Experiments in Cognitive Radio and Dynamic Spectrum Access using An Ontology-Rule Hybrid Architecture

Allen Ginsberg, Jeffrey D. Poston, and William D. Horne
The MITRE Corporation
McLean, VA
{aginsberg, jdposton, whorne}@MITRE.org

Abstract

Cognitive Radio and Dynamic Spectrum Access represent two complementary developments that will refashion the world of wireless communication. In order to investigate the roles of knowledge representation and reasoning technologies in this domain, we have developed, and continue to extend, an experimental cognitive radio simulation environment in which dynamic multi-device scenarios can be modeled. One of our goals has been to show how these technologies can be used to provide a "knowledge-driven differential-response" capability. That is, a conventional radio when operating in a particular communications mode always follows the same procedure and either succeeds or fails at a given task. A cognitive radio, by contrast, can use knowledge of radio technology and policy, representations of goals, and other contextual parameters to reason about a failed attempt to satisfy a goal and attempt alternate courses of action depending upon the circumstances. We talk about the roles ontologies and rules play in the various components of the simulation architecture that enables differential response.

1. Introduction

The terms "Cognitive Radio" and "Dynamic Spectrum Access" arise from recent technological and regulatory trends in the area of wireless communication. A "radio" is any device that can receive and transmit radio frequency signals, everything from cell phones to WiFi-enabled computers. Recent advances in radio technology now allow radios to be "agile" in the sense that they can reconfigure themselves to dynamically receive and transmit on a wide variety of channels using a variety of communication modes. A cognitive radio is a radio that has some ability to represent and reason about goals, policies, features of the external world, aspects of its internal state, etc., that relate to radio communication. For example a cognitive cell-phone, upon boarding an aircraft, could use its policy knowledge, context awareness, and reasoning capability to automatically disable its radio frequency functions.

Organizations that regulate radio frequency spectrum access are beginning to take account of these developments. For example, crowding on certain areas of the spectrum, has led regulators to propose and, in some regions, adopt, policies under which portions of licensed or regulated spectrum - spectrum that has been pre-allocated to certain users or uses - can be dynamically accessed by secondary users when not in use by the licensed users. For example, in some low-population areas of the US, reuse of idle analog TV spectrum for wireless internet access seems to be a good way of providing service to areas that are currently under-served.

A world of cognitive radios will be a world in which policy-makers and designers/implementers will need to share a common understanding of evolving radio technologies and cognitive radios themselves will need to be able to communicate their status, priorities and goals with each other in an unambiguous manner. In order to investigate and determine the knowledge representation and reasoning technologies and capabilities required to achieve this vision, we have developed, and continue to extend, an experimental cognitive radio simulation environment in which dynamic multi-device scenarios can be modeled. In this paper we discuss the problem area in more detail and present a detailed example scenario. We talk about the roles ontologies and rules play in the various components of the simulation architecture both from a semantic standpoint and from the point of view of orchestrating onotology-based reasoning with rule-execution, as well as the execution of methods defined in an object-oriented programming language. We close by considering what aspects of this hybrid architecture represent a valid generalizable approach to a principled framework for further development.

2. Background and Related Work

The advent of *Software Defined Radio* (SDR) [1] technology offers far greater processing resources than prior radio technology. With this increased capability comes the burden of defining architecture constructs and developing the software applicable for the various scenarios the SDR may encounter. One technology that promises to not only utilize this processing capability but to also provide an autonomous and flexible architecture that is applicable to a wide array of operational scenarios is the *Cognitive Radio* (CR) [2]. The ultimate vision of CR technology—denoted by Mitola as the "ideal cognitive radio (iCR)"—encompasses many facets of intelligent behavior such as context awareness, adaptation of action due to stimulus and prior information, inferring information not explicitly stated, learning, natural language processing, and planning. A growing research community is investigating the means for taking advantage of the processing resources in SDR platforms to develop the iCR; to date, most researchers choose to focus on one or a few of these facets of intelligence. As a result, literature on the subject defines the term CR in a variety ways, usually in a narrow, applications-specific manner.

