OWL-POLAR: Semantic Policies for Agent Reasoning*

Murat Şensoy[1], Timothy J. Norman[1], Wamberto W. Vasconcelos[1], and Katia Sycara[1,2]

[1] Department of Computing Science, University of Aberdeen, AB24 3UE, Aberdeen, UK
{m.sensoy,t.j.norman,w.w.vasconcelos}@abdn.ac.uk
[2] Carnegie Mellon University, Robotics Institute, Pittsburgh, PA 15213, USA
katia@cs.cmu.edu

Abstract. Policies are declarations of constraints on the behaviour of components within distributed systems, and are often used to capture norms within agent-based systems. A few machine-processable representations for policies have been proposed, but they tend to be either limited in the types of policies that can be expressed or limited by the complexity of associated reasoning mechanisms. In this paper, we argue for a language that sufficiently expresses the types of policies essential in practical systems, and which enables both policy-governed decision-making and policy analysis within the bounds of decidability. We then propose an OWL-based representation of policies that meets these criteria using and a reasoning mechanism that uses a novel combination of ontology consistency checking and query answering. In this way, agent-based systems can be developed that operate flexibly and effectively in policy-constrained environments.

1 Introduction

In this paper, we present a novel and powerful OWL 2.0 [7] knowledge representation and reasoning mechanism for policies: OWL-POLAR (an acronym for OWL-based POlicy Language for Agent Reasoning). Policies (aka. norms) are system-level principles of ideal activity that are binding upon the components of that system. Depending on the nature of the system itself, policies may serve to control, regulate or simply guide the activities of components. In systems security, for instance, the aim is typically to control behaviour such that the system complies with the policies [18]. In real socio-technical systems, however, there are important limits to this and the aim is to develop effective sets of policies along with incentives to regulate behaviour [2]. In systems of autonomous agents, the term norm is most prevalent, but the concept and issues remain the same [4]; for example, norms are used to regulate the behaviour of agents representing disparate interests in electronic institutions [6]. The objective of this research is to capture the essential requirements of policy representation and reasoning. In meeting this objective three key requirements must be met:

---

1. System/institutional policies must be machine understandable and underpinned by a clear interpretation.
2. The representation must be sufficiently expressive to capture the notion of a policy across domains.
3. Policies must be able to be effectively shared/interpreted at run-time.

The choice of OWL 2.0 as an underlying language addresses the first requirement, but in meeting the second two, we must clearly outline what is required of a policy language and what reasoning should be supported by it. The desiderata of a model of policies that motivates the language OWL-POLAR are as follows:

- **Representational adequacy.** Policies (or norms) must capture the distinction between activities that are required (obliged), restricted (prohibited) and, in some way, authorised but not necessarily expected (permitted) by some representational entity within the environment. It is essential to capture the authority from which the policy/norm comes, the subject (agent) to whom it applies, the object (activity) to which the policy/norm refers, and the circumstances within which it applies.
- **Supporting decisions.** Any reasoning mechanism that is driven/guided by policies must support both the determination of what policies/norms apply in a given situation, and what activities are warranted by the normative state of the agent if it were to comply with these policies.
- **Supporting analysis.** Any reasoning mechanism that is driven/guided by normative/policy constraints must support the assessment of policies in terms of: (i) whether a policy/norm is meaningful and (ii) whether norms conflict, and in what circumstances they do conflict.

With the introduction of data ranges within the OWL 2.0 specification [7], we believe that this desiderata of a model of policies can be met within the confines of OWL-DL. If this claim can be shown to be valid (as we aim to do within this paper), we believe that OWL-POLAR provides, for the first time, a sufficiently expressive policy language for which the key reasoning mechanisms required of such a language are decidable.

The paper is organised as follows: in Section 2 we formally specify the OWL-POLAR language within OWL-DL; in Section 3 we describe how a set of active policies may be computed, and how decisions about what activities are warranted by some set of policies may be made; then in Section 4 we present in detail the reasoning mechanisms that support the analysis of policies. OWL-POLAR is then compared to existing languages for policies in Section 5, and we present our conclusions in Section 6.

2   Semantic Representation of Policies

The proposed language for semantic representation of policies is based on OWL-DL [7]. An OWL-DL ontology $o = (TBox_o, ABox_o)$ consists of a set of axioms defining the classes and relations ($TBox_o$) as well as a set of assertional axioms about the individuals in the domain ($ABox_o$). Concept axioms have the form $C \sqsubseteq D$ where C and D are concept descriptions, and relation axioms are expressions of the form $R \sqsubseteq S$, where R and S are relation descriptions. The ABox contains concept assertions of the form $C(a)$ where C is a concept and a is an individual name, and relation assertions of the form $R(a, b)$, where R is a relation and a and b are individual names.

Conjunctive semantic formulas are used to express policies. A conjunctive semantic formula $F_{\boldsymbol{v}}^o = \bigwedge_{i=0}^{n} \phi_i$ over an ontology o is a conjunction of atomic assertions ϕ_i, where $\boldsymbol{v} = \langle ?x_0, \ldots, ?x_n \rangle$ represents a vector of variables used in these assertions. For the sake of convenience, we assume $\bigwedge_{i=0}^{n} \phi_i \equiv \{\phi_1, \ldots \phi_n\}$ in order to consider a conjunctive formula as a set of atomic assertions. Based on this, $F_{\boldsymbol{v}}^o$ can be considered as $T_{\boldsymbol{v}}^o \cup R_{\boldsymbol{v}}^o \cup C_{\boldsymbol{v}}^o$, where $T_{\boldsymbol{v}}^o$ is a set of type assertions using the concepts from o, e.g., $\{student(?x_i), nurse(?x_j)\}$; $R_{\boldsymbol{v}}^o$ is set of of relation assertions using the relations from o, e.g., $\{marriedTo(?x_i, ?x_j)\}$; $C_{\boldsymbol{v}}^o$ is a set of constraint assertions on variables. Each constraint assertion is of the form $?x_i \triangleleft \beta$, where β is a constant and $\triangleleft$ is any of the symbols $\{>, <, =, \neq, \geq, \leq\}$. A constant is either a data literal (e.g, a numerical value) or an individual defined in o.

