

Planning Approach to Dynamic Spectrum Access in Cognitive Radio Networks

T. Charles Clancy¹ Zhu Ji² Beibei Wang² K. J. Ray Liu²
clancy@LTSnet.net zhuji@umd.edu bebewang@umd.edu kjrlui@umd.edu

¹Laboratory for Telecommunications Sciences, US Department of Defense

²Electrical and Computer Engineering, University of Maryland, College Park

Abstract—Over the past few years, significant effort has been put into various treatments of dynamic spectrum access. This study has resulted in a many protocols and heuristic techniques for coexistence between primary and secondary users in a frequency band. However, none of these approaches has resulted in a description suitable for implementation on a general-purpose cognitive engine. As a result, devices implementing these algorithms cannot truly be called “cognitive radios” as they are not controlled by a cognitive engine.

This work focuses on translating the basic semantics of dynamic spectrum access into the Action Description Language (ADL), and in particular the primary-prioritized Markov approach to dynamic spectrum access. Consequently, this description is implemented within the Open-Source Cognitive Radio project, which is a research effort that interfaces the OSSIE SDR with the Soar cognitive engine.

I. INTRODUCTION

Dynamic Spectrum Access (DSA) is a fast-growing field in current cognitive radio research. There are a large number of protocols [1], [2], [3], [4], [5], [6], [7], [8] each with their own scheme – some centralized and some distributed, and each with their own heuristics to try and outperform others. However, these algorithms has been defined in such a way that it could be directly implemented on a *real* cognitive radio.

Here, we define a *real* cognitive radio is a wireless device governed by a *cognitive engine*, which provides the “brain”. Generally cognitive engines provide both the ability to reason (i.e. AI planning) and learn (i.e. machine learning) [9]. Reasoning is used by the engine to decide how best to act in a particular scenario given knowledge of how its action will affect its progress toward a goal. Learning is used when it’s unclear how a particular action will affect the overall system state, and that action must first be “tried out”. In more concrete terms, reasoning is an off-line search for the best system state, while learning is an on-line search.

There have been a couple efforts [10], [11] to develop a truly generic cognitive radio that can be programmed with any objective function, basic information about its environment, a set of actions, and how those actions affect its environment. Here we shall focus on the Open-Source Cognitive Radio (OSCR) [11]. We will extend the initial formulation of dynamic spectrum access presented in [12] to include the Primary-Prioritized Markov Approach (PPMA) [13], which will allow both spectrum sharing both in frequency and in time.

The final version of this paper will discuss results from actually implementing these algorithms within OSCR.

The remainder of this paper is organized as follows. Section 2 describes the basic cognitive radio architecture. Section 3 outlines how actions and goals are specified within our cognitive engine. Section 4 details how PPMA can be implemented within the description language. Section 5 discusses our experiences implementing PPMA within OSCR. Section 6 concludes.

II. COGNITIVE RADIO ARCHITECTURE

A software radio (SR) can be defined as a radio implemented with generic hardware that can be programmed to transmit and receive a variety of waveforms. Cognitive radio is often thought of as an extension to software radio, and here we treat it as such. A cognitive radio extends a software radio by adding an independent cognitive engine, composed of a knowledge-base, reasoning engine, and a learning engine, to drive software modifications. A well-defined API dictates communication between the cognitive engine and the SR. Figure 1 illustrates this architecture and the interaction between various components.

At any given time, the cognitive engine generates conclusions based on information defined in the knowledge base, the radio’s long-term memory. These conclusions are the result of extrapolations of this information based on reasoning or learning. The reasoning engine is what is often referred to in AI literature as an expert system. The learning engine is responsible for manipulating the knowledge base from experience. As lessons are learned, the learning engine stores them in the knowledge base for future reference by the reasoning engine. Depending on the application, the learning engine may only be run to train a newly initialized radio, or it could be run periodically as the radio operates. In this paper we focus on planning. For a more detailed description of how learning and planning interact within a cognitive radio, see [12].

The SR exports variables that are either read-only or read-write. The read-only parameters represent statistics maintained by the SR, such as signal to noise ratio or bit error rate. The read-write variables represent configurable parameters such as transmit power, coding rate, or symbol constellation.

