

## Article

# Towards Collaborative and Dynamic Spectrum Sharing via Interpretation of Spectrum Access Policies

Jakub Moskal <sup>1</sup>, Jae-Kark Choi <sup>2</sup>, Mieczyslaw M. Kokar <sup>1,3,\*</sup>, Soobin Um <sup>4</sup> and Jeung Won Choi <sup>4</sup><sup>1</sup> VISTology, Inc., Framingham, MA 0160, USA; jmoskal@vistology.com<sup>2</sup> Hanwha Systems Co. Ltd., Seoul 04541, Korea; jaekark.choi@hanwha.com<sup>3</sup> Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115, USA<sup>4</sup> Agency for Defense Development, Daejeon 34186, Korea; sbum1989@add.re.kr (S.U.); jwchoi@add.re.kr (J.W.C.)

\* Correspondence: mkokar@ece.neu.edu

**Abstract:** This paper describes some of the challenges that need to be addressed in order to develop collaborative spectrum-sharing systems. The importance of these challenges stems from the assumption that rules for spectrum sharing can change after the deployment of radio networks and that the whole system must be able to adapt to them. To address such a requirement, we used a policy-based approach in which transmissions are controlled by a policy interpreter system, and the policies can be modified during system operation. Our primary goal was to develop a prototype of such a system. In this paper, we outline the implementation of policy interpretation, automatic generation of transmission opportunities in case a request for transmission is denied by the policy reasoner, and the generation of rendezvous channels for the synchronization of otherwise asynchronously running software-defined radios.



**Citation:** Moskal, J.; Choi, J.-K.; Kokar, M.M.; Um, S.; Choi, J.W. Towards Collaborative and Dynamic Spectrum Sharing via Interpretation of Spectrum Access Policies. *Appl. Sci.* **2021**, *11*, 7056. <http://doi.org/10.3390/app11157056>

Academic Editors: Junsu Kim and Won Cheol Lee

Received: 19 June 2021  
Accepted: 27 July 2021  
Published: 30 July 2021

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

**Keywords:** spectrum sharing; collaborative spectrum sharing; policy-based sharing; transmission opportunities; rendezvous channels; policy interpreter

## 1. Introduction

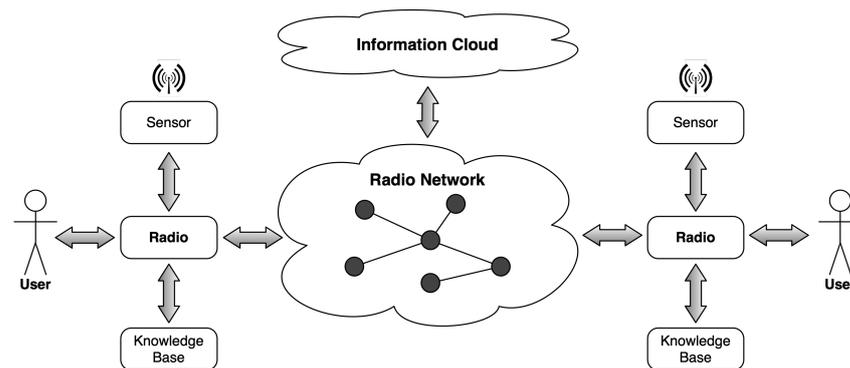
A number of definitions of a cognitive radio (CR) have been published in the literature. We are using here a definition that is provided by the IEEE P1900.1 standard [1]:

- (a) “A type of radio in which communication systems are aware of their environment and internal state and can make decisions about their radio operating behavior based on that information and predefined objectives.
- (b) Cognitive radio [as defined in item (a)] that uses software-defined radio, adaptive radio, and other technologies to adjust automatically its behavior or operations to achieve desired objectives.”

To realize the full potential of this definition, a number of capabilities need to be supported by a radio [2]. Among others, the capabilities should include: (1) Collecting *meters*—sensed information about the Radio Frequency (RF) environment; (2) Answering queries from users and other cognitive radios; (3) Situational awareness—understanding the various relations between the parameters; (4) Self-awareness—knowing own parameters and settings (*knobs*); (4) Sending queries and requests to other radios; (5) Command execution—accepting command requests from other radios, reasoning about them, and executing the commands; (6) Obtaining information relevant to the control of own capabilities (adaptation)—accessing and interpreting information stored in knowledge and data bases.

Figure 1 [2] shows the context for a cognitive radio. Two radios shown in the figure interact with the User, Sensor, Knowledge Base and other radios within a Radio Network using Application Programming Interfaces (APIs), depicted as bidirectional arrows. Since in this paper we are primarily focusing on the dynamic spectrum sharing (DSS) among the radios operating in the same RF environment—not necessarily being part of the same

Radio Network—we need to cover the issues of collaboration. “Collaboration is defined as the process of two or more entities or organizations working together to complete a task or achieve a goal” [3]. Collaboration is an inherent part of communications. Thus, the question is *how* to implement collaboration rather than *whether* to support it. Collaboration implies the willingness of the entities to interact and make trade-offs with each other. The organization of the collaboration process must follow some policies that all the entities should obey. In the use cases considered in this paper, the objective of collaboration is to make a collaborative decision about the use of the RF spectrum—a crucial step in Electromagnetic Spectrum Management (ESM).



**Figure 1.** A cognitive radio in context.

The next question is how policies should be encoded on a radio? The most typical approach is to encode each policy in an imperative language, e.g., C++, and then implement a control loop that matches the current ESM situation to an existing policy and executes it. While the hand-coded policies are likely to be computationally efficient, this approach has a number of disadvantages. The most difficult requirement to satisfy is the ability to modify the policies after deployment because the code developed in an imperative language, once modified, must be compiled, tested and deployed.

Another approach is to add a declarative meta-layer of code on top of the imperative code. It can be eXtensible Markup Language (XML) code that provides pre-conditions for a policy execution as well as the post-conditions that need to be interpreted and verified by the controller. While this approach is more flexible than the one just described above, it still requires re-coding of the XML-specified testing of the pre-conditions as well as the procedures implementing the post-conditions.

The cognitive radio requires that the language in which policies are coded must be machine processable and “understandable”. This means that the cognitive radio software must be capable of interpreting and executing dynamically added policies, without any modification of the procedural code. One way to achieve this is to express policies in a formal declarative language with formal syntax and semantics and incorporate a policy interpreter that can interpret policies for as long as they are expressed in that language. There are two kinds of knowledge that needs to be represented using such a formal declarative language: (1) the shared concepts between radios and networks; (2) the policies that are used to control the behavior of the radio.

The shared concepts between radios and network can be defined in common ontologies encoded in a (formal) ontology representation language, such as Web Ontology Language (OWL) [4]. In artificial intelligence, computer science and information science, the term “ontology” refers to a formal, explicit specification of a set of concepts in a specific domain and relationships between these concepts. The term “formal” means that the ontology is “machine processable for the purpose of knowledge reuse and sharing” [5]. In this paper, we describe policies that are represented in a language that uses ontological concepts.

## 2. Cognitive Policy-Based Radio

In this section, we describe the concept of a Cognitive Policy-Based Radio (CPBR), starting with brief definitions of related types of radios, based on the IEEE P1900.1 standard [1]. It should be noted that the P1900.1 definitions do not necessarily fully align with the regulatory definitions of the International Telecommunication Union (ITU) or the Federal Communications Commission (FCC).

*Hardware Radio* is “a type of radio whose all communications functions are entirely implemented in hardware”, but most importantly, changes in communications capabilities can only be achieved by changing the hardware. Thus, a radio is considered to be a Hardware Radio even if some of its functions are implemented in software, but the “regulated emission or reception parameters cannot be changed in the field, post manufacture, without physically modifying the device.”

*Software-Defined Radio (SDR)* is “a radio in which some or all of the physical layer functions are software defined”, i.e., implemented in software (which implies more than just controlling the functions implemented in hardware). Thus, unlike in a Hardware Radio, regulated parameters in SDR are changeable through software modification.

*Policy Based Radio (PBR)* is a type of radio whose behavior is governed by a policy-based control mechanism, with policy defined as a “set of rules, expressed in a machine-readable format that are independent of the radio implementation regardless of whether the radio implementation is in hardware or software”. Consequently, although PBRs are usually SDRs, this is not a logical requirement for being policy-based.

*Cognitive Policy-Based Radio (CPBR)* is a type of radio that is (1) cognitive, i.e., “aware of its environment and internal state and can make decisions about its radio operating behavior based on that information and predefined objectives” (c.f. Section 1), and (2) uses a policy-based control mechanism (akin to PBR). Policies in CPBR, expressed in a machine-readable format (referred to as a *policy representation language*), can not be only ingested and displayed but prescribe a behavior that is interpreted and executed by the radio.

Access to the Electromagnetic Spectrum (EMS) is one of the aspects that is controlled according to policies it has. In most cases, this means that they dynamically sense the spectrum for current usage and use the parts of it that are not occupied by the Primary Users (PU). As soon as a PU is sensed and there is evidence that transmitting in the PU’s band could result in interference to the PU, the CR vacates the current channel and tries to grab a different, unoccupied part of the spectrum. The decisions on whether to use a specific spectrum can be made via the interchange of messages between the collaborating radios, or it can be implemented autonomously by each node following the locally available policies.

Policies specify a number of factors that should be taken into consideration by the radios in order to arrive at a spectrum use decision, for example, the location of the device, the frequency and time range of interest or the maximum transmit power. Different policies pertain to different geographical areas and can change dynamically. Therefore, the CR design cannot account for all possible policies and must incorporate a mechanism for interaction with some authority that can dynamically grant or deny access to particular parts of the spectrum. Once a new policy is fetched from an authority, a CPBR radio designates it as an *active policy* and incorporates it into its control mechanism. In our approach, the Policy Manager (PM) component is in charge of interacting with an authority and maintaining the active policy, while the Policy Engine (PE) component interprets new policies and ensures that the radio operates within the constraints established by the PM.

