

**The Soundness and Completeness
of Margrave with Respect to a Subset of
XACML**

Michael Matthew Greenberg, Casey Marks,
Leo Alexander Meyerovich, and
Michael Carl Tschantz

Department of Computer Science
Brown University
Providence, Rhode Island 02912

CS-05-05

April 2005

The Soundness and Completeness of Margrave with Respect to a Subset of XACML

Michael Matthew Greenberg Casey Marks Leo Alexander Meyerovich
Michael Carl Tschantz

February 13, 2005

Abstract

Provides a natural semantics for a subset of XACML, a language for the specification of access-control policies. Presents the key algorithm of Margrave, a tool for analyzing XACML policies. Proves the soundness and completeness of the algorithm with respect to the subset of XACML.

1 Introduction

Multi-user and web applications necessitate controlling the actions a user may perform on a resource. Generally, some policy governs what each user is allowed to access. Applying *ad hoc* checks spread throughout the code to enforce such policies makes the policy hard to understand or modify. Thus, such informal policies should be formalized into a central access-control policy, which is enforced throughout the program. To allow for these policies to be shared across applications and for APIs to aid in enforcement, standardized languages are arising for the specification of access-control policies.

XACML is a language for the specification of such access-control policies [2]. XACML is an OASIS standard with backing from Sun Microsystems and others. In this paper we are concerned with the subset of XACML that is sufficient for role-based access control. In Section 3, we provide a natural semantics for this subset of XACML.

After specifying a policy in a language like XACML, the policy administrator should verify that the specification matches the intent of the policy. The policy verification tool Margrave aids in this task [1]. Margrave analyzes XACML policies belonging to the subset of XACML modeled in Section 3. Margrave uses an internal representation of an XACML policy called a multi-terminal binary decision diagram (MTBDD). In Section 4, we formalize the algorithm that produces a MTBDD from an XACML file. In Section 5, we prove the soundness and completeness of the results of this algorithm with respect to the subset of XACML presented in Section 3.

2 Syntax

For readability we will not use an XML style syntax for our subset of XACML, but rather a scheme-like syntax.

XACML has two syntaxes: one for specifications of policies and one for requests. Since the former has a key word “Policy”, we will call it the spec syntax. The latter is called the request syntax.

2.1 XACML Spec Syntax

The start nonterminal is S .

```

S ::= C
C ::= Ps | POL
Ps ::= (PolicySet CA T C*)
CA ::= First-Applicable | Deny-Overrides | Permit-Overrides
T ::= ( (SUB) (RES) (ACT) )
SUB ::= Any | ALLOW+
RES ::= Any | ALLOW+
ACT ::= Any | ALLOW+
ALLOW ::= (AVC+)
AVC ::= (ID VAL)
POL ::= (Policy CA T R*)
R ::= (Rule T EFFECT)
EFFECT ::= Permit | Deny

```

The elements of the syntax category R are called “Rules”; Ps, “PolicySets”; POL, “Policies”; and T, “targets”. The elements of syntax categories of SUB, RES, and ACT are called the “subtargets”. The elements of syntax categories of Ps, POL, and R are called the “laws”.

2.2 XACML request syntax

```

Q ::= ((S) (R) (A))
S ::= AVP*
R ::= AVP*
A ::= AVP*
AVP ::= ( ID VAL )

```

2.3 Parsing

To make the semantics more understandable, we assume the existence of a parser. Informally, we assume that is parser can take a string generated by the grammar and produce denotational object suggested by the strings form, mostly lists. Since the above syntax matches the list notation in Scheme, this parser is equivalent to the Scheme command quote.

3 Semantics

Let \mathcal{P} be the set of all policies (members of the syntactic category C of the language Spec). \mathcal{P} includes both XACML Policies and PolicySets.¹ Let \mathcal{Q} be the set of all requests (members of the syntactic category Q). Let \mathcal{D} be the set of all decisions ($\mathcal{D} = \{\text{Permit, Deny, NotApplicable}\}$). We will use the symbol used for the nonterminal in syntax in the natural semantics to refer to any element of the syntax category.

At the core of the semantics of XACML is the notation of a request matching the target of a rule, policy, or policy set. We will denote this relation by $q \in t$ where $q \in \mathcal{Q}$ and t is a target. We will define \in using a natural semantics given below in Table 1.

The semantics of the Margrave subset of XACML is given in a natural semantics that makes use of the relation \in . The results of evaluating a Rule is given in Table 2 in terms of the \models_r relationship. The \models_r relation is then used to define the \models relation, which gives the semantics of our subset of XACML. The default behavior of \models is given in Table 3. The behavior of \models under each of the combining algorithms are given in Table 4 for permit-overrides, in Table 5 for deny-overrides, and in Table 6 for first-applicable.

