Competitive Analysis of Opportunistic Spectrum Access Strategies

Nicholas B. Chang and Mingyan Liu
Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, MI 48109-2122
Email: nchang@ll.mit.edu, mingyan@eecs.umich.edu

Abstract—We consider opportunistic spectrum access (OSA) strategies for a transmitter in a multichannel wireless system, where a channel may or may not be available and the transmitter must sense/probe the channel to find out before transmission. Applications for this work include joint probing and transmission for a secondary user in a cognitive radio network. Limited by resources, e.g., energy and time, the transmitter must decide on a subset of a potentially very large number of channels to probe and can only use for transmission those that have been found to be available. In contrast to previous works, we do not assume the user has a priori knowledge regarding the statistics of channel states. The main goal of this work is to design strategies that decide, based only on knowledge of the channel bandwidths/data rates, which channels to probe. We derive optimal strategies that maximize the total expected bandwidth/data rate in the worst-case, via a performance measure in the form of a competitive regret (ratio) between the average performance of a strategy and a genie (or omniscient observer). We examine the performance of these optimal strategies under a wide range of system parameters and practical channel models via numerical studies.

I. INTRODUCTION

In this paper, we examine optimal channel sensing/probing strategies for opportunistic spectrum access (OSA), where a transmitter seeks to maximize its achievable data rate by opportunistically transmitting over a select subset of a potentially large number of channels. This is done by optimally constructing a channel sensing/probing strategy to find out which channels are available before transmitting.

This problem is motivated by wireless systems where a transmitter is supplied with more channels than needed for transmission. A channel, for example, could be a frequency in an FDMA system, a code in a CDMA system, etc. Software defined radio [1] and cognitive radio systems [2] may provide users with multiple channels (e.g. tunable frequency bands and modulation techniques) by means of a programmable hardware which is controlled by software. Specifically, a transmitter could be a secondary user in a cognitive radio network seeking to utilize portions of the spectrum currently not being used by a set of primary users. In doing so it needs to find out which

part of the spectrum is available for transmission. Furthermore, channel quality is in general time-varying due to multi-path fading caused by factors such as mobility and interference. Therefore the ability to probe or sense channel quality before transmission can help transmitters select the best one(s) to use. One method of channel probing is through the exchange of control packets between transmitter and receiver [3]. Another method, for example in a cognitive radio system, is to use a spectrum sensor [2], [4] at the physical layer to provide transmitters with information about channel quality.

Due to constraints on time, energy, and other resources, a transmitter may only be able to probe a limited number of channels. Therefore, it is imperative that it judiciously select which channels to probe. If the transmitter has *a priori* knowledge on the random distribution describing the states of different channels, then it can design strategies which maximize the expected transmission gain. Recent works such as [5], [6], [7] have studied the tradeoff between probing to gain more information about channel quality and maximizing transmission gain. Meanwhile, [2] examined the joint design of spectrum sensing at the physical layer and access policies at the medium access control (MAC) layer. All of these works require that the transmitter has *a priori* knowledge on the channel distribution in order to design optimal strategies.

On the other hand, if the transmitter cannot accurately model the distribution of channel states, these strategies may no longer be applicable. In addition, if the distribution of channel quality is constantly changing over time, e.g., in the case of a mobile node, then it becomes impractical for a transmitter to continually adjust its optimal strategy according to the varying distribution. It may also be possible that the transmitter's knowledge on the channel distributions is incorrect, which could lead to errors in determining the optimal strategy. Therefore, it is necessary in such cases to design channel probing and transmission strategies which are more robust and do not require *a priori* knowledge of the channel statistics.

Motivated by these practical considerations, in this paper we use competitive analysis [8] methods and seek to design strategies that perform well in the worst case (to be defined more precisely in Section II) for the following problem. A transmitter has to decide which subset (up to K channels) of N channels to probe, whereby it can only transmit using those channels that have been probed to be available, and its

performance is measured against a genie who knows exactly which channels are available. We will see that the resulting strategies do not require a priori knowledge on the statistics of channel state. To the best of our knowledge, this is the first work which analyzes opportunistic spectrum access policies using a competitive framework.

The main contributions of this paper include the following: First, under a variety of worst-case performance measures, we derive a class of optimal randomized strategies. These strategies are obtained as functions of a range of system parameters, e.g. number of channels, channel probes, etc. and are thus tunable to many practical scenarios. The worst-case performances of these strategies are also obtained, allowing us to study the optimal performance guarantee if a transmitter does not have knowledge on channel distributions. Second, we derive an algorithm that constructs a strategy belonging to this optimal class. Finally, we examine the performance of these strategies via numerical means.

The rest of this paper is organized as follows. Section II formulates the problem and presents the two primary worst-case performance metrics (a difference metric and a ratio metric) considered in this paper. Section III provides notation and preliminary results. Section IV presents a class of optimal strategies and describes an algorithm which sequentially constructs a subclass of the optimal strategies. Section V presents numerical results and Section VI concludes the paper.

Unless otherwise stated, all results are proven in the Appendix. We have stated *it can be shown* for a few results that are easy to verify and thus not included due to space constraints.

II. PROBLEM FORMULATION

In this section we formulate the problem considered in this paper. We then present various performance measures and corresponding objectives.

A. Formulation and Strategies

We consider a wireless system consisting of N channels, indexed by the set $\Omega = \{1, 2, \cdots, N\}$ and a transmitter who has access to these channels. The system works in discrete time, and at each time step the transmitter needs to determine which channels to use for transmission. To do so, the transmitter can probe up to K of the N channels, where $1 \leq K \leq N$.

A channel j has bandwidth or data rate denoted as r_j. We assume a binary channel model, where in each time step a channel may be in one of two states, available (on) or unavailable (off), with certain probability. If a channel j is probed to be available and the transmitter decides to transmit in that channel, then he receives a rate (also referred so as the channel reward) of r_j. If the channel is probed to be unavailable, then the transmitter is not allowed to use that channel for transmission. In addition, the transmitter is not allowed to transmit using unprobed channels.

We assume that the transmitter knows the vector $\mathbf{r} = [r_1, r_2, \cdots, r_N]$. Without loss of generality, we will assume

that $r_1 \geq r_2 \geq \cdots \geq r_N \geq 0$. However, in contrast to previous works, we assume that the transmitter does *not* have a priori information regarding the channel statistics, i.e., the probability of a channel being available. Thus the only information the transmitter has in making the probing decision is the channel bandwidths, but not the likelihood of a channel being available.

We describe the channel state by binary (random) vectors $\mathbf{x} = [x_1, x_2, \cdots, x_N]$ where $x_j = 1$ if channel j is on; otherwise, $x_j = 0$. We denote by X the set of all possible channel states (random or deterministic). When the transmitter does not have any knowledge regarding how many channels are available, we have $\mathbf{x} \in X$.

We will also consider the case when somehow the transmitter knows the average number of available channels. It will be seen that having such knowledge can help the transmitter design better strategies. For this case, we define the following set of channel states:

$$X_L = \left\{ \mathbf{x} \in X \ : E_{\mathbf{x}} \left[\sum_{j=1}^{N} I_{\{x_j=1\}} \right] \leq L \right\} , \qquad (1)$$

where $E_{\mathbf{x}}$ denotes the expectation with respect to $\mathbf{x}$, and $I_{\{\cdot\}}$ is the indicator function. X_L is the set of channel states such that, on average, at most L channels are on. Note that $X_N = X$. For generality, we will derive results in this paper for general L, noting that by using $L = N$ our results directly apply to the case where the transmitter has no knowledge on the average number of available channels.

