

Jointly optimal transmission and probing strategies for multichannel wireless systems

(Invited Paper)

Sudipto Guha*, Kamesh Munagala[†], and Saswati Sarkar[‡]

* Dept. of Computer and Information Sciences, UPenn, Philadelphia, PA 19104. Email: sudipto@cis.upenn.edu

[†] Dept. of Computer Science, Duke University, Durham, NC 27708. Email: kamesh@cs.duke.edu

[‡] Dept. of Electrical and Systems Engineering, UPenn, Philadelphia, PA 19104. Email: swati@seas.upenn.edu

Abstract—We consider a wireless system with multiple channels when each channel has several different transmission states. Different states are associated with different probabilities of successful transmissions. In such networks, we are faced with making transmission decisions in the presence of partial information about channel states. This (typically probabilistic) information about any channel can be refined by sending control packets in the channels. In presence of multiple alternative channels, this process of probing every channel to find the best one is onerous and resource consuming. There is a natural tradeoff between the resource consumed in probing and the estimate of channel state we can obtain. The desired tradeoff can be attained by judiciously determining which and how many channels to probe and also which channel to transmit. We present adaptive algorithms for provably approximating the desired tradeoffs within constant factors.

I. INTRODUCTION

Future wireless networks are likely to provide each node access to a large number of channels. A channel can for example be a frequency in a frequency division multiple access (FDMA) network, or a code in a code division multiple access (CDMA) network, or an antenna or a polarization state (vertical or horizontal) of an antenna in a device with multiple antennas (MIMO). Several existing wireless technologies, e.g., IEEE 802.11a [1], IEEE802.11b [8], IEEE802.11h [2] propose to use multiple frequencies. For example, IEEE 802.11a protocol has 8 channels for indoor use and 4 channels for outdoor use in the 5GHz band, while the IEEE 802.11b protocol has 3 channels in the 2.4GHz band. The potential deregulation of the wireless spectrum is likely to enable the use of a significantly larger number of frequencies. Due to significant advances in device technology, laptops with multiple antennas (antenna arrays) incorporated in the front lid, and devices with smart antennas have already been developed, and the number of such antennas are likely to significantly increase in near future. This increase in the number of channels is expected to significantly enhance network capacity and enable several new bandwidth-intensive applications as multiple transmissions can now proceed simultaneously in a vicinity using different channels, and the probability (at any given time) of existence of at least one channel with acceptable transmission quality significantly increases.

The main challenge, however, in exploiting multiple channels is that a node has only limited information about the trans-

mission quality of the individual channels which stochastically vary with time. Presumably, a node's transmission decisions will become closer to optimal as the available information about its channels increases. However, the bandwidth and the energy expended in acquiring such information also increases with the amount of information acquired. Note that a node usually probes in a channel by transmitting a control packet in the channel, and the receiver informs the sender about the quality of the channel in a response packet (e.g., the RTS and CTS packet exchange in IEEE 802.11). The exchange of control packets consumes additional energy, and prevents other neighboring users from simultaneously utilizing the channel. Thus, each probe is associated with a cost.

Owing to the probing costs, the amount of information a node acquires about its channels becomes an important decision variable. Before each transmission, a node needs to determine how many and which channels it will probe and also the sequence in which these channels will be probed (*probing policy*). Note that depending on the available hardware (e.g., availability, or lack thereof, of multiple network interface cards, or compatible transmission circuits to appropriately distribute the power across the antennas), a node may, or may not, be able to simultaneously transmit in multiple channels. In this paper, we consider the scenario where a node can transmit in only one channel in a time slot and transmits one packet in each slot. Based on the outcomes of the probes, a node must select one of the available channels (*channel selection policy*), which need not be those that it has probed.

An important performance goal in such networks is to design a jointly optimal probing and channel selection policy that maximizes a system utility which is the difference between the probability of successful transmission and the expected probing cost before each transmission. Loosely, this utility function represents the “gain” or the “profit” of the sender if the sender receives credit from the receiver for each packet it delivers successfully and needs to additionally compensate the wireless provider for each probe packet it transmits¹.