Variables are divided into two categories; data-type and object variables. A data-type variable refers to data values (e.g., integers) and can be used only once in $R_{\boldsymbol{v}}^o$. On the other hand, an object variable refers to individuals (e.g., University_of_Aberdeen) and can be used freely many times in $R_{\boldsymbol{v}}^o$. Equivalence and distinction between the values of object variables can be defined using OWL properties *sameAs* and *differentFrom* respectively, e.g., *owl:sameAs(?x, ?y)*. In the rest of the paper, we use the symbols α, ρ, φ, and e as a short hand for semantic formulas.

Given an ontology o, a conditional policy is defined as $\alpha \longrightarrow N_{\chi:\rho} (a : \varphi) / e$, where

1. α, a conjunctive semantic formula, is the activation condition of the policy.
2. $N \in \{O, P, F\}$ indicates if the policy is an obligation, permission or prohibition.
3. χ is the policy addressee and ρ describes χ using only the *role* concepts from the ontology (e.g., $?x : student(?x) \wedge female(?x)$, where *student* and *female* are defined as sub-concepts of the *role* concept in the ontology). That is, ρ is of the form $\bigwedge_{i=0}^{n} r_i(\chi)$, where $r_i \sqsubseteq role$. Note that χ may directly refer to a specific individual (e.g., *John*) in the ontology or a variable.
4. $a : \varphi$ describes what is prohibited, permitted or obliged by the policy. Specifically, a is a variable referring to the action to be regulated by the policy and φ describes a as an action instance using the concepts and properties from the ontology (e.g., $?a : SendFileAction(?a) \wedge hasReceiver(?a, John) \wedge hasFile(?a, TechReport218.pdf)$, where *SendFileAction* is an *action* concept). Each action concept has only a number of functional relations (aka. functional properties) [7] and these relations are used while describing an instance of that action.
5. e defines the expiration condition.

Table 1 illustrates how a conditional policy can be represented using the proposed approach. The policy in the table states that a person is obliged to leave a location when there is a fire risk.

Table 1. A person has to leave a location when there is a fire risk

α	$Place(?b) \wedge hasFireRisk(?b, true) \wedge in(?x, ?b)$
N	O
$\chi : \rho$	$?x : Person(?x)$
$a : \varphi$	$?a : LeavingAction(?a) \wedge about(?a, ?b) \wedge hasActor(?a, ?x)$
e	$hasFireRisk(?b, false)$

Given a semantic representation for the state of the world, policies are used to reason about actions that are permitted, obliged or prohibited. Let Δ_o be a semantic representation for a state of the world based on an ontology o. Each state of the world is partially observable; hence Δ_o is a partial representation of the world. Δ_o itself is represented as an ontology composed of $(TBox_o, ABox_\Delta)$ where $ABox_\Delta$ is an extension of $ABox_o$.

3   Reasoning with Policies

When its activation conditions are satisfied, a conditional policy leads to an activated policy. Definition 1 summarizes how a conditional policy is activated using ontological reasoning over a state of the world. Here we use query answering to determine activated policies and reason about actions. The query answering mechanism we use in this work is DL-safe; i.e. variables are bound only to the named individuals, to guarantee decidability [8]. In this section, we address some of the key issues in supporting decisions governed by policies: activation and expiration, and reasoning about interactions between policies and actions.

Definition 1. Let Δ_o be a state of the world represented based on a domain ontology o. If there is a substitution σ such that $\Delta_o \vdash (\alpha \wedge \rho) \cdot \sigma$, but there is no substitution σ' such that $\Delta_o \vdash (e \cdot \sigma) \cdot \sigma'$, then the policy $(N_\chi (a : \varphi)) \cdot \sigma$ becomes active. This policy expires when there exists a substitution σ' such that $\Delta_o \vdash (e \cdot \sigma) \cdot \sigma'$. ∎


Fig. 1. A partial state of the world represented based on a domain ontology

3.1   Policy Activation

A policy is activated for a specific agent when the world state is such that the activation condition holds for that agent and the expiration condition does not hold, and expires when this latter condition holds. The above definition is rather standard [11], but we now describe how this is implemented efficiently through query answering. A conjunctive semantic formula can be trivially converted to a SPARQL query [15] and can be evaluated by OWL-DL reasoners with SPARQL-DL [17] support such as Pellet [17] to

find a substitution for its variables satisfying a specific state of the world. Therefore, we can test $\Delta_o \vdash (\alpha \land \rho) \cdot \sigma$ by writing a query for $(\alpha \land \rho)$ and testing whether it is entailed by Δ_o or not. Consider the conditional policy in Table 1 and assume that we have the partially represented state of the world in Figure 1. We can write the semantic query in Figure 2 to find σ for the conditional policy. When we query the state of the world using SPARQL, each result in the result set provides a substitution σ; in our case, we have two σ values: {*?x/John, ?b/Room245*} and {*?x/Jane, ?b/Room245*}, representing that there is a fire risk in the room 245 of the Central Hospital and that John and Jane are in that room.

```
Query:

  q(?x, ?b):-
    Place(?b) ∧
    hasFireRisk(?b, true) ∧
    Person(?x) ∧
    in(?x,?b).
```

```
SPARQL SYNTAX:

  PREFIX example: <http://www.example.com/ns#>
  PREFIX rdf: <http://www.w3.org/...rdf-syntax-ns#>
  PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
  SELECT ?x ?b
  WHERE {
    ?b rdf:type example:Place.
    ?b example:hasFireRisk "true"^^xsd:boolean.
    ?x rdf:type example:Person.
    ?x example:in ?b.
  }
```

Fig. 2. Query for the activation of a policy

Now, using the computed σ values, we should try to find a σ' such that $\Delta_o \vdash (e \cdot \sigma) \cdot \sigma'$. In our case, for this purpose, we can use the semantic query "*q():- has-FireRisk(Room245, false)*". When the SPARQL representation of this query is executed over the state of the world shown in Figure 1, it returns *false*; that is the RDF graph pattern represented by the query could not be found in the ontology. This means that the policy in Table 1 should be activated using the variable bindings in σ. The result is activations of O_{John}(*?a : LeavingAction(?a) ∧ about(?a,Room245)*) and O_{Jane}(*?a : LeavingAction(?a) ∧ about(?a,Room245)*). These policies mean that *John* and *Jane* are obliged to leave the *room 245*; the obligation expires when the fire risk is removed.