These radio parameters are bound to predicates in the knowledge base. Knowledge bases are very common in AI planning. The one we describe here contains two basic data

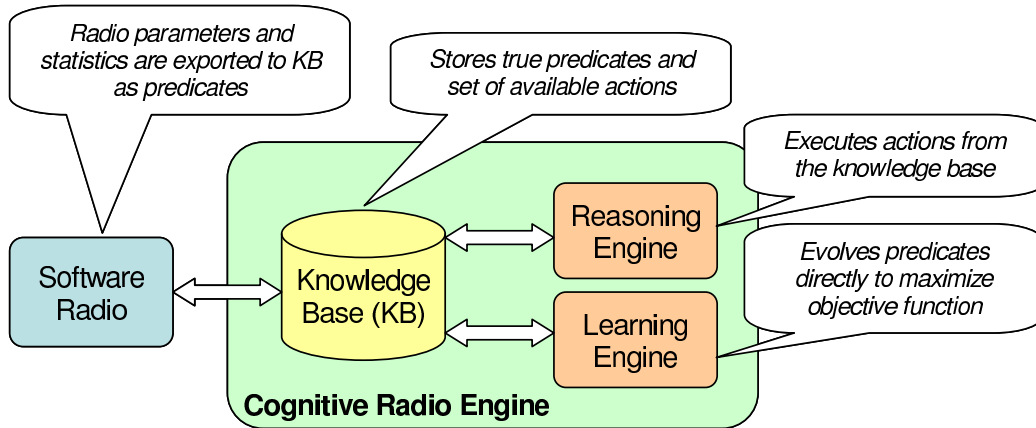


Fig. 1: Cognitive radio architecture showing the interactions between the software radio, knowledge-base, and policy and learning engines.

structures. The first is a logic expression made up of predicates that represents the state of the environment. Predicates are expressions in first-order logic that evaluate to either true or false.

The second set of data contained within the knowledge-base is actions. Actions define operations the reasoning engine could perform to change the state of its environment. Actions consist of a set of preconditions and postconditions. Preconditions must be inferable from the knowledge base and evaluate true for the action to be selected. An action's postconditions describe the modified state of the knowledge base.

To better illustrate the discussion, consider the following example, the objective of which is to decrease the modulation rate with a decrease in SNR.

The knowledge base contains the following predicates¹

$$\text{modRate}(QPSK) \wedge \text{snr}(5 \text{ dB}) \quad (1)$$

and the following action

action : decreaseModulationRate

$$\text{precond} : \text{modRate}(QPSK) \wedge \text{snr}(\leq 8 \text{ dB}) \quad (2)$$

$$\text{postcond} : \neg \text{modRate}(QPSK) \wedge \text{modRate}(BPSK)$$

The reasoning engine uses planning, which is a field of AI that works with logic². At any given time, it looks at the current state and determines which actions are executable in that state. All the possible resulting states are then evaluated to see which is optimal, where optimality is determined by an objective function $f(\cdot)$.

In our current example, we can successfully infer the preconditions from our knowledge-base. As a result, the **decreaseModulationRate** action is executed and the postcondi-

tions are applied to the knowledge-base, resulting in

$$KB' = KB \wedge \text{postcond} = \text{modRate}(BPSK) \wedge \text{snr}(5 \text{ dB}) \quad (3)$$

Observe how modulation was changed from QPSK to BPSK when the SNR drops below 8 dB. While this example may seem elementary, it provides the fundamentals for reasoning in our cognitive radio.

III. DYNAMIC SPECTRUM ACCESS

In this section, we first describe a dynamic spectrum access (DSA) system. Then we propose a Primary-Prioritized Markov Approach (PPMA) for dynamic spectrum access which can efficiently and fairly utilize the spectrum resources, and implement the PPMA using AI logic.

A. DSA System Description

Dynamic Spectrum Access (DSA) involves locating frequency bands and times when which a cognitive radio can transmit without causing harmful interference to other transceivers [14]. For example, consider a cognitive radio network operating in the UHF television bands, where transmission is permissible provided devices can guarantee they will not interfere with licensed broadcasts.