We first enumerate the challenges in designing the optimal policy. We consider a single node with access to n channels. The optimal policy needs to probe adaptively, i.e., the result

¹The sender may have to share with the provider part of the credit it receives from the receiver for each successfully delivered packet; then the credit we are considering here is the credit remaining after the sharing process.

of a probe determines the channels to be probed subsequently. For example, consider channels with 3 possible states (0, 1, 2), each of which is associated with a different transmission quality. Clearly, the probing terminates if a probed channel is in the highest state. Now, let a probed channel be in the intermediate state (state 1). Then the subsequent probes should be limited to channels that have high probabilities of being in the highest state. However, if all channels that have been probed in a slot are in the lowest state, then the channels that have high probabilities of being in the intermediate state may also be subsequently probed. Furthermore, the channel selection decision depends on the outcome of the probes and also the expectation and uncertainty of the transmission quality of the channels that have not been probed. The optimal policy is therefore a decision tree over n variables. The time to compute the optimum decision tree using a naive optimization which evaluates all the decisions trees over n variables is therefore clearly exponential in the problem size. Next, the space required to store the optimum tree will also be exponential in the problem space as this requires storage of decisions associated with all the branches.

In a companion paper [10], we have showed that for two state channels the optimum policy can be computed and stored in polynomial complexity. In this paper, we show that for an arbitrary number of states the optimal net gain can be approximated within a factor of $\frac{1}{2}$ using a simple approximation algorithm (Section IV), and when the number of states is 3 the approximation ratio can be improved to $\frac{2}{3}$ (Section V). The computation and storage complexities of our policies are polynomial in the number of channels. We review the related work in Section II and define the system model in Section III.

II. RELATED WORK

Opportunistic selection of channels with complete knowledge of channel states has been comprehensively investigated over the last decade (e.g., [18]). The joint optimization of the reward obtained from informed selections and the cost incurred in acquiring the required information however remains largely unexplored. Recently, Kanodia *et. al.* [13] and Ji *et. al.* [12] considered scenarios where a node probes multiple channels and selects a channel based on the outcomes of the probes. They consider only statistically identical channels with equal probing costs and assume that a node can transmit in only a channel that has been probed. Thus, their problem reduces to a decision of how many channels to probe which is equivalent to that of the well-investigated optimal stopping time problem [5]. Optimizing the order of evaluation of random variables so as to minimize the cost of evaluation (“pipelined filters”) has been investigated in several different contexts like diagnostic tests in fault detection and medical diagnosis, optimizing conjunctive query and joint ordering in data-stream systems, web services, and sensor networks [6], [14], [11], [7], [15], [3], [16], [17], [4]. However our work is different from all the above (including [13], [12]) in that, we allow a node to transmit in a channel even if the channel has not been

probed. Furthermore, we allow for channels with different distributions of the transmission qualities and different probing costs which is not considered in [13], [12], and consider multi-state channel models which pipeline filters seldom consider. These complications significantly alter the decision issues and the optimal solutions.

III. SYSTEM MODEL AND PROBLEM DEFINITION

A sender U has access to n channels which are denoted as channels $1, 2, \dots, n$, each of which has K possible states, $0, \dots, K - 1$. We assume that time is slotted. In any slot channel j is in state i with probability p_{ij} independent of its state in other slots and the states of other channels in any slot. In any slot, U transmits a data packet in one channel, and if the channel is in state j , the transmission is successful with probability r_j . Thus, r_j is the reward associated with state j . Without loss of generality we assume $0 \leq r_0 < r_1 < \dots < r_{K-1}$. For simplicity, we also assume that $r_0 = 0$; all analytical results can however be generalized to the scenario where $r_0 > 0$. Whenever U probes a channel i , it pays a cost of $c_i \geq 0$. Probing different channels may incur different costs as the probing process for different channels may interfere with the channel access of different number of users (based on geometry and allocation of channels). We now formally define the policies and the performance metrics.

Definition 3.1: A *probing policy* is a rule that, given the set of channels the sender has already probed in a slot (which would be empty at the beginning of the slot) and the states of the channels probed in the slot, determines (a) whether the sender should probe any more channels and (b) if the sender probes additional channels which channel it should probe next. The sender knows the state of a channel in a slot if and only if it probes the channel in the slot.

