leq \hat{P}_i^{(m^*)})$
    or $(\Delta P_i^{(m^*)} + \sum_{m=1}^{M} P_i^{(m)} \leq P_{\max,i}))$
      $violation\_flag \leftarrow 0$
    else
      $\mathcal{C} \leftarrow \mathcal{C} - \{m^*\}$
    end-if
  end-while

  /* Inter-link selection */
  exchange with neighbors the message:
  (link id i, channel id m^*, $\eta_i^{(m^*)}$)
  if ($\eta_i^{(m^*)}$ is the minimum among neighbors)
    increase $r_i^{(m^*)}$ from u_k to u_{k+1}
    if ($r_i^{(m^*)} == u_K$)
      $\mathcal{C} \leftarrow \mathcal{C} - \{m^*\}$
    end-if
    send rate-adjustment message
  end-if

  /* Collision elimination routine */
  if (rate-adjustment message is received from link j)
    calculate $h_{S(j)D(i)}$ based on received signal strength
    if ($h_{S(j)D(i)} \hat{P}_j^{(m^*)} > P_{I,CR}$)
      if ($r_i^{(m^*)} \leq r_j^{(m^*)}$)
        $r_i^{(m^*)} \leftarrow 0$ and $\mathcal{C} \leftarrow \mathcal{C} - \{m\}$
      else
        $S(i)$ sends a rate-adjustment message
      end-if
    end-if
  end-if
end-while
Output: $r_i^{(m)}$, for $m = 1, \ldots, M$

TABLE 2
Pseudo-code for the EF-based distributed algorithm.

by one level, i.e., from u_k to u_{k+1}. At the same time, the sender of this link, say $S(j)$, broadcasts at power $P_{\max,j}$ the following *rate-adjustment* message to its neighbors: (link id j, channel id m^*, $r_j^{(m^*)}$, $\hat{P}_j^{(m^*)}$, $P_{\max,j}$).

Whenever a CR link i receives a rate-adjustment message from link j, it performs a *collision elimination* routine. Specifically, the receiver of the ith link, $D(i)$, calculates the path loss from $S(j)$ to $D(i)$ based on the received signal strength. From the power mask information in the message, $D(i)$ can then decide whether $S(j)$'s transmission will interfere with the reception at $D(i)$ on channel m. If so, $D(i)$ compares $r_i^{(m^*)}$ with $r_j^{(m^*)}$. If $r_i^{(m^*)} \leq r_j^{(m^*)}$, $D(i)$ notifies $S(i)$ to set $r_i^{(m^*)}$ to zero and deletes m^* from $\mathcal{C}$. If $r_i^{(m^*)} > r_j^{(m^*)}$, $D(i)$ notifies $S(i)$ to send a rate-adjustment message that triggers link j to avoid using channel m^*.

A pseudo-code of the algorithm is given in Table 2. Because at least one $r_i^{(m)}$ will be increased by one level in each iteration in each interfering-link set, the rate adjustment will terminate in at most MK iterations. Theorem 2 quantifies the efficiency of this algorithm.

Theorem 2: The EF-based distributed algorithm can achieve at least $1/(\kappa^* + 1)$ of the optimal throughput, where $\kappa^* = \max_{i,m} |I_i^{(m)}|$ is the maximum interference degree of all CR links over all channels and $|\cdot|$ denotes the cardinality of a set.

Proof: The rate adjustment in the EF algorithm is analogous to the well-known single-user optimal Levin-Campello greedy algorithm [17] for bit loading in OFDM systems. In allocating each bit, this greedy algorithm calculates the cost to add one more bit in each sub-channel and chooses the sub-channel that requires the least cost, where the cost is the necessary power increment. It has been shown in [13] that for multi-user multi-carrier systems, if we assume no interference exists between users, then the same greedy algorithm achieves the optimal performance. Denote the optimal sum-rate of this idealized non-interfering multi-user system by $R_{tot,\max}^{(0)}$. This sum is calculated as:

$$R_{tot,\max}^{(0)} = \sum_{i=1}^{N} \sum_{m=1}^{M} R_i^{(m)} \tag{13}$$

where $R_i^{(m)}$ are the output of the greedy algorithm when interference between users is ignored. When interference is accounted for, the third "if" statement in the collision-elimination routine of the EF-based algorithm (see Table 2) guarantees that for every interfering link set, only the link that achieves the largest rate is kept (i.e., only this link can access the channel). All other interfering links are prohibited from accessing the channel (their rates on this channel are set to 0). Denote by $Z^{(m)}$ the set of links that can access channel m when the EF algorithm is used. For all $z \in Z^{(m)}$, it must be true that $R_z^{(m)} \geq R_j^{(m)}$, $\forall j \in I_z^{(m)}$. When interference is accounted for, denote the sum-rate of the EF algorithm by $R_{tot,EF}^{(1)}$. Then $R_{tot,EF}^{(1)} = \sum_{m=1}^{M} \sum_{z \in Z^{(m)}} R_z^{(m)}$. We further have the following relationship:

$$\begin{aligned} R_{tot,\max}^{(0)} &= \sum_{m=1}^{M} \sum_{i=1}^{N} R_i^{(m)} \\ &\leq \sum_{m=1}^{M} \sum_{z \in Z^{(m)}} (|I_z^{(m)}| + 1) R_z^{(m)} \\ &\leq \sum_{m=1}^{M} \sum_{z \in Z^{(m)}} (\kappa^* + 1) R_z^{(m)} \\ &= (\kappa^* + 1) R_{tot,EF}^{(1)}. \end{aligned} \tag{14}$$

When interference exists between users, we denote the optimal sum-rate by $R_{tot,\max}^{(1)}$. Obviously,

$$R_{tot,\max}^{(1)} \leq R_{tot,\max}^{(0)} \leq (\kappa^* + 1) R_{tot,EF}^{(1)}. \tag{15}$$

So it follows that $R^{(1)}_{tot,EF} \geq \frac{1}{\kappa^*+1} R^{(1)}_{tot,\max}$. Then Theorem 2 follows. ∎

Theorem 2 shows that the EF-based algorithm is optimal when $\kappa^* = 0$, e.g., when every two CR links are sufficiently separated such that they do not interfere with each other. When interference exists, the lower bound on the algorithm's performance decreases linearly with κ^*. The actual performance gap is evaluated later by simulations.

