

Reliable Open Spectrum Communications Through Proactive Spectrum Access

Prashanth Aravinda Kumar Acharya*, Sumit Singh†, Haitao Zheng*

* Department of Computer Science, University of California, Santa Barbara, CA 93106 U.S.A

† Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106 U.S.A

Email: acharya@cs.ucsb.edu, sumit@ece.ucsb.edu, htzheng@cs.ucsb.edu

Abstract—Open Spectrum systems offer an attractive solution to the reuse of under-utilized licensed spectrum. Existing proposals take a reactive sense-and-avoid approach to impulsively reconfigure spectrum usage without any knowledge of future dynamics. We propose a proactive spectrum access approach where secondary users utilize past observations to build predictive models on spectrum availability, and intelligently plan channel usage to maximize utilization and minimize disruptions to primary users. Based on the characteristics of TV-broadcast, we develop a simple availability metric and apply a usability filter to eliminate unreliable channels with heavy and frequent appearance of primary users. Our experimental results show that the proactive approach can significantly reduce the number of disruptions. We also observe a clear tradeoff between the disruption rate and the throughput at secondary users. By varying the usability filter threshold, we can control this tradeoff according to the constraints of primary users and the application requirements at secondary users.

I. INTRODUCTION

Existing spectrum management policies assign static spectrum ranges to different technologies in order to prevent interference. Over time, this has led to significant over-allocation and under-utilization of the licensed spectrum [1]. Based on the flexible technology of Cognitive Radios [2], [3], Open Spectrum systems promote *opportunistic spectrum access*, where new wireless devices (or *secondary users*) access spectrum bands currently unused by legacy or *primary users*. Successful deployments of Open Spectrum systems require secondary users to efficiently utilize unused spectrum while avoiding interference with primary users [4], [5].

Primary users can have spectrum usages that vary over time, hence, secondary users experience dynamic spectrum availability and need to intelligently adapt their spectrum usage. Current solutions are *reactive* wherein secondary users monitor local spectrum through individual or collaborative sensing [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], and reconfigure their own spectrum usage according to the spectrum usage by primary users. This adaptation is impulsive and it is based on the current observations of primary users. On detecting a primary user’s appearance on their current band in use, secondary users pause transmissions, relinquish the band and seek other opportunities to resume transmissions. Without any expectations or prior information about future spectrum availability to help schedule transmissions, secondary users can frequently encounter primary users. Due to non-negligible delay of sensing and radio reconfiguration, such a spectrum access approach results in frequent transmission disruptions

at both primary users and secondary users. Primary users can suffer from short-term interference until they are discovered, while secondary users experience transmission interruptions until they find other available channels and synchronize with their peers.

Disruptions occur when nodes in reactive systems can not foresee the occurrence of future events, and can only react to current observations. In this paper, we propose a *proactive* approach where secondary users proactively estimate future spectrum availability by building a predictive model of primary users’ spectrum usage based on past observations.

Our main contributions in this work are two-fold. We outline these next.

Proactive Spectrum Access We build a simple statistical model to characterize the future usability of each spectrum channel with an availability metric, and apply a usability filter to avoid unreliable channels with frequent and heavy appearance of primary users. Secondary users use the estimated channel availability to coordinate with peers to share available spectrum while minimizing the number of reconfigurations.

Experiments in broadcast-TV scenarios We also conduct simulation experiments on a simple TV-broadcast scenario to evaluate our proposed approach. Results show that the proactive approach can significantly reduce the occurrence of disruptions to primary users while achieving similar throughput performance. In addition, the use of usability filter allows secondary users to control the tradeoff between conservative spectrum access to reduce rate of disruptions and aggressive spectrum access to maximize spectrum utilization.

The rest of the paper is organized as follows. Section II provides a brief overview of Open Spectrum systems and existing solutions for spectrum sensing and coordination. Section III presents the proposed system architecture and the implementation assumptions. We discuss the detailed system design in Section IV. We describe our evaluation methodology in Section V, followed by experimental results in Section VI. Finally, we outline future work and conclude in Section VII.

II. BACKGROUND AND RELATED WORK

The fundamental challenge in Open Spectrum systems is to efficiently utilize and share unused spectrum among secondary users while avoiding interference to primary users. This section provides a brief review of prior work in this area. We focus on sensing algorithms to identify primary users, and coordination algorithms to enable fair sharing of available

among secondary users. We will describe a typical Open Spectrum system where TV-receivers are the primary users, and outline the related research challenges and problems.

A. Overview of Open Spectrum systems

A critical challenge in Open Spectrum systems is the need for secondary users to detect neighboring primary users to avoid interfering with them. Secondary users perform spectrum sensing, individually or through collaboration with neighboring peers. They can periodically monitor spectrum activity through energy detection [6], [7], [8], [9], [10], [11], [12], and interference temperature measurement [16]. Others apply cyclostationary feature detection [13] to identify periodic spectrum usage patterns [15], and statistically characterize spectrum usage intervals [14], [17] or available intervals [12]. To reduce sensing overhead, secondary users can leverage external sensors [18] to perform sensing and to broadcast the sensed spectrum usage information, *e.g.*, a sensor mounted next to primary receivers [19]. Using sensing information, secondary users must dynamically reconfigure their spectrum usage to avoid interference with primary users.