Definition 3.2: A *selection policy* is a rule that selects a channel for the transmission of a data packet in a slot on the basis of the states of the probed channels, after the completion of the probing process in the slot. The selection policy can select a channel even if it has not been probed in the slot, and in that case, the channel is referred to as a *backup* channel.

Definition 3.3: The *probing cost* is the sum of the costs of all channels probed in the slot. The probing cost is clearly a random variable that depends on the probing policy and the outcomes of the probes (as the sender may probe subsequent channels depending on the outcomes of the previous probes). The *expected probing cost* is the expectation of this random variable and depends on both the probing policy and the channel statistics.

Definition 3.4: In any slot, the *transmission reward* is 1 if the packet is successfully transmitted in the slot and 0 otherwise. Again, the *transmission reward* in any slot is a random variable that depends on the probing and selection policies and the states of the channels in the slot. The *expected transmission reward* depends on the probing and selection policies and the channel statistics.

Definition 3.5: The *expected net gain* of the sender, denoted simply as *gain*, is the difference between the expected transmission reward and expected probing cost. This depends on the probing and selection policies and the channel statistics.

Problem Definition: Given $\{c_i\}, \{r_i\}$ and $\{p_{ij}\}$ find a probing and selection policy so as to maximize the expected gain for independent channels. Let *OPT* denote the denote gain of the optimal policy.

Since we are considering the independent channel model, the optimal probing and selection policies in a slot need not depend on the decisions and the observations in other slots. Also, the optimal policies remain the same in all slots, though the specific choices made by each policy may be different in different slots depending on the outcome of the probes.

In [10], we showed that the optimal probing and channel selection policy for two state channels can be computed in polynomial time. In this paper we consider K state channels when $K \geq 3$.

IV. OPTIMAL POLICIES WHEN $K \geq 3$

We first show that the optimal policy in the class of policies that does not transmit on an unprobed (backup) channel can be computed in polynomial time (Subsection IV-A). We then obtain a policy that may transmit in a backup channel, but is guaranteed to attain at least $\frac{1}{2}$ the maximum gain in the class of all policies that may or may not use backup (Subsection IV-B).

A. Optimal Algorithm without Backups

We present an optimal polynomial time algorithm for multi-state channels when no backup channel is allowed to be used. We first introduce the following definition.

Definition 4.1: Define $\tilde{r}_i[u] = \frac{\sum_{v:u \leq v} p_{vi} r_{vi}}{\sum_{v:u \leq v} p_{vi}}$ and $\tilde{p}_i[u] = \frac{\sum_{v:u \leq v} p_{vi}}{\sum_{v:u \leq v} p_{vi}}$. Define $H_u = \Phi$ for all $u > K$. Recursively, starting from H_K , define $H_u = \left\{ i \mid i \notin \bigcup_{v:u > v} H_v \text{ and } \tilde{r}_i[u] - \frac{c_i}{\tilde{p}_i[u]} > r_{u-1} \right\}$. Assume $c_i/\tilde{p}_i[u] = +\infty$ when $\tilde{p}_i[u] = 0$.

OPTNOBKUP

Consider each H_u in decreasing order of u starting from $u = K$.
Within each H_u probe in non-increasing order of $\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]}$, and stop if any channel is found to be in state u or above.
Select the channel which is in the highest state among all probed channels.

We now present the intuition behind OPTNOBKUP. Note that once a sender observes that a probed channel is in state u it can not increase its gain any further by discovering another probed channel in state u or lower. Thus, subsequently it probes only the channels j for which the incremental gain $(\tilde{r}_j[u+1]\tilde{p}_j[u+1] - r_u)$ is less than the cost c_j , i.e., the

channels in $H_v, v > u$. The probing sequence in each H_u naturally follows an increasing order of the incremental gains.

Theorem 4.1: The expected gain of OPTNOBKUP is maximum among all strategies that do not use a backup.