More concretely, the goal is to locate center frequencies, bandwidths, and times when which a cognitive radio can transmit, while maximizing capacity and minimizing interference.

If we limit our DSA approach to multiplexing in frequency with licensed signals, then the problem simplifies significantly. Imagine our SR exports predicates to the knowledge-base regarding detected signals s_1, s_2, \dots, s_N that are of the form

$$\text{signalFreq}(s_i, f_i) \wedge \text{signalBW}(s_i, W_i) \quad (4)$$

Our goal is to find some f_c and W that does not overlap any detected signal, while maximizing W and consequently the radio's capacity. Let's also include a parameter $T \in [0, 1]$ representing the radio's duty cycle.

¹We use the conventional logic operators: \wedge AND; \vee OR; \neg NOT.

²We mostly consider scenarios that use forward chaining rather than backward chaining for inferencing, since we often don't have a particular goal state.

First, let's define a helper function

$$\begin{aligned} \text{notOverlap}(f_c, W, s_i) \\ = (f_i + W_i/2 < f_c - W/2) \vee (f_i - W_i/2 > f_c + W/2) \end{aligned} \quad (5)$$

Then we can define our predicate

$$\begin{aligned} \text{action : moveBand}(\mathbf{f}_{\text{old}}, \mathbf{W}_{\text{old}}, \mathbf{T}_{\text{old}}, \mathbf{f}_{\text{new}}, \mathbf{W}_{\text{new}}) \\ \text{precond : } \forall i \leq N : \text{notOverlap}(f_{\text{new}}, W_{\text{new}}, s_i) \\ \quad \wedge \text{centerFreq}(f_{\text{old}}) \wedge \text{bandwidth}(W_{\text{old}}) \\ \quad \wedge \text{dutyCycle}(T_{\text{old}}) \\ \text{postcond : } \neg(\text{centerFreq}(f_{\text{old}}) \wedge \text{bandwidth}(W_{\text{old}}) \\ \quad \wedge \text{dutyCycle}(T_{\text{old}})) \\ \quad \wedge \text{centerFreq}(f_{\text{new}}) \wedge \text{bandwidth}(W_{\text{new}}) \\ \quad \wedge \text{dutyCycle}(1) \end{aligned} \quad (6)$$

Then we define our objective function (with a maximization goal)

$$f(\text{bandwidth}(W)) = W \quad (7)$$

We now have a policy-based cognitive radio that will search out the largest continuous piece of bandwidth for communication.

However, imagine that the largest continuous piece of bandwidth was relatively small. Our best strategy may be to try and coexist in the time domain with a primary signal. When the primary signal is transmitting, we cease transmitting, and when it's not transmitting, we resume.

Let's say our signals s_i are extended to include duty cycle and are of the form

$$\text{signalFreq}(s_i, f_i) \wedge \text{signalBW}(s_i, W_i) \wedge \text{signalDC}(s_i, T_i) \quad (8)$$

We can now add an additional predicate to our knowledge base:

$$\begin{aligned} \text{action : moveSameBand}(\mathbf{f}_{\text{old}}, \mathbf{W}_{\text{old}}, \mathbf{T}_{\text{old}}, \mathbf{f}_{\text{new}}, \mathbf{W}_{\text{new}}) \\ \text{precond : } \exists i \leq N : \neg \text{notOverlap}(f_{\text{new}}, W_{\text{new}}, s_i) \\ \quad \wedge \text{centerFreq}(f_{\text{old}}) \wedge \text{bandwidth}(W_{\text{old}}) \\ \quad \wedge \text{dutyCycle}(T_{\text{old}}) \\ \text{postcond : } \neg(\text{centerFreq}(f_{\text{old}}) \wedge \text{bandwidth}(W_{\text{old}}) \\ \quad \wedge \text{dutyCycle}(T_{\text{old}})) \\ \quad \wedge \text{centerFreq}(f_{\text{new}}) \wedge \text{bandwidth}(W_{\text{new}}) \\ \quad \wedge \text{dutyCycle}(1 - T_i) \end{aligned} \quad (9)$$

This additional action allows us to occupy the same frequency band as a licensed signal with an appropriate duty cycle. We can adapt our objective function to be

$$f(\text{bandwidth}(W) \wedge \text{dutyCycle}(T)) = W \cdot T \quad (10)$$

Thus when using these actions together, we can evaluate the objective function over all possible choices.