Remark: Depending on the CR's hardware capabilities as well as other regulatory factors, additional constraints on the CRN may be imposed. These include:

C4: Number of Parallel Transmissions: The maximum number of channels a CR transmitter can use at one time may be bounded by M_t. In the BLP framework, this constraint is presented as

$$\tilde{C}4: \quad \sum_{m=1}^{M} \sum_{k=1}^{K} y_{k,i}^{(m)} \leq M_t, \quad \text{for } i = 1, \ldots, N. \tag{16}$$

C5: Transmission Bandwidth: The total bandwidth a CR can transmit over at any time is bounded by B_t. Formally,

$$\tilde{C}5: \quad \sum_{m=1}^{M} \sum_{k=1}^{K} B_m y_{k,i}^{(m)} \leq B_t, \quad \text{for } i = 1, \ldots, N. \tag{17}$$

C6: Forbidden Channels: A CR link i may be prohibited from using a certain set of channels, say $\mathbf{BF}_i \subseteq \{1, \ldots, M\}$. This constraint can be modeled as

$$\tilde{C}6: \quad y_{k,i}^{(m)} = 0, \quad \text{for } k = 1, \ldots, K, \text{and } m \in \mathbf{BF}_i. \tag{18}$$

An examination of (16) through (18) shows that the additional constraints are linear in the $y_{k,i}^{(m)}$'s. Thus, they do not fundamentally change the BLP formulation and its solutions discussed in the previous sections. Extensions of LPSF and the EF-based algorithm to handle such constraints are trivial, and thus are omitted due to space limitation.

5 Example Application and Practical Considerations

In this section, we illustrate through an application the main idea behind the multi-level spectrum opportunity. We consider a spectrum-leasing scenario, where a CRN shares the spectrum with an infrastructure PRN, as shown in Figure 2. The PRN consists of multiple static base stations (BSs) that are interconnected via a broadband wired network. We assume that the PRN operates using frequency division duplexing (FDD). At any given time, each BS tunes to some of the M uplink channels to receive signals from the PR mobile stations (MSs) (not shown in the figure). We only consider spectrum sharing on the uplink. To be consistent with the model in Section 3, a BS operating on multiple channels can be modeled as multiple virtual BSs that operate on individual non-overlapping channels.

We consider two different spectrum-opportunity discovery mechanisms: distributed sensing (DS) and

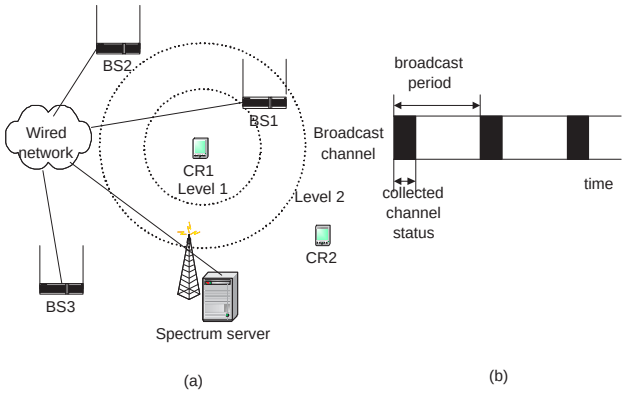

Fig. 2. Spectrum leasing application: (a) system architecture; (b) timing of the broadcasted channel-status information from the spectrum server.

subscription-based (SB). DS is what is commonly studied in most previous works. Each CR periodically senses channels and discovers binary-type spectrum opportunities. The power mask is either $P_{\max}$ or 0, depending on whether the channel is idle or busy. In SB, each PR BS periodically reports its status (receiving or idle) on each channel to the spectrum server via the wired network. Along with the location of each BS, the collected real-time channel-status information is broadcasted by the server to the CRN. By subscribing to the broadcast, each CR can calculate its multi-level spectrum opportunity, as described shortly. Note that the SB scheme fits within the CRN operational model recently advocated by the FCC [6], which calls for establishing a database that CR systems must first register with. This database also provides geo-location information of PRs, and assists the CR in identifying spectrum opportunities. In reality, the SB scheme may be used by CRs to share spectrum with an infrastructure-based PRN, such as a cellular IS-95 or 802.16 WiMAX system.

In contrast to the DS scheme, which assumes that the PRN is ambivalent to the existence of the CRN and no information exchange takes place between the two, the SB scheme assumes that the PRN collaborates with the CRN in identifying spectrum opportunities. This collaboration can be justified by *economic* considerations, where a PRN opens its spectrum to secondary reuse for a profit. The subscription component in SB is extremely suitable for implementing fee-based services, and thus provides a good incentive for the PRN to collaborate. Besides the economic reasons, other incentives to use the SB scheme include (1) providing the spectrum leaser more control over the pricing and control of the shared spectrum, and (2) lowering the CR's hardware complexity (and cost) because the sensing functionality can now be removed from the CR.