As an example CR application, the Federal Communications Commission (FCC) is considering an application where vacant portions of the TV broadcast bands could be shared with unlicensed devices with sufficient intelligence to detect the licensed users and avoid causing harmful interference to them [3]; a related effort by the IEEE 802.22 standards committee seeks to create a technical standard for a network of these devices [4]. While this definition of CR does include a radio with some awareness of the spectrum and some ability to adapt operating behavior based upon that information, this definition otherwise duplicates conventional radio technology with procedural-style specification of the radio's behavior.

The DARPA XG program [5] aims to demonstrate opportunistic spectrum access of otherwise idle spectrum under a range of conditions. An important component of that application is a policy checking entity that determines whether or not the dynamic spectrum access adheres to a policy. Their current approach employs a Prolog-based policy reasoner to evaluate such queries [6]. Other researchers such as Berlmann et al. [7] proposed policy-based reasoning to check a broader range of CR behaviors.

Work by Neel et al. [8] applied game theory principles to design distributed algorithms for adaptive behaviors. Rondeau et al. [9] proposed genetic algorithms for optimizing the settings of the many control parameters available to CRs. Both [8] and [9] addressed the problem of adaptation in CRs, and the latter could also be viewed as addressing a learning component.

The research group of Kokar investigated how to create CRs with self awareness of their own capabilities via an ontology framework [10] and how to replace procedural-style radio control constructs with machine reasoning techniques.

The use of an ontology as a knowledge representation mechanism is central to this paper's approach. On a related note, some of the authors of this paper [11] took an ontology-based approach to providing CRs with context awareness; for example, the term radio "channel" has many possible meanings, and in order to reason about the availability of a channel the CR must know what definition applies in a given context.

Recognizing the possible applications and approaches for introducing machine reasoning to radio systems as discussed in this section, this paper explores a CR incorporating a "differential-response" capability by augmenting the existing SDR processing paradigm using ontologies and rules. Sections that follow note the importance of a "knowledge-driven differential-response" capability not found in prior work and describe components necessary to instantiate it. Finally, the paper closes by describing a simulated prototype CR with this capability and how it can achieve goals despite facing conflicts that would have thwarted conventional radios.

2. Motivation and requirements for knowledge-driven differential-response

As a motivating example, this paper considers the problem of a CR attempting to gain access to a portion of the radio band governed by radio beacons at one or more locations. A number of beacon-based protocols are possible to facilitate dynamic spectrum access. This example assumes a policy regime in which *both* positive and negative control beacons are employed. Thus in order for the CR to be able to access a radio channel C, two conditions must be fulfilled: 1) the CR must be within radio range of a beacon station from which it receives a coded beacon message authorizing access to channel C [12], and 2) the CR must not simultaneously be in range of a beacon station from which it receives a coded beacon message *denying* access to C. Although there are a number of security, protocol, and radio engineering issues that must be addressed in the design and implementation of such beacons, it is not necessary to describe those aspects of the beacon in order to appreciate the value of differential-response capability.

For both conventional radios and the types of CRs described in the previous section, if the radio's location is such that conditions (1) and (2) are not satisfied, it cannot access the channel, and—more fundamentally—it cannot reason about why the goal of channel access has failed or what alternative conditions would permit overcoming the failure.

Note that there are essentially 2 ways in which the goal of using channel C can be thwarted: a) no positive beacon signal for access to C is received, and b) both positive and negative beacon signals for access to C are received. The cases in which no beacon signals are detected at all, or in which only negative beacon signals are received can be viewed as being subsumed under the other cases. For example, in the case where only negative signals are received, the radio needs to somehow get a positive signal, which is case (a). If it manages to solve that problem and it still is receiving one or more negative signals, then it is now in case (b) (otherwise it has solved the problem of gaining access).