Fig. 2: The rate diagram of PP-CTMC.

B. Primary-Prioritized Markov Approach

Once frequency-band selection is complete, we now need the ability to detect licensed users and halt transmission. In this subsection, we propose the Primary-Prioritized Markov Approach (PPMA) for dynamic spectrum access, since it works well on a basis of a fairly simple state machine that can be expressed through AI logic with relative ease.

We first describe how PPMA coordinates the access of multiple secondary users in the temporarily unused licensed band without conflicting with the primary spectrum holders' usage. Denote the primary user by P , and the secondary users by A and B . For each user γ , where $\gamma \in \{A, B, P\}$, its offered traffic is modeled with two independent Poisson processes, with the service-request rate λ_γ and the departure rate μ_γ .

Since the primary user's spectrum usage in its licensed band should not be affected by the operation of any other secondary user, we assume that once primary user P appears, any secondary user should stop transmission, buffer their interrupted traffic, continue scanning the licensed band, and immediately resume transmission once the licensed band becomes idle again. Moreover, if more than one secondary users are allowed to share the licensed frequency-band, the efficiency of spectrum usage can be further improved. However, too much coexistence of the secondary users may result in mutual interference, and we will describe later how to coordinate the access of secondary users to alleviate the interference level.

From the preceding assumptions, we model the interactions between the primary user and the secondary users as a Primary-Prioritized Continuous-Time Markov Chain (PP-CTMC), illustrated in Figure 2. The states of PP-CTMC are described in Table I.

Assume at first the licensed spectrum band is idle, i.e., PP-CTMC is in state $(0, 0)$. The two secondary users contend to operate in the spectrum. Upon the first access attempt of some user, say user A , PP-CTMC enters state $(0, A)$ with transition rate λ_A . If user A finishes its service duration before any other user requests spectrum access, PP-CTMC then transits to state $(0, 0)$ with rate μ_A . If user B 's service request arrives before A

TABLE I: The Eight States of PP-CTMC

Index	State	Description
0	(0, 0)	Spectrum is idle
1	(0, B)	Secondary user B is in service
2	(0, A)	Secondary user A is in service
3	(0, AB)	Both A and B are in service
4	(P, 0)	Primary user P is in service
5	(P, B _w)	P is in service; B is waiting
6	(P, A _w)	P is in service; A is waiting
7	(P, (AB) _w)	P is in service; A and B are waiting

completes its service, PP-CTMC transits to state (0, AB) with rate λ_B , where both secondary users share the spectrum usage using Code-Division Multiple Access (CDMA) techniques. Once user B (or A)'s service duration is completed, PP-CTMC transits from state (0, AB) to (0, A) (or (0, B)), with rate μ_B (or μ_A).

However, primary user P may, once in a while, appear during the service duration of the secondary users, i.e., when PP-CTMC is in state (0, A), (0, B) or (0, AB). Suppose the licensed band is being occupied by user A. If user A detects that primary user P needs to acquire the spectrum band, A ceases its transmission, buffers its interrupted traffic, and keeps sensing the band until P finishes operating in the band. Therefore, PP-CTMC transits from state (0, A) to (P, A_w) with rate λ_P . If primary user P finishes its service before B's access, A will continue its transmission, and PP-CTMC transits from state (P, A_w) to (0, A) with rate μ_P . In contrast, if the access request of B arrives before primary user P completes its service duration, B also buffers its traffic, and PP-CTMC transits to state (P, (AB)_w) with rate λ_B . In state (P, (AB)_w), both A and B keep sensing the spectrum. Once P is sensed to vacate, PP-CTMC transits to state (0, AB) with rate μ_P , where A and B share the spectrum band. Also, when PP-CTMC is in state (P, 0), if unlicensed users attempt to access the spectrum, they are kept sensing until P finishes its service, and PP-CTMC transits to state (P, A_w) or (P, B_w), with rate λ_A or λ_B , respectively.