4 Margrave

Margrave is a tool for the analysis of XACML policies. The key function in Margrave takes an XACML policy and produces a function from requests to decisions. We model this function as MAR:

$$\text{MAR} : \mathcal{P} \rightarrow (\mathcal{Q} \rightarrow \mathcal{D}).$$

Below we define MAR as pseudo-Scheme program that consumes a parsed XACML file and produces a data structure called a MTBDD that defines the returned function from requests to decisions. First we describe MTBDDs and then we give a definition of MAR.

4.1 MTBDDs

Let f be a function from $\{\text{T, F}\}^n$ to a finite set of terminals M with the formal arguments called \vec{x} . The function f can be represented as a binary tree called a binary decision tree. Each leaf of the tree is labeled with an element of M . Each internal node represents a question of the form “is the i th element of ordered n -tuple \vec{x} true?”. Given \vec{x} , one determines $f(\vec{x})$ by starting at the root node and if the answer to the question at the current node is “yes”, recur on the right child (the high branch); otherwise, recur on left child (the low branch). The base case is reached at a leaf. The returned value is the label at the reached leaf.

The same procedure succeeds when the tree requirement is lifted and directed acyclic graphs (DAGs) are allowed. Such a DAG is called a multi-terminal binary decision diagram (MTBDD). Given an MTBDD, one can easily convert it to be an MTBDD in which at most one leaf node is labeled by

¹The word “policy” when uncapitalized will refer to access-control policies in general. When “policy” is capitalized, it will refer to the XACML tag `Policy` and the associated structure.

$$\frac{S \in_S \text{SUB} \quad R \in_R \text{RES} \quad A \in_A \text{ACT}}{(S, R, A) \in (\text{SUB}, \text{RES}, \text{ACT})} \quad (1)$$

$$X \in \text{Any} \quad \forall X \in \{S, R, A\} \quad (2)$$

$$\frac{\exists i \text{ s.t. } S \in \text{ALLOW}_i}{S \in_S (\text{ALLOW}_1, \text{ALLOW}_2, \dots, \text{ALLOW}_n)} \quad (3)$$

$$\frac{\exists i \text{ s.t. } R \in \text{ALLOW}_i}{R \in_R (\text{ALLOW}_1, \text{ALLOW}_2, \dots, \text{ALLOW}_n)} \quad (4)$$

$$\frac{\exists i \text{ s.t. } A \in \text{ALLOW}_i}{A \in_A (\text{ALLOW}_1, \text{ALLOW}_2, \dots, \text{ALLOW}_n)} \quad (5)$$

$$\frac{\forall i \quad X \in_X \text{AVC}_i}{X \in_X (\text{AVC}_1, \text{AVC}_2, \dots, \text{AVC}_n)} \text{ where } X \in \{S, R, A\} \quad (6)$$

$$\frac{\exists j \text{ s.t. } \text{AVP}_j = \text{AVC}}{(\text{AVP}_1, \text{AVP}_2, \dots, \text{AVP}_n) \in_X \text{AVC}} \text{ where } X \in \{S, R, A\} \quad (7)$$

Table 1: The Match Relationship

$$\frac{Q \notin T}{((\text{Rule } T \text{ EFFECT}), Q) \models_r \text{NotApplicable}} \quad (8)$$

$$\frac{Q \in T}{((\text{Rule } T \text{ Permit}), Q) \models_r \text{Permit}} \quad (9)$$

$$\frac{Q \in T}{((\text{Rule } T \text{ Deny}), Q) \models_r \text{Deny}} \quad (10)$$

Table 2: The Rule Relationship \models_r

$$\langle (\text{Policy } \text{CA } T), Q \rangle \models \text{NotApplicable} \quad (11)$$

$$\langle (\text{PolicySet } \text{CA } T), Q \rangle \models \text{NotApplicable} \quad (12)$$

$$\frac{Q \notin T}{\langle (\text{Policy } \text{CA } T \text{ R}^*), Q \rangle \models \text{NotApplicable}} \quad (13)$$

$$\frac{Q \notin T}{\langle (\text{PolicySet } \text{CA } T \text{ C}^*), Q \rangle \models \text{NotApplicable}} \quad (14)$$

Table 3: Default NotApplicable Judgements

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models_r \text{Permit}}{\langle (\text{Policy } \text{Permit-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Permit}} \quad (15)$$

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models \text{Permit}}{\langle (\text{PolicySet } \text{Permit-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Permit}} \quad (16)$$

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models_r \text{Deny} \quad \forall j \neq i, \langle C_j, Q \rangle \not\models_r \text{Permit}}{\langle (\text{Policy } \text{Permit-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Deny}} \quad (17)$$

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models \text{Deny} \quad \forall j \neq i, \langle C_j, Q \rangle \not\models \text{Permit}}{\langle (\text{PolicySet } \text{Permit-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Deny}} \quad (18)$$

$$\frac{Q \in T \quad \forall i, \langle C_i, Q \rangle \models_r \text{NotApplicable}}{\langle (\text{Policy } \text{Permit-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{NotApplicable}} \quad (19)$$

$$\frac{Q \in T \quad \forall i, \langle C_i, Q \rangle \models \text{NotApplicable}}{\langle (\text{PolicySet } \text{Permit-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{NotApplicable}} \quad (20)$$