The interaction inside the radio between the main Controller and the PE requires a language for expressing the spectrum access requests and the responses to such requests. In case the access cannot be granted, the PE may respond with a list of opportunities, i.e., other channels, in which opportunistic access is possible. The above requirements cannot be achieved via static definitions of vocabularies (e.g., in XML or a relational database) since adding new terms to such vocabularies would require developing new software (for interpreting XML tags or relational schema and data). On the contrary, this

kind of capability can be achieved by an ontological approach in which new terms and their definitions can be introduced dynamically. This is the consequence of bounding ontologies to a highly expressive language (such as OWL [4]) in which new terms can be defined on the fly. These definitions are then interpreted by generic inference engines that are bound to the same language (such as OWL), thus avoiding the need to rewrite the underlying software.

While opportunistic access to the RF spectrum is reasonable for addressing the problem of no-interference to PUs, such an approach does not solve the problem of spectrum sharing among the radios within the same group, i.e., either PUs or the friendly secondary users (SUs). In this paper, we present parts of a policy-based solution to such cases. First, we briefly describe an architecture in which such scenarios can be investigated. Then, we describe our experiments with policy interpretation, policy-based generation of transmission opportunities, policy-based definition of rendezvous and policy-based optimization of radio parameters (knobs) based on the transmission feedback (meters). Finally, we analyze the consequences of a radio not having some of the capabilities considered in this paper.

### 3. Architecture

The main component of the cognitive radio is a Cognitive Engine, as shown in Figure 2. The architecture of the radio must support the functionality described above, plus it must support the aspects of policy management and policy interpretation and execution. Most of the cognitive system architectures are organized around an iterative processing loop (cf. SOAR [6], ACT-R [7]). Some others, e.g., the Endsley's model [8], are patterned after the OODA loop—Object, Orient, Decide and Act [9]. The OODA loop, in general, and Endsley's model, in particular, have been widely used in the domain of Information Fusion, where it serves as a basis for situational awareness.

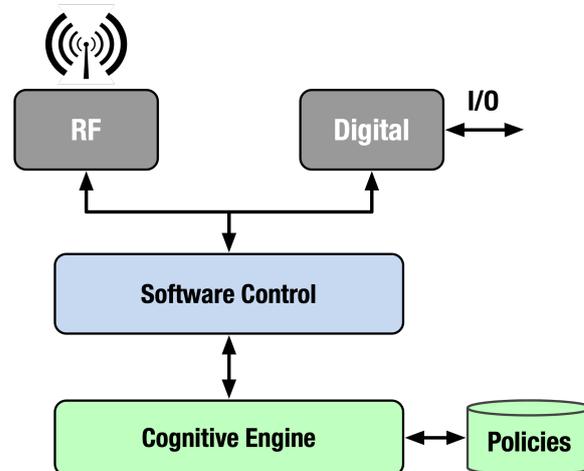


Figure 2. Cognitive radio—top level.

Mitola [10] proposed to extend the OODA loop and use it as the basic processing loop of SDR-based cognitive radio. He introduced two more processes to the control loop: Plan and Learn. This model, also known as ideal CR architecture (iCRA), is characterized by the Observe–Orient–Plan–Decide–Learn–Act (OOPDLA) loop (see Figure 3). The processes inherited from the original OODA model have the following interpretation in the context of cognitive radio: *Observe*—collect data about the environment using an array of sensors; *Orient*—integrate the observations with policies, goals and preferences, hardware and software limitations and past experience; *Decide*—identify the changes that need to be made in the SDR configuration in order to address the new situation; *Act*—reconfigure the SDR. Furthermore, the added processes are *Plan*—planning decisions for the long run—and *Learn*—machine learning and adaptation based on past experience.

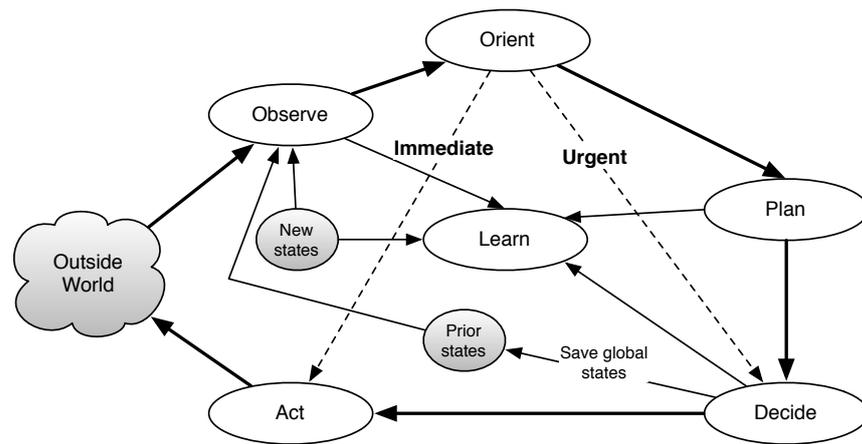


Figure 3. The extended OODA loop.

One of the possible passes through the OODA loop is as follows:

1. Receive a request (from Software Control) to transmit a packet of information; the request includes the metadata-frequency, time, power and other items of information, as needed.
2. Identify policies relevant to the current request.
3. Invoke the policy engine (PE) to verify that the transmission request is compliant with at least one of the policies and not forbidden by any of the relevant policies.
4. If the decision is “allow”, then convey this decision to Software Control.
5. If the decision is “disallow”, then invoke the function Compute Transmission Opportunities (planning).
6. Send transmission opportunities or the “disallow” decision (if no opportunities were found) to Software Control
7. Return to Step 1.

The architecture of the experimental system used in this work, termed the Cognitive Radio Engine (CRE), is shown in Figure 4. It consists of six components: Policy Manager, Policy Engine, Planner, Learner, Controller and Inference Store (IS). The Controller coordinates the interactions between the first four components and the radio’s Software Control. These four components also interact with the domain-independent IS, VISTology’s semantic storage. All data within CRE is represented in OWL (Web Ontology Language [4]). A general-purpose OWL Reasoner that can take full advantage of the semantic information encoded in OWL is embedded inside IS (not explicitly shown in the figure).

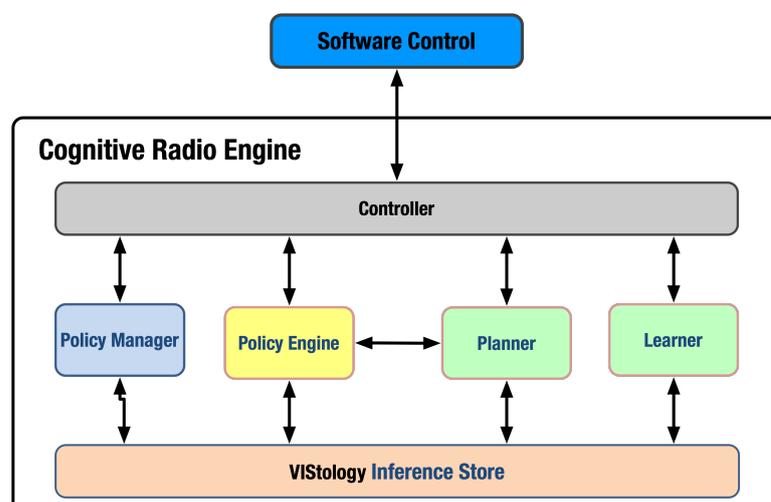


Figure 4. The Cognitive Radio Engine architecture described in this work.

The components of this architecture are in a direct relation to the OODA loop, as shown in Table 1. The table shows some functions that have not been discussed so far but are needed for the realization of collaboration among the CRs. In particular, we identify the operation of Schedule Rendezvous as a functionality that is needed for synchronizing collaboration among the radios.

**Table 1.** Mapping the OODA loop to the Cognitive Radio Engine architecture.

OODA	Function	Component
Observe	Get/describe transmit request	Controller
Orient	Identify matching active policies	Policy manager
Observe	Get spectrum measurement	Controller
Decide	Verify compliance	Policy Engine
Plan	Generate transmission opportunities Schedule rendezvous	Planner
Act	Issue transmit command	Controller
Learn	Update policy	Learner

The functionality of the particular components is described in more detail in the rest of the paper.

#### 4. Ontology

For the implementation of communications systems two parts are needed: a protocol to exchange messages and a language in which the messages are represented. In our work, we make use of *ontologies* as a representation language. First, we briefly introduce the notion of ontology since this concept is relatively new to the radio communications community. To put the notion of ontology in perspective, we overview other concepts that are close to ontologies.

*Vocabulary:* It is a list of terms, e.g., Radio, Receiver, Transmitter, Antenna, Battery, etc. All it contains is lexical representations of different concepts that a human reader can interpret just because the reader knows these names. A computer program would treat them simply as strings and thus would be able to perform operations that are applicable to the data type String, such as compare, insert, delete, assess the length of the name and so on.

*Dictionary:* It includes terms with natural language explanations of their meaning, where the explanations use the same language that includes the defined terms. Here is an example of dictionary term “radio” as can be found by Google if you enter the search term “radio”.

*Radio:* Noun: *the transmission and reception of electromagnetic waves of radio frequency, especially those carrying sound messages.* Use: “cellular phones are linked by radio rather than wires”; Verb: *communicate or send a message by radio.* Use: “the pilot radioed for help”.

The dictionary contains more information than vocabulary since it can be used by the human to understand what a term means in case the human does not know it. However, dictionaries are difficult to process by computers since computers are not too good (so far) at processing natural language text.

*Ontology:* An ontology includes a number of *classes, properties* and *individuals*, e.g., Radio could be one of the classes. Any class can have multiple individuals (aka *instances*), and thus, the Radio class can have a number of individuals—members of the Radio class, e.g., R-AFK2178, RD-FM-365M. An ontology can represent the fact that the Radio comprises components such as a Transmitter, a Receiver, an Antenna, a Modulator, a Demodulator, etc., (using the *aggregateOf* object property).

A small fragment of a radio ontology is shown in Figure 5, with classes represented as rounded rectangles with a circle icon, individuals as rounded rectangles with a diamond icon and properties as arrows (different colors denoting different types of properties).