The second fundamental challenge is how to efficiently utilize the dynamically available spectrum. Prior work follows a reactive approach where secondary users react to current observations and dynamically reconfigure upon each occurrence of spectrum variations, *i.e.*, appearance of a primary user. They pause existing transmissions, relinquish the channel and re-coordinate to find other available channels to resume transmissions. Along with non-negligible sensing and reconfiguration delay, this creates frequent disruptions to both primary and secondary users.

In addition to configuring spectrum usage to maximize utilization, secondary users also need to share available spectrum with other peers in order to coexist. Existing approaches can be divided into two categories with different assumptions on network infrastructure. Centralized approaches [18], [20], [21], [22], [23], [24] use a spectrum broker to determine allocations among devices with the knowledge of global topology. Distributed approaches [25], [26], [27], [28], [29], [30], [31] rely on device negotiations or competition to resolve conflicts and approach the results of global optimization. We can also leverage and extend conventional multi-channel MAC protocols that use periodic coordination to schedule transmissions [32], [33], [34], [35].

B. A Motivating Example: Broadcast-TV Spectrum

Recent policy development and technology advancement promotes deployment of Open Spectrum systems on broadcast-TV spectrum. To avoid interference with primary users (TV receivers), secondary users can detect primary transmitters and determine a no-talk zone through energy detectors [6], [7], [8]. Since TV-broadcast stations are always transmitting, this approach identifies regional spectrum availability rather than local availability around each primary TV receiver. An alternative is to mount a low-cost sensor on each TV receiver that determines spectrum usage by monitoring the

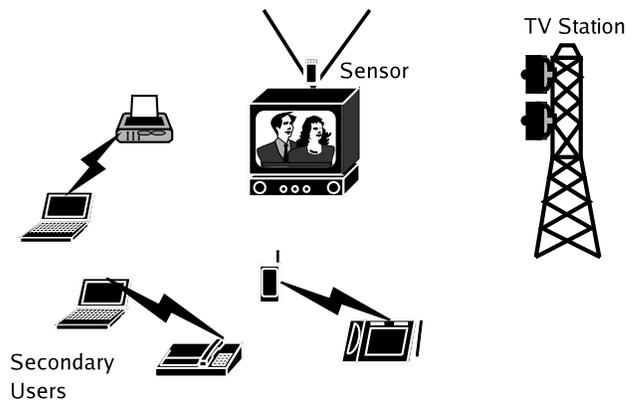


Fig. 1. A Broadcast-TV scenario where secondary users are laptops and PDAs who communicate with each other. The primary user is the TV receiver, and a sensor attached to it monitors its spectrum usage and broadcasts to secondary users.

local oscillator (LO) leakage power [19] and broadcasts usage information to neighboring secondary users (see Figure 1).

TV-viewing patterns are dynamic, which makes local spectrum availability fluctuate over time. Secondary users who apply the reactive approach to schedule channel usage suffer from frequent disruptions due to variations of spectrum availability. Because of lack of knowledge about future spectrum availability, secondary users search for channels that are currently unused, but these channels can face immediate appearance of primary users. Such blind and impulsive adaptation leads to frequent disruptions because of the non-negligible sensing and reconfiguration delay.

There have been studies to build models of spectrum usage in order to identify and track primary users with periodic usage or strong short-term temporal correlation [14], [15], [17]. These approaches model the duration of each primary channel usage, or the duration of each channel availability. However, they are mostly suited for periodic spectrum access with deterministic usage patterns, and do not address the frequency of spectrum being occupied and ways to select among multiple channels given the knowledge. On the other hand, TV-viewing patterns are much more complex, and vary significantly over time, and have strong temporal correlations of various scales, such as “time-of-the-day” and “day-of-the-week.” Hence, we need to develop general models of spectrum availability, and need to design adaptation mechanisms to take advantage of this knowledge.

III. PROACTIVE SPECTRUM ACCESS

In order to achieve the goals of maximizing spectrum utilization while minimizing disruptions, we propose a *proactive* spectrum access approach to build predictive models of spectrum availability and schedule channel usage to maintain reliable communications. Our approach consists of two components, (1) *building a predictive model of spectrum availability on each channel*, and (2) *scheduling channel usage to minimize disruptions and reconfiguration overhead*. We propose to characterize each spectrum channel’s usability with an availability metric, and apply a usability filter to eliminate

“unreliable” channels with frequent and prolonged primary user appearance. Secondary users then coordinate to share spectrum based on the knowledge of channel usability.

In this paper, we address the problem of Open Spectrum access in broadcast TV based systems described in Section II-B. However, this is not a limitation of our work. Our proposed approach can be easily extended to more general scenarios, and can be integrated with existing work on primary user characterization [15] to refine predictive models based on prior knowledge about primary users. We describe on-going work to expand this approach in Section VII.

Next, we describe the architecture of the proposed system on each secondary device, and discuss the assumptions and implementation issues.

A. System Architecture

Our proposed system is implemented on each secondary device. It consists of the following three modules, located at the MAC layer.

Characterization. This module takes input from spectrum sensing results, and builds statistical models of spectrum availability on each channel. Since primary users (TV-receivers) have temporal correlations over different time scales (hour/day/week), we propose to build models with different time granularity and integrate them together to derive a comprehensive characterization of spectrum availability.

Usability Filtering. This module applies a threshold based filter to examine spectrum availability and eliminate “unreliable” channels with frequent and heavy appearance of primary users. The filter threshold is reconfigurable, and offers the capability to control the tradeoff between aggressive spectrum access towards maximizing spectrum utilization, and conservative spectrum access for minimizing disruptions to primary users.