One of the most important goals in spectrum sharing is efficient spectrum utilization, i.e., high throughput achieved by each secondary user through successful acquisition of a spectrum band. From a statistical point of view, the secondary users want to maximize their average throughput. Given the rate diagram of PP-CTMC, we can obtain its stationary state probabilities, denoted by Π_{s_i} , where $s_i \in \{(0, A), (0, B), (0, AB)\}$. Since Π_{s_i} can be equivalently viewed as the ratio of allocation time to state s_i to the entire reference time, the product of Π_{s_i} and the capacity that secondary user γ achieves when operating in state s_i represents one average throughput component acquired by user γ in state s_i . Therefore, we can express the total average throughput for user γ as follows,

$$U_\gamma = \Pi_{(0,\gamma)} r_1^\gamma + \Pi_{(0,AB)} r_2^\gamma, \quad (11)$$

where r_1^γ and r_2^γ are channel capacities for user γ when it operates in the licensed band alone and with the other secondary user, respectively.

In order to alleviate the mutual interference among secondary users in PPMA, we introduce the spectrum access probability for user A and user B, denoted by a_A and a_B , respectively. Then, the resulting random access process can be approximated by slightly modifying the original CTMC. Because each secondary user γ 's traffic is admitted with probability a_γ , the actual arrival rate is approximated by $a_\gamma \lambda_\gamma$.

Then, the optimization goal is to determine a_A and a_B , such that the utility function can be maximized, i.e.,

$$\{a_\gamma\} = \arg \max_{0 \leq a_\gamma \leq 1} U(\{a_\gamma\}), \quad (12)$$

where $\forall \gamma \in \{A, B\}$.

According to different objectives, the utility function $U(\{a_\gamma\})$ can have different definitions. A good spectrum sharing scheme not only can efficiently utilize the spectrum resources, but also can provide fairness among different users, so we first propose to maximize the average throughput based on the proportional-fairness (PF) criterion. Thus, in (12), $U(a_A, a_B)$ can be written as

$$U(a_A, a_B) = \prod_{\gamma \in \{A, B\}} U_\gamma(a_A, a_B). \quad (13)$$

Other optimality criteria can also be employed, such as the maximal-throughput criterion

$$U(a_A, a_B) = \sum_{\gamma \in \{A, B\}} U_\gamma(a_A, a_B), \quad (14)$$

and the max-min fairness criterion

$$U(a_A, a_B) = \min_{\gamma \in \{A, B\}} U_\gamma(a_A, a_B). \quad (15)$$

C. Implementation of PPMA

According to the preceding discussions, we can implement the PPMA in the following steps.

Assume the frequency-band selection is complete, which is authorized to primary signal s_P . The knowledge base contains the following predicate

$$\begin{aligned} & \text{signalFreq}(s_P, f_P) \wedge \text{signalBW}(s_P, W_P) \\ & \wedge \text{signalDC}(s_P, T_P). \end{aligned} \quad (16)$$

The slave software radio is assumed to be able to measure the traffic statistics of different users, then we can add additional predicates to our knowledge base:

$$\text{arrivalRate}(\lambda_P) \wedge \text{serviceRate}(\mu_P), \quad (17)$$

and

$$\text{arrivalRate}(\lambda_\gamma) \wedge \text{serviceRate}(\mu_\gamma), \quad (18)$$

where $\gamma \in \{A, B\}$.

Having these measurements of the traffic statistics and the rate diagram in Figure 2, according to [13], the software radio can construct the equation arrays governing the DSA system, and compute the stationary state probabilities $\{\Pi_{s_i}\}$, where $s_i \in \{(0, A), (0, B), (0, AB)\}$. Since they are equivalent to the allocation time ratios to the secondary users, which

are functions of their service-request rates, we denote the corresponding predicates as follows

$$\begin{aligned} & \text{timeRatioA}(\{\lambda_\gamma\}) \wedge \text{timeRatioB}(\{\lambda_\gamma\}) \\ & \wedge \text{timeRatioAB}(\{\lambda_\gamma\}). \end{aligned} \quad (19)$$