Table 4: Permit-Overrides Judgements

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models_r \text{Deny}}{\langle (\text{Policy Deny-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Deny}} \quad (21)$$

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models \text{Deny}}{\langle (\text{PolicySet Deny-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Deny}} \quad (22)$$

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models_r \text{Permit} \quad \forall j \neq i, \langle C_j, Q \rangle \not\models_r \text{Deny}}{\langle (\text{Policy Deny-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Permit}} \quad (23)$$

$$\frac{Q \in T \quad \exists i \text{ s.t. } \langle C_i, Q \rangle \models \text{Permit} \quad \forall j \neq i, \langle C_j, Q \rangle \not\models \text{Deny}}{\langle (\text{PolicySet Deny-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Permit}} \quad (24)$$

$$\frac{Q \in T \quad \forall i, \langle C_i, Q \rangle \models_r \text{NotApplicable}}{\langle (\text{Policy Deny-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{NotApplicable}} \quad (25)$$

$$\frac{Q \in T \quad \forall i, \langle C_i, Q \rangle \models \text{NotApplicable}}{\langle (\text{PolicySet Deny-Overrides } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{NotApplicable}} \quad (26)$$

Table 5: Deny-Overrides Judgements

$$\frac{Q \in T \quad \langle R_1, Q \rangle \models \text{Permit}}{\langle (\text{Policy First-Applicable } T \ R_1, R_2, \dots, R_n), Q \rangle \models \text{Permit}} \quad (27)$$

$$\frac{Q \in T \quad \langle C_1, Q \rangle \models \text{Permit}}{\langle (\text{PolicySet First-Applicable } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Permit}} \quad (28)$$

$$\frac{Q \in T \quad \langle R_1, Q \rangle \models \text{Deny}}{\langle (\text{Policy First-Applicable } T \ R_1, R_2, \dots, R_n), Q \rangle \models \text{Deny}} \quad (29)$$

$$\frac{Q \in T \quad \langle C_1, Q \rangle \models \text{Deny}}{\langle (\text{PolicySet First-Applicable } T \ C_1, C_2, \dots, C_n), Q \rangle \models \text{Deny}} \quad (30)$$

$$\frac{\langle R_1, Q \rangle \models \text{NotApplicable} \quad \langle (\text{Policy First-Applicable } T \ R_2, \dots, R_n), Q \rangle \models D}{\langle (\text{Policy First-Applicable } T \ R_1, R_2, \dots, R_n), Q \rangle \models D} \quad \forall D \in \mathcal{D} \quad (31)$$

$$\frac{\langle C_1, Q \rangle \models \text{NotApplicable} \quad \langle (\text{PolicySet First-Applicable } T \ C_2, \dots, C_n), Q \rangle \models D}{\langle (\text{PolicySet First-Applicable } T \ C_1, C_2, \dots, C_n), Q \rangle \models D} \quad \forall D \in \mathcal{D} \quad (32)$$

Table 6: First-Applicable Judgements




Figure 1: An example MTBDD. The low branches are shown as dashed lines and the high branches as solid lines.

each terminal: for each $m \in M$, make a new node n labeled with m ; for each edge going from an internal node n' to a leaf labeled with m , make an edge going from n' to n ; remove all the leaves labeled with m except n and all the edges going to them. Henceforth, we assume that all MTBDDs discussed have this property. If \hat{f} is an MTBDD representation of f , we will let $\hat{f}(\vec{x}) = f(\vec{x})$.

A Binary Decision Diagram (BDD) is an MTBDD in which the set of terminals is $\{\mathsf{T}, \mathsf{F}\}$. A BDD may be viewed as a set where all elements that are mapped to T are in the set and all mapped to F are not in the set. Each element of the set formed from a BDD is itself a set of answers to the questions in the BDD. We let \in_{bdd} represent: an element is “in” the BDD. That is, if \hat{b} is a BDD representing a set B with a characteristic function of b , then

$$x \in_{\text{bdd}} \hat{b} \iff \hat{b}(x) \iff b(x) \iff x \in B \quad (33)$$

Given a policy, Margrave constructs a MTBDD that represents that policy. The terminals of this MTBDD are elements of the set $\{\text{Permit}, \text{Deny}, \text{NotApplicable}\}$. During this process Margrave uses BDDs. For both the three-terminal MTBDDs and the BDDs. \vec{x} is a representation of a request.

The algorithm makes use of the following constants: *deny-term*, *permit-term*, *na-term*, *true-term*, and *false-term* are the leaf node labeled the Deny, Permit, NotApplicable, T , and F , respectively.