Coordination. This module is responsible for coordinating spectrum usage among multiple co-located secondary users to avoid peer interference. Secondary users utilize the knowledge of spectrum availability and information of neighbor conflict to schedule and coordinate channel usage with their peers. In essence, each user tries to select a channel with highest spectrum availability but least usage from other peers. Our contribution is to integrate channel availability into the procedures of distributed coordination [33], [35].

Figure 2 illustrates the proposed architecture on each secondary device. We will discuss the detailed design of spectrum characterization, filtering and coordination in Section IV.

B. Assumptions and Implementation Issues

Before we proceed to the detailed system design, we briefly describe design assumptions and implementation issues. We argue that these assumptions arise naturally from the requirements of a cognitive radio based system. Having them helps to simplify our system design without affecting the generality of our proposed solutions and results.

Spectrum Sensing While sensing is not the focus of this paper, the sensing performance has direct impact on the effectiveness of spectrum access algorithms. In this paper,

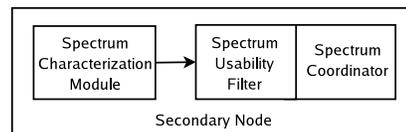


Fig. 2. Secondary Node Architecture: Each secondary node has a) Spectrum Characterization Module b) Spectrum Usability Filter and c) Spectrum Coordinator.

we assume the presence of an external sensor mounted next to the TV receiver that detects and announces the spectrum usage at the receiver [19]. The sensor sends an alarm message when the TV receiver switches channels. Announcements are sent over a dedicated control channel to secondary users with some propagation and processing delay, but without any transmission errors. Therefore, we assume that each secondary user has perfect knowledge of current channel in use.

We further assume that upon entering the network, each secondary user is required to collect statistics of spectrum availability through aggregating sensor broadcasts over a period of time. However, a significant amount of observation time is required to capture the long-term temporal correlation in spectrum availability. To address this issue, we assume the existence of an external (regional) knowledge base that constructs long-term models of spectrum availability and store them as policies. Secondary users can download these policies from regional servers periodically to schedule spectrum usage.

Common Control Channel We assume that there is a dedicated channel for secondary users to exchange control information [34], [36], such as neighbor discovery and coordination messages. To minimize control overhead, we assume an in-band solution where secondary users coordinate periodically by reconfiguring themselves into a low-power ultra-wide-band (UWB) radio to exchange messages without affecting primary users. Similar approaches are proposed in [34].

Communications in the presence of Primary Transmitters In TV broadcast scenarios, broadcast stations always transmit while individual TV receivers select channels to receive signals. We assume that secondary users apply sophisticated coding techniques to maintain reasonable communication capacity in the presence of high power interferers [19]. One example is the *dirty paper* type coding [37].

IV. DETAILED SYSTEM DESIGN

In this section, we present the detailed design of the proposed approach.

A. Spectrum Characterization

Secondary devices utilize past observations to build predictive models of spectrum availability. In order to characterize the “quality” of each spectrum channel, we use a channel *availability* metric that is defined as the expected channel free time after each primary user’s appearance. To capture temporal correlations over different time scales, we propose to build multi-tier models using data collected with different observation windows. For a given tier k with observation

window of length t_k , the availability of channel i is estimated as

$$\omega_i(t_k) = E\left\{\frac{\text{Total channel free duration over interval } t_k}{t_k \cdot \text{Number of primary user switches to channel } i}\right\}$$

and $0 \leq \omega_i(t) \leq 1$. This metric captures both the average duration of each channel availability and the frequency of interruptions from primary users.

Based on logical expectations of TV viewing patterns, we build a three-tier model including availability statistics from three types of observation windows.

(1) Immediate ($t_1=1$ minute) : This tier assumes strong immediate correlation in TV-viewing patterns. If a channel is viewed during the past 1 minute, the probability that it will be viewed in the immediate future is expected to be high.

(2) Short-term ($t_2= 10$ minutes) : This tier uses average availability over the last 10 minutes to predict the future availability.

(3) Long-term ($t_3= 30$ minutes over the past few weeks) : This tier captures the long-term time-of-day and day-of-week patterns. Secondary users can access this model from the knowledge base or through long-term sensing. We propose to integrate the availability metric at each tier to construct a single metric to characterize spectrum availability. In this paper, we define the overall *usability* of channel i u_i as

$$u_i = \omega_i(t_1) \cdot (\alpha \cdot \omega_i(t_2) + (1 - \alpha) \cdot \omega_i(t_3)), 0 \leq u_i \leq 1, (1)$$

where $u_i = 0$ if the channel has been used in the immediate past, *i.e.*, $\omega_i(t_1) = 0$. The parameters α and $1 - \alpha$ are weights assigned to short-term and long-term statistics respectively. The choice of weights depends on the relative level of correlation on each layer, and can be configured adaptively. In this paper, we determine the weights using simple heuristics and experimental training. In other complex scenarios, we can apply online learning to adapt the weights according to the behavior of primary users. We are currently investigating this approach and will include it in a future study.

We note that our predictive model is simplistic, and it does not capture all the statistics of TV-viewing patterns. This results in inevitable errors and model uncertainty, which could impact the efficiency of spectrum scheduling. In addition to exploring more advanced and complex prediction mechanisms [38] to improve the model accuracy, we also propose to develop recovery mechanisms to minimize the impact of errors. In this paper, secondary users recover by pausing their transmissions and selecting other usable channels. In the worst case scenario where spectrum availability on every channel is statistically random, our proposed approach will fall back to the conventional reactive approach.