At time T_P , primary user P is sensed to complete its service and vacate the frequency-band, then the secondary users begin to operate in the band. The corresponding predicates are defined as

$$\begin{aligned} & \text{action : startTraffic}(s_\gamma) \\ & \text{precond : } \neg(\text{centerFreq}(f_\gamma) \wedge \text{bandwidth}(W_\gamma) \\ & \quad \wedge \text{dutyCycle}(T_\gamma)) \\ & \text{postcond : } \text{centerFreq}(f_\gamma) \wedge \text{bandwidth}(W_\gamma) \\ & \quad \wedge \text{dutyCycle}(T_\gamma), \end{aligned} \quad (20)$$

where $\gamma \in \{A, B\}$.

Assume after the secondary users start their transmissions in the licensed band, our software radio can measure the capacity they can achieve. Then we can add the corresponding predicates as $COneUser(\gamma)$ and $CTwoUser(\gamma)$. The resulting average throughput for each secondary user is

$$\begin{aligned} U(\gamma) = & \text{timeRatioA}(\{\lambda_\gamma\})COneUser(\gamma) \\ & + \text{timeRatioAB}(\{\lambda_\gamma\})CTwoUser(\gamma), \end{aligned} \quad (21)$$

where $\gamma \in \{A, B\}$.

Therefore, the objective function is

$$\begin{aligned} & f(\text{timeRatioA}(\{\lambda_\gamma\}) \wedge \text{timeRatioB}(\{\lambda_\gamma\}) \\ & \wedge \text{timeRatioAB}(\{\lambda_\gamma\}) \wedge COneUser(\gamma) \wedge CTwoUser(\gamma)) \\ & = U(\{\lambda_\gamma\}), \end{aligned} \quad (22)$$

where $U(\{\lambda_\gamma\})$ can be selected from (13)-(15).

Since we have proved in [13] that the objective function for the PF criterion is concave in the access probabilities and the traffic arrival rates, we are then able to apply a gradient search algorithm to find the optimal solution. If we denote the optimal arrival rate as predicates $\lambda_\gamma^{\text{new}}$, the following actions should be executed in the knowledge base,

$$\begin{aligned} & \text{action : adaptArrivalRate}(s_\gamma, \lambda_\gamma, \lambda_\gamma^{\text{new}}) \\ & \text{precond : } \text{arrivalRate}(\lambda_\gamma) \wedge (\lambda_\gamma^{\text{new}} \neq \lambda_\gamma) \\ & \text{postcond : } \text{arrivalRate}(\lambda_\gamma^{\text{new}}). \end{aligned} \quad (23)$$

Consider the scenarios that primary user P is sensed to acquire its frequency-band with some new duty cycle T_P^{new} , however, the secondary users are still in their duty cycles. Then the secondary users should buffer their interrupted traffic, and the predicate is defined as

$$\begin{aligned} & \text{action : bufferTraffic}(s_\gamma) \\ & \text{precond : } (\text{centerFreq}(f_\gamma) \wedge \text{bandwidth}(W_\gamma) \\ & \quad \wedge \text{dutyCycle}(T_\gamma)) \wedge \neg(\text{centerFreq}(f_P) \\ & \quad \wedge \text{bandwidth}(W_P) \wedge \text{dutyCycle}(T_P^{\text{new}})) \\ & \text{postcond : } \neg(\text{centerFreq}(f_\gamma) \wedge \text{bandwidth}(W_\gamma) \\ & \quad \wedge \text{dutyCycle}(T_\gamma)) \wedge (\text{centerFreq}(f_P) \\ & \quad \wedge \text{bandwidth}(W_P) \wedge \text{dutyCycle}(T_P^{\text{new}})), \end{aligned} \quad (24)$$

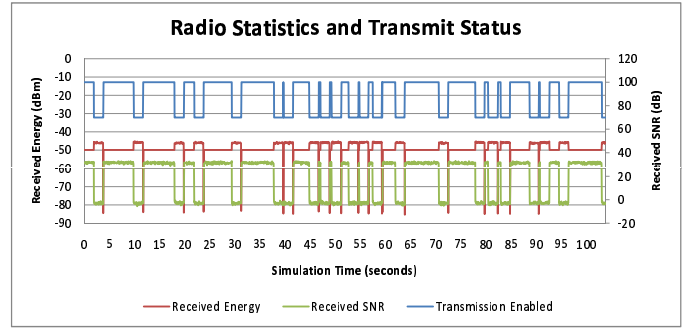


Fig. 3: Plot of the energy and SNR statistics measured during the experiment, along with the decision on whether or not to enable the secondary transmitter

where $\gamma \in \{A, B\}$.