Ontological classes are organized along the `subClassOf` relation, resulting in a hierarchy (a lattice). For instance, `Radio` is a `subClassOf` `System`, which in turn is a `subClassOf` `Object`. For the sake of clarity, the ontology fragment shown in the figure is limited to only a few classes and properties, and two individuals: `R-AFK2178` (an instance of class `Radio`) and `TR-FRK0999` (an instance of `Transmitter`). The superclass of every other class in an OWL ontology is called `Thing` (not shown in the figure). Thus, any individual is also an instance of `Thing` (via a chain of the `subClassOf` relations).

The radio ontology, developed with Protégé [11], was encoded (conceptually) as a Resource Description Framework (RDF) [12] triples  $\langle \text{subject}, \text{predicate}, \text{object} \rangle$ . Here, the subject would need to be an individual from a class, while the object is either an individual of a class or an instance of a data type, e.g., `Integer` or `String`.

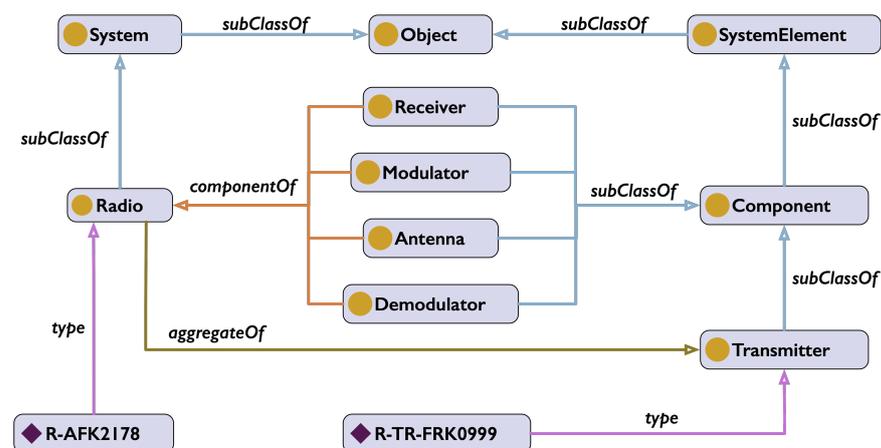


Figure 5. A fragment of the Radio Ontology.

The predicate is an individual of a special class called `Property`. Properties represent relations between pairs of individuals. For example, the fact that a specific radio has a transmitter would be represented as a triple  $\langle R-AFK2178, \text{aggregateOf}, TR-FRK0999 \rangle$ . Here, `aggregateOf` is a predicate. The facts that `R-AFK2178` is a radio and `TR-FRK0999` is a transmitter would be represented as two other triples:  $\langle R-AFK2178, \text{rdf:type}, Radio \rangle$  and  $\langle TR-FRK0999, \text{rdf:type}, Transmitter \rangle$ , in which a special OWL predicate `rdf:type` is used to associate individuals with classes. In Figure 5, `rdf:type` is shown simply as `type`.

Predicates represent either relations between individuals or attributes of individuals. Some of the characteristics are the “defining characteristics”. If an individual has those characteristics, it can be classified as an instance of that class. For example, a thing that has a transmitter and a receiver must be an instance of the class `Radio`.

Some other characteristics are just “descriptive”, e.g., characteristics that are common to various classes (various things have the property of weight). If it is known that a thing is an instance of a given class, then it is expected to have characteristics that are necessary of such a class. For example, if a thing is an instance of the class `Radio`, then it must have a transmitter and a receiver.

For ontologies to be processable by computers, they must be described in a language that has formal semantics. To be “understandable”, the language must also have formal, computer processable *semantics*. Roughly, it means that there needs to be a computer program that can answer queries about information encoded in the language. For example, in the context of relational databases, queries are expressed in the SQL language, and a database management system answers them by retrieving data stored in the database. The answers need to be *sound*, i.e., the answers must be correct with respect to what the database contains. For an ontological language, we require more than a database. An ontological language must have a program called “*reasoner* (or *inference engine*)”, which can do more than just retrieve information from its *knowledgebase*, but it should be able to extract information that is only implicit in the *knowledgebase*. e.g., if the knowledge base

describes the radio communications domain, it will encode facts about some of the domain objects as *axioms* expressed in the language. The reasoner then will be able to process the input facts and the axioms and answer queries based only on partial information about an object. For example, if the *knowledgebase* contains an object classified as jammer, the reasoner will be able to infer that it has a transmitter. The use of axioms has very powerful consequences because it provides means for extracting facts that are only implicit and for checking the logical consistency (lack of contradictions) of the *knowledgebase*.

#### 4.1. A DSA Ontology

Since our approach heavily relies on an ontology, we first provide a brief overview of the overall structure of the ontology.

Ontologies, similarly to other engineering artifacts, such as hardware or software, are built in a modular and hierarchical fashion. The structure of our Dynamic Spectrum Access (DSA) ontology (here, termed OntoDSA) is shown in Figure 6. The arrows in this figure represent the “import” relationships. OntoDSA imports SpectralSPARQL [13], which, in turn, imports GeoSPARQL [14], as well as ontologies to describe Lists and Uncertainty (both of them import the Nuvio foundational ontology [15]).

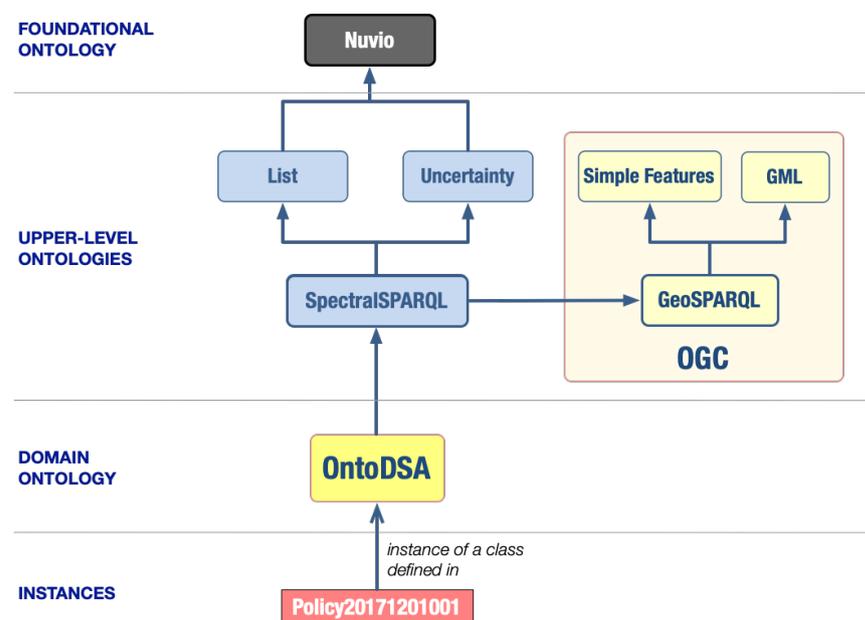


Figure 6. The Architecture of the DSA Ontology.

The structure of SpectralSPARQL is patterned upon GeoSPARQL [14]—an Open Geospatial Consortium (OGC) standard extension of SPARQL. GeoSPARQL provides a common representation of geospatial data in RDF [12] and facilitates querying and filtering of such data based on relationships between geospatial entities. The standard consists of an ontology of geospatial concepts, a set of SPARQL [16] functions and query transformation rules. With GeoSPARQL, one can express high-level queries that ask about specific geographic objects of features, their attributes and relations between them.

SpectralSPARQL is built on top of the base ontology, Nuvio [15], and thus includes its top-level classes: Attribute, InformationEntity, Object, Process, Situation, UnitOfMeasure and Value; all mutually exclusive. Each of the classes in SpectralSPARQL are subclasses of at least one of these concepts. The two top-level concepts in Nuvio are Object and Process—entities that are either wholly represented or that can only be partially represented at any given snapshot of time, respectively.

The three fundamental objects in SpectralSPARQL are Signal, Receiver and Emitter. The process subclasses are used to model different activities that the objects can participate in. For instance, an emitter can participate in a frequency hopping transmission.

The subclasses of the Attribute class model properties for describing both objects and processes. Some of them are simple quantities with associated units of measure, e.g., Frequency with FrequencyUnit, while others are more complex and require special-purpose properties to fully define them, e.g., FrequencyRange with object properties referring to the minimum and maximum values of the range. There are three major Attribute categories (subclasses): Spectral, Temporal and Signal. The first one is the central focus of the SpectralSPARQL functions. The SpectralAttribute class was modeled in a similar fashion as Feature in GeoSPARQL. It is used to define SpectralSPARQL functions that calculate relationships between two different spectral attributes, e.g., the intersection, regardless whether they are single frequencies, frequency sets or ranges, very similar to the *Intersects* query function in GeoSPARQL.

The OntoDSA ontology imports all of the above by simply importing SpectralSPARQL. The most significant classes in this ontology include Policy, Rule and the SAPHyperspace classes. These classes relate directly to the following Spectrum Access Policy (SAP) use cases implemented by the CRE:

- Each instance of the Policy class can be related to many instances of the Rule class by the property *hasRule*. It may also have several data property assertions, such as *hasID* or *hasRevision*.
- Instances of the Rule class can be related to instances of the SAPHyperspace by the property *governs*, denoting a relationship between a SAP rule and the spatio-spectro-temporal region that it applies to. Properties of each rule, such as its *ID*, *revision*, *type* (with the value of either *allow* or *deny*), or *maximum transmit power*, are asserted directly on the instance of the Rule class.
- Instances of the SAPHyperspace class represent spatio-spectro-temporal regions that are governed by some SAP rules. The term *spatial* is interpreted here in a more abstract sense, more than just the physical space. Any of the quantities can be considered as a dimension of such a space. Each hyperspace defines its frequency range, time interval and the geospatial region, to which the associated rule is applied and more. While the first two properties are defined using the SpectralSPARQL concepts, the latter is defined using GeoSPARQL.

An example of a policy is shown in Figure 7 later in the paper, after we discuss policies.