Proof: The proof follows immediately from Lemmas 4.2 and 4.3. ■

Lemma 4.2: The optimum policy probes only channels in $\bigcup_{v>u} H_v$, after it observes a channel to be in state u . Further if there is an un-probed channel in $\bigcup_{v>u} H_v$ and the best state seen so far is u , then probing that channel improves the expected gain.

Proof: The proof is immediate for $u = K$, where no further probing is needed. Consider $j \notin \bigcup_{v>u} H_v$. Since $j \notin H_{u+1}$ we know that $\tilde{r}_j[u+1]\tilde{p}_j[u+1] - r_u \leq c_j$. But $\tilde{r}_j[u+1]\tilde{p}_j[u+1] - r_u$ is the expected gain (over the already seen channel in state u), and this is less than the cost of probing the channel. Clearly it is suboptimal to probe such a channel after we have seen a channel at state u . Likewise consider j in H_v for some $v > u$. The expected gain (reward minus cost) of probing j is $\tilde{r}_j[v]\tilde{p}_j[v] - c_j - r_u > r_v - r_u > 0$. ■

The optimum cannot therefore “stop” in a state u if any channel from H_v for some $v > u$ is left unprobed.

Lemma 4.3: The optimum policy must probe the channels of H_u in non-increasing order of $\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]}$ provided it has not seen a channel in state u or better so far. Further if $v > u$ then the optimum policy must probe all the channels of H_v before probing any channel not in H_u .

Proof: We will prove by induction, first on u (starting from $u = K+1$) and then on the number of unprobed channels of H_u remaining in a particular sequence/run of the optimum policy. The base case is $u > K$ and there is nothing to prove.

Assume that we are in some inductive case u . We assume there is some channel in H_u which is unprobed and the best state seen so far is worse than u ; otherwise there is nothing to prove for u . Among those (unprobed) channels of H_u let j be the channel with the largest $\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]}$ value.

Suppose the optimum policy at the current point is to probe some $i \neq j$ contradicting the hypothesis.

If we find the channel i in state u or better, the optimum policy is to stop since by the induction hypothesis on $v > u$ all states in H_v have been probed and there is no further benefit (in expectation) possible by Lemma (4.2). If we observe any worse state, we probe j next by the induction hypothesis (since the number of unprobed channels in H_u decreases, we can apply the hypothesis).

The situation resembles a decision tree as in Figure (1a). The trees $T_1 \dots T_{u^2}$ correspond to observing the ordered pair $(i = u', j = u'')$ where $0 \leq u', u'' \leq u-1$. The square boxes denote that we will definitely not probe anything else.

Now consider an alternate scenario of probing as shown in Figure (1b) where j is probed first and then i . The tree T' corresponding to the ordered pair $(i = u', j = u'')$ is assigned appropriately, on the branch corresponding to observing j in




Fig. 1. The decision trees of the Optimal policy for $u = 3$

u'' and subsequently observing i in u' . The contributions to the gain from the trees T_1, \dots, T_{u^2} remain the same because in both the scenarios the probability of probing these trees are the same.

The expected gain from scenario (a) (from not considering T_1 etc.) is $\tilde{p}_i[u]\tilde{r}_i[u] - c_i + (1 - \tilde{p}_i[u])(\tilde{p}_j[u]\tilde{r}_j[u] - c_j)$. This accounts for stopping after probing i as well as stopping after probing i and then j . Note that the reward for probing j is $\tilde{r}_j[u]\tilde{p}_j[u]$ which is the weighted reward from observing j in states u or better.

The expected gain in scenario (b) is $\tilde{p}_j[u]\tilde{r}_j[u] - c_j + (1 - \tilde{p}_j[u])(\tilde{p}_i[u]\tilde{r}_i[u] - c_i)$. Now if $i \in H_u$ then we have $\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]} > \tilde{r}_i[u] - \frac{c_i}{\tilde{p}_i[u]}$ which is the condition that arises from violating the non-increasing order. Otherwise $i \notin H_u$ implies $\tilde{r}_i[u] - \frac{c_i}{\tilde{p}_i[u]} \leq r_{u-1}$. But $\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]} > r_{u-1}$ since $j \in H_u$. Therefore in both cases we have $\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]} \leq \tilde{r}_i[u] - \frac{c_i}{\tilde{p}_i[u]}$. But this implies that