B. Spectrum Usability Filtering

Spectrum usability estimates the quality of each channel. Secondary users can use this metric to schedule their spectrum usage, such as selecting the channel with the highest u_i metric value. However, there is no clear indication of whether a channel is safe to use without causing severe disruptions to

primary users. Towards this goal, we propose to utilize usability filtering to eliminate unreliable channels with frequent and prolonged appearance of primary users.

We define a configurable parameter called the filter threshold T ($T > 0$), where a channel can be considered for use only if its usability metric is above T . The choice of threshold T significantly impacts the system performance. Choosing a small T implies aggressive spectrum access where secondary users intend to maximize spectrum utilization by even using channels with low availability. Choosing a large T leads to conservative spectrum access that emphasizes on channel reliability and minimization of disruptions. The optimal choice of T depends on the tolerance of primary users and the application requirements of secondary users. In this paper, we explore the impact of T through experiments, which is summarized in Section V.

C. Spectrum Coordination Module

To support reliable communication in a multi-user system, secondary users in close proximity must share available spectrum and avoid interference with each other. Without infrastructure support, we propose the use of distributed coordination where secondary users inform competing peers about their spectrum usage, and dynamically adapt their own channel usage accordingly to reduce the level of interference. When two secondary users want to communicate with each other, they also exchange individual assessment of channel availability, and coordinate to select an appropriate channel.

While prior research [25], [33] has demonstrated the feasibility of distributed coordination in static spectrum availability scenarios, the use of this approach in dynamic spectrum systems needs to consider channel usability. We propose a usability-aware coordination approach where secondary users select channels based on the information of channel usability and the level of existing traffic load from other peers on each channel.

The proposed coordination procedure is similar to the approach of [33], [35]. Before starting a data transmission, a secondary transmitter sends a negotiation message TxRequest to its destination on the control channel¹. Each TxRequest message contains a prioritized channel list derived based on the channel availability, the level of existing traffic, and the intended traffic volume. The intended receiver combines the received statistics with its own assessment to select an appropriate channel and responds with a TxResponse message to the transmitter. The transmitter then sends a Confirm message and starts data packet transmissions. Figure 3 illustrates a sample message flow.

We also propose to reduce possible disruptions to primary users by limiting the transmission duration at secondary users. Each receiver estimates an *allowed transmission duration* (ATD) on the selected channel i based on the assessment of channel availability at the link pair,

$$ATD = \min(u_i \cdot \text{MaxTxDur}, \text{RTD}),$$

¹In this paper, secondary users periodically configure themselves into a UWB radio to perform coordination.

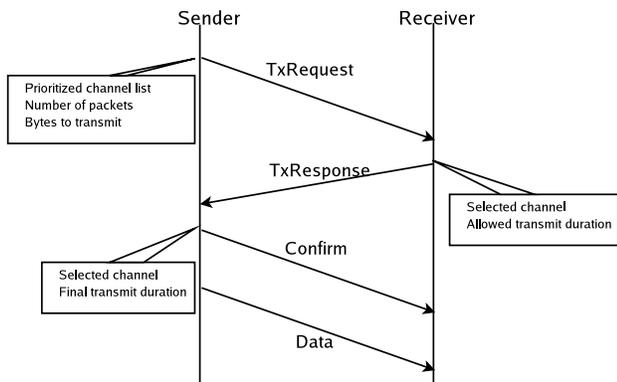


Fig. 3. Control message flow for data transfer negotiation.

where u_i is the minimum of the availability on channel i estimated by both the transmitter and the receiver, MaxTxDur represents the maximum allowed transmission duration of a secondary sender-receiver pair before the next coordination period², and RTD represents the estimated time to transmit all the packets requested by the transmitter. The receiver embeds ATD in the TxResponse message to inform the transmitter. By eavesdropping on the TxResponse and Confirm messages, neighboring secondary nodes can obtain estimates of existing traffic from other peers, and schedule transmission to avoid channels that are already occupied by other peers.

Upon being interrupted by an appearance of primary users, secondary users pause transmissions and wait till the next coordination window to re-negotiate and resume transmissions. Hence, having a large MaxTxDur might produce long interruptions at secondary users. To reduce disruptions, we can limit the value of MaxTxDur and even fall back to per-packet negotiation. However, this is at the cost of higher coordination overhead, and is only suitable for low-rate traffic. A more intelligent solution is to adapt MaxTxDur to the average duration between two consecutive traffic bursts at primary users. We plan address this issue in a future study.

V. EXPERIMENT SETUP AND EVALUATION METRICS

We implement the proposed approach using QualNet wireless network simulator [39]. We simulate a home network scenario - a living room with one TV representing the primary receiver, and a number of wireless devices as the secondary users (see Figure 1). We only consider single-hop communications because all the secondary users are assumed to be within the transmission range of each other. Performance study in a multi-hop environment is part of our ongoing research.

We assume that the spectrum space is divided into 10 orthogonal channels, each capable of providing 1Mbps raw data rate for secondary users. We only consider one primary user (the TV receiver) that accesses one channel at a time. We assume that each secondary user has knowledge of the long-term spectrum usage of the primary user, represented by $\omega(t_3)$ at each time instance. We assume that secondary

²This is equal to the length of the interval between subsequent configurations to UWB.

users are transmitting UDP packets. This allows us to examine the impact of spectrum dynamics and the performance of the proposed approach without the effect of higher-layer protocols (e.g., TCP). Our simulation results are averaged over 10 runs with different random seed values, and each simulation covers 3 hours of network activity.