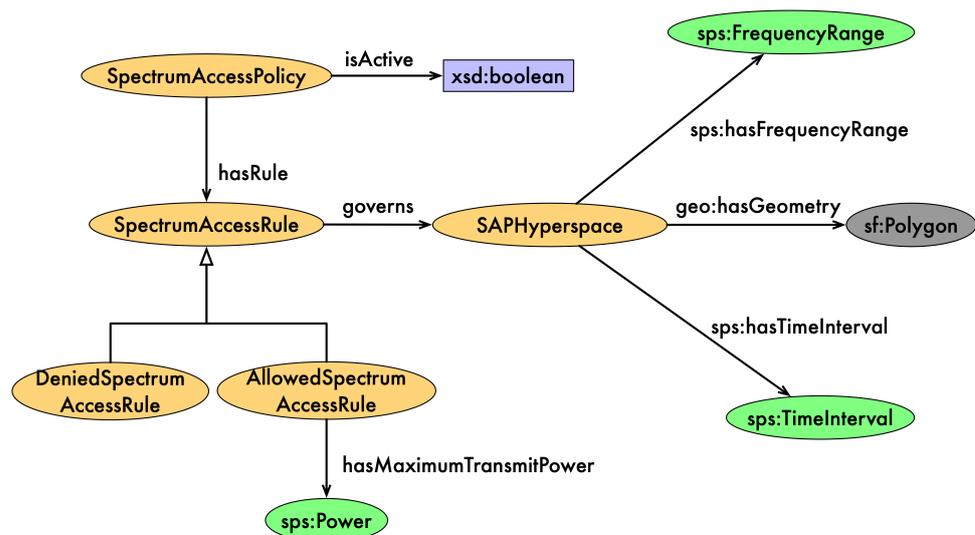


Figure 7. A snapshot of the SAP Ontology.

### 5. Policies

We distinguish two kinds of policies: (a) Hyper-cube based policies, which are constraints and requirements based on sets of states of the environment and the transmitter,

and (b) higher-level policies that are defined by logical conditions. First, we describe each policy kind in greater detail, then we show an example of a logic-based policy.

### 5.1. Hyper-Cube Based Policies

The policies are formalized by a number of rules,  $P_p$ . The rules are formally defined as functions that assign decisions depending on subsets of conditions,  $C$ , that are either satisfied or not. The conditions are expressed either as subsets of states,  $S$ , of the RF environment and radio characteristics or logical conditions of such states. A generic form of a rule can be formalized as:

$$P_p : 2^C \rightarrow D \tag{1}$$

where the conditions  $C$  are combinations of set-based and logical conditions:

$$C = S \cup L \tag{2}$$

where  $S$  is an explicitly specified set of states, and  $L$  is a set of logical values; expressions of the form  $2^L \rightarrow \{T, F, U\}$ , where  $T$  stands for logical *true*,  $F$  for *false* and  $U$  for *unknown*.

In a specific case, the sets can be intervals of dimensional values representing  $X$  and  $Y$  locations, with  $X_i, Y_i$ , time intervals  $T_i$ , frequency ranges  $F_i$ , transmit power  $W_i$  and other types of sets. An example of a spectrum access policy signature is:

$$P_p : X_i \times Y_i \times T_i \times F_i \times W_i \times C_i \rightarrow \{a, d\} \tag{3}$$

where  $\{a, d\}$  shows whether the policy is an *allowed policy* or *restrictive policy*.  $P$  allows transmitting when the value is  $a$ , while it disallows transmitting when it is  $d$ .

Examples of projections of this kind of policies are shown in Figure 8, where policy C is a restrictive policy, while the rest are allowed policies. This figure shows projections from 6D to 2D. In 6D, we refer to them as hypercubes or D-rectangles (in this case,  $D = 6$ ).

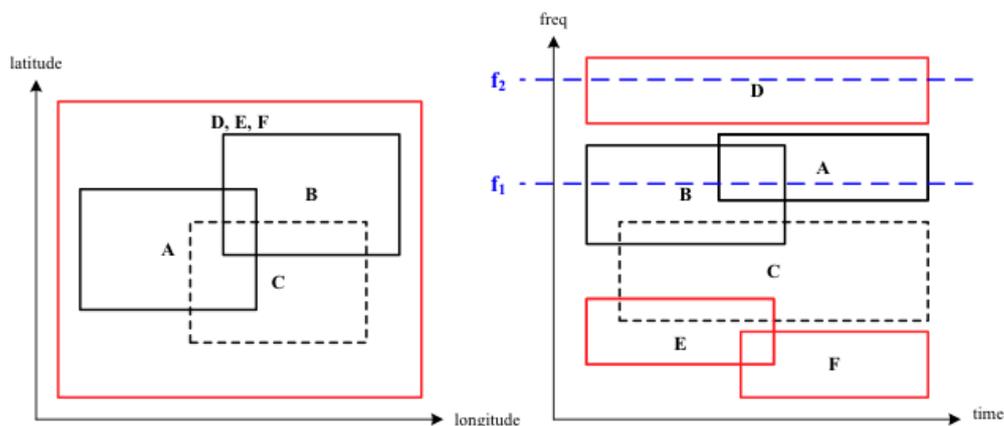


Figure 8. An Example of Policies.

### 5.2. Higher-Level Policies

Policies that are based on logical conditions cannot be represented as hypercubes. In our approach, we represent policies as SPARQL *Update* expressions augmented with the SpectralSPARQL extension (c.f. Section 4.1). The Update mechanism of SPARQL allows making conditional changes in the knowledge base, which fits well with the condition-action structure of the policies.

As an example, consider a policy expressed in natural language:

**Policy 1:**

Allow transmission for a transmitter  $T$  if one of the following conditions is met:

- $T$  is located at least 2 km from  $NodeX$ ; its Tx power is below  $P_1$ ; and its modulation and coding scheme is QPSK.
- $T$  is located at least 3 km from  $NodeX$ ; its Tx power is below  $P_2$ ; and its modulation and coding scheme is 16QAM.

This policy can be expressed as a SpectralSPARQL Update expression, as shown in Listing 1. Note how the two subexpressions (surrounded by curly brackets) of the *WHERE* clause directly correspond to the different conditions in the natural text.

Listing 1: A SPARQL representation of Policy 1.

```

INSERT {
  ?request :hasDecision :allow
}
WHERE {
  {
    ?request a :TransmissionRequest;
    :forNode: ?Node1 .
    ?Node1 a :Emitter;
    :hasMCS QPSK;
    :hasTxPower ?TxPower;
    :distanceToNodeX ?distance.
    FILTER (?distance > 2000 && TxPower < P1)
  }
  UNION
  {
    ?request a :TransmissionRequest;
    :forNode: ?Node1 .
    ?Node1 a Emitter;
    :hasMCS 16QAM;
    :hasTxPower ?TxPower;
    :distanceToNodeX ?distance.
    FILTER (?distance > 3000 && TxPower < P2)
  }
}

```

## 6. An Implementation of a Policy Based Spectrum Sharing System

In this section, we describe how the Cognitive Radio Engine architecture described in Section 3 was implemented in our experimental system.

### 6.1. Controller and SDR Platform

As was shown in Figure 2, the radio platform includes two kinds of processing: RF and Digital. While the RF processing is responsible for receiving and transmitting radio signals, the Digital processing is responsible for receiving and sending out digital information converted from and to the radio signals, respectively. Finally, the digital information is exchanged with the radio user via dedicated I/O streams.

In the case of SDR, the radio platform also has a Software Control module, which is used for setting RF and Digital hardware and software, as well as controlling the radio behavior. The Software Control module interacts with the Controller module (c.f. Figure 4) of the Cognitive Radio Engine through the Controller's API. This API handles both control messages, e.g., transmission requests, and data messages, e.g., packets, that the CRE needs to send to other radios. In addition to message handling, the Controller's API is also used for fetching policies from a Policy Authority, as needed.

One of the primary responsibilities of the Controller is the translation between the language used to capture information in Software Control and the language used in the CRE. Since on the CRE's side, everything is expressed in OWL and SPARQL, and on the Software Control side, either in JSON or CORBA IDL, the Controller must support this "adaptation". The Controller passes policies to the Policy Manager component and allows the CRE to get *meters* and set *knobs*, through which the radio platform is monitored and controlled.

### 6.2. Policy Manager

The PM manages policies and associated policy interpretation rules. In particular, it maintains a list of active policies, i.e., the policies that are applicable to the current circumstances as imposed by the Authoritative Policy Source. The function of the PM is to support the authoring of new policies, editing of the existing policies, storage and retrieval of policies in policy database and activating/deactivating policies. For this purpose, the Policy Manager supports functions that execute tasks such as Select Active Policies, Add Active Policy, Remove Active Policy, Select Allowed Policies, Select Disallowed Policies and so on.

Looking deeper into the structure of the system, the PM manages the policies that are stored locally in the Inference Store (see Section 6.3), based on actions initiated from the Controller. Each time a modification is made (rules are modified, policy is activated, etc.), the Inference Store automatically invokes BaseVISor [17]—an OWL 2 RL Reasoner—to infer any implicit facts about the policies. All facts, explicit and derived, are subject to querying when, for instance, the PE is determining the transmission opportunities, which requires access to the currently active policy definition.

### 6.3. Policy Engine and Inference Store

Policies are interpreted by the PE component. The most important function is to infer whether a request for transmission can be granted or not. The case of "allow" is rather simple; the Controller returns a permission (allow) to Software Control, which then can start transmitting. In the case it is denied, there may be a bigger issue that should be resolved by the PE. For instance, if the request cannot be granted for transmitting at this time and at this power, the question is: what should the Controller do? If the only response from PE is allow or deny, the Controller may need to attempt many different requests (assuming it knows how to generate them) before one is finally granted, if at all. To avoid this "long-loop" scenario, the PE can produce a list of transmission opportunities that would be allowed in the current state and given the current policy.