The system proceeds as follows. The transmitter decides which set of K channels in Ω to probe. After probing, the transmitter then determines which K_0 of these K probed channels to use for transmission. We describe a user's strategy as a (random) binary vector $\mathbf{u} = [u_1, u_2, \cdots, u_N]$, where $u_j = 1$ means the user probes channel j; otherwise, $u_j = 0$.

We denote the probability mass function of u_j by p_{u_j}. Thus, $p_{u_j}(k) = P(u_j = k)$ for $k \in \{0, 1\}$. Let U be the set of strategies, fixed or randomized. We define the following set of strategies:

$$U_K = \left\{ \mathbf{u} \in U \ : \ P \left(\sum_{j=1}^{N} u_j \leq K \right) = 1 \right\} . \qquad (2)$$

That is, U_K is the set of strategies that probe at most K channels. Note that any strategy in U_K must also satisfy: $\sum_{j=1}^{N} p_{u_j}(1) = E_{\mathbf{u}} \left[\sum_{j=1}^{N} I_{\{u_j=1\}} \right] \leq K$.

B. Competitive Regret

Our first worst-case performance measure is in the form of the difference with respect to a genie as we present below. For any integer k, we will let $g_k(\mathbf{z})$ denote the sum of the k largest elements of any vector $\mathbf{z}$. In addition, we let $\mathbf{z}_1 \cdot \mathbf{z}_2$ denote the dot product between any two vectors $\mathbf{z}_1$ and $\mathbf{z}_2$.

Consider any strategy $\mathbf{u} \in U_K$. Its expected reward for any $\mathbf{x}$ is given by:

$$V_{\mathbf{x}}^{\mathbf{u}} = E_{\mathbf{u}} E_{\mathbf{x}} \left[g_{K_0} \left(\mathbf{r} \cdot \mathbf{x} \cdot \mathbf{u} \right) \right] , \qquad (3)$$

because the transmitter will use the K_0 channels with highest rate that are probed and available. Now consider an omniscient observer or genie who knows the realization of channel states in advance and can use up to K_0 channels for transmission. For a given realization, the genie can simply probe and use the K_0 available channels with highest bandwidth. Its expected reward is thus given by:

$$V_{\mathbf{x}}^* = E_{\mathbf{x}}\left[g_{K_0}\left(\mathbf{r} \cdot \mathbf{x}\right)\right] \ . \tag{4}$$

Taking the difference between the genie's expected reward and that of strategy $\mathbf{u}$, we obtain a comparison between the relative performance of $\mathbf{u}$ and the best possible. As the transmitter does not have any *a priori* information regarding $\mathbf{x}$, except possibly the average number of available channels, we take a maximum of this difference over all $\mathbf{x}$ to obtain the following worst-case performance measure of $\mathbf{u}$:.

$$\delta^{\mathbf{u}} = \max_{\mathbf{x} \in X_L}\left\{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\right\}$$

This performance measure can also be interpreted as the *regret* [9], or *minimax regret* [8] of $\mathbf{u}$. The quantity $\delta^{\mathbf{u}}$ provides an upperbound on the performance of strategy $\mathbf{u}$ relative to the best possible. That is, for any arbitrary state $\mathbf{x} \in X_L$, strategy $\mathbf{u}$ will obtain expected reward within $\delta^{\mathbf{u}}$ of the omniscient observer. We will refer to $\delta^{\mathbf{u}}$ as the *worst-case reward (difference) of strategy* $\mathbf{u}$. We emphasize again that when the transmitter does not have knowledge on the average number of available channels, then we can set $L = N$ in the equation for $\delta^{\mathbf{u}}$ to obtain a performance measure with respect to all possible channel states.

One can also view $\delta^{\mathbf{u}}$ as the competitive regret of $\mathbf{u}$ versus an oblivious adversary [8] who knows the distribution (but not realization) of $\mathbf{u}$ in advance. Given $\mathbf{u}$, the adversary will choose an $\mathbf{x} \in X_L$ that maximizes the difference between the reward of $\mathbf{u}$ and the genie.

The objective for the first problem, which we call (A), is to find the strategy which minimizes the above worst-case measure, i.e. obtain the following minimum:

$$\delta^* = \min_{\mathbf{u} \in U_K} \max_{\mathbf{x} \in X_L}\left\{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\right\} \ , \tag{5}$$

where we refer to δ^* as the minimum worst-case difference.

C. Competitive Ratio

Our second worst-case performance measure, which is also commonly used, is in the form of the *ratio* between the reward of a strategy and the genie. Consider the following worst-case performance measure of any strategy $\mathbf{u}$:

$$\rho^{\mathbf{u}} = \min_{\mathbf{x} \in X_L} \frac{V_{\mathbf{x}}^{\mathbf{u}}}{V_{\mathbf{x}}^*} = \min_{\mathbf{x} \in X_L}\left\{\frac{E\left[g_{K_0}(\mathbf{r} \cdot \mathbf{x} \cdot \mathbf{u})\right]}{E\left[g_{K_0}(\mathbf{r} \cdot \mathbf{x})\right]}\right\} \ , \tag{6}$$

which is also known as the competitive ratio [8] of strategy $\mathbf{u}$. For any state $\mathbf{x}$, strategy $\mathbf{u}$ is guaranteed to obtain performance within a factor $\rho^{\mathbf{u}}$ of the best possible. Note that $\rho^{\mathbf{u}} \leq 1$ for any strategy $\mathbf{u}$, as it is impossible to do better than the genie. The corresponding objective for this problem, which we call (B), is to determine the strategy achieving maximum worst-case performance, $\max_{\mathbf{u} \in U_K} \rho^{\mathbf{u}}$.

III. PRELIMINARIES

Before presenting the optimal strategies under the objectives described in the previous section, we first describe some important preliminary notation and results. We will assume that $r_K > 0$. If this were not true, then it means there are less than K channels with positive rate, and the optimal strategy is trivial: probe channel j if $r_j > 0$.

We first present the following important result:

Lemma 1: For any K_0, L, let $\tilde{L} = \min\{L, K_0\}$. Then for any $\mathbf{u} \in U_K$:

$$\delta^{\mathbf{u}} = \max_{\mathbf{x} \in X_L}\left\{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\right\} = \max_{\mathbf{x} \in X_{\tilde{L}}}\left\{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\right\} \tag{7}$$

Thus in calculating $\delta^{\mathbf{u}}$ for any strategy, such as in (8), we can replace X_L with $X_{\tilde{L}}$. For the remainder of this paper, we will let $\tilde{L} = \min\{L, K_0\}$. This result implies the following:

$$\delta^{\mathbf{u}} = \max_{\mathbf{x} \in X_{\tilde{L}}} E_{\mathbf{u}} E_{\mathbf{x}}\left[\sum_{j=1}^{N} r_j I_{\{x_j=1, u_j=0\}}\right]$$

$$= \max_{\mathbf{x} \in X_{\tilde{L}}}\left\{\sum_{j=1}^{N} r_j \cdot p_{x_j}(1) \cdot p_{u_j}(0)\right\} \ , \tag{8}$$

where we have used the linearity of expectation and independence between $\mathbf{x}$ and $\mathbf{u}$. Note that $\delta^{\mathbf{u}}$ only depends on $\mathbf{u}$ and $\mathbf{x}$ through their marginal probabilities, and not their joint probabilities. Therefore throughout this paper we will describe strategies and channel states in terms of their marginal probabilities. In Section IV, we will demonstrate a method for constructing strategies based on these marginal probabilities.