$$\begin{aligned} & \tilde{p}_j[u]\tilde{r}_j[u] - c_j + (1 - \tilde{p}_j[u])\{\tilde{p}_i[u]\tilde{r}_i[u] - c_i\} \\ & - \tilde{p}_i[u]\tilde{r}_i[u] + c_i - (1 - \tilde{p}_i[u])\{\tilde{p}_j[u]\tilde{r}_j[u] - c_j\} = \\ & \tilde{p}_i[u]\tilde{p}_j[u]\left(\tilde{r}_j[u] - \frac{c_j}{\tilde{p}_j[u]} - \tilde{r}_i[u] + \frac{c_i}{\tilde{p}_i[u]}\right) > 0 \end{aligned}$$

Thus by considering the scenario (b), we increase the profit of the optimum solution, which is impossible and we arrive at a contradiction. Thus by induction the lemma is true. ■

B. Approximation Algorithm for the Backup Case

We now consider the case that the optimal policy can transmit in an unprobed channel, and present a policy which attains at least $\frac{1}{2}$ the optimal gain, and has a computation complexity which is polynomial in n .

APPROXBKUP

Let ℓ denote the channel with the highest expected reward, $\tilde{r}_\ell[0]$.
Compute the gain R of OPTNOBKUP.

If $R > \tilde{r}_\ell[0]$ **then** use OPTNOBKUP,
else do not probe any channels, and select ℓ .

Theorem 4.4: The gain of algorithm APPROXBKUP is at least half the optimal gain.

Proof: In the optimal policy, let the expected gain from using backups (given a backup is used) be z and let α denote the probability with which backups are used. Thus, the total gain from the backups is αz . Let ALG denote the expected gain of APPROXBKUP. We first have

$$ALG \geq z. \quad (1)$$

Now modify the optimal policy so that the backups are not used, but the rest of the policy remains the same. Let OPT' denote the expected gain of this policy, and let x denote the expected gain of this policy given that the optimal uses a backup. Then, $OPT - OPT' = \alpha(z - x)$. Thus, $OPT \leq OPT' + \alpha z$. In addition, since OPTNOBKUP returns a solution with gain at least OPT' , we have $ALG \geq OPT'$. Therefore,

$$ALG \geq OPT - \alpha z. \quad (2)$$

Combining Equations 1 and 2, we have $2ALG \geq OPT$. ■

Note that the gain of APPROXBKUP is at least $\max(\frac{OPT}{1+\alpha}, OPT - \max_i c_i)$, where α is the probability with which the optimal solution uses backups [9]. Thus, the approximation ratio is better than the worst case bound in many cases.

V. THREE STATE CHANNELS

We present an improved approximation for 3-state channels.

Definition 5.1: A 0/1-path in a decision tree is a sub-tree where the next action is the same irrespective of whether a probed channel is in state 0 or 1.

Note that 0/1 paths are not paths but behave like paths; hence henceforth we will not distinguish between a path and a 0/1 path. Recall that $H_2 = \{i | (r_2 - r_1) > \frac{c_i}{p_{2i}}\}$. and $\tilde{r}_i[0] = r_2 p_{2i} + r_1 p_{1i}$ which is the expected benefit of using the channel i as a backup.

Clearly, after having seen a channel in state 1 the subsequent sequence of actions are identical if a channel is observed in

state 0 1. Thus, after a channel is seen in a state 1, the optimal decision tree becomes a path.

The key technical lemma in this section is the following.

Lemma 5.1: If the optimum policy uses a backup after seeing some channel j in state 1, then there exists another optimal policy where the decision tree rooted at channel j is a 0/1-path ending in the backup and the two policies are otherwise the same.

Proof: Consider a node (say node m) closest to the decision tree at which some channel j is probed and the path which corresponds to the observation that j is in state 1 uses a backup channel. Since after a channel is observed in state 1, the decision tree becomes a 0/1 path, the backup channel will be used, unless a channel is observed in state 2 somewhere before. Note that the expected reward of this backup is at least r_1 . Let the decision tree that arises after probing j and observing it to be in state 0 be A . If we observe j to be in state 1 then the decisions form a path P . This is shown in Figure 2(a).