Internally, to generate opportunities, the Controller uses the Planner (see Section 6.4) and VISology's Inference Store. The Inference Store (IS) is a component that combines the power of BaseVISor's OWL inference capability [17] with an in-memory triple store that exposes a SPARQL 1.1 interface [16]. Upon changes to the triple store (assertion of new facts), BaseVISor is automatically triggered and injects implicit facts into the store. Thus, at query time, all facts (explicitly asserted and inferred) are immediately available.

IS stores the currently active policy (provided by the PM) in its ontological form. Each policy is represented in OWL according to the structure that is partially represented in Figure 7. The top-level class—SpectrumAccessPolicy—can be marked as *active*, and the PM ensures that exactly one policy is active at any given time. The policy comprises a number of SpectrumAccessRules that are related to it via the *hasRule* object property. Each rule *governs* exactly one SAPHyperspace, which represents a multi-dimensional space defined by frequency range, time interval, a geospatial region and other sets, as described in Section 5. The concrete definitions of these dimensions are imported from the pre-existing ontologies: SpectralSPARQL [13] and OGC GeoSPARQL [14]. Each SpectrumAccessRule either allows or denies access to the SAPHyperspace that it governs, which is modeled in the ontology with two subclasses: AllowedSpectrumAccessRule and DeniedSpectrumAccessRule. Instances of the AllowedSpectrumAccessRule can specify the maximum transmit power

(defined in SpectralSPARQL) that is allowed within the governed hyperspace. The ontology was designed using the best ontology engineering practices, i.e., modularity and reuse, which allow for a rapid development and support the integration of the policy data with other sources of information in future scenarios.

#### 6.4. Planner

The responsibilities of the Planner module include, among others, the generation of transmission opportunities (mentioned above) and generation plans for rendezvous. Other responsibilities may be added to this module as needed.

##### 6.4.1. Transmission Opportunities

While computing opportunities may seem to be a simple issue, the problem is that the number of possible opportunities may be quite large, especially for policies that comprise vast numbers of rules. One simplistic way to solve such a problem would be for the Controller to continuously generate requests (in a loop) in response to denials. This would lead to very many interactions between the Controller and the PE, which means it is clearly not a great solution from the point of view of the efficiency of the CRE. In our approach, the PE provides the Controller with an API for querying about transmission opportunities for any combination of location, time, frequency and power, under the currently active policy.

In search for possibly readily available solutions, we reviewed a number of published algorithms. One of the algorithms that relatively well-matched the requirements is called the *BSP—Binary Space Partitions* (cf. [18,19]). The research on BSP led us to the concept of a *k-d tree* (short for *k-dimensional tree*) [20], which is a space-partitioning data structure for organizing points in a *k-dimensional space*. The *k-d tree* data structure is a special case of a binary space partitioning tree that is used for representation and search that use multidimensional search keys (e.g., range searches and nearest neighbor searches). We have also looked at the graph-theoretic concept called *boxicity*. In graph theory, boxicity is a graph invariant. “The boxicity of a graph is the minimum dimension in which a given graph can be represented as an intersection graph of axis-parallel boxes”. In other words, “there must exist a one-to-one correspondence between the vertices of the graph and a set of boxes, such that two boxes intersect if and only if there is an edge connecting the corresponding vertices”. Our conclusion here was that this concept is not directly applicable to our problem. Finally, we analyzed the *R-trees* [21] and *R\*-trees* [22]. *R-tree* data structures are used for indexing multi-dimensional information that can be represented in terms of coordinates, e.g., geographical coordinates, rectangles or polygons. *R\*-trees* are a variant of *R-trees*. *R\*-trees* may have a higher computational complexity in the construction step than standard *R-trees*. The additional computational cost is due to the need to reinsert the data. However, this cost is compensated by a better query performance.

In our implementation, both policy rules and multi-criteria transmission requests are viewed as hyper-rectangles, which is a simplification assumption for hyperspaces, stemming from the fact that rules approximate geographical areas as rectangles as opposed to arbitrary polygons in a more general case. Edges in each hyper-rectangle dimension correspond to intervals along the geographical dimensions of longitude, latitude, elevation, the time interval, the frequency interval and the maximum transmission power. The problem of finding transmission opportunities for a given transmission request is then decomposed as:

- Retrieve all rules for the currently active policy.
- Discard the rules that representational hyperspaces do not intersect the hyperspace representing the query.
- For each remaining rule, determine the rule intersection with the query and insert the resulting hyperspace into either one of two lists, depending on the original rule type:
  - the list of “allow” rule intersections.
  - the list of “deny” rule intersections.

- In an effort to remove potential overlapping sections of hyperspaces from the results, if so required, for each new candidate hyperspace intersection with the query, one can also extract the subset of previously inserted hyperspaces that intersect the candidate to insertion, determine the complements of the candidate with the existing ones and insert those complements into the list in lieu of the whole candidate.
- Apply the disallowed rule intersections to the allowed list, effectively carving out their hyperspaces from the allowed hyperspace list wherever an overlap is found.
- Depending on the shapes to be carved out, this last step may result in multiple output hyperspaces for each input-allowed hyperspace.

Depending on the expected size of the active policy in terms of the number of rules, and given the simplification assumption introduced above, whereby geographical areas specified in the rules and the queries are rectangular, one can improve the algorithmic efficiency by indexing hyper-rectangles on each of their dimensions using Interval Trees. For that purpose, our implementation uses an augmented AVL implementation of a balanced binary search tree, which guarantees a worst-case complexity for the needed insertions and intersection lookups as  $\mathcal{O}(\log n)$ .

Once the opportunities for transmission are determined as stated above, one can simply derive a few useful functions, for example, get the allowed frequency ranges: given a multi-criteria transmission query that bounds the operational space—time interval and communication parameters—the algorithm returns an ordered list of disjoint frequency ranges that a radio may use to communicate with its peers. The processing follows the steps shown below.

- Evaluate the set of opportunities answering the query.
- Index the opportunities on frequency intervals using an interval tree.
- To ensure resulting intervals are disjoint, for each opportunity candidate for insertion in the tree, split the candidate into complements to readily inserted elements and insert those complements in lieu of the candidate.
- Traverse the tree in-order while merging adjacent intervals to create the ordered list of allowed frequency ranges.

We have also implemented a SPARQL-based version of searching for opportunities that is based on a rather complex SPARQL query associated with a couple of registered functions to determine whether hyper-rectangles intersect and what their intersection is.

#### 6.4.2. Rendezvous Channels

The objective here is to define the Rendezvous Channels Selection as an optimization problem. In other words, the idea is to first formalize the problem and then try to figure out whether this problem is possibly a case of a known problem already formulated (and possibly solved) somewhere else.

We are considering a scenario in which  $r$  wireless radio nodes distributed over some geographical 2D space are trying to find a number of channels on which they could communicate while not violating any of the policies that the radios must obey. For every two radios to communicate at the same time, assuming every pair needs a separate channel (i.e., every pair tuned to the same frequency, while all of the other pairs use different frequencies), we need  $f$  frequencies, where  $f$  is defined as:

$$f = r(r - 1) / 2$$

We can think of the value of  $f$  as the minimal requirement. The Controller will need more frequencies to choose from.

This requirement makes the problem complex since if there are  $r$  radios and  $f$  frequencies, there are  $f^r$  selection options, of which only  $f$  selections satisfy the requirement of tuning to the same frequency. While a full formalization of this problem would require a definition of the tuning precision, conceptually they must be on *the same frequency* in order to communicate.

The policies are known to all the radios, however, the radios know only their own locations but not the locations of the other nodes. Moreover, it is assumed that all the nodes must obey the policies, i.e., they can transmit only if there is a policy that allows them to transmit in a given location at a specified time on a given frequency, and there is no disallowing policy. The rendezvous algorithm will be executed on all the radios with the same input data (i.e., the same set of policies), and thus, the result should be the same for all the radios.

We are considering a scenario in which any two wireless radio nodes distributed over some geographical 2D space are trying to find common channels on which they could communicate. The objective here is to define the Rendezvous Channels Selection as an optimization problem. The rules (pertaining to the currently active policy) are known to all the radios; however, the radios know only their own locations but not the locations of the other nodes. Moreover, it is assumed that all the nodes must obey the rules, i.e., they can only transmit if there is a policy (rule) that allows them to transmit in a given location at a specified time on a given frequency. The rendezvous algorithm will be executed on all the radios with the same input data (i.e., the same set of rules) and is designed to guarantee some number of common channels between any two radios, even though the locations of radios are different.

#### Problem Formulation:

##### Given:

- A collection of “primary” policies,  $P$  (policies that cover all other “allow” policies geographically, e.g., D, E and F in Figure 8);
- A collection of “secondary” policies  $S$  (all “allow” rules, e.g., A, B, D, E and F in Figure 8);
- A ratio,  $r$ , of the number of selections of center frequencies from the policies  $F_P$  vs. the frequencies selected from the policies  $F_S$ ,  $r = |F_P|/|F|$ , where  $F = F_P \cup F_S$ ;
- The number of maximum communication channels,  $n$ ;
- One half of the width of each channel,  $b$ ;
- A set of sub-bands  $R = \{R_1, \dots, R_k\}$  (derived by the opportunities algorithm);
- Allowed frequency multiplier,  $\Delta$ .

##### Find:

- Center frequencies,  $F = \{f_1, \dots, f_n\}$

##### Constraints:

- $\forall_{f_i, f_j \in F} \bullet [f_i - b, f_i + b] \cap [f_j - b, f_j + b] = \emptyset$
- $\forall_{f_i \in F} \exists_{R_j \in R} \bullet [f_i - b, f_i + b] \subset R_j$
- $\forall_{f_i \in F} \exists_{m \in \mathcal{N}} \bullet f_i = m \cdot \Delta$
- $|F_P|/n \leq r$

The algorithm depends on the following three functions:

- **func<sub>1</sub>**—builds the list of opportunities with relaxed constraints, considering only the time-frequency map of the geographically all-encompassing rules of the active spectrum access policy and from there, carving out the time-frequency map of the union of disallowed rules, which is also expanded to the full geographical area covered by the active policy. The resulting set of opportunities is then passed through a “channel provider” filter that converts it to a list of “channels” determined by their center frequencies that obey parameterized constraints and limit the number of such channels returned to  $n \cdot r$ .
- **func<sub>2</sub>**—builds the list of opportunities with relaxed constraints, considering only the time-frequency map of “allow” rules of the active spectrum access policy and ignoring the disallowed rules of the said policy. The resulting set of opportunities is also passed through the “channel provider” filter while limiting the returned set of channels to  $n \cdot (1 - r)$

- ***func<sub>3</sub>***—consolidates the results by only retaining results from *func<sub>1</sub>* and *func<sub>2</sub>* that are within the fully constrained list of opportunities as computed from the active policy.