Next, we define the following parameter for any K:

$$M = \max\left\{n \in \{K, \cdots, N\} : r_n > 0, r_n \geq \frac{(n-K)}{\sum_{j=1}^{n} \frac{1}{r_j}}\right\} \tag{9}$$

Note that the above set is guaranteed to be nonempty, since $r_K > 0$ and therefore K lies in the above set. Also note that because $r_j \geq r_k$ whenever $j < k$, the above summation does not involve any division by 0. It will be seen that parameter M is crucial to describing the optimal strategy, and for some special cases denotes the channels which should be probed.

We will assume without loss of generality that $N \geq M + \tilde{L}$. If this condition does not hold, one can add extra channels with rates $r_{N+1} = \cdots = r_{N+\tilde{L}} = 0$. As a user has no incentive to use these extra channels and they do not contribute to the total reward, they do not affect the optimal strategy or its expected reward. This assumption is made to avoid boundary conditions in describing strategies and their worst-case rewards, thus facilitating the description.

We introduce additional notation as follows. For any L, K, N define M from (9). Then, for any integer $1 \leq \bar{M} \leq K$, define:

$$\eta(\bar{M}) = \max\left\{m \in \{0, \cdots, \tilde{L}\} : r_{\bar{M}+m} \sum_{j=1}^{\bar{M}} \frac{1}{r_j} \geq \bar{M} - K\right\}$$

where $\eta(\bar{M}) = 0$ if the above set is empty. This term $\eta(\bar{M})$ is introduced simply for notational purposes in describing the following class of algorithms and their worst-case rewards:

Definition 1: For any set of channels, define M from (9). Then for any integer K, strategy $\mathbf{u}[l, \bar{M}]$ is defined by configurable parameters (integers) l and $\bar{M}$, where $K \leq \bar{M} \leq M$ and either $l = 0$ or $\eta(\bar{M}) + 1 \leq l \leq \eta(\bar{M} - 1) - 1$. The marginal probabilities for $\mathbf{u}[l, \bar{M}]$ depend on l and $\bar{M}$ as follows:

If $l = 0$, the marginal probabilities for $\mathbf{u}[l, \bar{M}]$ are:

$$p_{u_i}(1) = \begin{cases} 1 - \dfrac{\bar{M} - K}{r_i \sum_{j=1}^{\bar{M}} \frac{1}{r_j}} & \text{if } 1 \leq i \leq \bar{M} \\ 0 & \text{if } \bar{M} + 1 \leq i \leq N \end{cases}$$

If $\eta(\bar{M}) + 1 \leq l \leq \eta(\bar{M} - 1) - 1$, the marginal probabilities for $\mathbf{u}[l, \bar{M}]$ are given by:

$$p_{u_i}(1) = \begin{cases} 1 - \dfrac{r_{\bar{M}+l}}{r_i} & \text{if } 1 \leq i \leq \bar{M} - 1 \\ 1 - \bar{M} + K + r_{\bar{M}+l} \sum_{j=1}^{\bar{M}-1} \frac{1}{r_j} & \text{if } i = \bar{M} \\ 0 & \text{if } \bar{M} + 1 \leq i \leq N \end{cases}$$

The parameters l and $\bar{M}$ essentially determine a randomized strategy. Note that from the definitions of M and η the strategy $\mathbf{u}[l, \bar{M}]$ has valid marginal probabilities for the range of l and $\bar{M}$ specified in Definition 1. For example, when $l > 0$ the condition $\eta(\bar{M}) + 1 \leq l \leq \eta(\bar{M} - 1) - 1$ is needed to ensure that $0 \leq p_{u_i}(1) \leq 1$ for all i.

We provide some intuition for considering the above class of strategies via an example. Suppose $M = N$ and consider strategy $\mathbf{u}[0, N]$. We see that the term $p_{u_j}(0)r_j$ is constant for all $1 \leq j \leq N$. From (8), this means that the worst-case difference for $\mathbf{u}$ only depends on the sum of the marginal probabilities of $\mathbf{x}$, but not the individual marginal probabilities. Thus, $\mathbf{u}[0, M]$ has constant performance for different $\mathbf{x}$. This type of constant performance with respect to $\mathbf{x}$ is a common trait of robust worst-case strategies [8]. Similar reasoning can be applied to $\mathbf{u}[l, \bar{M}]$ for other values of l and $\bar{M}$. For these strategies, the term $p_{u_j}(0)r_j$ is constant for some values of j.

This property is illustrated in Figure 1 through the following example. Suppose we have $N = 10$ channels, with rates given by $\mathbf{r} = [r_1, r_2, \cdots, r_9, r_{10}] = [1, 0.9, \cdots, 0.2, 0.1]$. If $K = K_0 = 6$, we can compute $M = 9$ and use Definition 1 to determine various strategies $\mathbf{u}[l, \bar{M}]$. In Figure 1 (TOP), we plot $p_{u_j}(0)r_j$ as j varies for strategies $\mathbf{u}[0, 7]$ and $\mathbf{u}[0, 9]$. Note that $\bar{M}$ essentially controls the region over which $p_{u_j}(0)r_j$ is constant, as well as its value in this region.

For fixed $\tilde{L}$, we know from (8) that the $\mathbf{x}$ which maximizes the regret satisfies $p_{x_j}(1) = 1$ for the $\tilde{L}$ highest values of $p_{u_j}(0)r_j$, and $p_{x_j}(1) = 0$ for other values. Thus, from Figure 1 we see a tradeoff between using $\mathbf{u}[0, 7]$ and $\mathbf{u}[0, 9]$. If $\tilde{L}$ is small, $\mathbf{u}[0, 9]$ will have better worst-case performance since its worst-case values of $p_{u_j}(0)r_j$ are smaller than that of $\mathbf{u}[0, 7]$. On the other hand, if $\tilde{L}$ is large, then $\mathbf{u}[0, 7]$ performs better in the worst-case since it has smaller $p_{u_j}(0)r_j$ for most values of j, e.g. $1 \leq j \leq 7$. A similar property holds for strategies $\mathbf{u}[l, \bar{M}]$ when $l > 0$, as shown in Figure 1 (BOTTOM). We see
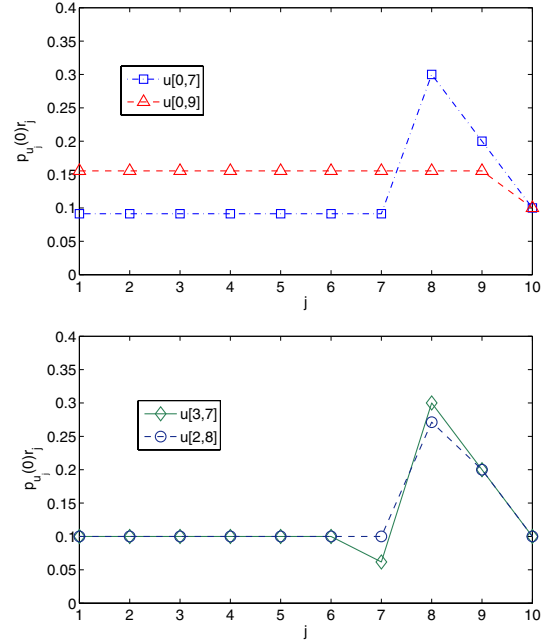

Fig. 1. For the scenario described in Section III, a plot of the term $p_{u_j}(0) \cdot r_j$ as a function of j, for the following strategies: (TOP): $\mathbf{u}[0, 7]$ and $\mathbf{u}[0, 9]$ (BOTTOM): $\mathbf{u}[3, 7]$ and $\mathbf{u}[2, 8]$.

again that l and $\bar{M}$ essentially control the region over which $p_{u_j}(0)r_j$ is constant.