The situations covered by cases (a) and (b) can be used to elaborate upon the notion of *knowledge-driven* differential-response. In both cases a radio will fail in its goal of getting access to channel C. Depending upon its functionality, a conventional radio might be able to distinguish between the two cases of failure, in the sense that it *goes* into a different internal state depending upon the circumstances. Through a user-interface it may be able to give an indication of its current state. However, even if the concrete indicators conveyed by the radio are different in the two cases, this is still not a knowledge-driven differential-response. First of all, one can say with some justice, that in both cases the radio is really doing exactly the same thing: upon failure to access a channel, convey the current internal state to the user. That is an accurate description of the radio's actions because that is exactly how the radio is programmed to behave. Secondly, even if, for the sake of argument, the responses are deemed to be different, they do not illustrate *knowledge-driven* differential-response. The reason is that the radio being in a distinct internal state is an irreducible and non-analyzable *cause* of its taking whatever action it takes. From a formal point of view, we may say that the radio behaves as a simple finite-state machine. The particulars of the internal state and the way in which those particulars relate to aspects of the external world do not enter into an explanation of the radio's behavior. The latter is at least part of what is required in order for any agent to be capable of cognition, and this is what we mean by *knowledge-driven differential-response.*

So, returning to the example scenario, what could a CR do in case (b) as opposed to case (a)? A CR would have representations of policy conditions (1) and (2) and it could also have a representation of a *beacon signal conflict*, that is, a situation in which two or more conflicting beacon signals are received. It would also have the knowledge, expressed in a *rule*, that when a goal cannot be achieved due to such a conflict, one can attempt to move to a location where only the desired beacon is in range. Depending on the radio's mobility capabilities, it could then act in a number of ways depending on the circumstances. For example, using its inherent signal strength detection capability it could guide its user to a region where only the desired beacon signal is received.

We have implemented a prototype simulation in which a CR exhibits this type of functionality by successfully guiding a user to a region where the desired channel is accessible. This is described later in the paper. First we need to extrapolate a general set of requirements that are entailed by the discussion of this example. We may distinguish three broad categories: knowledge and reasoning requirements, perceptual requirements, and action requirements.

2.1. Knowledge and reasoning requirements

As we shall see later, knowledge requirements for CR are of two sorts: conceptual and rule-based. Conceptual knowledge includes knowing the meanings of fundamental notions in a domain of interest as well as fundamental principles relating those concepts. In current practice, this kind of knowledge is said to be *ontological* and is formally encoded in knowledge representation frameworks know as *ontology* languages. *An* ontology is a formal representation of the key concepts and principles of a domain of interest. Rule-based knowledge, which is typically represented using some formal rule-language, can be thought of as the "bridge" that relates conceptual knowledge to the problem-solving needs of a particular application.

Knowledge and reasoning go hand-in-hand. A piece of knowledge that is not somehow related to other pieces of knowledge through inference (reasoning) is basically useless. How information and inference should function together effectively depends on the application and its goals. Since the knowledge and reasoning requirements are fundamental to the CR we discuss them in detail in a separate section below (see "Role of Knowledge Representation").

2.2. Perceptual requirements

It is clear from the example scenario that a CR must be able to recognize sensory inputs, or patterns thereof, as *being* or *representing* something in its environment. For example, a CR needs to have some way of translating the low-level events and processes occurring in its signal processing software and hardware into a construct that represents the fact that it perceives a signal with such-and-such properties.

Although beyond the scope of this paper, we also include *self-perception* in this category. Broadly speaking, a CR must be able to know what it is doing and why it is doing it. This implies that a CR must be able to recognize certain of its internal states and processes as representing certain facts about *itself*. For example, a CR should be able to perceive its current level of power consumption or its current memory disposition as properties belonging to itself just as a human perceives a bodily sensation such as pain as something that is internal.