The algorithm makes use of a few basic operations on MTBDDs given below²:

(define (*make-bdd* *subtarget-name* *attribute-id* *attribute-value*)

²Although many of the following algorithms are given a Scheme like notation, we intend for these algorithms to be given in the meta-language and not an additional object language. Thus, no distinction is made between variables in the Scheme notation and meta-variables.

creates a BDD with one internal node which corresponds to the question “Does the *attribute-id* in the *subtarget-name* have the value *attribute-value*?” and has the terminal T on its high bench and the terminal F on its low branch. Formally, $(make\text{-}bdd\ n\ i\ v)$ returns a BDD that represents a function f from \mathcal{Q} to $\{\text{T}, \text{F}\}$ such given a request $q = (S, R, A)$

$$f = \lambda(S, R, A) \in \mathcal{Q}[(i, v) \in X] \quad (34)$$

where $X = S$ if $n = \text{“Subject”}$, $X = R$ if $n = \text{“Resource”}$, $X = A$ if $n = \text{“Action”}$.

(bdd-and l-bdd r-bdd)

returns the BDD that represents the function $f = \lambda\vec{x}[f_l(\vec{x}) \wedge f_r(\vec{x})]$ where f_l is the function represented by *l-bdd*, and f_r by *r-bdd*. This function is extended to a list of BDDs as

```
(define (bdd-and* bdd-list)
  (if (empty? bdd-list)
      true-term
      (bdd-and (first bdd-list) (bdd-and* (rest bdd-list)))))
```

(bdd-or l-bdd r-bdd)

returns the BDD that represents the function $f = \lambda\vec{x}[f_l(\vec{x}) \vee f_r(\vec{x})]$ where f_l is the function represented by *l-bdd*, and f_r by *r-bdd*. This function is extended to a list of BDDs as

```
(define (bdd-or* bdd-list)
  (if (empty? bdd-list)
      false-term
      (bdd-or (first bdd-list) (bdd-or* (rest bdd-list)))))
```

(replace-terminal terminal replacing-mtbdd in-mtbdd)

returns a BDD that is identical to *in-mtbdd* except that the terminal *terminal* is replaced by the whole MTBDD *replacing-mtbdd*. (After the replacement all the leaf nodes labeled with the same value are combined as above.)

4.2 The Algorithm of Margrave

Rather than give the transition rules and axioms as in the normal semantics, we give pseudo-Scheme code so as to follow the actual implementation discussed in section 4.3 more closely.

The algorithm consists of two major parts: creating a BDD representations of targets and combining these to represent Rules, Policies, and PolicySets.

4.2.1 Representing Targets

The target of a policy may be viewed as a representation of the set of requests to which this policy applies. The Margrave algorithm creates these sets using BDDs. The algorithm *process-target* takes a (parsed) target and converts it to such a BDD.

The algorithm for creating and representing these sets as BDDs is given below. Note that the functions starting with *ast:* simply access the named child in the abstract syntax tree. For example, given a element of the syntax category of the nonterminal T called *target*, (*ast:target-subjects target*) returns the subject section of the target. If the abstract syntax tree is viewed as scheme list as explained in Section 2.3, this would be (*first target*).

```
(define (process-target target)
  (bdd-and (process-subtarget 'Subject (ast:target-subjects target))
           (process-subtarget 'Resource (ast:target-resources target))
           (process-subtarget 'Action (ast:target-actions target))))
```

```
(define (process-subtarget subtarget-name subtarget-body)
  (if (= (first subtarget-body) 'Any)
      true-term
      (bdd-or* (map process-allowance subtarget-body))))
```

```
(define (process-allowance subtarget-name match-list)
  (bdd-and* process-avc match-list))
```

```
(define (process-avc subtarget-name match)
  (make-bdd subtarget-name
            (ast:attribute-id match)
            (ast:attribute-value match)))
```

4.2.2 Forming Rules, Policies, and PolicySets

Each law can be viewed a function from requests to a decision (Permit, Deny, or NotApplicable). The function *ast->mtbdd* takes a (parsed) law and returns just such a function.

```
(define (ast->mtbdd law)
  (cond [(ast:rule? law)
         (augment-rule (process-target (ast:law-target law))
                       (cond [(eq? (ast:rule-effect law) 'Deny) deny-term]
                             [(eq? (ast:rule-effect law) 'Permit) permit-term]))]
        [(or (ast:policy? law) (ast:policySet? law))
         (let [ [ca (ast:get-ca law)]
                [children (map ast->mtbdd (ast:get-children law))] ]
           (augment-rule (process-target (ast:law-target law))
                         (cond [(equal? ca "first-applicable")
                                (build-first-applicable children)]
                               [(equal? ca "deny-overrides")
                                (build-deny-override children)]
                               [(equal? ca "permit-overrides")
                                (build-permit-override children)]))))))
```