In the next section, we will show the optimal strategy for (A) belongs to the class of strategies given by Definition 1.

IV. OPTIMAL STRATEGIES

In this section, we present optimal strategies for (A) and (B). We first focus on (A).

A. Optimal Strategies for (A)

As mentioned in Section III, (8) allows us to compute the optimal worst-case strategy by only considering the marginal probabilities for $\mathbf{u}$ and $\mathbf{x}$. Thus we first describe strategies in terms of their marginal probabilities, before presenting an algorithm for constructing such strategies.

For any fixed K, we let $W_K = \cup_{\bar{M}} \cup_l \mathbf{u}[l, \bar{M}]$ be the set of all strategies given by Definition 1. We now describe the optimal strategy for (A):

Theorem 1: *Optimal Strategy for* (A) For any set of channels, define M from (9). Then the optimal strategy $\mathbf{u} \in W_K$. That is, $\min_{\mathbf{u} \in U_K} \delta^{\mathbf{u}} = \min_{\mathbf{u} \in W_K} \delta^{\mathbf{u}}$. Moreover, the optimal strategy is determined as follows.

1) If there exists $K \leq M^* \leq M$ satisfying the following:

$$r_{M^*} \geq \frac{\tilde{L} - \eta(M^*)}{\sum_{j=1}^{M^*} \frac{1}{r_j}} \geq r_{M^*+1} . \qquad (10)$$

Then $\mathbf{u}[0, M^*]$ is the optimal strategy.

2) Otherwise, there must exist an M^*, where $K + 1 \leq M^* \leq M$, satisfying the following:

$$\tilde{L} - \eta(M^*) > r_{M^*} \sum_{j=1}^{M^*} \frac{1}{r_j} \geq \tilde{L} - \eta(M^* - 1) + 1 .$$

Then the optimal strategy is $\mathbf{u}[l^*, M^*]$, where:

$$l^* = \min \left\{ m \in \{0, \cdots, \tilde{L}\} : m \geq \tilde{L} - r_{M^*} \sum_{j=1}^{M^*} \frac{1}{r_j} \right\}$$

Theorem 1 is proved in Section IV-C. Also, the corresponding worst-case rewards (which are the minimum obtainable) are given in Lemma 3 in closed form. Note that Theorem 1 provides a method to determine the optimal strategy, and this procedure is a fairly simple one. In particular, cases 1 and 2 each require checking whether a condition holds, and there are only $M - K + 1$ possibilities for the value of M^*. The term $\eta(\cdot)$ is also not difficult to determine.

An alternative method for finding the optimal strategy is as follows. First, note that W_K contains at most $(M - K + 1)\tilde{L}$ strategies. For each $\mathbf{u}[l, \bar{M}] \in W_K$, the worst-case difference $\delta^{\mathbf{u}[l,\bar{M}]}$ can be easily determined, as we will show in Lemma 3 of Section IV-C. Thus combining this with $|W_K| \leq (M - K + 1)\tilde{L}$, the best strategy in W_K can be determined with low computational complexity.

An interesting special case of the above result is the following. When $K + \tilde{L} \leq M$, then from the definition of M, case 1 of Theorem 1 holds with $M^* = M$. Thus we have the following corollary:

Corollary 1: For any set of channels, define M from (9). Then if $K + \tilde{L} \leq M$, the optimal strategy $\mathbf{u}^*$ has the same marginal probabilities as strategy $\mathbf{u}[0, M]$ given by Definition 1.

Note that Corollary 1 and Theorem 1 describe the optimal marginal probabilities, but they do not immediately or uniquely determine an optimal strategy. By definition of U_K, at most K channels can be probed for any realization of $\mathbf{u}$. Thus, it cannot be possible that channel selection is done independently for each channel, and therefore we cannot simply multiply the marginal probabilities to obtain joint probabilities, or come up with a strategy by flipping independent coins to determine the inclusion/exclusion of each channel. Additionally, suppose that using these marginal probabilities one can construct the optimal strategy by appropriately defining the probability mass function of $\mathbf{u}^*$ over the set of all possible binary vectors of length N. This set contains 2^N elements, however, which grows very quickly as N increases.

Below we present an procedure that computes $\mathbf{u}[0, M]$ in a sequential manner without considering the entire set of binary vectors of length N. Consider the following algorithm which takes inputs N, K, K_0 and sequentially generates $\mathbf{u}[0, \bar{M}]$.

Algorithm 1: (*Sequential Method to Construct* $\mathbf{u}[0, \bar{M}]$)
For any N and K, define M as in (9) and consider any $K \leq \bar{M} \leq M$.

Initially: $\bar{M}, K$ are defined. Set $j = 1$ and $u_l = 0$ for all $1 \leq l \leq N$.

Step 1: Let $p = 1 - \frac{\bar{M} - K}{r_j \sum_{l=1}^{\bar{M}} \frac{1}{r_l}}$. Update K and $\mathbf{u}$ as follows:
- With probability p, set $u_j = 1$ and $K = K - 1$.
- Otherwise (with probability $1 - p$), do not change $\mathbf{u}$, K.

Step 2: Go to step 3 if $K = 0$. Otherwise, repeat Step 1 with $j = j + 1$, $\bar{M} = \bar{M} - 1$, and using the updated $\mathbf{u}$ and K.

Step 3: Stop and use strategy $\mathbf{u}$.

We note from Step 2 that Algorithm 1 constructs a realization of strategy $\mathbf{u}$ which, with probability 1, probes at most K channels. Thus, $\mathbf{u} \in U_K$. We have the following result which shows that Algorithm 1 constructs the optimal strategy for case 1 of Theorem 1.

Lemma 2: Algorithm 1 constructs a strategy with marginal probabilities which match the marginal probabilities of $\mathbf{u}[0, \bar{M}]$ given by Definition 1.

Therefore, by using $\bar{M} = M$ in Algorithm 1, we can construct the optimal strategy given in Corollary 1. In Section V, we will show that $\mathbf{u}[0, M]$ performs very well compared to the optimal strategy given by Theorem 1. Thus, strategy $\mathbf{u}[0, M]$ is a practical alternative if the optimal strategy of Theorem 1 is difficult to construct. Algorithm 1 also constructs strategy $\mathbf{u}[0, \bar{M}]$ for any $K \leq \bar{M} \leq M$; thus it generates the optimal strategy whenever case 1 of Theorem 1 holds. For $l > 0$, it is not immediately clear whether there exists a sequential algorithm to easily construct $\mathbf{u}[l, \bar{M}]$. It may be possible that constructing such a strategy requires defining a probability mass function over all vectors of length N.

B. Optimal Strategies for (B)

For Problem (B) we have the following result:

Theorem 2: Consider any set of channels and integers L, K. We have the following result for (B): $\max_{\mathbf{u} \in U_K} \rho^{\mathbf{u}} = K/N$, where the optimal strategy $\mathbf{u}^*$ achieving this maximum reward ratio has marginal probabilities given by: $p_{u_j^*}(1) = K/N$ for all $1 \leq j \leq N$.

We see from Theorem 2 there exists an optimal strategy which uniformly chooses from all length-N binary vectors with exactly K elements equal to 1. In Section V, we compare the performance of this uniform strategy to the optimal strategies of Theorem 1.

C. Deriving Optimal Strategies

To derive Theorem 1, we first prove the following preliminary results on the set of strategies given by Definition 1.

