Sharing Spectrum the Smart Way

Introduction

Cognitive radio is the next-generation technology that will enable wireless systems to address and manage critical issues in the realms of spectrum efficiency, interference, and coexisting heterogeneous networks. This emerging technology allows primary and secondary users in a hierarchical access structure to interoperate and use spectrum efficiently. In such a network, secondary users use cognitive radios to exploit instantaneous spectrum opportunities. Cognitive radios actively sense and learn about the communication environment, thus enabling secondary users to identify and take advantage of these spectrum opportunities while avoiding interference with primary users.

(Credit: FCC)

Nearly all of the spectrum frequencies have already been allocated by the FCC.

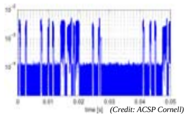

(Credit: ACSP Cornell)

However, in reality much of the spectrum is idle at any given time.

Challenges

The complex interaction between distributed secondary users is currently one of the most pressing issues in cognitive radio systems. Without a central controller in a distributed secondary network, a secondary user can only use its own local observations of the network to decide which of the spectrum channels to sense and access. Because secondary users cannot exchange information with one another, it is a standing challenge to develop an optimal policy for individual secondary users to rely on their own observations to achieve a tradeoff between avoiding collisions with other secondary users and selecting the channel that is most likely to be idle.

Problem Formulation

A spectrum communication environment can be modeled as X independently distributed channels, where secondary users can select one channel to sense in each time slot. If the channel is sensed to be idle "1" in a specific slot, the user is able to access the channel for that time slot. If the channel is sensed to be busy "0", the user cannot access the channel for that time slot and must wait for the next slot to select another channel to sense. A partially observable Markov decision process (POMDP) can be used to design sensing policies, whereby users attempt to increase their average throughput, the percentage of sensing success, over time. In a Markov Decision Process, decisions are made based only on information given by the present state, rather than taking the history of previous states into account.

pij = Pr (next state = j | current state = i)

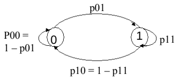

Here, i and j are represented by 0 and 1, where 0 represents a busy slot and 1 represents an idle slot.

A POMDP is a special case of the Markov Decision Process, applicable in this case because users can only "partially observe" the state of the various channels. Because of their limited sensing capabilities, users cannot know what the state of every channel is for a given slot. However, users must update their beliefs for each channel after each slot in order to reflect updated information received in the most recent slot. The belief vector, ω, which represents how much the user believes each channel will be idle in the next slot, is therefore updated using only the user's observations or beliefs from the most recent time slot.

$$\omega_i(t+1) = \begin{cases} p11, & \text{if channel } i \text{ is sensed to be idle in slot t} \\ p01, & \text{if channel } i \text{ is sensed to be busy in slot t} \\ \omega_i(t)p11 + (1 - \omega_i(t))p01, & \text{if channel } i \text{ is not sensed in slot t} \end{cases}$$

If a channel has not yet been observed by the user, its probability of being idle is assumed to be the static distribution:

$$q = p01/(1 + p01 - p11)$$

Ideally, a successful sensing policy for a secondary user will maximize the percentage of rewards over time. For a single user environment, a myopic policy, where only immediate, and not future, reward is maximized, has been established as an optimal policy in terms of both throughput and simplicity. In this myopic policy, a user simply selects the channel that it believes is most likely to be idle to sense. The result is improved average throughput over time, as shown in Fig. 1.

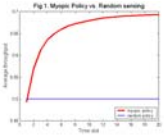

Fig. 1 shows the improved throughput of a myopic policy over a random policy, where the user randomly selects a channel to sense in each slot.

However, the myopic single user approach leads to poor network performance in a multi-user setting; if all secondary users greedily choose to sense the channel they believe is most likely to be idle, they will collide with one another and opportunities in other channels will be left unexplored. Even if users start off in different channels, whenever they sense the same channel during the same slot, they receive the same information, resulting in their belief vectors becoming increasingly synchronized.

To alleviate this problem, a randomized sensing policy has been proposed, where a user normalizes its belief vector by dividing each belief vector element by the sum of all the elements in its belief vector. Each belief vector element is now considered the probability of the user going to that specific channel. Thus, even when users have similar belief vectors, they will not become synchronized in their channel selections. The effects of these two policies in a multi-user setting can be seen in Fig. 2.