2.2. Action requirements

A CR needs to be able to initiate action based upon the conclusions it reaches. For example, if a CR decides that it should attempt to communicate with another trusted host on behalf of its user in some emergency, then it must be able to "translate that thought into action," so to speak. Action usually requires a certain amount of knowledge but also, at some point, requires the basic ability to *do* something. For example, engaging in a protocol with another CR would involve knowledge of the protocol, but is also means that when a message is to be sent, the CR has the ability to cause that to happen.

Requirements in categories (2) and (3) together delineate what we call the *World Model* and *Self Model* of a CR. A CR's world model is a time-varying construct that represents the state of the world as perceived by the CR. For example, at a certain time a particular signal may be in the world model of a CR. That "same signal" may also exist in the world model of another CR located nearby. That is because objects in the world can be perceived and acted upon by multiple agents.

A CR's self model is a time-varying construct that represents the internal state of that CR as well as its basic capabilities and properties. Something that is in the self model of one CR cannot literally represent the same item as anything in another CR's self model. For example, suppose two CRs, CR-1 and CR-2, both have the ability to perform a certain basic action, e.g., transmit a signal on a certain channel. The self models of CR-1 and CR-2 will both contain a representation of that action. However, these cannot be viewed as literally representing the same action. Otherwise, there would be no difference at all between CR-1 and CR-2 transmitting on that channel.

3. Architecture

The discussion thus far can be encapsulated in a proposed architecture for a CR device, as shown in Figure 1. In the left side of the figure we see that the goal of augmenting SDR processing structures is accomplished in this architecture by means of the *Perception & Action Abstraction Layer* (PAAL). The Perception and Action Abstraction Layer (PAAL) is defined in terms of certain *standard* radio concepts and is used to characterize device observables and actions in a *platform-independent* knowledge representation.

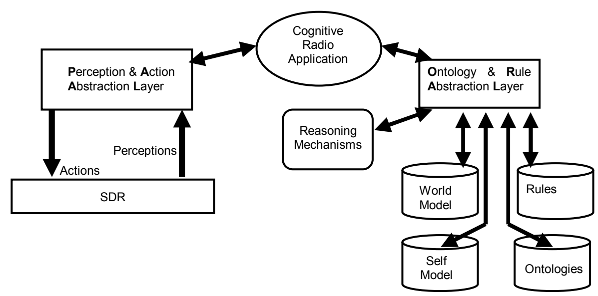

Figure 1. CR Architecture

This is a key layer if one wants to allow for reuse of the cognitive portion of the architecture with different conventional radio implementations. That is, different radios could use different signal processing algorithms at a very low level that have no bearing on whether or not something is an instance of, for example, a certain kind of waveform. The PAAL makes it possible for a device to interpret its sensory input in perceptual terms that can be used to drive a CRs world and self models. Going in the other direction, it also makes it possible for a CR to *do* things by exporting SDR primitive actions in a platform independent format.

The right side of the figure shows the components involved in augmenting an SDR architecture to allow for cognitive capabilities. The World and Self model components were discussed above. The Rules, Concepts, and Reasoning Mechanisms, will be discussed below. The remaining component is another abstraction layer, the Ontology & Rule Abstraction Layer. This layer serves a purpose that is symmetric to

PAAL. It allows ontology and rule concepts to be represented in a *platform-independent standard.* This is important if one wants to allow the same radio implementation to be used with alternate ontology and rule reasoning platforms. Just as radio notions such as *signal* and *waveform* should have meaning independent of any particular radio implementation, so too notions such as *concept*, and *rule* should have meaning independent of any particular implementations.

As an overview of how the proposed architecture works to augment existing SDR implementations, suppose the radio senses some waveform. We assume that the radio's SDR interface can be used to program a *wrapper* around its existing methods so that such an event triggers a method defined in terms of PAAL that allows an appropriate instance of a *signal* object (as defined in the ontology) to be constructed and deposited into the world model. Conversely, suppose the reasoner concludes that a certain action, such as evacuating a channel, should be taken. From its self model it knows that it is capable of taking such an action. Then, by virtue of the PAAL layer, the ontological element that represents that action will be linked to a method that can invoke the radio's native interface with a call to perform that action (or perform some procedure).