3.2   Reasoning about Actions

Let us assume that a specific action $a' : \varphi'$ will be performed by x, where a' is a URI referring to the action instance and φ' is a conjunctive semantic formula describing a' without using any variables. Let Δ_o be the current state of the world. We can test if the action a' is permitted, forbidden or prohibited in Δ_o. For this purpose, based on Δ_o, we create a "sandbox" (hypothetical) state of the world Δ'_o to make *what-if* reasoning [21], i.e., Δ'_o shows what happens if the action is performed. This is achieved by simply adding the described action instance to Δ_o, i.e., $\Delta'_o = \Delta_o \cup \varphi'$. For example, the state of the world in Figure 1 is extended using action instance *LeaveAct_1: LeavingAction(LeaveAct_1) ∧ hasActor(LeaveAct_1,John) ∧ about(LeaveAct_1,room245)*. The resulting state of the world is shown in Figure 3.
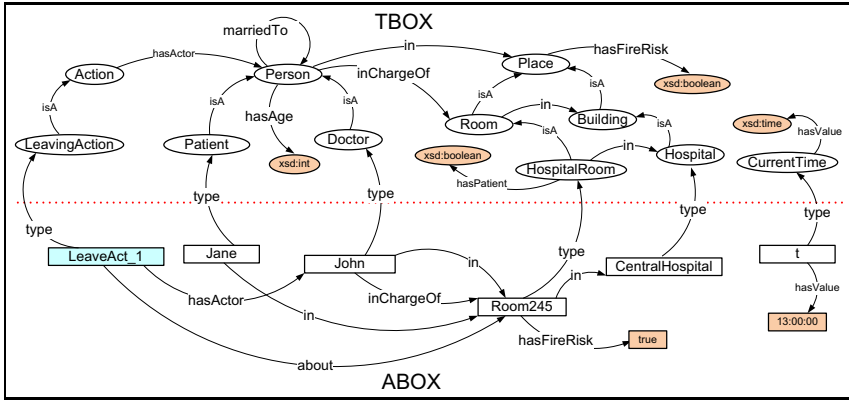
Fig. 3. The "sandbox" (hypothetical) state of the world

For each active policy $N_x(y : \varphi_y)$, we test the expiration conditions on Δ'_o as explained before. If the policy's expiration conditions are satisfied, we can conclude that the action $a' : \varphi'$ leads to the expiration of the policy. Otherwise, a semantic query Q of the form $q(\boldsymbol{v}_{\varphi_y})$:- φ_y is created, where $\boldsymbol{v}_{\varphi_y}$ is the vector of variables in φ_y. Then, Δ'_o is queried with Q. Let the query return a result set rs; each result $r \in rs$ is a substitution such that $\Delta'_o \vdash \varphi_y \cdot r$. If $y \cdot r = a'$ for any such r, then a' is regulated by the policy. In this case, we can interpret the policy based on its modality as follows:

1. $N_x = O$**:** In this case, the policy represents an obligation; that is, x is obliged to perform a'. Performing a' will remove this obligation.
2. $N_x = P$**:** Performing a' is explicitly permitted.
3. $N_x = F$**:** Performing a' is prohibited.

After examining the active policies as described above, we can identify a number of possible normative positions with respect to the action instance a': (i) doing a' may be explicitly permitted if there is a policy permitting it; (ii) doing a' may be obligatory if there exists a policy obliging it; (iii) doing a' may be prohibited if there is a policy prohibiting it; and (iv) there may be a conflict in the normative position with respect to a' if it is either both prohibited and explicitly permitted, or both prohibited and obliged.

4   Reasoning about Policies

In this section, we demonstrate reasoning techniques to support the analysis of policies in terms of their meaningfulness (Section 4.2) and possibility of conflict (Section 4.3), and hence address our third desideratum. Prior to this, however, we propose methods for reasoning about semantic formulas to underpin our mechanisms for policy analysis.

4.1   Reasoning about Semantic Formulas

Here, we introduce methods for reasoning about semantic conjunctive formulas using query freezing and constraint transformation.

Conjunctive Queries. There is a relation between conjunctive formulas and conjunctive queries. A conjunctive semantic formula can trivially be converted into a conjunctive semantic query. For example, $A^o_{v_1}$ can be converted into the query $q_A():- A^o_{v_1}$. Therefore, we can use query reasoning techniques to reason about semantic formulas. For instance, in order to reason about the subsumption between semantic formulas, we can use query subsumption (containment).

In conjunctive query literature, in order to test whether q_A subsumes q_B, the standard technique of *query freezing* is used to reduce query containment problem to query answering in Description Logics [13,20]. For this purpose, we build a canonical ABox Φ_{q_B} from the query $q_B():- B^o_{v_2}$ in three steps. First, for each variable in v_2, we put a fresh individual into Φ_{q_B} using the type assertions about the variable. Note that this individual should not exist in o. Second, we add each individual appearing in q_B into Φ_{q_B}. This is done using the information about the individual from the $ABox_o$ (e.g., type assertions). Third, relationships between individuals and constants defined in q_B are inserted into Φ_{q_B}. As a result of this process, Φ_{q_B} contains a pattern that exists only in ontologies that satisfy q_B. We combine Φ_{q_B} and our $TBox_o$ to create a new canonical ontology, $o' = (TBox_o, \Phi_{q_B})$. Example 1 demonstrates a simple case. Based on [20,13], we conclude that $o \vdash q_B \sqsubseteq q_A$ if and only if o' entails q_A. In order to test whether o' entails q_A or not, we query o'. That is, o' entails q_A if there exists at least one match for q_A in o'. This can easily be achieved by converting q_A to SPARQL syntax and use Pellet's SPARQL-DL query engine to answer q_A on o' [17].

Example 1. Let query q_A be *q():- Person(?p) ∧ marriedTo(?p,?x) ∧ Patient(?x)* and query q_B be *q():- Doctor(?x) ∧ marriedTo(?x,Jane) ∧ hasChild(?x,?c)*. Then, Φ_{q_B} contains an individual x, which is created for the variable ?x. The individual x is defined as of type $Doctor$. In Φ_{q_B}, we also have another individual $Jane$, which is defined in the original $ABox_o$ as an instance of the *Patient* class; we get all of its type assertions from the $ABox_o$. Then, we insert the object property *marriedTo* between the individuals x and $Jane$. Lastly, we create another individual c for the variable ?c in Φ_{q_B} and insert the $hasChild$ object property between x and c. The resulting ontology is shown in Figure 4.