We implement the proposed proactive approach, and a baseline approach (*reactive, no prediction*), where nodes randomly select channels to use, without any knowledge about future dynamics. This approach applies the same three-way handshaking procedure to negotiate channel usage, but does not use any usability information to prioritize channels.

A. TV Viewing Patterns

To the best of our knowledge, there are no simple TV-viewing patterns publicly available. Hence, we generate TV-viewing patterns based on logical expectations and our daily life experience. We use a sitcom model where time is divided into half hour bins, with day-of-week and hour-of-day patterns. The viewer has a *preferred channel set* for each half hour bin, and selects channel from this set with high priority. Within each bin, the viewer randomly switches across channels, selecting a channel from the preferred set with probability p_1 and the rest with probability $1 - p_1$.

The channel viewing time on the preferred and non-preferred channels E_1 , and E_2 are exponentially distributed with means of 90 and 15 seconds, respectively. If the same channel is chosen over two consecutive viewing intervals, we disregard the channel switch and consider it to be a single viewing interval. We have evaluated the system performance with different values of p_1 and have observed similar performance trends. Hence, in this paper, we will use $p_1 = 0.8$ as an illustrative example. Similar observations apply to E_1 and E_2 , and we will use 90 and 15 seconds, respectively.

We note that the TV-viewing model is an empirical model to provide preliminary evaluation of our proposed approach, but it does not limit the scope of the proposed approach.

B. Evaluation Metrics

We evaluate the system performance in terms of two metrics: *primary user channel hit ratio* and *secondary user throughput*. These metrics examine the tradeoff between conservative spectrum access to minimize occurrence of disruptions, and aggressive spectrum access to utilize as much available spectrum as possible to maximize throughput and channel utilization.

Primary Channel Hit Ratio. This metric is defined as the ratio of channel hits experienced by the primary user to the total number of primary user channel switches. A channel hit occurs if the primary user switches to a channel that is currently occupied by some secondary users, and suffers from short-term disruptions due to the non-negligible sensing, broadcast and evacuation delays. In addition to primary user disruptions, this metric also represents the number of interruptions to secondary user packet transmissions since they need to immediately evacuate the channel and resume transmissions in the next coordination interval.

Throughput. The throughput measures the system throughput obtained by all secondary users. We only measure the long-term throughput averaged over the simulation duration (3 hours). Since the coordination duration MaxTxDur is small compared to the average viewing time, the impact of interruptions is negligible. However, these interruptions cause packet losses, which is reflected by the primary channel hit ratio.

VI. EXPERIMENTAL RESULTS AND DISCUSSIONS

A. Comparison to the Baseline Approach

Figure 4 illustrates the primary channel hit ratio and the aggregate throughput of all secondary users versus the number of secondary user pairs. Our results show that the proposed proactive approach can significantly reduce the number of hits over the baseline (reactive) approach. By selecting channels with high availability, secondary users can avoid using channels with frequent appearance of primary users. Hence, even without any usability filter, the proactive approach can reduce primary hit ratio by 20-30%.

We also observe a strong dependency of the performance on the filter threshold T . The use of large T can effectively reduce the hit ratio to 4%, particularly at high system load. When a large number of users compete for spectrum usage, they spread selections across channels to reduce peer interference, and even select channels that are frequently accessed by the primary users. Hence the use of higher threshold to restrict the use of “unreliable” channels can effectively control the hit ratios. Reduced throughput from higher threshold values is due to multiple secondary pairs sharing fewer usable channels.

We also examine the impact of T with a fixed number of 7 secondary user pairs under different levels of traffic load. From Figure 5, we see that increasing per-link load has no effect on the hit ratio but it improves the system throughput. The use of high T confines secondary users to “reliable” channels, and system throughput saturates to the channel capacity as the number of links increases. We also examine the performance under a constant total traffic load of 800Kbps but with varying number of communicating secondary user pairs. The results in Figure 6 show a similar trend as that in Figure 4, where the primary channel hit ratio increases linearly with the number of secondary user pairs at small T but remains constant at large T .

We observe that the primary hit ratio is non-zero even in a lightly loaded system with a single secondary user link. This is because of the inherent random nature of our TV-viewing model, and the limitations of the prediction accuracy of the simple availability metric. In addition to triggering recovery mechanisms at secondary users to quickly evacuate the channels and resume communications, we can also apply strict filtering to avoid the use of unpredictable channels.

B. Adaptive Filtering

From Figures 4 and 5, we observe a tradeoff between the disruption rate (primary channel hit ratio) and the secondary user throughput. The use of large T can significantly reduce the occurrence of primary user hits but at the cost of reduced throughput. Figure 7 plots this tradeoff as the channel hit

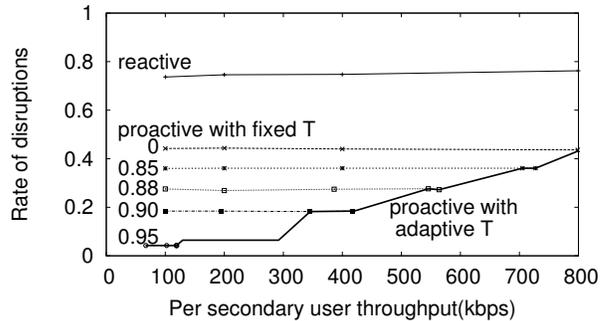


Fig. 7. Channel hit ratio as a function of sustainable throughput using different threshold (T) values.