4. Role of knowledge representation

The discussion of knowledge-driven differential-response touched upon various pieces and forms of knowledge that a CR would need in order to exhibit such a capability. This section explores this requirement in more detail.

We appealed to two types of knowledge in our account. The first type was *conceptual.* Conceptual knowledge is the kind of knowledge that ontologies are intended to represent. Conceptual knowledge is typically either *analytic* or *axiomatic* in nature. These types of knowledge are both thought of as representing necessary truths, but for different reasons. A piece of knowledge is analytic if it expresses or follows from the *meaning* of concepts. For example, it is useful to talk about radios that can be moved from place to place (without impairing their operational capabilities). The concept of a *mobile-radio* would therefore be defined as a radio that has this property. Representing this definition in an ontology would make it possible for it to be applied in a formal reasoning system.

Axiomatic conceptual knowledge, on the other hand, expresses fundamental conceptual relationships that are *not* based on meaning alone. For example, the fundamental principles that radio waves are a form of electromagnetic energy and that they travel at the speed of light might be considered axioms within an ontology of radio knowledge. Note that what might be considered an axiom from the point of view of one ontological domain, might be considered a "derived" piece of knowledge, or a "theorem," from the point of view of a more fundamental domain. This "relativity" of what is an axiom should be a familiar theme and has been demonstrated many times in the history of science. Kepler's laws of planetary motion, for example, had the status of independent axioms when initially formulated, but were later shown to be consequences of Newton's general laws of motion.

This brings us to *rules*. In terms of the distinctions we have drawn, rules may be thought of as theorems that are worth "committing to memory," so to speak, because 1) they are useful in an application of interest, and 2) because the computational cost of deriving them from axioms "on demand" is prohibitive. For example, it is known that certain frequencies of radio signals are likely to degrade because of atmospheric conditions and mathematical laws governing this phenomenon can be derived from first principles. However, for any application in which this kind of knowledge is critical, it is highly likely that even a human expert would depend upon known rules for calculating such attenuation rather than performing an analysis based on the fundamental laws of electromagnetism and meteorology.

As we saw in our discussion of related work and in the discussion of our CR scenario, there is a need to represent *policies* in a way that a CR can both be guided by them and reason about them. A *policy* itself is a convention or a norm that *ought* to be followed and is not something that is, strictly speaking, true or false. However, *that* a particular policy is in force in some region at some time is something that is true or false. Knowing what behavior is required in order to be in compliance with a policy is also factual information, but may sometimes require a complex reasoning process to derive. Therefore, statements that relate the existence of a policy in a region to actions that need to be taken (or avoided) in order to be in compliance with that policy might often be worth "committing to memory" in the form of rules.

As stated previously, a piece of knowledge, such as the definition of a concept or a rule that cannot be used to derive other knowledge is essentially useless from an application point of view. *Reasoning* is basically the process whereby one or more pieces of knowledge are related to another piece. Numerous rule-languages or rule-based systems have been developed and deployed in various applications. These tools allow rules to be represented and typically provide an *inference* mechanism whereby rules are automatically invoked and applied in an application environment.

Ontologies also enable reasoning. From a theoretical point of view the kind of reasoning afforded by ontologies differs from rule-based reasoning. *Subsumption* reasoning is one example. Thus, as discussed above, from the fact that r is a radio and has the property of being mobile, and the definition of *mobile radio*, one can infer that r is a mobile radio. In practice, subsumption reasoning can be implemented using an underlying rule-based approach, but that is not necessary.

As the ensuing discussion will show, ontological and rule-based knowledge complement one another. This is especially true in complex applications in which modular design is a necessity.