Fig. 4. The ontology created for q_B in Example 1

```
<owl:Class rdf:about="#AgeConst1">
 <rdfs:subClassOf>
    <owl:Restriction>
      <owl:onProperty rdf:resource="#hasAge"/>
        <owl:allValuesFrom>
         <rdfs:Datatype>
            <owl:onDataRange rdf:resource="&xsd;nonNegativeInteger"/>
            <xsd:minInclusive rdf:datatype="&xsd;int">10</xsd:minInclusive>
            <xsd:maxExclusive rdf:datatype="&xsd;int">20</xsd:maxExclusive>
         </rdfs:Datatype>
        </owl:allValuesFrom>
    </owl:Restriction>
 </rdfs:subClassOf>
</owl:Class>
```

Fig. 5. A concept named **AgeConst1** is created for *hasAge(?c, ?a) ∧ ?a ≥ 10 ∧ ?a ≤ 20*

The query freezing method described above enables us to create a canonical ABox for a semantic conjunctive formula; this ABox represents a pattern which only exists in ontologies satisfying the semantic formula. On the other hand, this method assumes that variables in queries can be assigned fresh individuals in a canonical ABox. However, in OWL-DL, individuals can refer to objects, but not data values [19]. Therefore, the proposed query freezing method can be used to test for subsumption between q_A and q_B only if the variables in q_A and q_B refer to objects. A variable can refer to an object if it is used as the domain of an object or datatype property (e.g., *hasAge(?x,10)*) or if it is used as the range of an object property (e.g., *marriedTo(Jack,?x)*). Unfortunately, in many real-life settings, queries may have variables referring to data values with various constraints, which we refer to here as *datatype variables*. In these settings, the query freezing described above cannot be used to test subsumption. Example 2 illustrates a simple scenario.

Example 2. Let query q_A be *q():- Person(?p) ∧ hasChild(?p,?c) ∧ hasAge(?c,?y) ∧ ?y ≥ 12 ∧ ?y ≤ 16* and query q_B be *q():- Doctor(?x) ∧ marriedTo(?x,Jane) ∧ hasChild(?x, ?c) ∧ hasAge(?c,?a) ∧ ?a ≥ 10 ∧ ?a ≤ 20*. In this example, the query freezing method cannot be used directly to test subsumption between q_A and q_B, because the variables *?y* and *?a* refer to data values, which cannot be represented by individuals in an OWL-DL ontology.

Constraint Transformation. Here, we propose *constraint transformation*. It is a pre-processing step which enables us to create a canonical ABox for semantic formulas with datatype variables. Note that a datatype variable is used in a semantic formula to constrain one datatype property, e.g., *?y* is used to constrain the $hasAge$ datatype property in q_A of Example 2. Constraint transformation in contrast uses *data-ranges* introduced in OWL 2.0 [7] to transform each constrained datatype property to a named OWL class. As a result, datatype variables and related datatype properties and constraints are replaced with type assertions. This procedure is detailed in Algorithm 1.

The algorithm takes a conjunctive semantic formula F_v^o and the ontology o as inputs (line 1). F_v^o is of the form $T_v^o \cup R_v^o \cup C_v^o$, where T_v^o, R_v^o, and C_v^o are sets of type, relation and constraint assertions respectively. The outputs of the algorithm are the transformed

Algorithm 1. Constraint transformation

1: **Inputs:**   Formula $F_v^o \equiv T_v^o \cup R_v^o \cup C_v^o$,
                    Ontology $o \equiv (ABox_o, TBox_o)$
2: **Outputs:**   Formula F_u^ϕ,
                    Ontology $\phi \equiv (ABox_o, TBox_\phi)$
3: **Initialization:** $F_u^\phi = T_v^o, \ TBox_\phi = TBox_o$
4: **for all** $(r(a, b) \in R_v^o)$ **do**
5:     **if** $(isDatatypeVariable(b))$ **then**
6:         $\gamma_d = getConstraints(b, C_v^o)$
7:         $c = createConcept(r, \gamma_d, TBox_\phi)$
8:         $\tau = createTypeAssertion(a, c)$
9:         $F_u^\phi = F_u^\phi \cup \tau$
10:    **else**
11:        $F_u^\phi = F_u^\phi \cup r(a, b)$
12:        $\gamma_b = getConstraints(b, C_v^o)$
13:        **if** $(\gamma_b \neq \emptyset \ \& \ \neg(\gamma_b \subset F_u^\phi))$ **then**
14:            $F_u^\phi = F_u^\phi \cup \gamma_b$
15:        **end if**
16:    **end if**
17: **end for**

semantic formula F_u^ϕ (containing no datatype variables) and the updated ontology ϕ (line 2). Initially, F_u^ϕ is set as equal to T_v^o and ϕ is the same as o (line 3). For each relation assertion $r(a, b)$ in R_v^o, we do the following (line 4). First, we check if b is a datatype variable (line 5). If so, this means that r is a datatype property with a variable in its range. In this case, we extract the set of constraints related to b from C_v^o, which is referred by γ_d (line 6). Based on r and γ_d, we create a concept c in $TBox_\phi$ using the *createConcept* function (line 7). This function works as follows:

1. If $\gamma_d \neq \emptyset$, then b implies some restrictions on the range of r. In this case, c should refer to objects that have the property r with the restrictions defined in γ_d on its range. While creating c in $TBox_\phi$, we use *data-ranges*[1] introduced in OWL 2.0 to restrict the range of r accordingly. For example, if $r(a, b)$ corresponds to *hasAge(?c, ?a)* and γ_d = {*?a ≥ 10, ?a ≤ 20*}, then a concept named *AgeConst1* can be described as shown in Figure 5. For more sophisticated constraints, we create more complex class expressions using the OWL constructors *owl:unionOf*, *owl:intersectionOf*, and *owl:complementOf*.
2. If $\gamma_d = \emptyset$, then b has no constraints, which means that the data-range of b is equivalent to the range of its data-type (i.e., for *xsd:int*, the range is min inclusive -2147483648 and max inclusive 2147483647).