The SB scheme requires calculating the power mask. The basic idea behind this calculation is to adapt the CR's interference range to the activity of neighboring BSs. The interference range is defined as the distance d_I for which

$\hat{P}_i^{(m)} h^{(m)}(d_I) \leq P_I$, where $h^{(m)}(d_I)$ is the channel gain at distance d_I on channel m and P_I is the interference tolerance, below which the interference can be deemed as harmless to the PR. We also assume that each CR has knowledge of its location, and thus can calculate its distance to neighboring PR BSs. The determination of the interference-range is illustrated in Figure 2: If the channel gain is dictated by the propagation distance, then when BS1 is receiving on channel m, CR1's power mask should be such that its interference range is smaller than the distance between CR1 and BS1, denoted by the smallest dotted circle (Level 1) in the figure. When BS1 is not receiving but BS2 is receiving, then the power mask can be increased such that CR1's interference range reaches the larger dotted circle (Level 2), and so on. Although this basic idea seems straightforward, the calculation needs to take into account the following two random factors.

5.1 Randomness of PR Activity

This randomness impacts the determination of the the power mask. For example, in Figure 2, even if BS1 is not currently receiving on channel m, there is a chance that it subsequently starts receiving data before the next reporting time. This status change cannot be conveyed to CRs until the next report. So if a CR transmits based on the power mask of Level 2, which is calculated according to the current status report, it will cause unacceptable interference to BS1. To account for this impact, we impose a soft guarantee, $\alpha^{(m)}$ for channel m, such that the ratio of the time the CR interferes with the PRN on channel m is smaller than $\alpha^{(m)}$. This constraint requires us to take into account the accumulated possibility of status-flipping (from not-receiving to receiving) of all idle BSs that are closer to the target CR than its closest active BS neighbor. As a result, it might not always be appropriate to use a power mask that corresponds to the closest active BS neighbor. For example, in Figure 2, even if BS2 is the closet active neighbor of CR1 in the current report, the CR should not use the power mask of Level 2, if the likelihood of BS1 flipping to receiving is greater than $\alpha^{(m)}$. The detailed mathematical treatment is given in the appendix.

5.2 Randomness of the Channel Gain

This randomness impacts the value of each power mask level. Given $\hat{P}_i^{(m)}$, the random channel fluctuation means that the received signal strength at distance d_I is a random variable: $\hat{p}_i^{(m)} = \hat{P}_i^{(m)} \bar{h}(d_I) \chi^{(m)}$, where $\chi^{(m)}$ is a unit-mean r.v. denoting the random fluctuation of the channel, $\bar{h}(d_I) = A_0 d_I^{-\mu}$ is the distance-related component of the path loss, A_0 is the close-in constant, and μ is the path loss exponent. To counter this random effect, we impose a second soft guarantee, $\beta^{(m)}$ for channel m, which requires $\Pr\{\hat{P}^{(m)} \bar{h}(d_I) \chi^{(m)} \geq P_I\} < \beta^{(m)}$. Since d_I is fixed (this corresponds to the interference range of the level selected in the previous section), $\hat{P}_i^{(m)}$ is calculated as $\hat{P}_i^{(m)} = \frac{P_I}{\bar{h}(d_I) Q^{(m)}(\beta^{(m)})}$, where $Q^{(m)}(\beta^{(m)})$ is the $(1-\beta^{(m)})$-quantile of the fluctuation $\chi^{(m)}$, i.e., $\Pr\{\chi^{(m)} \leq Q^{(m)}(\beta^{(m)})\} = 1 - \beta^{(m)}$.

Obviously, the rate at which each BS reports to the spectrum server impacts the calculated power mask. The lower this rate is, the larger the uncertainty in the BS' status between two consecutive reports, and therefore the more conservative the power mask has to be in order to guarantee the given PRN violation constraint. On the other hand, the bandwidth of the channel-status information broadcast channel also influences the throughput of the CRN: Because CRs update their power masks according to the periodic broadcast, the higher the broadcast bandwidth, the quicker each CR can acquire the channel-status information, thus more time left between two consecutive updates for a CR to deliver data. An interesting question is how much gain the multi-level scheme can attain when accounting for the broadcast overhead. We will answer this question using simulations.

6 Performance Evaluation

6.1 Accuracy of the Approximate Algorithms

We consider a 1000 meter $\times$ 1000 meter region, in which 5 PRNs (5 channels) coexist with 5 CR links. The parameter values used in our simulations are listed in Table 3. We assume the following rate-SINR relationship: $R_i^{(m)} = B_m \log_2(1 + SINR/8)$, and $r_i^{(m)} \in \{0, 1/2, 1, 3/2, 2\}$ bits/second/Hz for all i and m. The locations of the PR and CR transmitters and receivers are randomly assigned within the simulation region. A simple path loss model with exponent of 4 is assumed for the channel gain between any two points (i.e., $h_{ij} = d_{ij}^{-4}$). We assume the PRs on all channels follow the same 2-state Markov activity model, i.e., durations of ON/OFF states are exponentially distributed, with the average ON and OFF periods as given in Table 3. The power masks of all CRs are calculated periodically according to the SB scheme. A CR is capable of using all 5 channels at once. We compare the achieved sum-rate of all CR links in each reporting period under 3 different algorithms: an exhaustive-search algorithm that finds the optimal solution, our polynomial-time LPSF algorithm, and the EF algorithm.

A trace of the CRN sum-rate is plotted in Figure 3 for 50 consecutive reporting periods. The upper bound generated in the first iteration of the LPSF algorithm is also shown. It is clear that the LPSF and the EF algorithms give near-optimal solutions (within 5% from the optimal solution). In most of the cases, they give the optimal solution. The upper bound provided by the LPSF algorithm is reasonably tight. In all simulations, the gap between this bound and the optimal solution does not exceed 10%. So this bound provides a useful reference to evaluate the accuracy of the approximate solutions in large networks when the optimal solution is computationally difficult to obtain.

channel bandwidth	1 MHz
number of PRs over various channels	25, 10, 15, 20, 25
average PR ON period	1 second
average PR OFF period	10 second
transmission power of a PR	500 mW
total transmission power of a CR $P_{\max}$	1 W
PR interference tolerance P_I	0.12346 μW
sensitivity of CR receiver $P_{I,CR}$	0.06173 μW
PR status report period	100 ms
CR-to-PR violation bound $\alpha^{(m)}$	2% for all channel m

TABLE 3
Simulation parameters.