Lemma 3: For any set of channels, consider any strategy $\mathbf{u}[l, \bar{M}]$ given by Definition 1.

If $l = 0$,

$$\delta^{\mathbf{u}[0,\bar{M}]} = \sum_{j=1}^{\eta(\bar{M})} r_{\bar{M}+j} + \frac{\bar{M} - K}{\sum_{j=1}^{\bar{M}} \frac{1}{r_j}} \left(\tilde{L} - \eta(\bar{M}) \right) .$$

Otherwise, for $\eta(\bar{M}) + 1 \leq l \leq \eta(\bar{M} - 1) - 1$, then:

$$\delta^{\mathbf{u}[l,\bar{M}]} = \left(\bar{M} - K - r_{\bar{M}+l} \sum_{j=1}^{\bar{M}-1} \frac{1}{r_j} \right) r_{\bar{M}} + \sum_{k=1}^{l-1} r_{\bar{M}+k}$$
$$+ \left(\tilde{L} - l \right) r_{\bar{M}+l} .$$

Now we prove that these strategies are optimal. It can be shown (see [10]) that by interchanging maximum and minimum, the following is always true:

$$\min_{\mathbf{u} \in U_K} \max_{\mathbf{x} \in X_L} \{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\} \geq \max_{\mathbf{x} \in X_L} \min_{\mathbf{u} \in U_K} \{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\} \quad (11)$$

The inequality in (11) is useful as it allows us to derive lowerbounds to δ^* as follows. Note that for any particular $\mathbf{x}$, the righthand-side of (11) is an optimization problem with objective to minimize an average-reward criterion. By choosing $\mathbf{x} \in X_L$ such that the optimal average-reward is high, we can obtain a useful lowerbound to the minimax regret. We then show that these lower bounds match the quantities given in Lemma 3, thereby proving the optimality of $\mathbf{u}[0, M^*]$ for some M^*.

We first prove case 1 of Theorem 1. To derive a lowerbound for (11), we consider the following problem:

Problem 1: For any given L, K and set of channels, define M as in (9). Suppose case 1 of Theorem 1 holds for some M^*, and let state $\mathbf{x}$ have the following marginal probabilities. If $1 \leq i \leq M^*$:

$$p_{x_i}(1) = \frac{\tilde{L} - \eta(M^*)}{r_i \sum_{j=1}^{M^*} \frac{1}{r_j}} . \quad (12)$$

If $M^* + 1 \leq i \leq M^* + \eta(M^*)$, let $p_{x_i}(1) = 1$. Otherwise, for all other values of i, let $p_{x_i}(1) = 0$.

Solution: The optimal strategy for this problem will probe any K of the channels in set $\{1, \cdots, M^*\}$ (either randomly or deterministically), and satisfies: $\min_{\mathbf{u} \in U_K} \{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\} = \delta^{\mathbf{u}[0, M^*]}$

These results are proven in Appendix E. Combining the result of Problem 1 with (11), we see that $\mathbf{u}[0, M^*]$ must be the optimal worst-case strategy.

Now we show that when case 1 of Theorem 1 does not hold, then case 2 must be true. When case 1 of this theorem is not true for all $K \leq \bar{M} \leq M$, then for each $\bar{M}$ one of the following is true:

$$r_{\bar{M}} < \frac{\tilde{L} - \eta(\bar{M})}{\sum_{j=1}^{\bar{M}} \frac{1}{r_j}} \quad \text{or} \quad r_{\bar{M}+1} > \frac{\tilde{L} - \eta(\bar{M})}{\sum_{j=1}^{\bar{M}} \frac{1}{r_j}}$$

Combining this with $\eta(K) = \tilde{L}$ from the definition of η, it can be shown that case 2 must be satisfied for some M^*.

We note that for l^* and M^* described by case 2 of Theorem 1, it can be shown that $\eta(M^*) + 1 \leq l^* \leq \eta(M^* - 1) - 1$. Thus, $\mathbf{u}[l^*, M^*]$ is a valid strategy in Definition 1.

To prove that $\mathbf{u}[l^*, M^*]$ as described by case 2 is optimal, we consider the following problem.

Problem 2: For any given L, K and set of channels, define M as in (9). Suppose case 2 of Theorem 1 holds for some M^* and l. Let state $\mathbf{x}$ have the following marginal probabilities: For $1 \leq i \leq M^* - 1$, let $p_{x_i}(1) = \frac{r_{M^*}}{r_i}$. For $M^* \leq i \leq M^* + l - 1$, let $p_{x_i}(1) = 1$. When $i = M^* + l$, then let $p_{x_i}(1) = \tilde{L} - l - r_{M^*} \sum_{j=1}^{M^*-1} \frac{1}{r_j}$. Otherwise, for all other values of i, let $p_{x_i}(1) = 0$.

Solution: As shown in Appendix F, the optimal strategy for this problem probes any K of the channels in set $\{1, \cdots, M^*\}$ and satisfies: $\min_{\mathbf{u} \in U_K} \{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\} = \delta^{\mathbf{u}[l, M^*]}$.

Combining this with (11) proves that for case 2 of Theorem 1, $\mathbf{u}[l, M^*]$ must be the optimal strategy.

This concludes the proof of Theorem 1, since we have shown that either case 1 or 2 must hold, and we have derived the optimal strategy for each case.

V. NUMERICAL RESULTS

In this section, we examine the performance of the proposed algorithms under a range of system parameters. We first compare the performance between the optimal strategies under the different metrics described in Section II. Then we compare the worst-case strategies to other heuristics or algorithms.

A. Comparison of Metrics

For the numerical results, we chose parameters r_j as follows. For a given N, the r_j were uniformly distributed in $[0, 1]$. For each realization of the rates $\{r_j\}_{j \in S}$, we computed the average reward obtained by the following strategies:

1) genie, or omniscient observer: knows in advance the $\tilde{L}$ available channels with highest bandwidth.
2) $\mathbf{u}^*$: the optimal strategy within the class given by Definition 1, and shown to be optimal in Theorem 1. We used Lemma 3 to determine the best worst-case strategy among this class for a given L, K, N.
3) $\mathbf{u}[0, M]$: the optimal strategy under the conditions given by Corollary 1
4) uniform: the strategy which uniformly picks K out of the N channels to probe, regardless of their rate/bandwidth, shown to be optimal for (B) in Theorem 2.

Figure 2 (TOP) shows the performance of these two strategies when we fix $K_0 = K = L = 3$ and N varies from 5 to 10. In order to achieve $L = 3$, we set $p_{x_j}(1) = 3/N$ for all $1 \leq j \leq N$. We see that $\mathbf{u}^*$ outperforms $\mathbf{u}[0, M]$ and uniform for most N. Note that for large N, $\mathbf{u}[0, M]$ slightly outperforms $\mathbf{u}^*$. This is because even though $\mathbf{u}^*$ may have a lower worst-case difference (bound over all $\mathbf{x}$), it does not necessarily outperform $\mathbf{u}[0, M]$ for some given $\mathbf{x}$. In general, the performance of these two strategies is similar for most N, especially as N increases. The latter can be explained by the fact that as N (and M) increase while K stays fixed, it becomes more likely that the conditions of Corollary 1 are satisfied. Hence, as N increases then it is more likely that $\mathbf{u}^*$ and $\mathbf{u}[0, M]$ are the same strategy. The performance of the genie varies for each realization of the $\{r_j\}$ and $\mathbf{x}$, but its average performance is relatively constant with respect to N since $L = 3$ does not change and the genie can only use the $K_0 = 3$ available channels with highest rate.