After creating the concept c in $TBox_\phi$, we create a type assertion τ to declare a as an instance of c (e.g., *AgeConst1(?c)*) (line 8). This type assertion is added to F_u^ϕ in order to substitute $r(a, b)$ and γ_d in F_v^o (line 9). On the other hand, if b is not a datatype variable (line 10), there are two possibilities: (1) r is a datatype property but b is not a variable, or (2) r is an object property. In both cases, we directly add $r(a, b)$ to F_u^ϕ (line 11). If b

[1] http://www.w3.org/TR/2008/WD-owl2-syntax-20081008/#Data_Ranges

Fig. 6. The Ontology created for q_B in Example 2

has constraints defined in C_v^o, we extract these constraints and add them to F_u^ϕ if they are not already added (lines 12-15).

In order to test subsumption between q_A and q_B in Example 2, we should transform the bodies of these queries and update the ontology they are based on. For this purpose, we use constraint transformation twice. That is, we first update the ontology by adding the concept $AgeConst1$ to handle *hasAge(?c,?y)* $\land$ *?y* $\geq$ *10* $\land$ *?y* $\leq$ *20* and transform q_B to *q():- Doctor(?x)* $\land$ *marriedTo(?x,Jane)* $\land$ *hasChild(?x,?c)* $\land$ *AgeConst1(?c)*. Then, we add concept $AgeConst2$ to the ontology to handle *hasAge(?c,?y)* $\land$ *?y* $\geq$ *12* $\land$ *?y* $\leq$ *16* and transform q_A to *q():- Person(?p)* $\land$ *hasChild(?p,?c)* $\land$ *AgeConst2(?c)*. After this preprocessing step, we use query freezing to test $q_B \sqsubseteq q_A$; the ontology with a canonical ABox created during query freezing is shown in Figure 6.

With these techniques in place, we are now in a position to address the issue of policy analysis supported by OWL-POLAR. It is descried in the following sections.

4.2 Idle Policies

A policy is *idle* if it is never activated or the policy's expiration condition is satisfied whenever the policy is activated. This condition is formally described in Definition 2. If a policy is idle, it cannot be used to regulate any action, because either it never activates or whenever it activates an obligation, permission, or prohibition about an action, the activated policy expires. While designing policies, we may take domain knowledge into account to avoid idle policies.

Definition 2. A policy $\alpha \longrightarrow N_{\chi:\rho}(a:\varphi)/e$ is an idle policy if it does not activate for any state of the world Δ_o or there is a substitution σ' such that $\Delta_o \vdash (e \cdot \sigma) \cdot \sigma'$, whenever there is a substitution σ such that $\Delta_o \vdash (\alpha \land \rho) \cdot \sigma$. ∎

Let us demonstrate idle policies with a simple example. Assume that object property *hasParent* is an inverse property of *hasChild*. Also, let us assume in the domain ontology, we have a SWRL rule such as $hasSponsor(?c, true) \leftarrow hasParent(?c, ?p) \land hasAge(?c, ?age) \land ?age < 18$, which means that children under 18 have a sponsor if they have a parent. Now, consider the policy in Table 2. This policy is activated when a person $?p$ has a child $?c$, which is a student under 18. The activated policy expires when

Table 2. A simple idle policy example

α	$hasChild(?p, ?c) \wedge Student(?c) \wedge hasAge(?c, ?age) \wedge ?age < 18$
N	O
$\chi : \rho$	$?p : Person(?p)$
$a : \varphi$	$?a : PayTuitionsOfStudent(?a) \wedge about(?a, ?c) \wedge hasActor(?a, ?p)$
e	$hasSponsor(?c, true)$

$?c$ has a sponsor. Interestingly, whenever the policy is activated, the domain knowledge implies that $?c$ has a sponsor. That is, whenever the policy is activated, it expires.

In order to detect idle policies, we reason about the activation and expiration conditions of policies. Specifically, a policy $\alpha \longrightarrow N_{\chi:\rho} (a : \varphi) / e$ is an idle policy if $(\alpha \wedge \rho)$ is unrealistic or implies e using the knowledge in the domain ontology. More formally, we can show that the policy is idle if we show $(\alpha \wedge \rho)$ never holds or $(\alpha \wedge \rho) \rightarrow e$. This can be achieved as follows. First, we freeze $(\alpha \wedge \rho)$ and create a canonical ontology o'. If the resulting o' is not a consistent ontology, then we can conclude that the policy is an idle policy, because $(\alpha \wedge \rho)$ never holds. Let o' be consistent and σ be a substitution denoting the mapping of variables in $(\alpha \wedge \rho)$ to the fresh individuals in o'. If there exists a substitution σ' such that $o' \vdash (e \cdot \sigma) \cdot \sigma'$, we conclude that $(\alpha \wedge \rho) \rightarrow e$. We can test $o' \vdash (e \cdot \sigma) \cdot \sigma'$ by querying o' with $q() : - (e \cdot \sigma)$.

4.3   Anticipating Conflicts between Policies

In many settings, policies may conflict. In the simplest case, one policy may prohibit an action while another requires it. There are, however, many less obvious interactions between policies that may lead to logical conflicts [9,16,12,5]. Further developing our earlier example, consider the policy presented in Table 3 that states that a doctor cannot leave a room with patients if he is in charge of the room. This policy conflicts with the policy in Table 1 under some specific conditions. For example, in the scenario described Figure 1, *room 245* of Central Hospital has a fire risk and *Dr. John* is in charge of the room, in which there are some patients. In this setting, the policy in Table 1 obligates Dr. John to leave the room while the policy in Table 3 prohibits this action until the room has no patient.

Table 3. A doctor cannot leave a room containing patients if he is in charge of the room

α	$Room(?r) \wedge hasPatient(?r, true) \wedge inChargeOf(?d, ?r)$
N	F
$\chi : \rho$	$?d : Doctor(?d)$
$a : \varphi$	$?x : LeavingAction(?x) \wedge about(?x, ?r) \wedge hasActor(?x, ?d)$
e	$hasPatient(?r, false)$

If we can determine possible logical conflicts while designing policies, we can create better policies that are less likely to raise conflicts at run time. Furthermore, we can use various conflict resolution strategies such as setting a priority ordering between the policies to solve conflicts [11,21,22], once we determine that two policies may conflict.