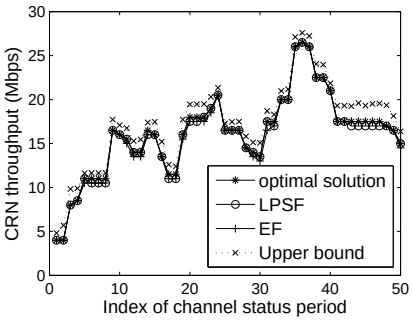

Fig. 3. Trace of the CRN throughput.

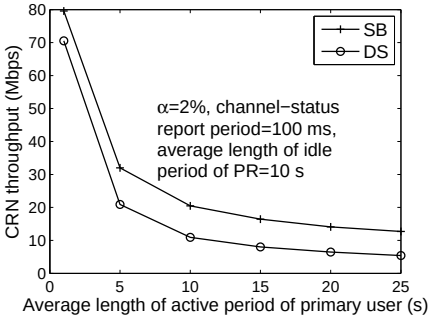

Fig. 4. CRN throughput vs. PR activity.

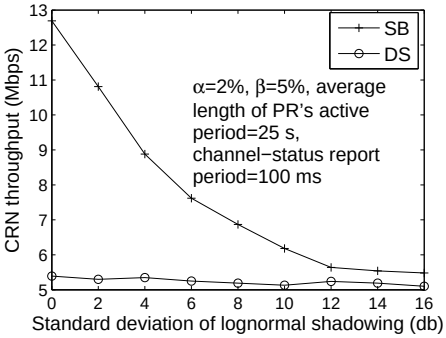

Fig. 5. CRN throughput vs. channel shadowing.

6.2 Comparison between Binary and Multi-level Opportunity

In this section, we simulate a large network and apply the EF algorithm for channel access. We consider 10 channels and 10 CR links over the same square area. The numbers of PRs operating on each channel are 25, 10, 15, 20, 25, 10, 5, 15, 20, and 25, respectively. The set of rates supported by a CR is now given by $\{0, 1/2, 1, 3/2, 2, 5/2, 3, 7/2, 4\}$ b/s/Hz. So the number of binary variables in the BLP is 800. Unless indicated otherwise, the other parameters are kept the same. The results presented below are averaged over 20 randomly generated topologies, with a simulation time of 1000 sensing/status-report periods for each topology.

For the DS scheme, we assume a channel-sensing period of 100 ms. We denote the status-report period of the SB scheme by T. The performance metric of interest is the CRN throughput, defined as the average number of data bits that can be transmitted by all CR links in one period divided by the duration of the period. Because under the SB scheme, a fraction of the period, denoted by T_B, is used to receive broadcast information at each CR, the actual data transmission time in each period is $T - T_B$. The overhead is given by $T_B = \frac{V_B}{B_B}$, where V_B is the number of bits of the collected channel-status information in one report period and B_B is the bandwidth of the broadcast channel. For our simulations, V_B is loosely upper bounded as follows. We assume that the channel-status information for one PR has the format (PR id $\|$ channel id $\|$ channel status). The total number of PRs is less than 200, so an 8-bit id field is enough to identify them. The 10 channels can be identified by a 4-bit field, and 1 bit is used to identify the status of the PR (ON/OFF). So $V_B < 200 \times (8+4+1) = 2.6$ Kbits. We use this value in the following overhead calculation. To give a conservative estimate of the gain attained by the SB scheme, we assume that channel sensing under the DS scheme takes zero time. Thus, the throughput of the DS scheme plotted below represents an upper bound on any channel access scheme that is based on the binary spectrum opportunity. We ignore the computation time of the EF algorithm in both schemes.

In Figure 4, we study the CRN throughput as a function of the average length of a PR's active period. Here, we fix $B_B = 260$ Kbps, corresponding to $T_B = 10$ ms. We observe that SB always achieves higher throughput than DS. This is because SB is able to exploit the microscopic multi-level spectrum opportunities, while DS can only utilize binary spectrum opportunities. More specifically, when a CR is close to an active PR, this CR cannot transmit on the same channel under the DS scheme. However, under the SB scheme, such a transmission can take place if a reasonably small transmission power is used. Therefore, more CR transmission opportunities are found under SB, leading to better CRN throughput performance. It can also be observed from Figure 4 that at low PR activity, the throughput of SB exceeds DS slightly

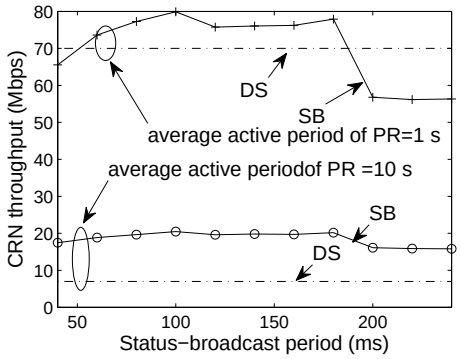

Fig. 6. CRN throughput vs. period of the status broadcast.