T	reactive	0.0	0.3	0.6	0.9
$\alpha = 1.0, \beta = 0.0$	0.74	0.56	0.56	0.55	0.52
$\alpha = 0.5, \beta = 0.5$	0.74	0.47	0.47	0.42	0.39
$\alpha = 0.0, \beta = 1.0$	0.74	0.44	0.43	0.42	0.18

TABLE I

CHANNEL HIT RATIOS FOR PREDICTION BY SHORT-TERM AND LONG-TERM USAGE CORRELATIONS.

ratio versus the throughput on each secondary link assuming 7 secondary user pairs. The bold line represents this performance tradeoff by adapting the value of T , where the throughput grows almost exponentially at small hit ratios (3-10%) but slowly converges to linear growth as the hit ratio increases. This shows that if primary users can tolerate even a low degree of interference, it can lead to substantial improvement of spectrum utilization at the secondary users.

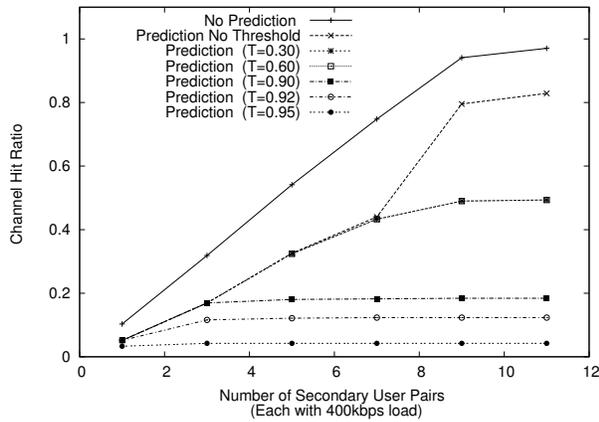
We also observe from these results a non-linear relationship between the threshold T and the disruption rate. Secondary users can clearly benefit by using a detailed mapping of the two values as a guideline to traverse this cost-performance tradeoff.

C. Impact of Predictive Model Parameters

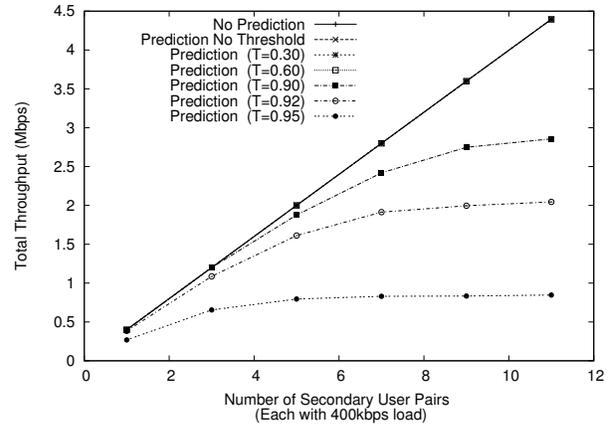
The performance of the availability model depends heavily on the choice of α . We examine three cases with $\alpha = 1.0, 0.5$ and 0.0 , corresponding to short-term, mixed, and long-term based predictions. Note that the immediate ($t_1 = 1s$) observation is directly embedded in the model since it will lead to an availability of 0. Table I summarizes the primary channel hit ratios under an example scenario of 7 secondary user pairs with 400kbps traffic per link. As expected, given the strong long-term temporal correlation present in TV-viewing patterns, the use of large α leads to significantly improved performance.

VII. CONCLUSION AND FUTURE WORK

In this paper, we examine the performance of Open Spectrum systems in a TV-broadcast scenario. We observe from this work that the spectrum usage patterns of primary users have great impact on the performance of any Open Spectrum system, and should be incorporated when scheduling spectrum

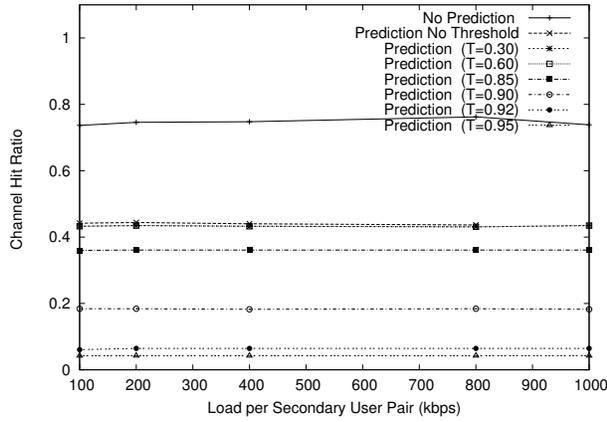


(a) Primary channel hit ratio

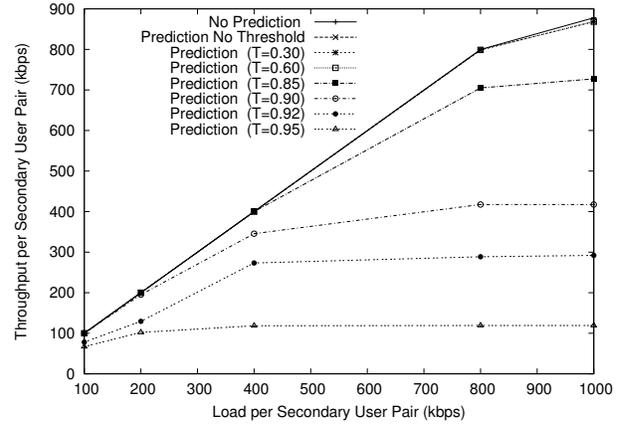


(b) Throughput

Fig. 4. Effect of increasing aggregate secondary user load by varying the number of secondary user pairs with CBR load of 400Kbps each. No prediction represents the baseline reactive approach.