Fig. 2 shows the synchronization between users when the single user myopic policy is directly applied to a multi user setting, as well as the randomized policy's ability to avoid synchronization.

Objective and Method

The objective of this project was to gain insights into the interaction between distributed secondary users and to understand the impact of temporal correlations of spectrum opportunities. The goal was to design high performance networking policies of low complexity. Specifically, a class of distributed randomized policies was investigated, and questions such as how to choose the parameters of the randomized policy, how to implement the policy efficiently, and how to improve the performance of the policy were addressed.

Using the Markov Decision Process explained in the Problem Formulation, MATLAB was used to simulate a proposed multi-user policy in various network settings. The data from the simulations were then analyzed to draw conclusions about the policy and to eventually improve the performance and implementation of the policy.

Effect of Secondary Network Parameters on Policy Design

The standard multi-user policy outlined in Fig. 2 (see Problem Formulation) normalizes all the channels in a system to promote channel selecting diversity. It was hypothesized that the same throughput may be achieved with greater efficiency through normalizing fewer channels, however, and the effect of various secondary network parameters on policy design – specifically the optimum number of channels to be normalized – was studied.

Fig. 3 illustrates the general relationship found between the number of channels normalized and the average throughput for different numbers of users in a system. When only a few users are present, the difference in average throughput when different numbers of channels are normalized is minimal. As the number of users increases, it is clear that normalizing a higher percentage of the channels becomes more optimal. However, Fig. 3 also shows that the average throughput plateaus as the number of normalized channels increases: the marginal benefit of normalizing more channels drops off substantially. This observation suggests that there is an "optimum k" for any given scenario, where k is the number of channels normalized, using the assumption that normalizing fewer channels is more efficient.

Under this reasoning, and choosing optimum k as the lowest number of channels normalized that will achieve an average throughput less than 5% away from the maximum throughput, an optimum k was determined for simulated systems with varying user-to-channel ratios "ratio" and probabilities of packet arrival "p." Fig. 4 shows the results, and it can be clearly seen that the optimum number of channels normalized increases directly with respect to both the user-to-channel ratio and the packet probability.

Fig 3. Effect of Normalizing Different Numbers of Channels

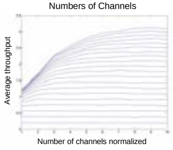

Number of normalized channels vs. average throughput in a 10 channel system. Shown for 1-20 users (lowest line is 1 user, highest line is 20 users).

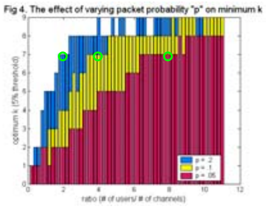

User-to-channel ratio vs. optimum k for p=.2, .1, and .05. The simulated system has 10 channels, and optimum k is calculated after 20 time-slots.

From this relationship, the following theorem can be derived:

THEOREM: For any p, M, X, and k^* in a multi-user network, where p is packet arrival probability, M is the number of secondary users, X is the number of channels, and k* is the optimum k,

$$\text{If} \quad p\frac{M}{X} = p'\frac{M'}{X'}$$

$$\text{Then} \quad k^*(p, M, X) = k^*(p', M', X')$$

When tested in practice, the theorem works quite well, and several points from Fig. 4 (circled in green) demonstrating its application are shown below:

Table 1. Application of Theorem 1 to values from Fig. 4

p	M/X	constant	Optimum k
.05	8	.4	7
.1	4	.4	7
.2	2	.4	7

Here, since the system has 10 channels, it follows that 70% of the channels should be normalized for maximum or near-maximum throughput.

The theorem was tested by varying secondary network parameters in system simulations, with the following properties verified:

Invariance Across Packet Arrival Probability

Experimental trials showed that while the optimum k predicted by the theorem does not exactly equal the experimental value of the optimum k in all cases, it is in general an accurate representation of the optimum k. As an example, Table 2 shows the accuracy of the theorem when packet arrival p = .2 vs. p = .05.

Table 2. Evaluating Theorem Accuracy by Varying P (based on data from Fig. 4)

P = .2													
Ratio	.2	.4	.6	.8	1	1.2	1.4	1.6	1.8	2.0	2.2	2.4	2.6
k*	1	2	2	4	5	5	5	6	7	7	7	8	8
P = .05													
Ratio	.8	1.6	2.4	3.2	4	4.8	5.6	6.4	7.2	8	8.8	9.6	10.4
k*	2	2	3	4	5	5	6	7	7	7	7	8	8

From Table 2 it can be seen that the optimum k was predicted exactly correctly by the theorem approximately 70% of the time. 100% of the predictions were within one channel (or 10%) of the actual optimum k as determined through simulation.