Figure 2 (BOTTOM) shows the performance of these strategies when $N = 10$, $L = 5$ and $K = K_0$ increases from 3 to 9. As expected, the average rewards of $\mathbf{u}^*$, $\mathbf{u}[0, M]$, and uniform all approach the genie's average reward as K increases. This holds because it becomes more likely that these strategies will

Fig. 3. Comparison of optimal average-reward strategies $\mathbf{v}_0$ and $\mathbf{v}_2$ to the optimal worst-case strategies $\mathbf{u}^*$, $\mathbf{u}[0, M]$ and uniform as the channel state distribution is varied, as described in Section V-B.
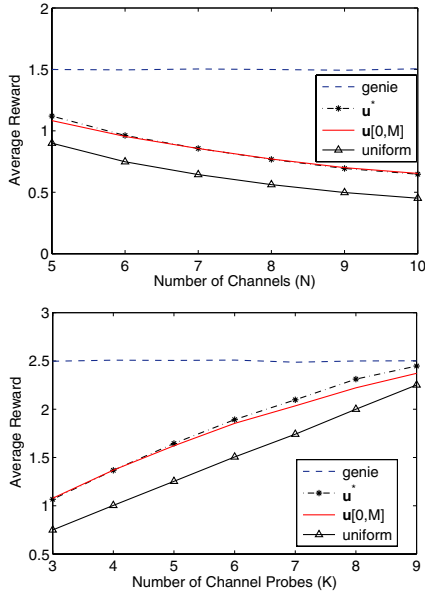
Fig. 2. (TOP): Average performance of the genie, algorithms $\mathbf{u}^*$, $\mathbf{u}[0, M]$, and uniform as described in Section V, when $K = L = 3$ and N varies from 5 to 10 (BOTTOM): When $N = 10$ and $L = 5$, average performance of these strategies as K increases from 3 to 9.

probe the best channels as K increases. On the other hand, we also see that $\mathbf{u}^*$ strictly outperforms uniform for all K.

These results indicate that the optimal strategies under the worst-case difference performance measure considered in Section II-B appear to outperform the optimal algorithm (uniform) arising from the worst-case ratio performance measure. In addition, if the procedure of determining $\mathbf{u}^*$ becomes impractical, then one can instead consider the strategy $\mathbf{u}[0, M]$ which has comparable performance for the scenario described in this section but is very easy to construct using the proposed Algorithm 1.

B. Sensitivity Analysis of Strategies

The optimal strategies in Section IV are derived under worst-case performance measures. If the transmitter knows the distribution of $\mathbf{x}$, i.e. channel states, then it can probe and use the K channels with highest values of $p_{x_i}(1)r_i$ in order to maximize $V_{\mathbf{x}}^{\mathbf{u}}$ (we call this the optimal average-reward strategy) and outperform the optimal worst-case strategies. On the other hand, if this knowledge of the distribution is incorrect then the optimal average-reward strategy may have a poor performance. The optimal worst-case strategies, on the other hand, are guaranteed to perform well over all $\mathbf{x}$ and therefore could be more robust to changes in $\mathbf{x}$.

To illustrate this, we compare various strategies for the following scenario, where $N = 10$ and $L = K = K_0 = 5$. For any given set of rewards $\{r_j\}$, $P_{x_j}(1) = j^\alpha/\beta$ for all $1 \le j \le N$ and some $-5 \le \alpha \le 5$. Here, β is a normalizing constant to ensure that $\mathbf{x} \in X_L$. Note that when $\alpha = 0$, then $\mathbf{x}$ takes on the uniform distribution. Negative (positive) α means the channels with higher (lower) rate have a higher probability
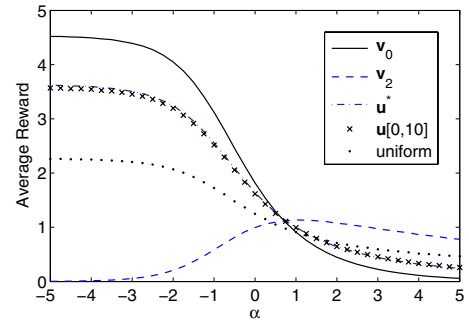
of being available. We let $\mathbf{v}_0$ and $\mathbf{v}_2$ to be the optimal average-reward strategies when the transmitter *believes* $\alpha = 0$, 2, respectively. Note that strategy $\mathbf{v}_0$ (or $\mathbf{v}_2$) probes the K channels with the highest values of r_j (or $r_j j^2/\beta$). Thus, these strategies can be seen as simple heuristics which probe based on weighted values of the channel rates.

As shown in Figure 3, we examined the performance of these strategies by varying the real value of α. For a given α, the $\{r_j\}$ were uniformly distributed in $[0, 1]$ and the average performance of $\mathbf{v}_0$, $\mathbf{v}_2$, $\mathbf{u}^*$, $\mathbf{u}[0, M]$, and uniform was determined for 10^4 realizations. This process was repeated for $-5 \le \alpha \le 5$. Note that performance of $\mathbf{v}_0$ and $\mathbf{v}_2$ under various α indicates how it is affected when there are errors in the transmitter's belief of the true α.

We see that for negative (positive) α, $\mathbf{v}_0$ ($\mathbf{v}_2$) performs very well. On the other hand, we see that as α increases (decreases), the expected reward of $\mathbf{v}_0$ ($\mathbf{v}_2$) approaches zero. This is because $\mathbf{v}_0$ probes the channels with the K highest rewards, thus it performs better when the channels with higher rate have a higher probability of being available. Similar reasoning can be applied for $\mathbf{v}_2$. On the other hand, $\mathbf{u}^*$ and $\mathbf{u}[0, M]$ are relatively robust to changes in the value of α, as their total reward does not approach 0 if α varies. This example illustrates the robustness of worst-case strategies compared to optimal average-reward strategies under varying distributions.

VI. CONCLUSION

In this paper, we examined optimal competitive algorithms for joint channel probing and transmission. We formulated multiple worst-case performance measures and derived a class of optimal strategies. We presented an algorithm which sequentially generates a subclass of these strategies with low computational complexity. The performance of these strategies were also examined via numerical studies. These results and algorithms are applicable to many practical scenarios, particularly when the channel quality is changing unpredictably and the transmitter cannot model it stochastically.

REFERENCES

[1] J. Kennedy and M. Sullivan, "Direction finding and "smart antennas" using software radio architectures," *IEEE Communications Magazine*, pp. 62–68, May 1995.

[2] Y. Chen, Q. Zhao, and A. Swami, "Joint design and separation principle for opportunistic spectrum access," *Proceedings of IEEE Asilomar Conference on Signals, Systems, and Computers*, November 2006.

[3] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive mac protocol for multi-hop wireless networks," *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2001, Rome, Italy.

[4] G. Ganesan and Y. Li, "Cooperative spectrum sensing in cognitive radio networks," *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pp. 137–143, November 2005.

[5] A. Sabharwal, A. Khoshnevis, and E. Knightly, "Opportunistic spectral usage: Bounds and a multi-band CSMA/CA protocol," *IEEE/ACM Transactions on Networking*, pp. 533–545, June 2007.

[6] S. Guha, K. Munagala, and S. Sarkar, "Approximation schemes for information acquisition and exploitation in multichannel wireless networks," *Proceedings of 44th Annual Allerton Conference on Communication, Control and Computing*, September 2006, Monticello, IL.

[7] N. Chang and M. Liu, "Optimal channel probing and transmission scheduling for opportunistic spectrum access," *Proceedings of the Thirteenth Annual International Conference on Mobile Computing and Networking (MobiCom)*, September 2007, Montreal, CN.