In this section, we propose techniques to anticipate possible conflicts between policies at design time. Suppose we have two non-idle policies $P_i = \alpha^i \longrightarrow A_{\chi^i:\rho^i} \left(a^i : \varphi^i\right) / e^i$ and $P_j = \alpha^j \longrightarrow B_{\chi^j:\rho^j} \left(a^j : \varphi^j\right) / e^j$. These policies are active for the same policy addressee in the same state of the world Δ if the following requirements are satisfied:

(i) $\Delta \vdash \left(\alpha^i \wedge \rho^i\right) \cdot \sigma_i$, but no σ_i' such that $\Delta \vdash \left(e^i \cdot \sigma_i\right) \cdot \sigma_i'$

(ii) $\Delta \vdash \left(\alpha^j \wedge \rho^j\right) \cdot \sigma_j$, but no σ_j' such that $\Delta \vdash \left(e^j \cdot \sigma_j\right) \cdot \sigma_j'$

(iii) $\chi^i \cdot \sigma_i = \chi^j \cdot \sigma_j$

The policies P_i and P_j conflict if the following requirements are also satisfied:

(iv) $(\varphi^i \cdot \sigma_i) \sqsubseteq (\varphi^j \cdot \sigma_j)$ or $(\varphi^j \cdot \sigma_j) \sqsubseteq (\varphi^i \cdot \sigma_i)$

(v) A conflicts with B. That is, $A \in \{P, O\}$ while $B \in \{F\}$ or vice versa.

We can use Algorithm 2 to test if it is possible to have such a state of the world where P_i conflicts with P_j. The first step of the algorithm is to test if A conflicts with B (line 2). If they are conflicting, we continue with testing the other requirements. We create a canonical state of the world Δ in which P_i is active by freezing $\left(\alpha^i \wedge \rho^i\right)$ with a substitution σ_i mapping the variables in $\left(\alpha^i \wedge \rho^i\right)$ to the fresh individuals in Δ. Given that $(\varphi^j \cdot \sigma) \sqsubseteq \varphi^j$ for any substitution σ mapping variables into individuals, the requirement (iv) implies that $(\varphi^i \cdot \sigma_i) \sqsubseteq \varphi^j$. We test this as follows. First, we create a canonical ontology o' by freezing $(\varphi^i \cdot \sigma_i)$ (line 4) and then query o' with φ^j (line 5). Each answer to this query defines a substitution σ_k mapping variables in φ^j into the terms in $(\varphi^i \cdot \sigma_i)$, so that $(\varphi^i \cdot \sigma_i) \sqsubseteq (\varphi^j \cdot \sigma_k)$. If φ^j does not have any variable but it repeats in o' as a pattern, the result set contains only one empty substitution. If the query fails, the result set is an empty set ($\emptyset$), which means that it is not possible to have a σ_k such that $(\varphi^i \cdot \sigma_i) \sqsubseteq (\varphi^j \cdot \sigma_k)$. For each σ_k satisfying $(\varphi^i \cdot \sigma_i) \sqsubseteq (\varphi^j \cdot \sigma_k)$, we test

---

Algorithm 2. An algorithm to anticipate if P_i may conflict with P_j

```
 1: Inputs:    Policy P_i = α^i ⟶ A_{χ^i:ρ^i} (a^i : φ^i) /e^i,
                Policy P_j = α^j ⟶ B_{χ^j:ρ^j} (a^j : φ^j) /e^j
 2: if ( (A ∈ {O, P} and B ∈ {F}) or (A ∈ {F} and B ∈ {O, P}) ) then
 3:     ⟨Δ, σ_i⟩ = freeze(α^i ∧ ρ^i)
 4:     ⟨o', _⟩  = freeze(φ^i · σ_i)
 5:     rs       = query(o', φ^j)
 6:     for all (σ_k ∈ rs) do
 7:         ⟨Δ, σ_j⟩ = update(Δ, (α^j ∧ ρ^j) · σ_k)
 8:         if (isConsistent(Δ)) then
 9:             if (query(Δ, e^i · σ_i) = ∅ and query(Δ, (e^j · σ_k) · σ_j) = ∅) then
10:                 return true
11:             end if
12:         end if
13:     end for
14: end if
15: return false
```

the other requirements as follows. First, we update Δ by freezing $(\alpha^j \wedge \rho^j) \cdot \sigma_k$ without removing any individual from its existing $ABox$ (line 7). Note that as a result of this process, σ_j is the substitution mapping the variables in $(\alpha^j \wedge \rho^j) \cdot \sigma_k$ to the new fresh individuals in the updated Δ, so that $\chi^i \cdot \sigma_i = (\chi^j \cdot \sigma_k) \cdot \sigma_j$. We test the consistency of the resulting state of the world Δ (line 8). If this is not consistent, we can conclude that it is not possible to have a state of the world satisfying the requirements. If the resulting Δ is consistent, we check the expiration conditions of the policies. If both are active in the resulting state of the world (line 9), the algorithm returns *true* (line 10). If any of these requirements do not hold, the algorithm returns *false* (line 15).

As described above, the algorithm transforms the problem of anticipating conflict between two policies into an ontology consistency checking problem. To check the consistency of the constructed canonical state of the world Δ, we have used the Pellet [17] reasoner. This reasoner adopts the *open world assumption* and does not have Unique Name Assumption (UNA). Hence, it searches for a model[2] of Δ, also considering the possible overlapping between the individuals (i.e., individuals referring the same object). If there is no model of Δ, it is not possible to have a state of the world satisfying the requirements stated above. We should also note that, while anticipating the conflict, Algorithm 2 tests only the case $(\varphi^i \cdot \sigma_i) \sqsubseteq (\varphi^j \cdot \sigma_j)$. However, we also need to test $(\varphi^j \cdot \sigma_j) \sqsubseteq (\varphi^i \cdot \sigma_i)$ to capture the possibility of conflict. Therefore, if the algorithm returns *false*, we should swap the policies and run the algorithm again. If it returns *true* with the swapped policies, we can conclude that there is a state of the world where these policies may conflict.