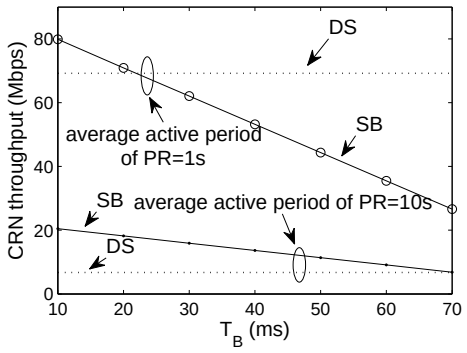

Fig. 7. CRN throughput vs. bandwidth of the broadcast channel.

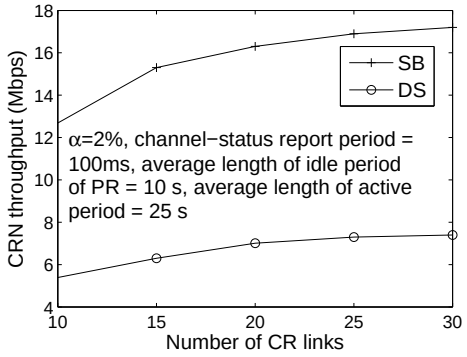

Fig. 8. CRN throughput vs. number of CR links.

(15% gain), but at high PR activity, SB exceeds DS significantly (150% gain). So, although the broadcast channel consumes about 2.6% of the total system bandwidth, it leads to at least 15% throughput gain in the worst case and 150% gain in the best case. The difference in gain is because when the PR activity is low, all neighboring PRs are often in the OFF state. The outcome of SB approaches that of DS, in the sense that most of the time a CR can transmit at power level $P_{\max}$. With increased PR activity, low-power CR transmission opportunities occur more frequently under the SB scheme, which cannot be exploited under DS, and thus the gap between the two schemes keeps growing.

We study the impact of channel fluctuations in Figure 5, where a channel is subject to log-normal shadowing. The channel gain is taken as $g_{ij} = d_{ij}^{-4} 10^{\frac{\chi}{10}}$, where χ is a zero-mean Gaussian random variable that denotes the channel fluctuation in db. The standard deviation of χ represents the severity of shadowing. For each channel, we require a soft guarantee $\beta = 5\%$. We first note that the average throughput under DS barely changes with std(χ) because of the fixed power mask set $(0, P_{\max})$. It is also observed that with the increase in channel fluctuation, the throughput under SB will decrease, and eventually it approaches that of DS. However, when the standard deviation is 6 db, which represents a typical shadowing environment, SB still achieves about 50% throughput gain over DS.

In Figure 6, we fix B_B to 260 Kbps ($T_B = 10$ ms) and vary the status-broadcast period. It can be observed that in general, a shorter broadcast period leads to higher throughput because of the increased certainty in the PR's activity between two consecutive reporting instances. However, when the broadcast period is very small, e.g., $T = 40$ ms, the throughput under SB is low. This is because the broadcast of status information occupies a significant portion of each broadcast period, thus less time is left for data transmissions.

In Figure 7, we fix the status-broadcast period at $T = 100$ ms and plot the throughput under SB for various T_B values (corresponding to various broadcast channel bandwidths). It is observed that the throughput degrades linearly with T_B (or equivalently, with the decrease in the broadcast bandwidth), because less time in each period is left for data transmissions. At low PR activity, the throughput of SB crosses that of DS once T_B is greater than 20 ms or B_B is smaller than 130 Kbps, which is about 1.3% of the total system bandwidth. At high PR activity, the crossing point is $T_B = 50$ ms. This corresponds to $B_B \approx 50$ Kbps (0.5% of the total bandwidth). The extremely small bandwidth at the crossing point in both situations indicate that the overhead of SB is basically negligible.

In Figure 8, we study the CRN throughput as a function of the number of CR links under the SB and DS schemes. It can be observed that for both schemes, throughput increases with the number of CR links. Interestingly, the throughput gap between SB and DS grows with the number of CR links. This phenomenon can be explained by noting that DS has a much larger keep-out region than SB. Therefore, under DS, increasing the number of CR links in the area does not lead to a significant increase in the number of active CR links. In contrast, under the SB scheme, because a CR that is close to an active PR can still transmit data by adjusting its transmission power, an increase in the number of CR links can be directly translated into more transmitting CRs, leading to higher throughput gain.

7 Conclusions

In this paper, we developed centralized and distributed algorithms for joint power/rate control and channel assignment in CRNs. The problem was formulated under a multi-level spectrum opportunity framework, which reflects the microscopic spatial opportunity available to CRs. We also applied our algorithms to study the achieved throughput gain over the conventional binary spectrum opportunity while taking its overhead into account. We showed that a significant gain can be achieved under the SB scheme with the assistance of a narrowband channel, which periodically broadcasts channel-status information. Our work only considered single-hop ad hoc CRNs. Our future efforts will focus on the routing problem in a multi-hop CRN.

References

[1] L. Cao and H. Zheng. Distributed spectrum allocation via local bargaining. In *Proceedings of IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, 2005.

[2] L. Cao and H. Zheng. On the efficiency and complexity of distributed spectrum allocation. In *Proceedings of International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM)*, 2007.

[3] N. B. Chang and M. Liu. Optimal channel probing and transmission scheduling for opportunistic spectrum access. In *the Proceedings of the ACM MobiCom Conference*, 2007.

[4] T. Chen, H. Zhang, G. M. Maggio, and I. Chlamtac. CogMesh: A cluster-based cognitive radio network. In *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 168–178, 2007.

[5] D. R. Cox. *Renewal Theory*. Butler & Tanner Ltd., 1967.

[6] FCC Document 08-260. Second report and order and memorandum opinion and order: In the matter of unlicensed operation in the TV broadcast bands and additional spectrum for unlicensed devices below 900 MHz and in the 3 GHz band, Nov. 2008.

[7] G. Ganesan and Y. Li. Cooperative spectrum sensing in cognitive radio networks. In *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pages 137–143, 2005.