The three functions are leveraged in the main control flow of the rendezvous algorithm as follows:

- Retrieve the time-frequency map of the geographically all-encompassing rules of the active spectrum access policy (assuming there are such rules).
- Retrieve the set of “disallowed” rules from the active policy.
- Append the ordered list of frequencies returned by *func<sub>2</sub>* to the list of frequencies returned by *func<sub>1</sub>*.
- Then, apply the filter *func<sub>3</sub>* to the list of channels obtained from the previous steps.
- Return the filtered list, ordered and freed of any duplicates.

### 6.5. Learner

One of the main responsibilities of the PE is to automatically and periodically determine optimal values for the Modulation and Coding Scheme (MCS) and the Transmit Power (PTx). The inputs to this process are the current settings and the history of previously used values for these parameters. The proposed transmissions are checked by the Policy Reasoner for the purpose of checking conformance with the active policies. The policy rules are executed by the BaseVISor reasoning engine, which is part of the VISTology Inference Store component.

The conceptual view of the optimization and learning algorithm while selecting proposed values of transmit power and MCS is shown in Figure 9. According to this figure, the PE maintains a representation (a table) of possible transmit states (*knobs*) and relations to other neighboring knobs. Periodically, the algorithm selects a knob and sends it to the radio to use in the next transmissions. The radio then sends feedback to the PE summarizing the performance of the radio using the selected knobs in the current environment parameters. The feedback is computed based on the dependency of the Bit Error Rate (BER) on the ratio of Energy per Bit (Eb) to the Spectral Noise Density (No), Eb/No. After receiving the feedback, the PE updates its state table accordingly. This is the learning step. Then, the algorithm selects the next knob to be sent to the radio.

In order to analyze the performance of the PE, we implemented the initialization and the main loop of the algorithm in Matlab. The conceptual model of the state transitions (Figure 9) was implemented as a table. It is a  $5 \times 3$  matrix, in which each row corresponds to a specific value of the transmit power, while each column corresponds to a specific MCS. There are two matrices associated with this model; we call them PTable and KTable. The entries in the PTable represent the expected quality of the transmission using the specific knobs. These values are updated by the learning procedure. The KTable, on the other hand, keeps the probabilities of success weighted by the MCS. The entries of this table were calculated by multiplying the PTable by 2, 3 and 4 for QPSK, 8-PSK and 16-QAM, respectively.

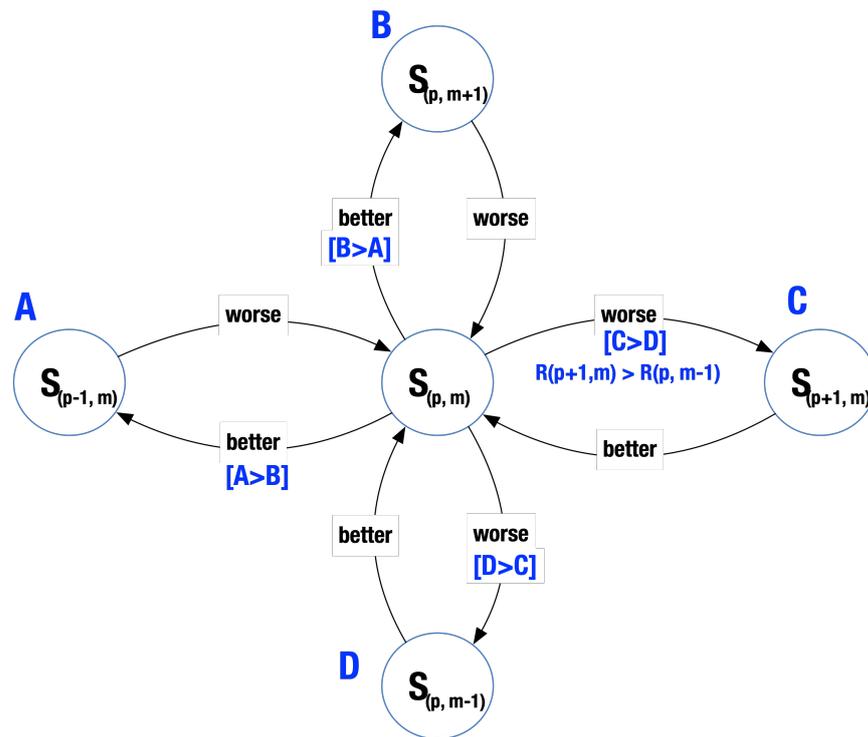


Figure 9. Modification of transmission parameters.

To compute the feedback, we first simulated the Bit Error Rates (BER) of various modulation schemes over an additive white Gaussian noise (AWGN) channel. The input argument for this simulation was,  $E_b/N_0$ , i.e., the ratio of bit energy to noise power spectral density, in dB, corresponding to the different  $E_b/N_0$  levels. The feedback calculations based on these models are shown in Figure 10.

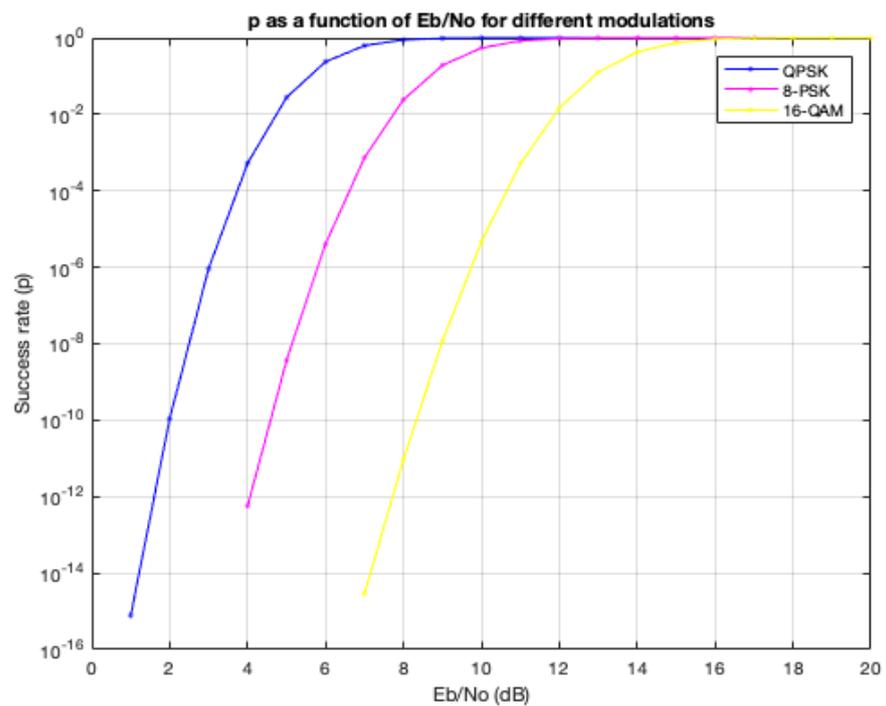


Figure 10. Success rate based on the BER models for different modulations.

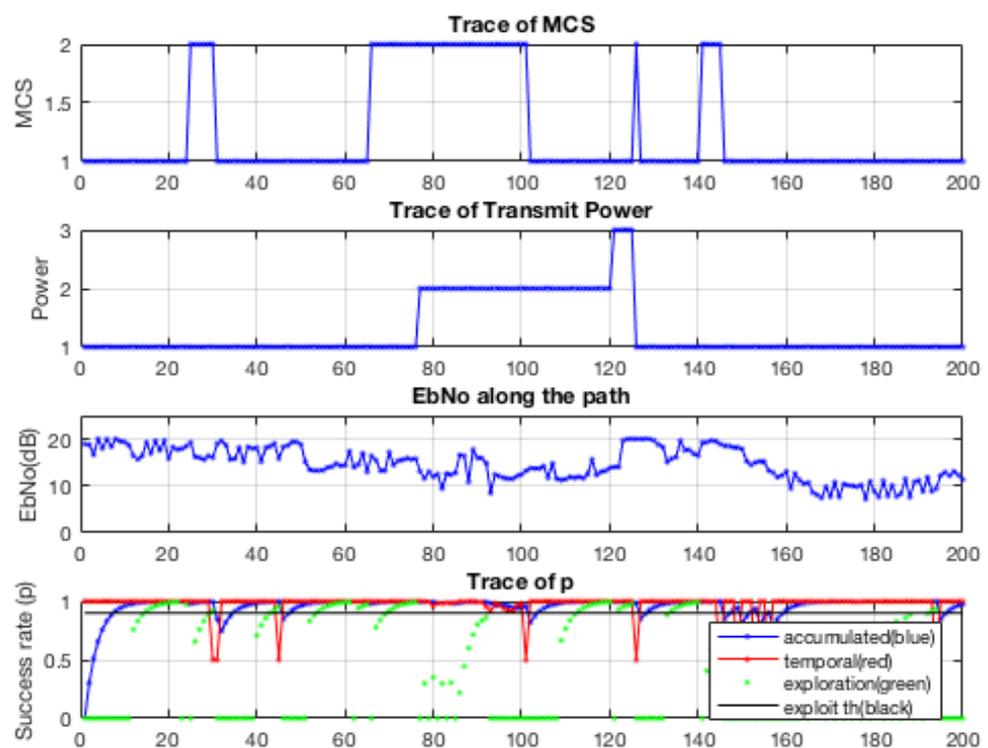
We used a simplified reinforcement learning procedure. First, noise was estimated according to the noise model described above. Based on the noise and the power level,  $P_{tx}$ , the values of  $E_b/N_0$  were computed and then used to find BER, Packet Error Rate (PER) and  $p$  (where  $N$  was set to 800):

$$PER = 1 - (1 - BER)^N$$

This was followed by the application of the value of  $p$  to the update rule of the  $PTable$ :

$$PTable(row, col) = p_{cur} + (p_{new} - p_{cur}) \cdot r$$

This was then followed by updating the  $KTable$ , as described earlier, and the selection of the maximal value and then using its matrix coordinates to select Power and MCS. To date, the above learning rule was used only for a single value update. In the next version of this algorithm, we plan to propagate the updates to other neighbors of the current node. The result of simulations are shown in Figure 11.