To demonstrate the algorithm, let us use the policies presented in Tables 1 and 3 and refer to them as P_i and P_j respectively. In this example, P_j is a prohibition while P_i is an obligation, so the algorithm proceeds as follows (line 2). We create a canonical state of the world Δ by freezing $Person(?x) \wedge Place(?b) \wedge hasFireRisk(?b, true) \wedge in(?x, ?b)$ with a substitution $\sigma_i = \{?x/x, ?b/b\}$ (line 3). Now we create a canonical ontology a' by freezing $\varphi^i \cdot \sigma_i$ with substitution $\{?a/a\}$ (line 4). This ontology has the following $ABox$ assertions: $LeavingAction(a), about(a, b), hasActor(a, x)$. We query o' with $LeavingAction(?x) \wedge about(?x, ?r) \wedge hasActor(?x, ?d)$ (line 5). The result set is composed of only one substitution: $\sigma_k = \{?x/a, ?r/b, ?d/x\}$. The next step is to update Δ by freezing $Doctor(x) \wedge Room(b) \wedge hasPatient(b, true) \wedge inChargeOf(x, b)$ without removing the current $ABox$ of Δ (line 7). The resulting canonical state of the world is shown in Figure 7. Lastly, we check whether both policies remain in effect by checking their expiration conditions (line 9). In this example, we query Δ with $hasFireRisk(b, false)$ and $hasPatient(b, false)$. Both of these queries return $\emptyset$, hence we conclude that there is a state of the world where these policies conflict (line 10).

5   Related Work and Discussion

There have been several policy languages proposed that are built upon Semantic Web technologies. Rei [10] is a policy language based on OWL-Lite and Prolog. It allows

---

[2] A model of an ontology *o* is an interpretation of *o* satisfying all of its axioms [1].

Fig. 7. The canonical state of the world where the policies of Table 1 and Table 3 conflict

logic-like variables to be used while describing policies. This gives it the flexibility to specify relations like *role value maps* that are not directly possible in OWL. The use of these variables, however, makes DL reasoning services (e.g., static conflict detection between policies) unavailable for Rei policies. KAoS [21] is, probably, the most developed language for describing policies that are built upon OWL. KAoS was originally designed to use OWL-DL to define actions and policies. This, however, restricts the expressive power to DL and prevents KAoS from defining policies in which one element of an action's context depends on the value of another part of the current context. For example, KAoS cannot be used to represent a policy like *two soldiers are allowed to communicate only if they are in the same team*. To handle such situations, KAoS has been enhanced with *role-value maps* using Stanford JTP, a general purpose theorem prover [21]. Unfortunately, subsumption reasoning is undecidable in the presence of arbitrary role-value-maps [1].

KAoS distinguishes between (positive and negative) obligation policies and (positive and negative) authorization policies. Authorization policies permit (positive) or forbid (negative) actions, whereas obligation policies require (positive) or do not require (negative) action. Thus the general types of policies that can be described are similar to those that we have discussed in this paper. Actions are also the object of a KAoS policy, and conditions on the application of policies can be described (context), although the subject (individual/role) of the policy is not explicit (it is, however, in Rei). In common with OWL-POLAR in its present form, KAoS does not capture the notion of the authority from which/whom a policy comes, but there is a notion of the priority of a policy which partially (although far from adequately) addresses this issue. Unlike OWL-POLAR, Rei and KAoS do not provide means to explicitly define expiration conditions of the policies.

Policy analysis within both KAoS and Rei is restricted to subsumption. A policy in KAoS is expressed as an OWL-DL class regulating an action, which is expressed as an OWL-DL class expression (e.g., using restrictions on properties such as *performedBy* and *hasDestination*). Two policies are regarded in conflict if their actions overlap (one subsumes another) while the modality of these policies conflict (e.g., negative vs. positive authorization). Similarly, if there exist two policies within Rei that overlap with

respect to the agent and action concerned and they are obligued and prohibited, then a conflict is recognised. In such a situation, meta-policies are used to resolve the conflict. Policy conflicts can also be detected within the Ponder2 framework [18,23], where analysis is far more sophisticated than that developed for either KAoS or Rei, but analysis is restricted to design time. In general, different methods can be used to resolve conflicts between policies. This issue has been explored in detail elsewhere [11].

The expressiveness of OWL-POLAR is not restricted to DL. Using semantic conjunctive formulas, it allows variables to be used while defining policies. However, in semantic formulas, OWL-POLAR allows only object-type variables to be compared using *owl:sameAs* and *owl:differentFrom* properties. On the other hand, data-type variables can be used to define constraints on the datatype properties. In other words, semantic formulas are restricted to describe states of the world, each of which can be represented as an OWL-DL ontology. Therefore, when a semantic formula is frozen, the result is a canonical OWL-DL ontology. OWL-POLAR converts problems of reasoning with and about policies into query answering and ontology consistency checking problems. Then, it uses an off-the-shelf reasoner (Pellet) to solve these problems. It is known that consistency checking in OWL-DL is decidable [17], and query answering in OWL-DL has also been shown to be decidable under DL-safety restrictions [8].

Ontology languages like KAoS are built on OWL 1.0, which does not support dataranges. Therefore, while defining policies, they either do not allow complex constraints to be defined on datatype properties or use non-standard representations for these constraints, which prevents them from using the off-the-shelf reasoning technologies. The clear distinctions between OWL-POLAR and KAoS, however, are manifest in the fact that data ranges are exploited in OWL-POLAR to enable the expression of more complex constraints on policies, and the sophistication of the reasoning mechansims described in this paper.

To the best of our knowledge, OWL-POLAR is the first policy framework that formally defines and detects idle policies. Existing approaches like KAoS and Rei analyse policies only to detect some type of conflict, considering only subsumption between policies. On the other hand, OWL-POLAR provides advanced policy analysis support that is not limited to subsumption checking. Consider the following policies: (i) *Dogs are prohibited from entering to a restaurant*, and (ii) *A member of CSI team is permitted to enter a crime scene*. There is no subsumption relationship between these policies, and so KAoS and Rei could not detect a conflict. However, OWL-POLAR anticipates a conflict by composing a state of the world where these policies are in conflict, e.g., the crime scene is a restaurant and there is a dog in the CSI team.