5. Scenario using prototype simulation

We have implemented a prototype simulation environment capable of handling the beacon signal conflict scenario we outlined above. The ontological knowledge is expressed in OWL. We use Jena [13] as our ontology API and we also use the rule-language provided with Jena for representing rules. The inference mechanisms are also Jena-based. The PAAL is implemented by linking the ontology-API with our own interface to a very simple "software defined" radio implementation in Java.

The simulation enables one or more CRs and one or more beacons to be represented in a 2-dimensional space. The CRs can be mobile. Currently this means that they are associated with a user who can move around in the simulation environment. As a radio is moved and as the various components of the environment change, an environment handler and a simulation manager ensure that the necessary events are propagated to the various elements of the simulation. See figure 2 for a screen shot of the current system display and CR user interface.

A beacon signal conflict situation occurs when two beacons with opposing policies for the same channel overlap in some region. In figure 2, we see a CR that is positioned more or less equidistant from two such beacons. Such a conflict will matter to a CR only if it causes a problem with respect to a one of its goals. Suppose that the CR user has indicated a desire to use channel C and that a beacon signal conflict exists for C. The CR knows that its user wants to use channel C directly from user input. In terms of our architecture this knowledge is encoded in the self model of the CR. How does the CR know that a beacon signal conflict with respect to channel C exists? It knows this due to a series of inferences which are enabled by its ontological knowledge. In terms of the architecture in figure 1, the sequence of events will follow this pattern. The basic SDR component will process the two incoming signals. Thanks to the PAAL, two distinct instances of type signal will be added to the world model of the CR. Each signal is known to be associated with a certain logical channel, and certain logical channels are known to be reserved for beacons. Therefore, using its ontological knowledge, the CR will be able to conclude that the two signals it is receiving are two distinct beacon signals. Once a signal is known to be from a beacon the CR is able to interpret the "content" of the signal based on properties of the

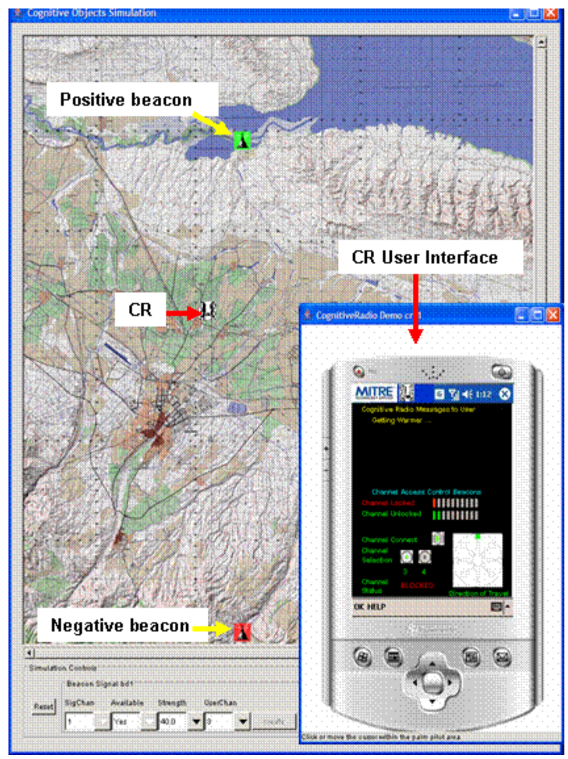

Figure 2. Screen Capture of Prototype.

signal. So the CR is now in a position to know that it is receiving a signal saying that it is allowed to use channel C and a signal that saying that it is not allowed to use channel C. The CR also knows the strength of each of these signals.

Formally, the kind of ontological reasoning just described relies upon the use of well-defined frameworks, such as OWL [14], in which definitions such as the following (schematic) definition can be encoded:

Beacon-Signal-Conflict-For-Channel-Use
subclass-of Radio-Policy-Conflict
GIVEN:
Logical-Channel c;
Beacon-Signal b1;
Beacon-Signal b2;
SUCH THAT:
b1 NOT-EQUAL b2;
b1 signal-content IS "c is available";
b2 signal-content IS "c is unavailable";

This definition provides a sufficient condition for determining when a beacon signal conflict exists. In our scenario, the world model of the CR contains two beacon signals that satisfy the conditions of this definition. This will automatically cause an instance of a Beacon-Signal-Conflict-For-Channel-Use conceptual object to also be inserted in CR's world model. This instance is parameterized with the references to other objects, such as logical channel c, that caused it to be inserted in the world model.