[8] A. Borodin and R. El-Yaniv, *Online Computation and Competitive Analysis*. Cambridge University Press, 1998, Cambridge, UK.

[9] D. Foster and R. Vohra, "Regret in the on-line decision problem," *Games and Economic Behavior*, vol. 29, pp. 1084–1090, 1999.

[10] S. Nash and A. Sofer, *Linear and Nonlinear Programming*. McGraw-Hill, 1996, New York, USA.

APPENDIX

A. Proof of Lemma 1

For notation, let $\delta_{\mathbf{x}}^{\mathbf{u}} = V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}$ for any strategy $\mathbf{u}$. First note that because $\tilde{L} \leq L$, we have $X_{\tilde{L}} \subseteq X_L$. Thus, $\max_{\mathbf{x} \in X_L} \{\delta_{\mathbf{x}}^{\mathbf{u}}\} \geq \max_{\mathbf{x} \in X_{\tilde{L}}} \{\delta_{\mathbf{x}}^{\mathbf{u}}\}$. Thus to complete the proof, we only need to show that the righthandside of this inequality is less than or equal to the lefthandside. To prove this, we show that for any $\mathbf{x} \in X_L$ there exists corresponding $\mathbf{y} \in X_{\tilde{L}}$ such that $V_{\mathbf{y}}^* - V_{\mathbf{y}}^{\mathbf{u}} \geq V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}$. Consider any $\mathbf{x} \in X_L$. Note from the equation for $V_{\mathbf{x}}^*$ that for any realization $\mathbf{x}(\omega)$, where $\mathbf{x} = [x_1(\omega), \cdots, x_N(\omega)]$, only the $\min\{K_0, |\mathbf{x}(\omega)|\}$ available channels with highest rate determine the expected reward of the genie. We can define a new channel state $\mathbf{y}$ such that $y_j(\omega) = 1$ if and only if $x_j(\omega) = 1$ and j is one of the $\min\{K_0, |\mathbf{x}(\omega)|\}$ available channels with the highest rate.

We see that because $|\mathbf{y}(\omega)| \leq \tilde{L}$ for all realizations ω, then $\mathbf{y} \in X_{\tilde{L}}$ and $V_{\mathbf{x}}^* = V_{\mathbf{y}}^*$. At the same time, clearly $V_{\mathbf{x}}^{\mathbf{u}} \geq V_{\mathbf{y}}^{\mathbf{u}}$ for any strategy $\mathbf{u}$. Thus, $V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}} \leq V_{\mathbf{y}}^* - V_{\mathbf{y}}^{\mathbf{u}}$. ∎

B. Proof of Lemma 2

Note that since Algorithm 1 randomly chooses which channels to use, then the value for K after each step is a random variable. Let $\bar{K}_i$ denote the value for K after i iterations of Step 1. The value of $\bar{M}$ also changes after each iteration. We let $\bar{K}_0$ and M_0 denote the initial (fixed) values for K and $\bar{M}$, respectively. Thus we need to show that any strategy generated by Algorithm 1 has the same marginal probabilities as $\mathbf{u}[0, M_0]$ of Definition 1, with $\bar{K}_0$ replacing K. We first show this holds for $1 \leq i \leq M_0$ by induction.

Induction Basis: As described in Algorithm 1, initially $p = 1 - \frac{M_0 - \bar{K}_0}{r_1 \sum_{j=1}^{M_0} \frac{1}{r_j}}$ and we set $u_1 = 1$ with probability p. Therefore, the result holds for $i = 1$.

Induction Hypothesis: First let's consider any $2 \leq l \leq M_0$. Suppose the result holds for $i = 1, \cdots, l-1$. We will show it also holds for $i = l$. The probability that we use channel l is a random variable given by: $p = 1 - [(\bar{M} - i + 1) - \bar{K}_{i-1}]/(r_i \sum_{j=i}^{M_0} \frac{1}{r_j})$. Taking the expectation of p gives the marginal probability that channel l is used:

$$p_{u_l}(1) = 1 - \frac{(M_0 - l + 1) - E[\bar{K}_{l-1}]}{r_l \sum_{j=l}^{M_0} \frac{1}{r_j}} \tag{13}$$

We note that $E[\bar{K}_{l-1}]$ can be calculated as follows by considering Step 2 of Algorithm 1:

$$E[\bar{K}_{l-1}] = \bar{K}_0 - \sum_{j=1}^{l-1} p_{u_j}(1) = \bar{K}_0 - \sum_{j=1}^{l-1} \left[1 - \frac{M_0 - \bar{K}_0}{r_j \sum_{l=1}^{M_0} \frac{1}{r_l}} \right],$$

where we have used the induction hypothesis which gave us $p_{u_j}(1)$ for all $j \leq l-1$. Plugging this result back into (13) and rearranging, we obtain: $p_{u_l}(1) = 1 - [M_0 - \bar{K}_0]/[r_l \sum_{j=1}^{M_0} \frac{1}{r_j}]$, which proves that the result holds.

Now consider any $M_0 + 1 \leq l \leq N$. We need to prove that $p_{u_l}(1) = 0$. By step 3 of the algorithm, this will always be true if $P(\bar{K}_{M_0} = 0) = 1$. We prove this as follows. Initially, $\bar{K}_0 \leq M_0$, and whenever $\bar{K}_{j-1} = M_0 - j + 1$ then from Step 1 we have $p = 1$ which means that in Step 2 we will obtain $\bar{K}_j = \bar{K}_{j-1} - 1 = M_0 - (j+1) + 1 = M_0 - j$. Thus, it is impossible to have $\bar{K}_j > M_0 - j$ for any j. Therefore, $\bar{K}_{M_0} \leq M_0 - M_0 = 0$ with probability 1, which proves the result. ∎

C. Proof of Theorem 2

First let's consider any strategy $\mathbf{u}^*$ with marginal probabilities described in Theorem 2. We have by using (6):

$$\rho^{\mathbf{u}^*} = \min_{\mathbf{x} \in X_L} \left\{ \frac{\frac{K}{N} \sum_{j=1}^{N} r_j p_{x_j}(1)}{\sum_{j=1}^{N} r_j p_{x_j}(1)} \right\} = \min_{\mathbf{x} \in X_L} \left\{ \frac{K}{N} \right\} = \frac{K}{N}$$

Now consider any other strategy $\mathbf{u}$ with marginal probabilities that differ from those given by Theorem 2, i.e. strategies which do not probe channels uniformly. This implies there exists $1 \leq j \leq N$ such that $p_{u_j}(1) \neq K/N$. On the other hand, we know from definition of U_K in (2) that the sum of the marginal probabilities is at most K. Thus, there exists $1 \leq i \leq N$ such that $p_{u_i}(1) < \frac{K}{N}$. Then we obtain the following for this i, by setting $p_{x_i}(1) = 1$ and $p_{x_j}(1) = 0$ for all $j \neq i$:

$$\rho^{\mathbf{u}} = \min_{\mathbf{x} \in X_L} \left\{ \frac{\sum_{j=1}^{N} r_j p_{x_j}(1) p_{u_j}(1)}{\sum_{j=1}^{N} r_j p_{x_j}(1)} \right\} \leq \frac{r_i p_{u_i}(1)}{r_i} < \frac{K}{N} .$$

Thus we see that $\rho^{\mathbf{u}^*} \geq \rho^{\mathbf{u}}$ for all $\mathbf{u} \in U_K$, and therefore $\mathbf{u}^*$ must be optimal. ∎