When the primary user finishes its service, the secondary users immediately resume their originally interrupted traffic, and the predicate is

$$\begin{aligned} & \text{action : resumeTraffic}(s_\gamma) \\ & \text{precond : } \neg(\text{centerFreq}(f_\gamma) \wedge \text{bandwidth}(W_\gamma) \\ & \quad \wedge \text{dutyCycle}(T_\gamma)) \wedge (\text{centerFreq}(f_P) \\ & \quad \wedge \text{bandwidth}(W_P) \wedge \text{dutyCycle}(T_P^{\text{new}})) \\ & \text{postcond : } (\text{centerFreq}(f_\gamma) \wedge \text{bandwidth}(W_\gamma) \\ & \quad \wedge \text{dutyCycle}(T_\gamma)) \wedge \neg(\text{centerFreq}(f_P) \\ & \quad \wedge \text{bandwidth}(W_P) \wedge \text{dutyCycle}(T_P^{\text{new}})). \end{aligned} \quad (25)$$

If the secondary users request the spectrum access during the primary user's duty cycle, they can buffer and resume the traffic using similar actions as defined in (24) and (25).

IV. IMPLEMENTATION RESULTS

The PPMA algorithms were implemented within the Open-Source Cognitive Radio (OSCR) [11]. In order to perform dynamic spectrum access, the radio needs to know whether the primary and/or secondary users are transmitting. The existing OSCR implementation already provided us with an SNR statistic that could be used to determine if a secondary user was transmitting. An additional statistic was added that computes the energy at the secondary receiver. Using both these statistics, the radio can determine if anyone was transmitting, and if so whether it was the primary or secondary user.

The radio then used this data to decide whether or not to activate the secondary transmitter. The statistics and ability to start and stop the transmitter was exported to the Soar cognitive engine, which then applied the PPMA algorithms to control the transmitter.

Figure 3 shows the experimental results. A received energy of -50 dBm indicates the secondary user is transmitting, while a received energy of -45 dBm indicates a primary user is transmitting. Once the primary user starts transmitting, the received SNR drops due to the interference. The cognitive engine's logic then triggers a stop in the secondary transmitter.




Fig. 4: Screen-shot of the initial state, where the primary user is transmitting and the secondary AP is idle.

Once the received energy drops below -80 dBm, the secondary user knows the primary user has vacated the channel, and resumes transmission. In our experimentation, the cognitive engine was able to adapt in 5 milliseconds to a change in the state of the primary user. This reaction time could be reduced by more frequent polling of the radio statistics.

Additionally, an implementation was created within GNU-radio to demonstrate RF over-the-air functionality with the Universal Software Radio Peripheral (USRP). Two USRP boards were used, one as a primary user and another as a secondary user access point (AP). The secondary user AP used energy detection to sense the presence of the primary user. If a primary user was detected, the AP prompts the secondary users to stop transmitting. Otherwise the AP will coordinate secondary user transmissions in the available spectrum.

The energy detection algorithm implemented within GNU-radio is an FFT block fed into a squaring and averaging block. The output is compared to a decision threshold to determine if a primary user's signal is present. Figures 4 and 5 show the power spectra for both states $(P, 0)$ and $(0, A)$ in the PPMA state machine.

V. CONCLUSION

Since the introduction of cognitive radio, there have been many high-level discussions on proposed capabilities of cognitive radios. In this article, we have formalized some of the architecture behind dynamic spectrum access-based applications.