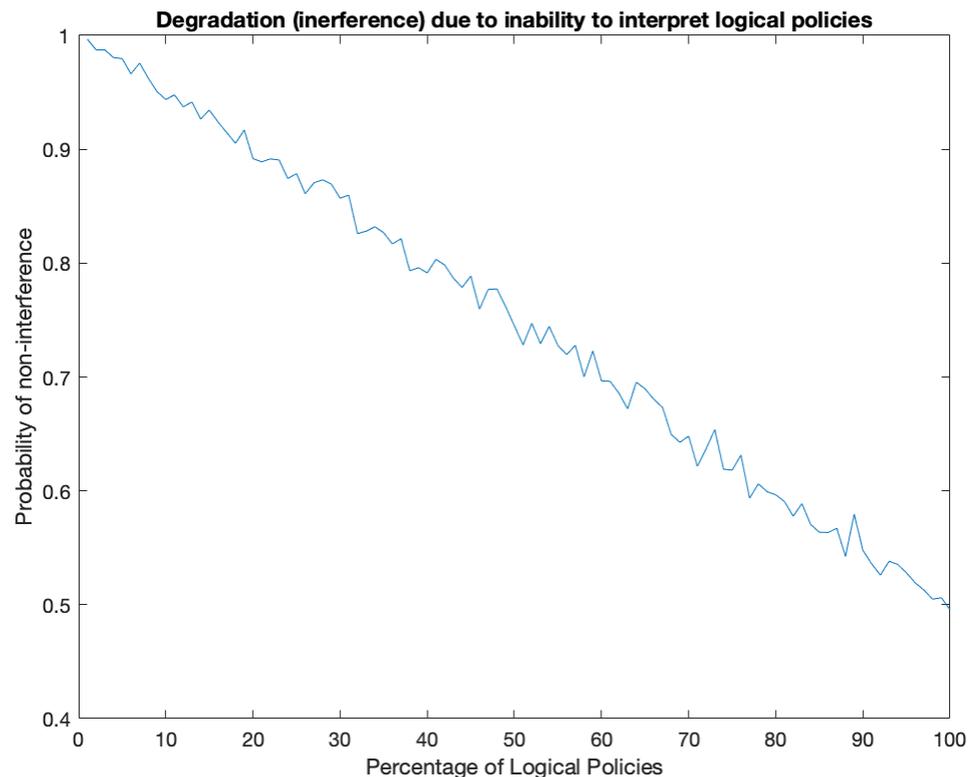
**Figure 11.** Simulation results indicating the behavior of the Learner in different conditions.

As we can see from this figure, in most of the steps, the algorithm selects values that guarantee the highest success rate. However, in some steps, the success rate falls very rapidly. Although it recovers in two steps, this is still a rather unexpected behavior, and such a drop is highly undesirable. The explanation of this behavior can be traced to the noise model. The drops coincide with the peaks of noise. There are two issues here. First of all, the modifications of the  $PTable$  by the learning algorithm make the table less than optimal. This pattern can be classified as some “instability” of the learning algorithm that results in the selection of less reliable MCS, e.g., 8-PSK instead of QPSK. Secondly, the algorithm does not have means to control the transmitter since the available transmit power is not sufficient.

## 7. Discussion

As stated in Section 1, the main objective described in this paper was to experiment with an ontology and policy-based system for dynamic spectrum sharing. Towards this

objective, we designed and implemented a system whose components were described above. In this section, we ask the question of how we can analyze, and possibly quantify, the impact of not having the functionalities of these modules on the overall performance of spectrum sharing. In short, we are interested in the question of what happens if spectrum management systems are not capable to interpret policies that involve logical rules (higher-level policies). Towards this end, we simulated scenarios in which policies were a mixture of both hypercube-based and logical rule-based policies. The results of these simulations are shown in Figure 12.



**Figure 12.** The impact of the missing capabilities on interpreting higher-level policies: simulation results.

This figure shows the dependence of non-interference on the percentage of the logical policies as opposed to the hypercube-based (quantitative) policies. In the simulations, it was assumed that at first, the percentage of the logical policies is negligible (on the left) and was increasing to 100% in the end (on the right). The second important assumption was that even if the interpretation is wrong, still, the interference would occur only randomly following the uniform distribution. For this reason, the right-hand side of the plot shows the degradation going down to only 0.5. This assumption is somewhat arbitrary, and its validity should be verified experimentally.

#### *Related Work*

There are many ways in which one could develop policy and ontology-based DSA systems. While the research in this area includes example implementations and discussions regarding the expressiveness of policy representation languages, more research is needed to develop guidelines for DSA practitioners. In particular, we argue that such research should provide quantitative metrics for system performance. To this end, we contributed two recent journal publications. In [23], we described a program to generate large volumes of device descriptions in OWL/RDF. The intent was to use this program for testing implementations of automatic reasoning systems for DSA. In [24], we described a program for automatic generation of queries against the repositories of device descriptions.

Currently, we are in the process of publishing another paper that compares XML with OWL as candidate languages for representing and querying cognitive radio capabilities.

The decision regarding the use of ontologies, in general, and for policy-based cognitive radios, in particular, is subject to a number of competing requirements:

- Expressiveness of the language for representing domain ontologies and policies;
- Computational complexity of reasoning;
- Capabilities of the ontology reasoning and policy reasoning engines;
- Behavioral characteristics of the engines;
- Use of open standards.

While a number of proposals and prototype implementations have been described in the literature, we briefly describe some of the most significant contributions. Each of the described approaches emphasizes some trade-offs between the expressiveness and the computational complexity. It is clear that the implementers have investigated various approaches, and none of them subsumes all the others. Our work is no exception to this.

Policy-based DSA systems have been promoted by the DARPA XG program. One of the outcomes was the Cognitive Radio Language (CoRaL) developed by SRI [25,26]. CoRaL is an expressive language; it provides both generic constructs such as types, functions and constraints, but it also defines various concepts that are specific to the cognitive radio domain. In particular, it uses constraints to specify the transmission opportunities and a CoRaL reasoner would return opportunities as simplifications of DSA policies. However, as the authors state, “A novel mix of reasoning techniques is required to implement such a reasoner.” Descriptions of two implementations of CoRaL were published: one in Prolog and one in Maude [27]—a language and theorem proving system with a highly desirable expressiveness. The first implementation did not support the inference of transmission opportunities. The second one, to the best of our knowledge, has not been developed as a standalone system that would address the undecidability issues of Maude.

Another attempt at a formal approach to DSA has been described in [28], in which Athena—a programming language and an interactive theorem proving environment—was used for experiments. It is based on first-order logic with arithmetic and algebraic data types and is a framework for working with constraint satisfaction problems (CSP) and satisfiability modulo theories (SMT), in which constraint solvers are used. According to the description of Athena available on its website, it is a theorem-proving environment that one can use for representing proofs and then verifying them. The formulation of the proofs is expected to be done by a human, or the human would have to specify some proof tactics that could be then used by the Athena processor. Clearly, Athena is a very good tool for experimenting with formal representations of policies and investigation of strategies for addressing issues of policy execution; however, this is not a standalone policy engine. The authors of [28] demonstrated that they can implement policy reasoning that includes the ability of returning transmission opportunities in case a request for transmission is not allowed.

Bahrak and their colleagues have developed a DSA system [29,30] that used a modified Pellet reasoner for checking the consistency of ontologies and reasoning about policies and leveraged SWRL as the policy representation language. The focus of their work was on the derivation of opportunity constraints, to which end, they developed four different algorithms. To the best of our understanding, their goal was to address the issue of undecidability of reasoning in SWRL by limiting its expressiveness. The policy reasoner (BRESAP) frames the policy reasoning problem as a graph-based Boolean function manipulation problem and represents policies as Multi-Terminal Binary Decision Diagrams. While this apparently addresses some issues related to the computation complexity, it is also a limitation with regard to the expressiveness; it is limited to policies that can be adequately represented as finite Boolean functions, which it encodes as binary decision diagrams (BDDs).

The last work that we selected comes from the Spectrum Sharing Company (SSC), one of the performers on the DARPA XG program. SSC was first developed to provide

a proof of feasibility to the policy-based DSA [31]. Their first paper did not provide any details about the reasoner that was used in the implementation (“the Policy Reasoner itself is implemented using numerous approaches”). The second paper [32] stated that OWL Lite (a less expressive profile of OWL) was used for representing ontologies, and SWRL was used for representing policies. However, it seems that the reasoner was a proprietary implementation (“Based on results from our initial Prolog-based implementation, we developed a C++ Policy Conformance Reasoner for meeting processor and main memory requirements”). Finally, in [33], SSC summarized their experiences with the work on the XG program. Overall, SSC has made great contributions towards greater acceptance of the policy-based DSA, but their products are not easily accessible to a wider community.

In summary, various ideas of how to address the trade-offs between the multiple competing criteria have been investigated. The main trade-off is between the expressiveness of the policy language and the computational cost of policy interpretation and execution. While the expressiveness ranges between the languages based on first-order logic and those based on propositional logic, the computational complexity ranges between linear time and undecidable. The expressiveness of our approach presented in this paper is based on W3C’s OWL 2 RL (less than OWL DL but significantly more than propositional logic), while the complexity is within the range of the decidable. The policy engine relies on procedural attachments, a mechanism for invoking procedural code from within rules, supported by VISTology’s BaseVISor OWL Reasoner. A comparison of BaseVISor with other OWL inference engines was presented in [17].