Let the optimum policy traverse node m with probability p . Let G_1 be the conditional expected gain of the optimum policy if it does not traverse node m , G_A be the conditional expected gain of the optimum policy if it traverses node m and j is observed to be in state 0, and G_P be the conditional expected gain of the optimum policy if it traverses node m and j is observed to be in state 1. Clearly, $OPT = (1 - p)G_1 + p(p_{0j}G_A + p_{1j}G_P + p_{2j}r_2)$. Now, consider a modified policy where the tree A is used in place of the path P if j is observed in state 1 at node m . We refer to the gain of this policy as OPT' . Clearly, $OPT' = (1 - p)G_1 + p(p_{0j}G_A + p_{1j}G_A + p_{2j}r_2)$. Since $OPT' \leq OPT$, $G_A \leq G_P$. Now, consider another policy which is obtained by modifying the optimal policy as follows: path P is used instead of tree A when j is observed to be in state 0 in node m . We refer to the gain of this policy as OPT'' . Since unless a channel is observed in state 2, P uses a backup, G_P is the optimal gain of this policy given that it traverses node m and j is observed to be in state 0. Thus, $OPT'' = (1 - p)G_1 + p(p_{0j}G_P + p_{1j}G_P + p_{2j}r_2)$. Since $G_P \geq G_A$, $OPT'' \geq OPT$. Thus, the second modification corresponds to an optimal policy as well. Note that the second modification is otherwise similar to the original optimum, but its decision sub-tree rooted at node m is a 0/1 path. The result follows. ■




Fig. 2. The first two figures show how the paths are formed, (c) shows the consequence of Lemma 5.1

Applying the above lemma bottom-up on the optimal decision tree yields the following structure theorem, whose proof we omit.

Theorem 5.2 (Structure Theorem): For three-state channels, there exists an optimum policy that uses a unique backup channel (if at all) on only one path.

The structure theorem implies that the choice of the backup does not depend on the outcomes of the probes. Note that the uniqueness of the path on which a backup is used implies that the probability of using a backup is likely to be small. Furthermore, this theorem allows us to improve the approximation guarantee to 2/3 by combining the policies APPROXBKUP, OPTNOBKUP and another policy RESERVEBKUP, which we describe next.

Definition 5.2: Let $P(\ell)$ denote the class of policies, each of which (a) never probes ℓ and (b) never use any channel other than ℓ as a backup.

The best algorithm in $P(\ell)$ (over all choices of ℓ) may still be suboptimal, but will give us the desired approximation. Consider the following algorithm.

RESERVEBKUP(ℓ)

- 1) If $\tilde{r}_\ell[0] \geq r_1$ then use the policy which is optimal under the two state model among all policies that use ℓ as backup (the two state model is obtained by treating state 1 the same as state 0, and the optimal policy in this case has been obtained in [10]).
- 2) Otherwise (for the remainder of the algorithm, $\tilde{r}_\ell[0] < r_1$) sort the channels in $H_2 \setminus \{\ell\}$ in decreasing order of $r_2 - \frac{c_{2j}}{p_{2j}}$.
- 3) Probe the channels in $H_2 \setminus \{\ell\}$ in the above order. Stop if a channel is found to be in state 2, and select the channel.
- 4) If all channels in $H_2 \setminus \{\ell\}$ have been probed and if a channel has been observed in state 1, select that channel.
- 5) Otherwise let $\tilde{H}_1(l) = \{j | j \notin H_2 \cup \{\ell\} \text{ and } \frac{\tilde{r}_j[0] - c_j}{1 - p_{0j}} > \tilde{r}_\ell[0]\}$. Probe channels in $\tilde{H}_1(l)$ in order of decreasing $\frac{\tilde{r}_j[0] - c_j}{1 - p_{0j}}$ and stop if any channel j is observed in states 2 or 1, and select channel j .
- 6) If all the channels probed so far are in state 0, use ℓ as a backup.