From our description thus far we can see that ontologies are the key logical device for providing "understanding" or "interpretation" of the lower-level factual inputs deposited into the world and self models of a CR through the PAAL. This is in keeping with our characterization of ontologies as providing knowledge concerning the analytic and axiomatic foundations of a domain. *Rules* comes into play as a more targeted form of knowledge relative to the application environment and typically encode knowledge about actions necessary to achieve goals, including cases where the normal course of action is not available. Rules are also useful for implementing constraints on action imposed by policy. The Reasoning Mechanisms component shown in figure 2 is responsible for making sure that any rule that is applicable in the current world and self model states is evaluated.

There are a number of rules relevant to the beacon-signal-conflict scenario. The rules implementing policy constraints come into play because the CR has the goal of using channel C. The rules state that a channel can be accessed only if there is a beacon signaling that the channel is available and that there is no conflicting beacon signal in the world model. Since the current situation includes a beacon signal conflict, the later rule will not be satisfied. Rather, a complementary rule, disallowing the use of C due to the conflict will be satisfied. When evaluated, this rule will have the effect of modifying the CR's self model to include the fact that the goal of using channel c is currently blocked due to the presence of a conflict. This change will cause the Reasoning Mechanism to review the set of rules again. At that point a *Beacon-Signal-Conflict-Rule* which can be given the following English rendering will be satisfied:

If
the goal of using channel C is blocked by a
Beacon-Signal-Conflict-For-Channel-Use
and b1 is the beacon-signal allowing use of C
and b2 is the beacon-signal disallowing use of C
Then
Attempt to move to a region in which b1 is still received but b2 is not received.

The "then" part of the actual rule in our simulation attempts to invoke a procedure (defined in terms of basic actions included in the CR's self model), that will cause the CR to ask the user to move in a direction. The user picks a direction and moves. The CR, following the aforementioned procedure, lets the user know whether he is getting "Warmer," "Colder", or neither. To be "Warmer" means that the signal strength of b1 is increased and the signal strength of b2 is decreased. This interactive procedure is repeated until signal b2 is no longer received. At that point, a rule allowing use of C will be satisfied and the CR will be able to take appropriate action to access the channel.

Besides illustrating the notion of knowledge-driven differential-response, this scenario also serves to illustrate important benefits of hybrid ontology-rule knowledge representation. One of the long-standing problems of knowledge representation is known as the issue of *qualification* which is related to the impossibility of stating rules in such a way that satisfaction of their conditions *guarantees* that the actions they recommend will have the desired result [15]. For example, if a CR is not mobile, then the above rule cannot succeed because the radio cannot be moved. The "if" part of Beacon-Signal-Conflict-Rule could be augmented with conditions stating that the CR must be mobile, and that it must actually have some means of locomotion, such as a user who is capable of self-locomotion. But this process degrades into an impossible situation if the goal is to state absolutely precise circumstances in which the rule can achieve the desired result. Thus, it need not be the case that the CR user is self-moving, rather it could also be the case that the CR user is merely mobile (not self-moving) but in control of another device that is self-moving. Or it could be the case that the latter device is also not self-moving but in control of another device and so on. In practical terms, the issue is how precise does one have to be in stating the conditions required for a rule to be applicable? Stated thusly, the problem is similar to any software validation problem: how does one know that a piece of code does the “right thing” for all possible input cases?