Invariance Across Time

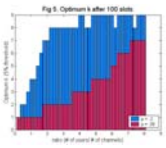

Fig. 5 shows the same system that was in simulated in Fig. 4, except with optimum k calculated after 100 time slots instead of 20. The accuracy results shown in Table 3 (rows 1 and 2) demonstrate that the percentage of channels that need to be normalized remains relatively constant over time.

Invariance Across Number of Channels

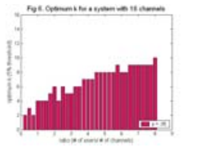

The theorem also accurately predicted the optimum k across systems with different numbers of channels. Fig. 6 illustrates a system with 16 channels instead of the 10 in Fig. 4, and the accuracy results are shown in Table 3 (row 3).

Table 3. Evaluating accuracy based on data from Figures 4, 5, and 6

		Accuracy		
Graphs Compared	**# of points tested**	**Perfect**	**Within 10%**	**Within 20%**
P = .2 (Figs 4 & 6)	40	70%	100%	100%
P = .05 (Figs 4 & 6)	40	30%	80%	100%
P = .1 (Figs 4 & 6)	8	37.5%	75%	100%

Effect of Primary Network Parameters on Policy Design

Although the discussion of secondary network parameters above shows the optimum number of normalized channels to be quite predictable from the theorem, it was discovered that one primary network parameter, the positive or negative correlation of p01 and p11 (defined in Problem Formulation), does affect the optimum k holding other parameters constant. The case where p11 > p01 is defined as positively correlated, while the case where p01 > p11 is defined as negatively correlated.

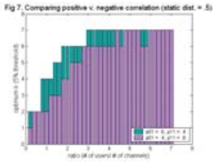

In Fig. 7, while both cases have the same static distribution, the optimum k is generally greater when p01 > p11 than when p11 > p01. A proposed explanation for this effect follows. It is necessary to return to differences in positively v. negatively correlated behavior in both single and multi user situations to find an explanation.

In a single user simulation, shown in Fig. 8, p11 > p01 produces a greater average throughput than p01 > p11 over time. This implies that in a p11 > p01 scenario, users are better at finding and accessing idle channels.

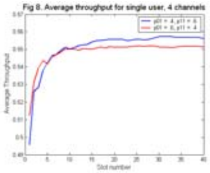

However, when multiple users employ the single user approach simultaneously, the p11 > p01 case actually produces a lower net throughput than the p01 > p11 case. Because users in a p11> p01 case are, as we know from the single user scenario, better at seeking out idle channels, they also end up synchronizing more easily as all the users select the same channels slot after slot. In Figures 9 and 10, the difference in the average throughputs clearly becomes more pronounced as the difference between p01 and p11 increases.

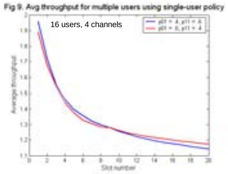

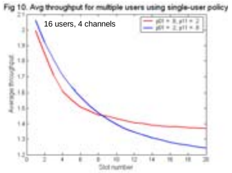

Considering this, a possible explanation for the effect of p01 and p11 correlation on the optimum k is as follows: from Figures 8, 9, and 10, it is clear that users in a p11 > p01 setting can indeed seek out idle channels with greater accuracy when a single-user policy is used. The reason for this can be seen by taking a closer look at the myopic policy. If a user is currently in an idle channel and p11 is high, the user should stay in that channel since its state is less likely to change; if p01 is high, however, the user should switch channels since the channel is more likely to become busy in the next slot. Since a user is typically in idle channels more than half the time (see Fig. 8), a user in a negatively correlated setting must therefore switch channels more frequently than a user in a positively correlated setting. When a user switches channels, it is forced to sense a channel whose belief has been evolving independently without direct observation. Intuitively, we can see that such a channel's belief will be less accurate than that of a channel that has just been sensed; thus, in a negatively correlated setting, a user must make decisions off of more inaccurate information, leading to lower average throughput. Because users are more sure of their beliefs in a p11 > p01 setting, there is a larger difference between the beliefs of the good vs. the bad channels. Thus, users are more likely to go to the best channels regardless of k; different values of k don't make much of a difference in average throughput, leading to a smaller optimum k. More investigation is required in order to validate this theory, however.