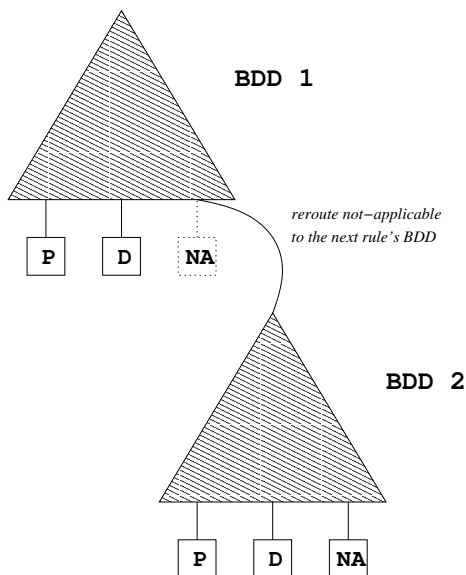


Figure 2: First-applicable.

```
(define (augment-rule target-bdd effect-term)
  (replace-terminal true-term effect-term
    (replace-terminal false-term na-term target-bdd)))
```

```
(define (build-first-applicable children)
  (if (empty? children)
      na-term
      (replace-terminal na-term (build-first-applicable (rest children))
        (first children))))
```

Deny-overrides can be seen in Figure 2.

```
(define (build-permit-overrides children)
  (if (empty? children)
      na-term
      (replace-terminal deny-term
        (replace-terminal na-term deny-term (build-permit-overrides (rest children)))
        (replace-terminal na-term (build-permit-overrides (rest children))
          (first children))))))
```

Permit-overrides can be seen in Figure 3.

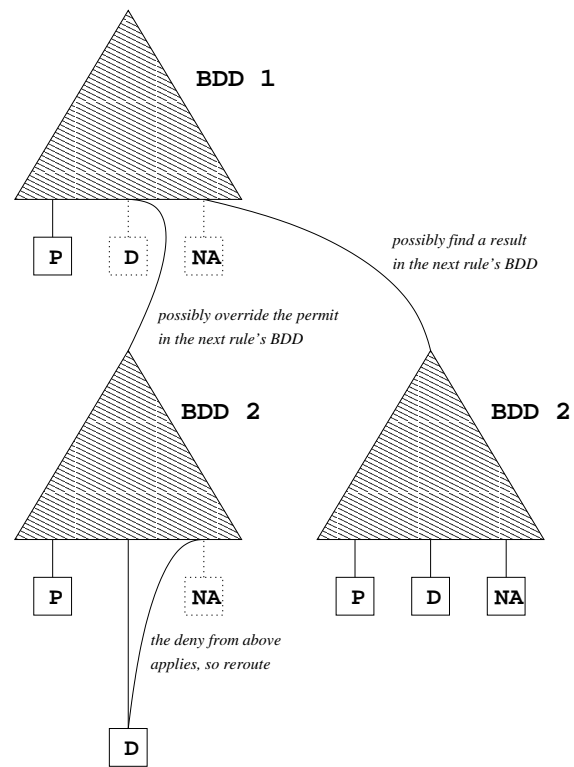


Figure 3: Permit-override.

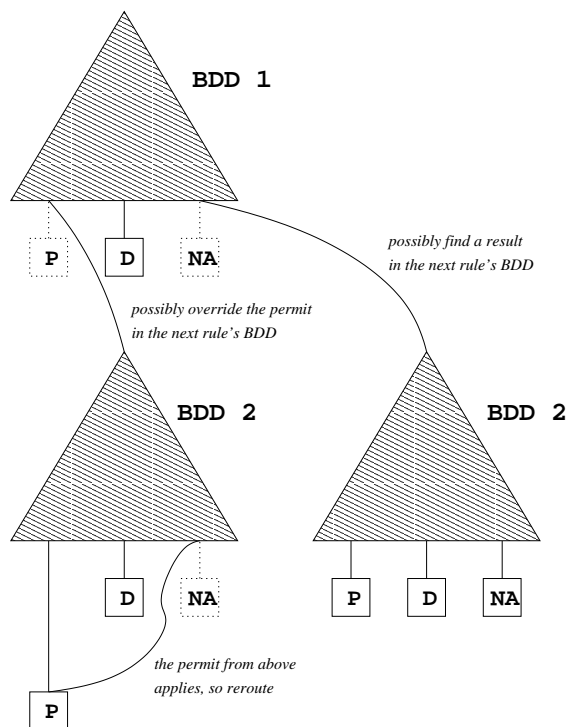


Figure 4: Deny-override.

```

(define (build-deny-overrides children)
  (if (empty? children)
      na-term
      (replace-terminal permit-term
        (replace-terminal na-term permit-term (build-deny-overrides (rest children)))
        (replace-terminal na-term (build-deny-overrides (rest children))
          (first children))))))

```

Deny-overrides can be seen in Figure 4.

The function MAR is (`lambda (policy) (ast->mtbdd (parse policy))`).

4.3 An Implementation of Margrave

An implementation of Margrave exists and is available at www.cs.brown.edu/research/plt/software/margrave. The implementation differs from the above algorithm in that the implementation uses reduced ordered multi-terminal BDDs instead of just MTBDDs [1]. Since the relationship between ROMTBDDs and general MTBDDs are well understood and the algorithms to maintain reduced ordering complicates the proof considerably, the following proofs are not for ROMTBDDs.