Lemma 5.3: The algorithm RESERVEBKUP(ℓ) is optimal for the class of policies $P(\ell)$.

Proof: First note if $\tilde{r}_\ell[0] \geq r_1$ then the best algorithm in the class $P(\ell)$ will use the backup as long as no state is observed in state 2. In effect the algorithm will simply ignore state 1. This reduces this case to the two state problem with backup (see [10]), and RESERVEBKUP uses this solution.

Thus it suffices to consider $\tilde{r}_\ell[0] < r_1$ in the rest of the proof.

Using arguments similar to those in the proofs of Lemma 4.2 and Lemma 4.3, we can show the following.

- 1) All channels in $H_2 \setminus \{\ell\}$ must be probed unless a channel is observed to be in state 2.

- 2) The optimal policy probes the channels in $H_2 \setminus \{\ell\}$ in decreasing order of $r_2 - \frac{c_{2j}}{p_{2j}}$.
- 3) If some channel in $H_2 \setminus \{\ell\}$ is found to be in state 1, channels outside $H_2 \setminus \{\ell\}$ need not be probed.

Parts (1) and (2) prove that the actions in step (3) of the algorithm are optimal for algorithms in $P(\ell)$.

Now, consider step (4) of the algorithm. From part (3) and since $\tilde{r}_\ell[0] \geq r_1$, after all channels in $H_2 \setminus \{\ell\}$ are probed, if any channel has been found in state 1, the probing will stop and the channel must be selected. Thus, it follows that the actions in step (4) of the algorithm are optimal for algorithms in $P(\ell)$.

We now outline the proof that the actions in step (5) are optimal for algorithms in $P(\ell)$. Note that the algorithm executes step (5) only when all channels in $H_2 \setminus \{\ell\}$ are observed to be in state 0. Again note that at this stage if any channel is found to be in state 1 or 2, the probing must stop and the channel must be selected. Thus, from this point onwards, the best algorithm in $P(\ell)$ must treat states 1 and 2 as the same. Thus, the channels in effect have two states, but the reward in state 1 depends on the channel. We can show that in this case the optimum algorithm in $P(\ell)$ must probe a channel in $\tilde{H}_1(l)$ before using the backup (as then the gain increases). Subsequently we show that two consequently probed channels must be in non-increasing order of $\frac{\tilde{r}_j[0] - c_j}{1 - p_{0j}}$, or the net gain can be increased by switching their order. The result follows.

Finally, we prove that if a channel j is probed just before using the backup and $\frac{\tilde{r}_j[0] - c_j}{1 - p_{0j}} \leq \tilde{r}_\ell[0]$ then the net gain does not decrease by eliminating the probe of j . Applying this condition recursively, we establish that only the channels in $\tilde{H}_1(l)$ need to be probed before the backup is used. Thus, the actions in step (6) are optimal for all algorithms in $\tilde{H}_1(l)$. ■

We now present the main approximation algorithm CHOICE:

CHOICE
 G_1 = Gain of using the best backup channel ($G_1 = \max_i \tilde{r}_i[0]$).
 G_2 = Gain of OPTNoBKUP.
 G_3 = Gain of RESERVEBKUP(ℓ) for the best choice of ℓ .
 Select the best of the above three solutions.

Theorem 5.4: The CHOICE algorithm gives a $2/3$ approximation to the best adaptive probing policy for 3-state channels.

Proof: By the structure Theorem (Theorem 5.2), the optimum policy uses a unique backup (if at all). Let this backup be ℓ . Recall that the reward of using the backup is $\tilde{r}_\ell[0]$. Let p_ℓ denote the probability with which the backup is used. Recall that OPT denotes the optimal gain. Let ALG denote the expected gain of CHOICE. Equations 1 and 2 hold (setting $z = \tilde{r}_\ell[0]$) just as in Theorem 4.4.