Certainly there is a great deal of future work in the field of cognitive radio. The architecture described here is flexible enough to address many different DSA protocols, provided they can be expressed in predicates, actions, and objective functions.

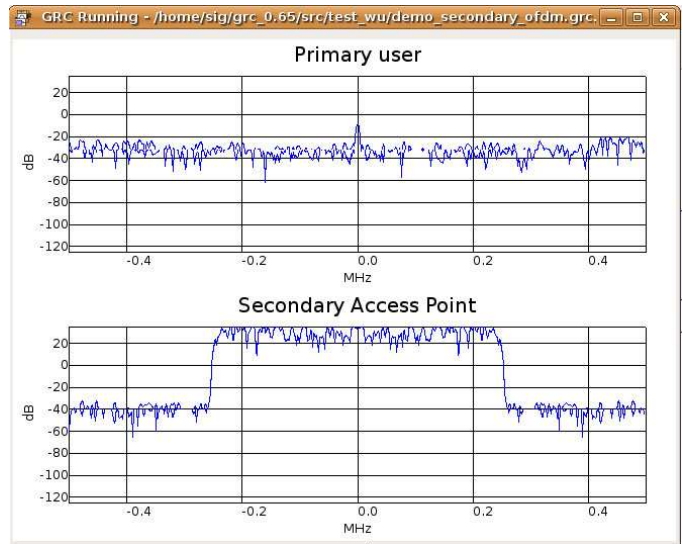


Fig. 5: Screen-shot of the transition state where the primary user has vacated the channel and the secondary user accesses it opportunistically.

VI. ACKNOWLEDGMENTS

The authors would like to thank Mahmoud AbdulRehem, Yongle Wu, and Tim O'Shea for their valuable implementation assistance.

REFERENCES

- [1] K. Challapali, S. Mangold, and Z. Zhong, "Spectrum agile radio: Detecting spectrum opportunities," IEEE GlobeCOM 2004.
- [2] D. Cabric, S. Mishra, D. Willkomm, R. Brodersen, and A. Wolisz, "A cognitive radio approach for usage of virtual unlicensed spectrum," IST Mobile Summit 2005.
- [3] C. Raman, R. Yates, and N. Mandayam, "Scheduling variable rate links via a spectrum server," IEEE DySPAN 2005.
- [4] N. Nie and C. Comaniciu, "Adaptive channel allocation spectrum etiquette for cognitive radio networks," IEEE DySPAN 2005.
- [5] L. Ma, X. Han, and C. Shen, "Dynamic open spectrum sharing MAC protocol for wireless ad hoc networks," IEEE DySPAN 2005.
- [6] H. Zheng and L. Cao, "Device-centric spectrum management," IEEE DySPAN 2005.
- [7] R. Etkin, A. Parekh, and D. Tse, "Spectrum sharing for unlicensed bands," IEEE DySPAN 2005.
- [8] Z. Ji and R. Liu, "Dynamic spectrum sharing: a game theoretical overview," to appear in *IEEE Communications Magazine (Feature Topic on Cognitive Radios for Dynamic Spectrum Access)*, 2007.
- [9] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2002.
- [10] T. Rondeau, B. Le, D. Maldonado, D. Scaperth, and C. Bostian, "Cognitive radio formulation and implementation," IEEE CROWNCOM 2006.
- [11] E. Stuntebeck, T. O'Shea, J. Hecker, and T. Clancy, "Architecture for an open-source cognitive radio," SDR Forum Conference 2006.
- [12] T. Clancy, J. Hecker, E. Stuntebeck, and T. O'Shea, "Applications of machine learning to cognitive radio networks," *IEEE Wireless Communications Magazine*, 2007.
- [13] B. Wang, Z. Ji, K. J. R. Liu, and T. Clancy, "Primary-prioritized markov approach for dynamic spectrum allocation," *ACM Mobile Networks and Applications Journal*, under submission.
- [14] P. Leaves, S. Ghaferi-Niri, R. Tafazolli, L. Christodoulides, T. Sammur, W. Stahl, and J. Huschke, "Dynamic spectrum allocation in a multi-radio environment: Concept and algorithm," Conference on 3G Mobile Communication Technologies 2001.