[8] Y. T. Hou, Y. Shi, and H. D. Sherali. Optimal spectrum sharing for multi-hop software defined radio networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 1–9, 2007.

[9] S. Huang, X. Liu, and Z. Ding. Opportunistic spectrum access in cognitive radio networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 1427–1435, Apr. 2008.

[10] J. Jia, Q. Zhang, and X. Shen. HC-MAC: A hardware-constrained cognitive MAC for efficient spectrum management. *IEEE Journal on Selected Areas in Communications*, 26(1):106–117, Jan. 2008.

[11] H. Kim and K. G. Shin. Efficient discovery of spectrum opportunities with MAC-layer sensing in cognitive radio networks. *IEEE Transactions on Mobile Computing*, 7(5):533–545, May 2008.

[12] H. Kim and K. G. Shin. In-band spectrum sensing in cognitive radio networks: energy detection or feature detection? In *the Proceedings of the ACM MobiCom Conference*, pages 14–25, 2008.

[13] J. Lee, R. V. Sonalkar, and J. M. Cioffi. Multiuser bit loading for multicarrier systems. *IEEE Transactions on Communications*, 54(7):1170–1174, July 2006.

[14] Y. Shi and Y. T. Hou. Optimal power control for multi-hop software defined radio networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 1694–1702, 2007.

[15] Y. Shi and Y. T. Hou. A distributed optimization algorithm for multi-hop cognitive radio networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 1966–1974, Apr. 2008.

[16] Y. Shi, Y. T. Hou, H. D. Sherali, and S. F. Midkiff. Optimal routing for UWB-based sensor networks. *IEEE Journal on Selected Areas in Communications*, 24(4):857–863, Apr. 2006.

[17] T. Starr, M. Sorbara, J. M. Cioffi, and P. J. Silverman. *DSL Advances*. Prentice-Hall, 2003.

[18] W. Wang and X. Liu. List-coloring based channel allocation for open-spectrum wireless networks. In *Proceedings of the IEEE Vehicular Tech. Conference*, pages 690–694, 2005 Fall.

[19] Y. Xing, R. Chandramouli, S. Mangold, and S. N. Shankar. Dynamic spectrum access in open spectrum wireless networks. *IEEE Journal on Selected Areas in Communications*, 24(3):626–637, Mar. 2006.

[20] Y. Xing, C. N. Mathur, M. A. Haleem, R. Chandramouli, and K. P. Subbalakshmi. Dynamic spectrum access with QoS and interference temperature constraints. *IEEE Transactions on Mobile Computing*, 6(4):423–433, Apr. 2007.

[21] Y. Yuan, P. Bahl, R. Chandra, T. Moscibroda, and Y. Wu. Allocating dynamic time-spectrum blocks in cognitive radio networks. In *Proceedings of the ACM MobiHoc Conference*, pages 130–139, 2007.

[22] J. Zhao, H. Zheng, and G. Yang. Distributed coordination in dynamic spectrum allocation networks. In *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2005.

[23] Q. Zhao, S. Geirhofer, L. Tong, and B. M. Sadler. Optimal dynamic spectrum access via periodic channel sensing. In *Proceedings of the IEEE WCNC Conference*, 2007.

[24] Q. Zhao and B. M. Sadler. A survey of dynamic spectrum access: signal processing, networking, and regulatory policy. *IEEE Signal Processing Magazine*, 24(3):79–89, 2007.

[25] Q. Zhao, L. Tong, A. Swami, and Y. Chen. Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: A POMDP framework. *IEEE Journal on Selected Areas in Communications*, 25(3):589–600, Apr. 2007.

[26] H. Zheng and L. Cao. Device-centric spectrum management. In *Proceedings of the IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2005.

[27] H. Zheng and C. Peng. Collaboration and fairness in opportunistic spectrum access. In *Proceedings of the IEEE ICC Conference*, 2005.

Appendix A
Calculation of the Power Mask under Random PR Activity

We use the calculation on channel m as an example. For a target CR, label its PR BS neighbors from closest to farthest as BS1, BS2, ... Denote its power mask set on channel m by $\mathbf{p}_{mask}^{(m)} = (p_{mask}^{(m)}(1), p_{mask}^{(m)}(2), \ldots, p_{mask}^{(m)}(N_m + 1))$, where $p_{mask}^{(m)}(j) = P_I / h_j^{(m)}$ for $j = 1, \ldots, N_m$, P_I is the PR's interference tolerance, $h_j^{(m)}$ is the path loss from the target CR to BS j with $h_1^{(m)} \geq \ldots \geq h_{N_m}^{(m)}$, and N_m is the BS beyond which the CR-to-PR interference is always smaller than P_I even if the CR is transmitting at its full power $P_{\max}$. We let $p_{mask}^{(m)}(N_m+1) \stackrel{\text{def}}{=} P_{\max}$. Define the channel-usage profile at the nth reporting time as an N_m-dimensional vector $\mathbf{S}^{(m)}(n) = (s_1^{(m)}(n), \ldots, s_{N_m}^{(m)}(n))$. For $i = 1, \ldots, N_m$:

$$s_i^{(m)}(n) = \begin{cases} 1, & \text{if BS } i \text{ is receiving at the } n\text{th report} \\ 0, & \text{if BS } i \text{ is not receiving at the } n\text{th report} \end{cases} \tag{19}$$