Now, modify the optimum policy so that ℓ is removed from all places where it is probed (but it may still be used as the backup). The probability with which it is probed is at most $1 - p_\ell$, and the gain from probing it is at most $\tilde{r}_\ell[0]$. Therefore, the expected gain from probing it is at most $(1 - p_\ell)\tilde{r}_\ell[0]$. Let

the new expected gain be OPT'' . We have $OPT'' \geq OPT - (1 - p_\ell)\tilde{r}_\ell[0]$. Since RESERVEBKUP(ℓ) returns a solution with at least this value, we have

$$ALG \geq OPT - (1 - p_\ell)\tilde{r}_\ell[0] \quad (3)$$

Adding Equations 1, 2, and 3, we have $3ALG \geq 2OPT$, implying a $\frac{2}{3}$ approximation. ■

VI. ACKNOWLEDGEMENTS

The contribution of the authors were supported in part by NSF grants NCR 02 – 38340, CCF 04 – 30376, CNS 04 – 543264 and CNS 05 – 40347.

REFERENCES

- [1] I. 802.11a Working Group, *Wireless LAN Medium Access Control(MAC) and Physical Layer(PHY) Specifications-Amendment 1:High speed Physical Layer in the 5GHz Band*, 1999.
- [2] I. 802.11h Working Group, *Wireless LAN Medium Access Control(MAC) and Physical Layer(PHY) Specifications-Amendment 5:Spectrum and Transmit Power Management Extensions in the 5GHz Band in Europe*, 2003.
- [3] S. Babu, R. Motwani, K. Munagala, I. Nishizawa, and J. Widom, "Adaptive ordering of pipelined stream filters," in *Proc. of the 2004 ACM SIGMOD Intl. Conf. on Management of Data*, June 2004, pp. 407–418.
- [4] J. Burge, K. Munagala, and U. Srivastava, "Ordering pipelined query operators with precedence constraints," *Submitted*, 2005.
- [5] Y. S. Chow, H. Robbins, and D. Siegmund, "Great expectations: The theory of optimal stopping," *Houghton Mifflin Company*, 1971.
- [6] E. Cohen, A. Fiat, and H. Kaplan, "Efficient sequences of trials," in *Proc. of the 2003 Annual ACM-SIAM Symp. on Discrete Algorithms*, 2003.
- [7] U. Feige, L. Lovász, and P. Tetali, "Approximating min-sum set cover," *Algorithmica*, 2004.
- [8] I. . W. Group, *Wireless LAN Medium Access Control(MAC) and Physical Layer(PHY) Specifications*, 1997.
- [9] S. Guha, K. Munagala, and S. Sarkar, "Adaptive probes and wireless channels," *Technical Report*, 2006.
- [10] —, "Optimizing transmission rate in wireless channels using adaptive probes," *Proc. of SIGMETRICS/Performance*, 2006.
- [11] J. Hellerstein, "Optimization techniques for queries with expensive methods," *ACM Trans. on Database Systems*, vol. 23, no. 2, pp. 113–157, 1998.
- [12] Z. Ji, Y. Yang, J. Zhou, M. Takai, and R. Bagrodia, "Exploiting medium access diversity in rate adaptive wireless lans," *ACM MOBICOM*, 2004.
- [13] V. Kanodia, A. Sabharwal, and E. Knightly, "Moar: A multi-channel opportunistic auto-rate media access protocol for ad hoc networks," *Proceedings of Broadnets*, October 2004.
- [14] H. Kaplan, E. Kushilevitz, and Y. Mansour, "Learning with attribute costs," in *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, 2005, pp. 356–365.
- [15] M. S. Kodialam, "The throughput of sequential testing," *Lecture Notes in Computer Science*, vol. 2081, pp. 280–292, 2001.
- [16] K. Munagala, S. Babu, R. Motwani, and J. Widom, "The pipelined set cover problem," *Proc. Intl. Conf. Database Theory*, 2005.
- [17] K. Munagala, U. Srivastava, and J. Widom, "Optimization of continuous queries with shared expensive filters," *Submitted*, 2005.
- [18] A. Stolyar, S. Shakkottai, and R. Srikant, "Pathwise optimality of the exponential scheduling rule for wireless channels," *Advances in Applied Probability*, vol. 36, no. 4, pp. 1021–1045, 2004.