5 Soundness and Completeness of Margrave Analysis

Margrave is equivalent to XACML evaluation in the following sense

$$\forall p \in \mathcal{P}, \forall q \in \mathcal{Q}, [\langle p, q \rangle \models d \iff (\text{MAR}(p))(q) = d]. \quad (35)$$

The proof of this equivalence follows from the soundness and completeness of Margrave with respect to XACML evaluation as shown in the next two sections. Below we will use the symbol that represents a nonterminal in the syntax to refer to the set of all strings in that syntax category.

5.1 The Match Relationship

Lemma 1.

$$\forall (i, v) \in AVC, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{bdd} (\textit{process-avc } 'Subject (i, v)) \iff s \in_S (i, v)] \quad (36)$$

$$\forall (i, v) \in AVC, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{bdd} (\textit{process-avc } 'Resource (i, v)) \iff r \in_R (i, v)] \quad (37)$$

$$\forall (i, v) \in AVC, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{bdd} (\textit{process-avc } 'Action (i, v)) \iff a \in_A (i, v)] \quad (38)$$

Proof. For the subject case, $(\textit{process-avc } 'Subject (i, v))$ produces a BDD with one internal node. The internal node represents the question “Does the request include in the Subject subrequest the id-value pair (i, v) ?”. The high branch goes to **T** and the low branch to **F**. Thus, $(s, r, a) \in_{bdd} (\textit{process-ac } 'Subject (i, v))$ if and only if (s, r, a) includes attribute-value pair (i, v) in the Subject subtarget. Formally, $(s, r, a) \in_{bdd} (\textit{process-ac } 'Subject (i, v))$ if and only if $\exists (i', v') \in S$ s.t. $(i', v') = (i, v)$. Such a request (s, r, a) will make $s \in_S (i, v)$ true since this is the same requirement to for Judgement 7.

For the other direction for the subject case, if $q \in (i, v)$, then $\exists (i', v') \in s$ s.t. $(i', v') = (i, v)$. This will allow the high branch of the single node in the BDD created to be taken yielding **T**. Thus, $q \in (i, v)$ implies $(s, r, a) \in_{bdd} (\textit{process-avc } 'Subject (i, v))$.

The resource and action cases follow the same reasoning as the subject case. \square

Lemma 2.

$$\forall \ell \in Allow, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{bdd} (\textit{process-allowance } 'Subject \ell) \iff q \in \ell] \quad (39)$$

$$\forall \ell \in Allow, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{bdd} (\textit{process-allowance } 'Resource \ell) \iff q \in \ell] \quad (40)$$

$$\forall \ell \in Allow, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{bdd} (\textit{process-allowance } 'Action \ell) \iff q \in \ell] \quad (41)$$

Proof. For subject case: Let $(i_j, v_j) \in AC$ be the j th element of $\ell \in Allow$ where j ranges from 1 to $|\ell|$. Let b be $(\textit{process-allowance } 'Subject \ell)$. $\textit{process-allowance}$ produces b as the $bdd\text{-and}$ of $(\textit{process-avc } 'Subject (i_j, v_j))$ for all j . Thus, $(s, r, a) \in_{bdd} b$ if and only if

$$(s, r, a) \in_{bdd} (\textit{process-avc } 'Subject (i_j, v_j)) \quad (42)$$

for all j . Thus,

$$(s, r, a) \in_{\text{bdd}} (\text{process-allowance } \text{'Subject } \ell) \iff \forall j, (s, r, a) \in_{\text{bdd}} (\text{process-avc } \text{'Subject } (i_j, v_j)) \quad (43)$$

$$\iff \forall j, (s, r, a) \in_S (i_j, v_j) \quad (44)$$

$$\iff (s, r, a) \in_S \ell. \quad (45)$$

where middle step follows from Lemma 1 and the last from Judgement 6.

The resource and action cases follows the same steps. \square

Lemma 3.

$$\forall s' \in \text{Sub}, \forall (s, r, a) \in \mathcal{Q}, \quad [(s, r, a) \in_{\text{bdd}} (\text{process-subtarget } \text{'Subject } s') \iff s \in s'] \quad (46)$$

$$\forall r' \in \text{Res}, \forall (s, r, a) \in \mathcal{Q}, \quad [(s, r, a) \in_{\text{bdd}} (\text{process-subtarget } \text{'Resource } r') \iff r \in r'] \quad (47)$$

$$\forall a' \in \text{Act}, \forall (s, r, a) \in \mathcal{Q}, \quad [(s, r, a) \in_{\text{bdd}} (\text{process-subtarget } \text{'Action } a') \iff a \in a'] \quad (48)$$

Proof. For the subject case: There are two cases: (1) s' is either a list of elements from *Allow* or (2) s' is the singleton list holding the keyword “Any”.