## 8. Conclusions

This paper addresses some of the issues associated with collaborative spectrum sharing. It describes some of the aspects we identified as relevant and important for any implementation of a spectrum sharing system. The main assumption was that rules for spectrum sharing could change after the deployment of radio networks, and the whole system must be able to adapt either on the fly or with as short interruptions as possible. To address such a requirement, we focused on a policy-based approach, in which transmissions are controlled by a policy interpreter system, while policies are represented in a formal language (OWL).

One of our primary goals was to develop a prototype of such a system. It is important to notice that three aspects of spectrum sharing, and not just dynamic spectrum access, were addressed. First of all, transmission opportunities were generated based on the active policies. They were all specific to radio locations. This insures that each radio can transmit according to the local policies. Second, rendezvous opportunities were computed locally, but all of them were based on the set of active policies. Finally, the adaptation mechanisms were implemented in each radio, allowing the incorporation of spectrum selection based on the propagation conditions at the radios’ locations. In summary, this paper outlines a number of issues relevant to dynamic spectrum sharing. However, due to the space limitations, the details of the implementation have not been presented.

**Author Contributions:** Conceptualization, M.M.K., J.M., J.-K.C., S.U. and J.W.C.; methodology, M.M.K., J.M. and J.-K.C.; software, J.M., J.-K.C. and M.M.K.; validation, J.-K.C. and J.M.; formal analysis, M.M.K., J.-K.C. and J.M.; investigation, J.M., M.M.K. and J.-K.C.; resources, S.U., J.W.C. and J.-K.C.; data curation, J.-K.C., S.U. and J.W.C.; writing—original draft preparation, M.M.K.; writing—review and editing, J.M. and J.-K.C.; visualization, J.M. and J.-K.C.; supervision, J.-K.C., M.M.K., S.U. and J.W.C.; project administration, S.U., J.W.C., J.-K.C. and M.M.K.; funding acquisition, J.-K.C., S.U. and J.W.C. All authors have read and agreed to the published version of the manuscript.

**Funding:** This work was supported by a grant-in-aid of Hanwha Systems and the Agency for Defense Development (ADD) in the Republic of Korea as part of the Contract UC160007ED.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Not applicable.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. DYSPAN P1900.1 Working Group. IEEE Standard Definitions and Concepts for Dynamic Spectrum Access: Terminology Relating to Emerging Wireless Networks, System Functionality, and Spectrum Management. In *Technical Report*; IEEE: Piscataway, NJ, USA, 2008.
2. Li, S.; Kokar, M.M. *Flexible Adaptation in Cognitive Radios*; Springer: Berlin/Heidelberg, Germany, 2012.
3. Marinez-Moyano, I.J. Exploring the Dynamics of Collaboration in Interorganizational Settings. In *Creating a Culture of Collaboration*; Schuman, S., Ed.; Wiley: Hoboken, NJ, USA, 2006; pp. 69–86.
4. W3C. OWL 2 Web Ontology Language Document Overview. 2009. Available online: <http://www.w3.org/TR/owl2-overview/> (accessed on 29 July 2021).
5. Studer, V.R.; Benjamins, D.F. Knowledge Engineering: Principles and Methods. *Data Knowl. Eng.* **2010**, *25*, 161–197. [[CrossRef](#)]
6. Laird, A.N.; Rosenbloom, P.S. Soar: An architecture for general intelligence. *J. Artif. Intell.* **1987**, *33*, 1–64. [[CrossRef](#)]
7. Anderson, J.R.; Bothell, D.; Byrne, S.D.; Lebiere, C.; Qin, Y. An integrated theory of the mind. *Psychol. Rev.* **2004**, *111*, 1036–1060. [[CrossRef](#)] [[PubMed](#)]
8. Endsley, M.R. Bringing cognitive engineering to the information fusion problem: Creating systems that understand situations. In *Proceedings of the Fusion'2011: International Conference on Information Fusion*, Chicago, IL, USA, 5–8 July 2011.
9. Boyd, J. A Discourse on Winning and Losing. In *Technical Report*; Maxwell AFB: Montgomery, AL, USA, 1987.
10. Mitola, J.J. Cognitive radio for flexible mobile multimedia communications. In *Proceedings of the IEEE International Workshop on Mobile Multimedia Communications (MoMuC'99)*, San Diego, CA, USA, 15–17 November 1999; pp. 3–10.
11. Musen, M.A. The protégé project: A look back and a look forward. *AI Matters* **2015**, *1*, 4–12. [[CrossRef](#)] [[PubMed](#)]
12. RDF Core Working Group. Resource Description Format (RDF) Primer. W3C Recommendation. 10 February 2004. Available online: <http://www.w3.org/TR/rdf-primer/> (accessed on 29 July 2021).
13. Moskal, J.; Kokar, M.; Roman, V.; Normoyle, R.; Guseman, R. Towards a SpectralSPARQL Standard for Exchanging EMS Knowledge. In *Proceedings of the Military Communications Conference*, Baltimore, MD, USA, 23–25 October 2017.
14. OGC. OGC GeoSPARQL-A Geographic Query Language for RDF Data. Technical Report, Open Geospatial Consortium. 2012. Available online: <http://www.opengis.net/doc/IS/geosparql/1.0> (accessed on 29 July 2021).
15. Cognitive Radio Ontology. Available online: <https://cogradio.org/> (accessed on 30 May 2021).
16. W3C. SPARQL Query Language for RDF. W3C Candidate Recommendation. 2006. Available online: <http://www.w3.org/TR/2006/CR-rdf-sparql-query-20060406/> (accessed on 29 July 2021).
17. Matheus, C.; Baclawski, K.; Kokar, M.M. BaseVISor: A Triples-Based Inference Engine Outfitted to Process RuleML and R-Entailment Rules. In *Proceedings of the 2nd International Conference on Rules and Rule Languages for the Semantic Web*. ISWC, Athens, GA, USA, 10–11 November 2006.
18. Dumitrescu, A.; Mitchell, J.S.B.; Sharir, M. Binary Space Partitions for Axis-Parallel Segments, Rectangles, and Hyperrectangles. *Discret. Comput. Geom.* **2004**, *31*, 207–227. [[CrossRef](#)]
19. Nguyen, V.H.; Widmayer, P. Binary space partitions for sets of hyperrectangles. In *Algorithms, Concurrency and Knowledge*; Kanchanasut, K.; Lévy, J.J., Eds.; Springer: Berlin/Heidelberg, Germany, 1995; pp. 59–72.
20. De Berg, M.; Cheong, O.; Van Kreveld, M.; Overmars, M., Orthogonal Range Searching. In *Computational Geometry: Algorithms and Applications*; Springer: Berlin/Heidelberg, Germany, 2008; pp. 95–120. [[CrossRef](#)]
21. Guttman, A. R-trees: A Dynamic Index Structure for Spatial Searching. In *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, New York, NY, USA, 1 June 1984; Volume SIGMOD '84, pp. 47–57. [[CrossRef](#)]
22. Beckmann, N.; Kriegel, H.P.; Schneider, R.; Seeger, B. The R\*-tree: An Efficient and Robust Access Method for Points and Rectangles. In *Proceedings of the 1990 ACM SIGMOD International Conference on Management of Data*, New York, NY, USA, 1 May 1990; Volume SIGMOD '90, pp. 322–331. [[CrossRef](#)]
23. Chen, Y.; Kokar, M.M.; Moskal, J.J. RDF object description generator. *Int. J. Web Eng. Technol.* **2020**, *15*, 140–169. [[CrossRef](#)]
24. Chen, Y.; Kokar, M.M.; Moskal, J.J. SPARQL Query Generator (SQG). *J. Data Semant.* **2021**, 1–17. [[CrossRef](#)]
25. Elenius, D.; Denker, G.; Stehr, M.O.; Senanayake, R.; Talcott, C.; Wilkins, D. CoRaL-Policy Language and Reasoning Techniques for Spectrum Policies. In *Proceedings of the Eighth IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'07)*, Bologna, Italy, 13–15 June 2007; pp. 261–265.
26. Wilkins, D.; Denker, G.; Stehr, M.O.; Elenius, D.; Senanayake, R.; Talcott, C. Policy-Based Cognitive Radios. *IEEE Wirel. Commun.* **2007**, *14*, 41–46. [[CrossRef](#)]
27. Maude Software Language. Available online: <https://www.sri.com/hoi/maude-software-language/> (accessed on 12 July 2021).
28. Arkoudas, K.; Chadha, R.; Chiang, C. An Application of Formal Methods to Cognitive Radios. In *DIFTS@FMCAD*; RWTH Aachen University: Aachen, Germany, 2011.
29. Bahrak, B.; Deshpande, A.; Jerry Park, J.M. Spectrum access policy reasoning for policy-based cognitive radios. *Comput. Net.* **2012**, *56*, 2649–2663. [[CrossRef](#)]

30. Bahrak, B.; Park, J.M.; Wu, H. Ontology-based spectrum access policies for policy-based cognitive radios. In Proceedings of the 2012 IEEE International Symposium on Dynamic Spectrum Access Networks, Bellevue, WA, USA, 16–19 October 2012; pp. 489–500.
31. Perich, F.; Foster, R.; Tenhula, P.; McHenry, M. Experimental Field Test Results on Feasibility of Declarative Spectrum Management. In Proceedings of the 2008 3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks, Chicago, IL, USA, 14–17 October 2008; pp. 1–10.
32. Perich, F.; McHenry, M. Policy-based spectrum access control for dynamic spectrum access network radios. *J. Web Semant.* **2009**, *7*, 21–27. [[CrossRef](#)]
33. Perich, F.; Morgan, E.; Ritterbush, O.; McHenry, M.; D'Itri, S. Efficient dynamic spectrum access implementation. In Proceedings of the 2010–Milcom 2010 Military Communications Conference, San Jose, CA, USA, 31 October–3 November 2010; pp. 1887–1892.