The combination of ontologies with rules provides a mechanism for dealing with such problems in a modular fashion. Instead of worrying about qualifying a rule with absolute precision, the rule can be stated in a simpler general form and the qualifications can be embedded in the appropriate ontological concepts. Thus, for example, a CR that does not have a means of locomotion at its disposal would know that it does not have any way of attempting to move. It would know this because the analytic knowledge contained in its ontology together with the knowledge contained in its self-model implies that it cannot move. Therefore it will not attempt to apply the Beacon-Signal-Conflict-Rule even if the “IF” part is satisfied by its current circumstances.

6. Closing remarks

The work presented in this paper is based on the premise that the path to CR is sure to be incremental. We have presented a high-level architecture that accommodates an incremental approach towards augmenting the SDR architecture with components required for CR. We have shown how the components we have discussed can work together to provide a system with a more robust form of goal-directed behavior, namely, knowledge-driven differential-response. We have built a simulation environment within which a scenario illustrating this cognitive capability has been successfully executed. Currently we are working to expand the simulation so that a wider range of scenarios can be accommodated.

7. Acknowledgment

This work was supported under MITRE Sponsored Research project 07MSR215.

8. Disclaimer

The affiliation of the authors with The MITRE Corporation is provided for identification purposes only and is not intended to imply MITRE’s support for the viewpoints expressed by the authors.

9. References

[1] J. Mitola III, “The Software Radio Architecture,” *IEEE Commun. Mag.*, vol. 33, no. 5, May 1995, pp. 26—38.

[2] J. Mitola III and G. Maguire, “Cognitive Radio: Making Software Radios More Personal,” *IEEE Personal Comm. Mag.*, Aug. 1999, pp.13—18.

[3] FCC, “Unlicensed operation in the TV broadcast bands,” ET Docket No. 04-186, Notice of Proposed Rulemaking (NPRM), May 2004.

[4] [On-line]. grouper.ieee.org/groups/802/22

[5] P. Marshall, "XG Communications Program Information Briefing," *Semantic Web Applications for National Security (SWANS) Conference*, Rosslyn, VA, Apr 7, 2005, [On-line]. http://www.daml.org/meetings/2005/04/pi/DARPA_XG.pdf

[6] G. Denkar, "Policy-Defined Cognitive Radios:What Are The Issues At Hand?", *Cognitive Radio Workshop*, San Francisco, CA, Apr. 10, 2006.

[7] L. Berlemann, S. Mangold, and B. H. Walke, “Policy-based Reasoning for Spectrum Sharing in Cognitive Radio Networks,” *Proc. IEEE DySPAN*, Baltimore, MD, Nov. 8—11, 2005, pp. 1—10.

[8] J. Neel, J. Reed, R. Gilles. “Game Models for Cognitive Radio Algorithm Analysis," *Proc. SDR’04*, Phoenix, AZ, Nov. 15—18, 2004.

[9] T. W. Rondeau, C. J. Rieser, B. Le, and C. W. Bostian, "Cognitive Radios with Genetic Algorithms: Intelligent Control of Software Defined Radios," *Proc. SDR’04*, Phoenix, AZ, Nov. 15—18, 2004.

[10] J. Wang, D. Brady, Baclawski K., M. Kokar, and L. Lechowicz, “The Use of Ontologies for the Self-Awareness of the Communication Nodes," *Proc. SDR’03*, Orlando, FL, 2003.

[11] J. D. Poston, W. D. Horne, M. G. Taylor, and F. Z. Zhu, Ontology based reasoning for context-aware radios: Insights and findings from prototype development,” *Proc. IEEE DySPAN*, Baltimore, MD, Nov. 8—11, 2005, pp. 634—637.

[12] M. J. Marcus, “Real time spectrum markets and interruptible spectrum: new concepts of spectrum use enabled by cognitive radio,” *Proc. IEEE DySPAN*, Baltimore, MD, Nov. 8—11, 2005, pp. 512—517.

[13] [On-line]. http://jena.sourceforge.net/.

[14] [On-line]. http://www.w3.org/2004/OWL.

[15] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach.* Prentice-Hall, 2002.