The status of BS i is characterized by a random process $\{x_i^{(m)}(t): t \geq 0\}$, where $x_i^{(m)}(t) = 1/0$ denotes the receiver is receiving (ON)/not-receiving (OFF) at time t. Let τ_n and τ_{n+1} denote the nth and the $(n+1)$th status reporting times, respectively. So $s_i^{(m)}(n)$ denote the value of $x_i^{(m)}(t)$ at time τ_n. Define the random variable $\xi_i^{(m)}$ as follows:

$$\xi_i^{(m)} = \begin{cases} 1, & \text{if } x_i^{(m)}(t) = 1 \text{ for some } t \in [\tau_n, \tau_{n+1}] \\ 0, & \text{if } x_i^{(m)}(t) = 0 \ \forall t \in [\tau_n, \tau_{n+1}]. \end{cases} \tag{20}$$

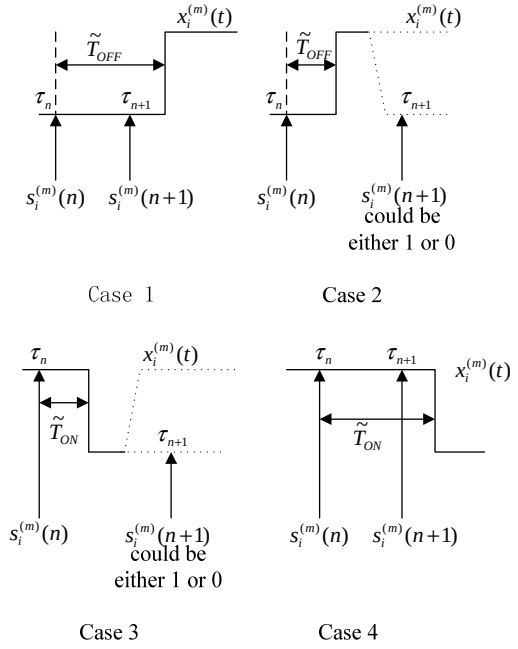

Fig. 9. Randomness of the PR activity.

Between τ_n and τ_{n+1}, $x_i^{(m)}(t)$ fluctuates according to one of the four possible cases in Figure 9. Using the current sensing time τ_n as a reference point, we denote the forward recurrence time of the OFF state by $\tilde{T}_{OFF}$ and the forward recurrence time of the ON state by $\tilde{T}_{ON}$. Following standard renewal theory results [5], the probability of each case can be calculated as follows:

$$\begin{aligned}\Pr\left\{\xi_i^{(m)}=0|s_i^{(m)}(n)=0\right\} &= \Pr\left\{\tilde{T}_{OFF}>T\right\}\\ &= 1-\int_0^T f_{\tilde{T}_{OFF}}(t)dt \quad (21)\end{aligned}$$

$$\begin{aligned}\Pr\left\{\xi_i^{(m)}=1|s_i^{(m)}(n)=0\right\} &= \Pr\left\{\tilde{T}_{OFF}\le T\right\}\\ &= \int_0^T f_{\tilde{T}_{OFF}}(t)dt \quad (22)\end{aligned}$$

$$\Pr\left\{\xi_i^{(m)}=0|s_i^{(m)}(n)=1\right\}=0 \quad (23)$$

$$\begin{aligned}\Pr\left\{\xi_i^{(m)}=1|s_i^{(m)}(n)=1\right\} &= 1-\Pr\left\{\xi_i^{(m)}=0|s_i^{(m)}(n)=1\right\}\\ &= 1 \quad (24)\end{aligned}$$

where $T \stackrel{\text{def}}{=} \tau_{n+1}-\tau_n$ is the sensing period and $f_{\tilde{T}_{OFF}}(t)$ is the pdf of $\tilde{T}_{OFF}$, which can be derived from the pdf of the OFF period $f_0^{(m)}$:

$$f_{\tilde{T}_{Off}}(t)=\frac{1-\int_0^t f_0^{(m)}(\tau)d\tau}{\int_0^\infty \tau f_0^{(m)}(\tau)d\tau}. \quad (25)$$

Note that (23) can be explained from the definition of the r.v. $\xi_i^{(m)}$ in (20). At time τ_n, given a sensing output $\mathbf{S}^{(m)}(n)$, the violation probability of the CR transmitter on channel m is conditioned on the power mask to be used. Because the neighboring PR receivers $v_i^{(m)}$'s are labeled in a descending order of their channel gains, at

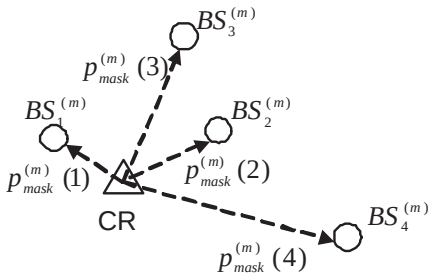

Fig. 10. An example of run-time spectrum opportunity detection.

a given power-mask level l ($1 \le l \le N_m+1$), a violation will happen if and only if there exists at least one PR neighbor i, $i<l$, for which $\xi_i^{(m)}=1$. Mathematically, the violation rate for a given level of power mask can be calculated as follows:

$$\begin{aligned}&\Pr\{\text{violation of PRN } m|\mathbf{S}^{(m)}, p_{mask}^{(m)}(l)\}\\ =\ &\Pr\{\xi_1^{(m)}=1|\mathbf{S}^{(m)}\}+\Pr\{\xi_1^{(m)}=0,\xi_2^{(m)}=1|\mathbf{S}^{(m)}\}\\ &+\ldots+\Pr\{\xi_1^{(m)}=0,\ldots,\xi_{l-2}^{(m)}=0,\xi_{l-1}^{(m)}=1|\mathbf{S}^{(m)}\}\\ =\ &\Pr\{\xi_1^{(m)}=1|s_1^{(m)}\}+\Pr\{\xi_1^{(m)}=0|s_1^{(m)}\}\Pr\{\xi_2^{(m)}=1|s_2^{(m)}\}\\ &+\ldots+\Pr\{\xi_1^{(m)}=0|s_1^{(m)}\}\Pr\{\xi_2^{(m)}=0|s_2^{(m)}\}\\ &\times\ldots\times\Pr\{\xi_{l-2}^{(m)}=0|s_{l-2}^{(m)}\}\Pr\{\xi_{l-1}^{(m)}=1|s_{l-1}^{(m)}\}. \quad (26)\end{aligned}$$