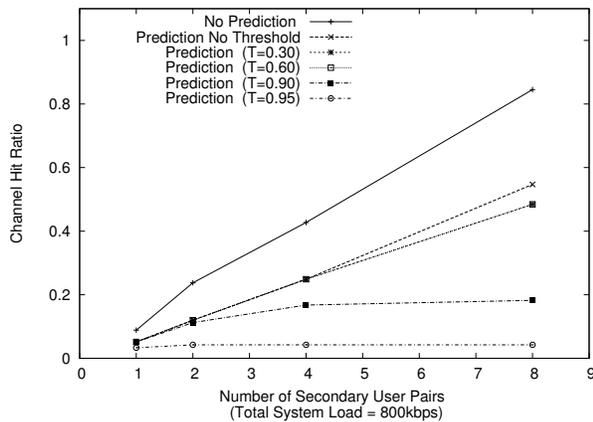


(a) Primary channel hit ratio

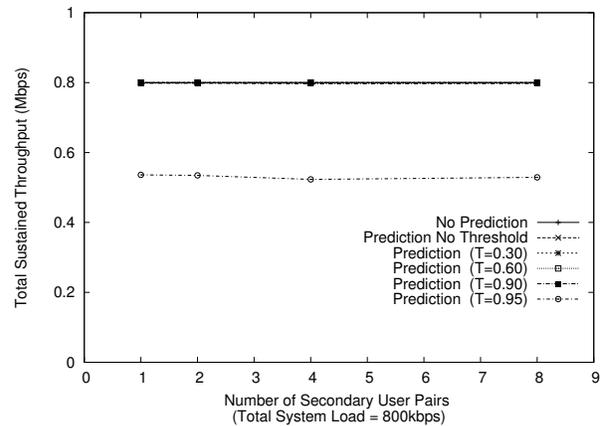


(b) Throughput

Fig. 5. Effect of increasing aggregate secondary user load by varying the CBR load for a fixed number (7) of secondary user pairs.



(a) Primary channel hit ratio



(b) Throughput

Fig. 6. Effect of number of secondary users with a constant aggregate load of 800Kbps.

usage. Conventional approaches apply a reactive sense-and-adapt solution based solely on current observations of primary users. Disruptions occur when nodes using reactive spectrum access can not foresee the occurrence of future events, and can only react to current observations. We propose a proactive spectrum access approach where secondary users make use of past observations to build predictive models of spectrum availability, eliminate unusable channels through a usability filter, and schedule their spectrum usage to maximize spectrum utilization while minimizing the occurrence of disruptions to primary users. We observe a clear tradeoff between the rate of disruptions to primary users and the throughput achieved by secondary users. By varying the usability filter characteristics, we observe that proactive spectrum access significantly reduces disruptions to primary users while providing similar throughput as the conventional reactive approach. The threshold parameter allows us to control the tradeoff between the disruption rate and the desired throughput.

While our initial results demonstrate the feasibility of proactive spectrum access in a simple scenario, the effectiveness of this approach depends heavily on the accuracy of the spectrum availability models. Our proposed statistical model is simple and can incorporate multiple tiers of different time granularities. However, in more complex scenarios, we need to investigate more sophisticated predictive models and apply intelligent self-tuning of system parameters to improve prediction accuracy. We are currently expanding our research in this area.