Building upon this research, we plan to explore various extensions to OWL-POLAR. We will explore extending the representation of policies to include deadlines and penalties associated with their violation, along the lines of [3]. Another issue we would like to investigate concerns how policing mechanisms [14] could make use of our representation and associated mechanisms to foster welfare in societies of self-interested components/agents. We plan to enhance our representation so as to allow constraints over arbitrary terms (and not just $?x \triangleleft \beta$, β being a constant), possibly using constraint satisfaction mechanisms to deal with these. Two further extensions should address policies over many actions (as in, for instance, "ξ is obliged to perform φ_1 and φ_2") and

disjunctions (as in, for instance, "ξ is obliged to perform φ_1 or φ_2"). Finally, we are exploring the use of OWL-POLAR in support of human decision-making, including joint planning activities in hybrid human-software agent teams.

6   Conclusions

Policies provide useful abstractions to constrain and control the behaviour of components in loosely coupled distributed systems. Policies, also called norms, help designers of large-scale, open, and heterogenous distributed systems (including multi-agent systems) to specify, in a concise fashion, acceptable (or policy-compliant) global and individual computational behaviours, thus providing guarantees for the system as a whole.

   In this paper, we have presented a semantically-rich representation for policies as well as efficient mechanisms to reason with/about them. OWL-POLAR meets all the essential requirements of policies, as well as achieving an effective balance between expressiveness (realistic policies can be adequately represented) and computational complexity of associated reasoning for decision-making and analysis (reasoning with and about policies operate in feasible time).

References

1. Baader, F., Calvanese, D., McGuinness, D.L., Nardi, D., Patel-Schneider, P.F. (eds.): The Description Logic Handbook: Theory, Implementation and Applications. Cambridge University Press, Cambridge (2003)
2. Beautement, A., Pym, D.: Structured systems economics for security management. In: Proceedings of the Ninth Workshop on the Economics of Information Security, Harvard, USA (June 2010)
3. Boella, G., Broersen, J., Torre, L.: Reasoning about constitutive norms, counts-as conditionals, institutions, deadlines and violations. In: Bui, T.D., Ho, T.V., Ha, Q.T. (eds.) PRIMA 2008. LNCS (LNAI), vol. 5357, pp. 86–97. Springer, Heidelberg (2008)
4. Castelfranchi, C.: Modelling social action for AI agents. Artificial Intelligence 103, 157–182 (1998)
5. Elhag, A., Breuker, J., Brouwer, P.: On the formal analysis of normative conflicts. Information & Communications Technology Law 9(3), 207–217 (2000)
6. García-Camino, A., Rodríguez-Aguilar, J.A., Sierra, C., Vasconcelos, W.: Constraint rule-based programming of norms for electronic institutions. Autonomous Agents and Multi-Agent Systems 18(1), 186–217 (2009)
7. W. O. W. Group. OWL 2 web ontology language: Document overview (October 2009), `http://www.w3.org/TR/owl2-overview`
8. Haase, P., Motik, B.: A mapping system for the integration of owl-dl ontologies. In: Proceedings of the First International Workshop on Interoperability of Heterogeneous Information Systems, IHIS 2005, pp. 9–16. ACM, New York (2005)
9. Hill, H.: A functional taxonomy of normative conflict. Law and Philosophy 6(2), 227–247 (1987)
10. Kagal, L., Finin, T., Joshi, A.: A policy language for a pervasive computing environment. In: Proceedings of the 4th IEEE International Workshop on Policies for Distributed Systems and Networks, POLICY 2003, pp. 63–74 (2003)

11. Kollingbaum, M.J., Norman, T.J.: Norm adoption and consistency in the NoA agent archi-
    tecture. In: Dastani, M.M., Dix, J., El Fallah-Seghrouchni, A. (eds.) PROMAS 2003. LNCS
    (LNAI), vol. 3067, pp. 169–186. Springer, Heidelberg (2004)
12. Lupu, E., Sloman, M.: Conflicts in policy-based distributed systems management. IEEE
    Transactions on Software Engineering 25(6), 852–869 (1999)
13. Motik, B.: Reasoning in Description Logics using Resolution and Deductive Databases. PhD
    thesis, Universitt Karlsruhe (TH), Karlsruhe, Germany (January 2006)
14. Patel, J., et al.: Agent-based virtual organisations for the grid. Int. Journal of Multi-Agent
    and Grid Systems 1(4), 237–249 (2005)
15. Prud'hommeaux, E., Seaborne, A.: SPARQL Query Language for RDF. Technical report,
    W3C (2006), `http://www.w3.org/TR/rdf-sparql-query/`
16. Sartor, G.: Normative conflicts in legal reasoning. Artificial Intelligence and Law 1(2), 209–
    235 (1992)
17. Sirin, E., Parsia, B., Grau, B.C., Kalyanpur, A., Katz, Y.: Pellet: A practical OWL-DL rea-
    soner. Web Semant. 5(2), 51–53 (2007)
18. Sloman, M., Lupu, E.: Policy specification for programmable networks. In: Covaci, S. (ed.)
    IWAN 1999. LNCS, vol. 1653, pp. 73–84. Springer, Heidelberg (1999)
19. Smith, M.K., Welty, C., McGuinness, D.L.: OWL: Web ontology language guide (February
    2004), `http://www.w3.org/TR/owl-guide`
20. Ullman, J.D.: Information integration using logical views. Theoretical Computer Sci-
    ence 239(2), 189–210 (2000)
21. Uszok, A., Bradshaw, J.M., Lott, J., Breedy, M., Bunch, L., Feltovich, P., Johnson, M., Jung,
    H.: New developments in ontology-based policy management: Increasing the practicality and
    comprehensiveness of KAoS. In: Proceedings of the 2008 IEEE Workshop on Policies for
    Distributed Systems and Networks, POLICY 2008, pp. 145–152 (2008)
22. Vasconcelos, W.W., Kollingbaum, M.J., Norman, T.J.: Normative conflict resolution in multi-
    agent systems. Autonomous Agents and Multi-Agent Systems 19(2), 124–152 (2009)
23. Zhao, H., Lobo, J., Bellovin, S.M.: An algebra for integration and analysis of ponder2 poli-
    cies. In: Proceedings of the 2008 IEEE Workshop on Policies for Distributed Systems and
    Networks, POLICY 2008, Washington, DC, USA, pp. 74–77. IEEE Computer Society, Los
    Alamitos (2008)