The multiplicative form in the last step of (26) is due to the assumed independence between different PR links. We can rewrite (26) in a compact form as

$$V(l,\mathbf{S}^{(m)}) \stackrel{\text{def}}{=} \sum_{i=1}^{l-1}\Pr\{\xi_i^{(m)}=1|s_i^{(m)}\}\prod_{j=1}^{i-1}\Pr\{\xi_j^{(m)}=0|s_j^{(m)}\} \quad (27)$$

for $l=1,\ldots,N_m+1$. In addition, we also define $V(N_m+2,\mathbf{S}^{(m)})=1$. The conditional probabilities $\Pr\{\xi_i^{(m)}|s_i^{(m)}\}$ in (27) can be computed according to (21) through (24). The level of spectrum opportunity l^* over channel m should be selected such that $V(l^*,\mathbf{S}^{(m)}) \le \alpha^{(m)}$ and $V(l^*+1,\mathbf{S}^{(m)})>\alpha^{(m)}$. It is easy to verify that $V(l,\mathbf{S}^{(m)})$ defined in (27) is a mono-increasing function of l. Thus, this selection is valid in the sense that some l^* can always be found. Then, the power mask of a given CR i is simply $\hat{P}_i^{(m)}=p_{mask}^{(m)}(l^*)$.

As an example of the above calculation, consider the scenario in Figure 10, where the CR has four PR BS neighbors. Assume the pdf function f_0 is exponential with a mean of 10 seconds, and the sensing period is 100 ms. The translation from a sensing output vector ($s_i^{(m)}$ corresponds to $BS_i^{(m)}$) to the level of maximum allowable power mask l^* is calculated in Table 4 for $\alpha=1\%$ and $\alpha=2\%$. From this table, it is clear that the microscopic spectrum opportunities are more aggressively exploited under a larger tolerance (i.e., a larger α) of the randomness of PR activities, which is in line with our intuition.

$(s_1^{(m)} s_2^{(m)} s_3^{(m)} s_4^{(m)})$	l^* $(\alpha = 1\%)$	l^* $(\alpha = 2\%)$
(0000)	1	3
(0001)	1	3
(0010)	1	3
(0011)	1	3
(0100)	1	2
(0101)	1	2
(0110)	1	2
(0111)	1	2
(1000)	1	1
(1001)	1	1
(1010)	1	1
(1011)	1	1
(1100)	1	1
(1101)	1	1
(1110)	1	1
(1111)	1	1

TABLE 4
Mapping between sensing output and spectrum opportunity.

PLACE PHOTO HERE

Tao Shu received the B.S. and M.S. degrees in electronic engineering from the South China University of Technology, Guangzhou, China, in 1996 and 1999, respectively, and the Ph.D. degree in electronic engineering from Tsinghua University, Beijing, China, in 2003. Currently he is a Ph.D. student at the electrical and computer engineering department at the University of Arizona, Tucson, USA. His research aims at addressing the security and performance issues in wireless networking systems, with strong emphasis on system architecture, protocol design, and performance optimization. His research interests span a broad range of problems, including applied cryptography, information assurance, network security and privacy, truthful mechanism design, dynamic spectrum access, radio resource management, and routing/MAC protocol design. His research covers a wide variety of wireless systems, including cognitive radio networks, wireless sensor networks, mobile ad hoc networks, and cellular networks. He is a student member of IEEE.

PLACE PHOTO HERE

Marwan Krunz is a professor of ECE at the University of Arizona. He also holds a joint appointment at the same rank in the Department of Computer Science. He directs the wireless and networking group and is also the UA site director for Connection One, a joint NSF/state/industry IUCRC cooperative center that focuses on RF and wireless communication systems and networks. Dr. Krunz received his Ph.D. degree in electrical engineering from Michigan State University in 1995. He joined the University of Arizona in January 1997, after a brief postdoctoral stint at the University of Maryland, College Park. He previously held visiting research positions at INRIA, HP Labs, University of Paris VI, and US West (now Qwest) Advanced Technologies. Currently, he is a visiting researcher at the University of Carlos III, Madrid, Spain, where he holds a Chair of Excellence ("Ctedra de Excelencia") position. Dr. Krunz's research interests lie in the fields of computer networking and wireless communications, with recent focus on cognitive radios and SDRs; distributed radio resource management in wireless networks; channel access and protocol design; MIMO and smart-antenna systems; UWB-based personal area networks; energy management and clustering in sensor networks; media streaming; QoS routing; and fault monitoring/detection in optical networks. He has published more than 160 journal articles and refereed conference papers, and is a co-inventor on three US patents. M. Krunz is a recipient of the National Science Foundation CAREER Award (1998). He currently serves on the editorial boards for the IEEE Transactions on Mobile Computing and the Computer Communications Journal. He previously served on the editorial board for the IEEE/ACM Transactions on Networking (2001-2008). He was a guest co-editor for special issues in IEEE Micro and IEEE Communications magazines. He served as a technical program chair for various international conferences, including the IEEE WoWMoM 2006, the IEEE SECON 2005, the IEEE INFOCOM 2004, and the 9th Hot Interconnects Symposium (2001). He has served and continues to serve on the executive and technical program committees of many international conferences and on the panels of several NSF directorates. He gave keynotes and tutorials, and participated in various panels at premier wireless networking conferences. He is a consultant for a number of companies in the telecommunications sector. He is an IEEE Fellow.