REFERENCES

- [1] M. McHenry, "Spectrum white space measurements," *New America Foundation Broadband Forum*, June 2003.
- [2] R. W. Brodersen, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, "Corvus: a cognitive radio approach for usage of virtual unlicensed spectrum," July 2004, whitepaper.
- [3] S. Mangold, Z. Zhong, K. Challapali, and C. T. Chou, "Spectrum agile radio: Radio resource measurements for opportunistic spectrum usage," in *Proc. of Globecom*, November-December 2004.
- [4] (2006) Spectrum Policy Task Force. [Online]. Available: <http://www.fcc.gov/sptf/>
- [5] P. J. Kolodzy, "Interference temperature: a metric for dynamic spectrum utilization," *Int. J. Netw. Manag.*, vol. 16, no. 2, pp. 103–113, 2006.
- [6] A. Sahai, N. Hoven, and R. Tandra, "Some fundamental limits on cognitive radio," in *Forty-second Allerton Conference on Communication, Control and Computing*, Monticello, IL, October 2004.
- [7] A. Sahai, R. Tandra, and N. Hoven, "Opportunistic spectrum use for sensor networks: the need for local cooperation," 2006, submitted to Information Processing in Sensor Networks (IPSN).
- [8] S. Mishra, A. Sahai, and R. Brodersen, "Cooperative sensing among cognitive radios," 2006, submitted to IEEE International Conference on Communications (ICC).
- [9] G. Ganesan and Y. G. Li, "Cooperative spectrum sensing in cognitive radio networks," in *Proc. of IEEE DySPAN*, November 2005.
- [10] A. Ghasemi and E. S. Sousa, "Collaborative spectrum sensing for opportunistic access in fading environments," in *Proc. of IEEE DySPAN*, November 2005.
- [11] D. Cabric, S. M. Mishra, and R. W. Brodersen, "Implementation issues in spectrum sensing for cognitive radios," in *Proceedings of Asilomar conference on signals, systems and computers*, 2004.
- [12] X. Liu and S. Shankar, "Sensing-based opportunistic channel access," *Accepted for publication by ACM MONET*, 2006.
- [13] A. Fehske, J. D. Gaeddert, and J. H. Reed, "A new approach to signal classification using spectral correlation and neural networks," in *Proc. of IEEE DySPAN*, November 2005.
- [14] H. Tang, "Some physical layer issues of wide-band cognitive radio system," in *Proc. of IEEE DySPAN*, November 2005.
- [15] K. Challapali, S. Mangold, and Z. Zhong, "Spectrum agile radio: Detecting spectrum opportunities," in *International Symposium on Advanced Radio Technologies (ISART)*, Boulder, CO, March 2004.
- [16] FCC, "Et docket no 03-237 notice of inquiry and notice of proposed rulemaking," ET Docket No. 03-237, Tech. Rep., November 2003.
- [17] S. D. Jones, N. Merheb, and I.-J. Wang, "An experiment for sensing-based opportunistic spectrum access in CSMA/CA networks," in *Proceedings of DySPAN 2005*, Baltimore, MD, Nov 2005.
- [18] V. Brik, E. Rozner, S. Banarjee, and P. Bahl, "Dsap: A protocol for coordinated spectrum access," in *Proc. of IEEE DySPAN*, November 2005.
- [19] B. Wild and K. Ramchandran, "Detecting primary receivers for cognitive radio applications," in *Proc. of IEEE DySPAN*, November 2005.
- [20] O. Ileri, D. Samardzija, and N. B. Mandayam, "Demand responsive pricing and competitive spectrum allocation via spectrum server," in *Proc. of IEEE DySPAN*, November 2005.
- [21] M. M. Buddhikot, P. Kolody, S. Miller, K. Ryan, and J. Evans, "Dimsumnet: New directions in wireless networking using coordinated dynamic spectrum access," in *Proc. of IEEE WoWMoM 2005*, June 2005.
- [22] M. M. Buddhikot and K. Ryan, "Spectrum management in coordinated dynamic spectrum access based cellular networks," in *Proc. of IEEE DySPAN*, November 2005.
- [23] C. Raman, R. D. Yates, and N. B. Mandayam, "Scheduling variable rate links via a spectrum server," in *Proc. of IEEE DySPAN*, November 2005.
- [24] T. A. Weiss, J. Hillenbrand, A. Krohn, and F. K. Jondral, "Efficient signaling of spectral resources in spectrum pooling systems," in *Proc. of Symposium on Communications and Vehicular Technology (SCVT)*, November 2003.
- [25] L. Cao and H. Zheng, "Spectrum allocation in ad hoc networks via local bargaining," in *Proc. of SECON*, September 2005.
- [26] H. Zheng and L. Cao, "Device-centric spectrum management," in *Proc. of IEEE DySPAN*, November 2005.
- [27] M. L. H. J. Huang, R. A. Berry, "Spectrum sharing with distributed interference compensation," in *Proc. of IEEE DySPAN*, November 2005.
- [28] L. Ma, X. Han, and C. C. Shen, "Dynamic open spectrum sharing mac protocol for wireless ad hoc network," in *Proc. of IEEE DySPAN*, November 2005.
- [29] Q. Zhao, L. Tong, and A. Swami, "Decentralized cognitive mac for dynamic spectrum access," in *Proc. of IEEE DySPAN*, November 2005.
- [30] N. Nie and C. Comaniciu, "Adaptive channel allocation spectrum etiquette for cognitive radio networks," in *Proc. of IEEE DySPAN*, November 2005.
- [31] S. Sankaranarayanan, P. Papadimitratos, A. Mishra, and S. Hershey, "A bandwidth sharing approach to improve licensed spectrum utilization," in *Proc. of IEEE DySPAN*, November 2005.
- [32] S. L. Wu, C. Y. Lin, Y. C. Tseng, and J. P. Sheu, "A new multi-channel (mac) protocol with on-demand channel assignment for multi-hop mobile ad hoc networks," in *Proc. of I-SPAN*, 2000, pp. 232–237.
- [33] J. So and N. Vaidya, "Multi-channel MAC for ad hoc networks: Handling multi-channel hidden terminals using a single transceiver," in *Proc. of MobiHoc*. ACM, May 2004.
- [34] R. W. Brodersen, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, "Corvus: A cognitive radio approach for usage of virtual unlicensed spectrum," UC Berkeley Whitepaper, July 2004.
- [35] J. Zhao, H. Zheng, and G. Yang, "Distributed coordination in dynamic spectrum allocation networks," in *Proc. of IEEE DySPAN*, November 2005.
- [36] X. Jing and D. Raychaudhuri, "Spectrum coexistence of ieee 802.11b and 802.16a networks using cscs etiquette protocol," in *Proc. of IEEE DySPAN*, November 2005.
- [37] M. H. M. Costa, "Writing on Dirty Paper," *IEEE Transactions on Information Theory*, vol. 29, no. 3, pp. 439–441, 1983.
- [38] C. M. Bishop, *Neural Networks for Pattern Recognition*. Oxford University Press, 1996.
- [39] (2006) Qualnet Network Simulator, version 3.9. [Online]. Available: <http://www.scalable-networks.com>