1. $s' \in \text{Allow}^+$: Let $\ell_j \in \text{Allow}$ be the j th element of $s' \in \text{Sub}$ where j ranges from 1 to $|s'|$. (*process-subtarget* 'Subject s') produces the *bdd-or* of (*process-allowance* 'Subject ℓ_j) over all j . Thus, $(s, r, a) \in_{\text{bdd}} (\text{process-subtarget } \text{'Subject } s')$ if and only if $(s, r, a) \in_{\text{bdd}} (\text{process-allowance } \text{'Subject } \ell_j)$ for at least one j . Thus,

$$(s, r, a) \in_{\text{bdd}} (\text{process-subtarget } \text{'Subject } s') \iff \exists j, (s, r, a) \in_{\text{bdd}} (\text{process-allowance } \text{'Subject } \ell_j) \quad (49)$$

$$\iff \exists j, (s, r, a) \in_S \ell_j \quad (50)$$

$$\iff (s, r, a) \in_S s'. \quad (51)$$

where middle step follows from Lemma 2 and the last from Judgement 3.

2. $s' = (\text{Any})$: In this case, *process-subtarget* will return *true-term*. Thus, for any request (s, r, a) , $(s, r, a) \in_{\text{bdd}} (\text{process-subtarget } \text{'Subject } s')$. Equivalently, by Judgement 2, all requests $(s, r, a) \in \mathcal{Q}$, $s \in (\text{Any})$.

The resource and action cases follows the same steps. \square

Lemma 4.

$$\forall (s', r', a') \in \text{Target}, \forall (s, r, a) \in \mathcal{Q}, [(s, r, a) \in_{\text{bdd}} (\text{process-target } (s', r', a')) \iff (s, r, a) \in (s', r', a')] \quad (52)$$

Proof. (*process-target* (s', r', a')) constructs the *bdd-and* of

- (*process-subtarget* 'Subject s'),

- (*process-subtarget* 'Resource r'), and
- (*process-subtarget* 'Action a').

Thus,

$$(s, r, a) \in_{\text{bdd}} (\text{process-target } (s', r', a')) \quad (53)$$

$$\iff (\text{process-subtarget 'Subject } s') \wedge (\text{process-subtarget 'Resource } r') \wedge (\text{process-subtarget 'Action } a') \quad (54)$$

$$\iff (s, r, a) \in_S s' \wedge (s, r, a) \in_R r' \wedge (s, r, a) \in_A a' \quad (55)$$

$$\iff (s, r, a) \in (s', r', a'). \quad (56)$$

where middle step follows from Lemma 3 and the last from Judgement 1. \square

5.2 Rules

Lemma 5.

$$\forall r \in \text{Rule}, \forall q \in \mathcal{Q}, [(\text{MAR}(r))(q) = d \iff \langle p, q \rangle \models_r d]. \quad (57)$$

Proof. The relative branch of *ast->add* is as follows:

$$\begin{aligned} & (\text{augment-rule } (\text{process-target } (\text{ast:law-target } \text{law}))) \\ & \quad (\mathbf{cond} [(eq? (\text{ast:rule-effect } \text{law}) \text{'Deny}) \text{deny-term}] \\ & \quad \quad [(eq? (\text{ast:rule-effect } \text{law}) \text{'Permit}) \text{permit-term}]]) \end{aligned}$$

The function *augment-rule* first replaces the false terminal with *na-term*, which means that any request that yields false in BDD (was not \in_{bdd} the BDD) will now yield **NotApplicable**. Second, *augment-rule* replaces the true terminal of the BDD produced by (*process-target* (*ast:law-target* *law*)) with the effect of the rule (either a *deny-term* or *permit-term* as selected the effected list in *law*).

Consider two cases: (1) where the request does not match the target and (2) where the request matches the target:

1. From Lemma 4, we know that the only requests that will yield F from BDD produced by (*process-target* (*ast:law-target* *law*)) are those that are not in the match relationship with the target. Thus, since the terminal for F was replaced by the terminal for **NotApplicable** and none of the internal nodes were altered, the MTBDD that is produced by *augment-rule* will yield **NotApplicable** if and only if the request was not in the match relationship \in with the target. This is the same behavior of the rule relationship \models_r found in Table 2 (note esp. Judgement 8).
2. Now consider the case where the match relationship does hold (as in Judgements 9 and 10). Consider the sub-case where the effect of the rule is **Permit**. Since the terminal for T was replaced by the terminal for **Permit** and none of the internal nodes were altered, the MTBDD that is produced by *augment-rule* will yield **Permit** if and only if the request is in the match

relationship \in with the target. Thus, the same behavior as found in Table 2 will be observed (note esp. Judgement 9). The proof of equivalence for the deny case follows in the same way.

Thus, in both cases MAR and \models_r behaves the same making the above equivalence true. \square

5.3 Policy

Lemma 6.

$$\forall p \in Pol, \forall q \in \mathcal{Q}, [(MAR(p))(q) = d \implies \langle p, q \rangle \models d]. \quad (58)$$

Proof. The case where the request does not match the target of the policy is almost identical to the corresponding case in the proof of Lemma 5. For the case where the request does match the target, we must proof that result returned by rule-combining function (either *build-first-applicable*, *build-permit-overrides*, or *build-deny-overrides*) is correct. We use proof by induction over the number of rules in p .