D. Proof of Lemma 3

First consider any strategy $\mathbf{u}[0, \bar{M}]$. We know from (8) that $\delta^{\mathbf{u}}$ is given by the following:

$$\delta^{\mathbf{u}} = \max_{\mathbf{x} \in X_{\tilde{L}}} \left\{ \frac{(\bar{M} - K)}{\sum_{j=1}^{\bar{M}} \frac{1}{r_j}} \sum_{i=1}^{\bar{M}} p_{x_i}(1) + \sum_{i=\bar{M}+1}^{N} r_i \cdot p_{x_i}(1) \right\} .$$

Meanwhile, by definition of $\eta(\bar{M})$, we have the following for all $1 \leq i \leq \eta(\bar{M}) < k$, where i, k are integers:

$$r_{\bar{M}+i} \sum_{j=1}^{\bar{M}} \frac{1}{r_j} \geq \bar{M} - K > r_{\bar{M}+k} \sum_{j=1}^{\bar{M}} \frac{1}{r_j}. \tag{14}$$

From these inequalities, we see there exists state $\mathbf{x}$ which achieves the maximum in $\delta^{\mathbf{u}}$ by setting $p_{x_{\bar{M}+j}}(1) = 0$ for $j \geq \eta(\bar{M}) + 1$, and $p_{x_{\bar{M}+j}}(1) = 1$ for $1 \leq j \leq \eta(\bar{M})$, and $p_{x_j}(1) = 1$ for any $\tilde{L} - \eta(\bar{M})$ values of j in $1 \leq j \leq \bar{M}$. Note that this is a valid $\mathbf{x}$ as $\eta(\bar{M}) \leq \tilde{L}$ by definition of $\eta(\bar{M})$. Thus $\delta^{\mathbf{u}}$ reduces to the equation given in Lemma 3.

Now consider $\mathbf{u}[l, \bar{M}]$ when $l > 0$. In this case $\delta^{\mathbf{u}}$ becomes

$$\delta^{\mathbf{u}} = \max_{\mathbf{x} \in X_{\tilde{L}}} \left\{ \sum_{i=1}^{\bar{M}-1} r_i p_{x_i}(1) \frac{r_{\bar{M}+l}}{r_i} + \sum_{i=\bar{M}+1}^{N} r_i p_{x_i}(1) \right.$$
$$\left. + r_{\bar{M}} p_{x_{\bar{M}}}(1) \left(\bar{M} - K - r_{\bar{M}+l} \sum_{j=1}^{\bar{M}-1} \frac{1}{r_j} \right) \right\} \tag{15}$$

By the definition of $\eta(\bar{M} - 1)$, we have the following: $r_{\bar{M}+\eta(\bar{M}-1)+1} \sum_{j=1}^{\bar{M}-1} \frac{1}{r_j} < \bar{M} - 1 - K$. Combining this with $l \geq \eta(\bar{M} - 1)$ and using $r_j \leq r_k$ for any $j \geq k$, we obtain the inequality: $r_{\bar{M}-1+l+1} \left(\sum_{j=1}^{\bar{M}} \frac{1}{r_j} \right) < (\bar{M} - K)$. Rearranging yields: $r_{\bar{M}+l} < r_{\bar{M}} \left(\bar{M} - K - \sum_{j=1}^{\bar{M}-1} (r_{\bar{M}+l}/r_j) \right)$.

Thus we see that (15) can be maximized by setting $p_{x_{\bar{M}+k}}(1) = 1$ for $0 \leq k \leq \tilde{l} - 1$, and $p_{x_j}(1) = 1$ for any $\tilde{L} - \tilde{l}$ values of j in $1 \leq j \leq \bar{M} - 1$. Plugging this maximization into (15) completes the proof. ∎

E. Optimal Strategy for Problem 1

We first note from the definition of M^* from case 1 of Theorem 1 that $0 \leq p_{x_j}(k) \leq 1$, i.e. the marginal probabilities for $\mathbf{x}$ are valid. Since $V_{\mathbf{x}}^*$ does not depend on $\mathbf{u}$, then we have the following:

$$\min_{\mathbf{u} \in U_K} \{V_{\mathbf{x}}^* - V_{\mathbf{x}}^{\mathbf{u}}\} = V_{\mathbf{x}}^* - \max_{\mathbf{u} \in U_K} V_{\mathbf{x}}^{\mathbf{u}} .$$

Therefore, the optimal strategy for this problem is the one which maximizes $V_{\mathbf{x}}^{\mathbf{u}}$.

Plugging the marginal probabilities for $\mathbf{x}$ into the defining equation for $V_{\mathbf{x}}^{\mathbf{u}}$, we obtain:

$$V_{\mathbf{x}}^{\mathbf{u}} = \sum_{j=1}^{M^*} p_{u_j}(1) \frac{\tilde{L} - \eta(M^*)}{\sum_{j=1}^{M} \frac{1}{r_j}} + \sum_{j=1}^{\eta(M^*)} p_{u_j}(1) r_{M^*+j}.$$

We see from the definition of M^* in case 1 of Theorem 1 that there exists a strategy which maximizes $V_{\mathbf{x}}^{\mathbf{u}}$ by setting $p_{u_j}(1) = 0$ for all $M + 1 \leq j \leq N$ and probes any K of the channels in $\{1, \cdots, M^*\}$. Thus, the reward of using optimal strategy $\mathbf{u}^*$ is given by: $V_{\mathbf{x}}^{\mathbf{u}^*} = [\tilde{L} - \eta(M^*)] K / (\sum_{j=1}^{M^*} \frac{1}{r_j})$, which holds regardless of which K of the channels are probed.

Meanwhile, we have:

$$V_{\mathbf{x}}^* = \sum_{l=1}^{N} r_l p_{x_l}(1) = \frac{\tilde{L} - \eta(M^*)}{\sum_{j=1}^{M^*} \frac{1}{r_j}} M^* .$$

Taking the difference between $V_{\mathbf{x}}^*$ and $V_{\mathbf{x}}^{\mathbf{u}^*}$, and comparing with $\delta^{\mathbf{u}[0,M^*]}$ given by Lemma 3, we obtain the result. ∎

F. Optimal Strategy for Problem 2

We follow a similar technique used in Appendix E for Problem 1. First, the optimal strategy maximizes $V_{\mathbf{x}}^{\mathbf{u}}$, which can be written as follows:

$$V_{\mathbf{x}}^{\mathbf{u}} = \sum_{j=1}^{M^*} \frac{r_{M^*}}{r_j} p_{u_j}(1) r_j + \sum_{j=1}^{l-1} p_{u_{M+j}}(1) r_{M+j}$$
$$+ \left(\tilde{L} - l - r_{M^*} \sum_{j=1}^{M^*-1} \frac{1}{r_j} \right) p_{u_{M^*+l}}(1) r_{M^*+l},$$

We see that because $r_j > r_k$ for all $j < k$, then $V_{\mathbf{x}}^{\mathbf{u}}$ is maximized by probing any K of the channels in $\{0, \cdots, M^*\}$. Thus its reward is simply $r_{M^*} K$.

Meanwhile, from (4) we have:

$$V_{\mathbf{x}}^* = r_{M^*} M^* + \sum_{j=1}^{l-1} r_{M+j} + \left(\tilde{L} - l - r_{M^*} \sum_{j=1}^{M^*-1} \frac{1}{r_j} \right) r_{M^*+l} .$$

Taking the difference between this and the reward of the optimal strategy, and comparing with $\delta^{\mathbf{u}[l,M^*]}$ given by Lemma 3, proves the result.