Conclusion and Discussion

Simulations in MATLAB show that not all channels in a system need to be normalized and considered for access at any given time slot to achieve optimal or near-optimal throughput. If users normalize only enough channels to achieve a throughput that is within 5% of the maximum throughput that can be achieved through the normalization method, a tradeoff between optimal throughput and user efficiency is reached.

In addition, a relationship between the optimum percentage of channels to normalize, the packet arrival probability, and the ratio of users to channels was found. This relationship, described in the Theorem in Experiments and Results, can be applied between systems with similar temporal correlations and remains relatively constant over time. In fact, predictions should get increasingly accurate for large numbers of slots, as the static distribution of each channel converges. The slight discrepancies between the theorem's predictions and actual simulated results may be due to normal variations in simulation or, in the case of the 16 channel comparison, different p11 and p01 settings. The relationship is beneficial because users can use a previously determined set of k^*(p,M,X) values to determine their own k^*(p,M,X) values, instead of having to simulate each setting individually.

Results also show that whether a system is positively or negatively correlated affects the optimum k. In general, it appears that the optimum k is generally larger when a system is negatively correlated than when it is positively correlated. The reason for this may go back to the difference in user behavior in each environment, which is explored in the results section. Clearly, however, when the correlations of two networks are very different, users must use separate sets of k^*(p,M,X) values to determine their optimum k.

Applications

The policy detailed in this experiment can have a wide range of applications including cell phone networks, portable devices, and wireless internet streaming. From the results, we can see that, depending on the number of users and channels, each user can save up to 90% of the energy and time used to normalize channels per slot. This is significant because each time slot generally lasts only a fraction of a second, and it is necessary for users to make good decisions quickly and efficiently. In addition, the policy is able to adapt to the ratio of users to channels in a system at a given time. For example, if a cell phone network is typically nearly unused in the morning, but very busy at night, users can normalize a small percentage of channels in the morning and more in the evening.

(Credit: Pop Wuping)

(Credit: Digital tech news)

(Credit: Alibaba.com)

Future Work

•Generalization of the Problem

Currently, parameters such as the behavior of primary users, the number of secondary users, the frequency of packet arrival, and the number of channels are taken into account during policy development. In the future, further parameters such as the location of users in relation to one another can also be addressed in policy development. However, because secondary users cannot communicate with one another, such factors will make modeling the problem much more complex.

•Generalization of Technical Approach

In this project, a partially observable Markov decision process is employed by secondary users to make channel sensing decisions. Other approaches of policy construction can also be tested in the future to see whether the average throughput in a multi user setting can be further improved. For example, methods such as Game Theory, a powerful mathematical tool that models distributed selfish but rational users, can also be applied to networking policies.

.

References

J. G. Kemeny and J. L. Snell, *Finite Markov Chains*. D. Van Nostrand Company, Inc., 1960.

S. Ashley, "Cognitive Radio," *Scientific American magazine*, vol. 294, no. 3, pp. 46-50, March 2006.

K. Liu, Q. Zhao, and Y. Chen, "Distributed Sensing and Access in Cognitive Radio Networks," in *Proc. of 10th International Symposium on Spread Spectrum Techniques and Applications (ISSSTA)*, August, 2008 (invited).

P. Mannion, "Sharing spectrum the smarter way", *EE Times*, pp.18-20, April, 2004.

Q. Zhao, B. Krishnamachari, and K. Liu, "On Myopic Sensing for Multi-Channel Opportunistic Access: Structure, Optimality, and Performance," in *IEEE Transactions on Wireless Communications*, vol. 7, no. 12, pp. 5431-5440, December, 2008.

Q. Zhao and B. Sadler, "A Survey of Dynamic Spectrum Access: Signal Processing, Networking, and Regulatory Policy," *IEEE Signal Processing magazine*, vol. 24, pp. 79-89, May 2007.

Acknowledgements

I would like to express my deep appreciation to Prof. Qing Zhao at UC Davis for her unwavering guidance throughout this research, and for all the time she spent teaching me. I am also grateful to all the graduate students at her lab, especially Keqin Liu, for their constant support and help. Finally, I would like to thank my family for always supporting me and encouraging me in all my pursuits.