Base Case: Zero rules. The functions *build-first-applicable*, *build-permit-overrides*, and *build-deny-overrides* all return *na-term* for the rule-less policy. $\forall p \in \mathcal{P}, \forall q \in \mathcal{Q}, \langle p, q \rangle \models \text{NotApplicable}$ holds when p is rule-less but not for either $\langle p, q \rangle \models \text{Permit}$ or $\langle p, q \rangle \models \text{Deny}$.

Inductive Case: Assume for $n \geq 0$ rules that $\forall p \in Pol, \forall q \in \mathcal{Q}, [q \in MAR(p, d) \implies \langle p, q \rangle \models d]$. Let p be in Pol and have $n + 1$ rules.

1. First-Applicable: Since the number of child rules is greater than zero, the applicable branch in *build-first-applicable* is

$$(\text{replace-terminal na-term } (\text{build-first-applicable } (\text{rest children})) \\ (\text{first children})).$$

Since (*first children*) is just a rule, it is correct by Lemma 5. Since the number rules in (*rest children*) is n , by the inductive hypothesis, (*build-first-applicable* (*rest children*)) is correct. These other rules only matter if the first rule yields **NotApplicable** for given request. Otherwise, the result of the first rule is the result for the given request. The results of *replace-terminal* ensures this behavior. If the first rule yields **NotApplicable**, then the other rules are consulted since the terminal for **NotApplicable** gets replaced by the correct MTBDD that results from the other rules. The other two terminals are unchanged and thus correctly returned.

2. Permit-Overrides: Since the number of child rules is greater than zero, the applicable branch in *build-permit-overrides* is

$$(\text{replace-terminal deny-term} \\ (\text{replace-terminal na-term deny-term } (\text{build-permit-overrides } (\text{rest children}))) \\ (\text{replace-terminal na-term } (\text{build-permit-overrides } (\text{rest children})) \\ (\text{first children}))))).$$

Lets decompose this as follows:

$$\hat{b} = (\textit{first children}) \quad (59)$$

$$\hat{c} = (\textit{build-permit-overrides (rest children)}) \quad (60)$$

$$\hat{d} = (\textit{replace-terminal na-term } \hat{c} \hat{b}) \quad (61)$$

$$\hat{e} = (\textit{replace-terminal na-term deny-term } \hat{c}) \quad (62)$$

$$\hat{f} = (\textit{replace-terminal deny-term } \hat{e} \hat{d}) \quad (63)$$

First we will prove soundness and then completeness:

Soundness:

- (a) $\hat{f}(q) = \text{Permit}$: For $\hat{f}(q)$ to equal Permit, either $\hat{d}(q) = \text{Deny}$ and $\hat{e}(q) = \text{Permit}$, or $\hat{d}(q) = \text{Permit}$.

If $\hat{e}(q) = \text{Permit}$, then $\hat{c}(q) = \text{Permit}$. Since \hat{c} is produced by n rules, by the inductive hypothesis, $\hat{c}(q) = \text{Permit}$ is sound. That is, if p' is a policy created out the last n rules of the given policy p , then $\langle p', q \rangle \models \text{Permit}$. Thus, since the only applicable judgement that implies that Permit is yielded is Judgement 15, the antecedents of Judgement 15 must hold. Thus, there exists a rule in (*rest children*) that yields Permit. This makes Judgement 15 applicable for policy p as well.

If $\hat{d}(q) = \text{Permit}$, then either $\hat{b}(q) = \text{Permit}$, or $\hat{b}(q) = \text{NotApplicable}$ and $\hat{c}(q) = \text{Permit}$. If $\hat{b}(q) = \text{Permit}$, then by Lemma 5, rule $n + 1$ must yield Permit making Judgement 15 applicable. If $\hat{c}(q) = \text{Permit}$, then as above, Judgement 15 holds.

Thus, in either case, Judgement 15 holds, and thus, $\langle p, q \rangle \models \text{Permit}$.

- (b) $\hat{f}(q) = \text{Deny}$: For $\hat{f}(q)$ to equal Deny, then both $\hat{e}(q)$ and $\hat{d}(q)$ must be Deny since \hat{f} is \hat{d} with its Deny terminal replaced by \hat{e} .

When $\hat{e}(q) = \text{Deny}$, $\hat{c}(q)$ must equal either Deny or NotApplicable since \hat{e} is \hat{c} with its NotApplicable terminal changed to be a Deny terminal. Since \hat{c} is created by n rules, by the inductive hypothesis, \hat{c} must be sound. Thus, since both applicable judgements that yield Deny or NotApplicable (Judgements 17 and 19) require that no rule in the Policy yields Permit, no rule in (*rest children*) may yield Permit.

When $\hat{d}(q) = \text{Deny}$, either $\hat{b}(q) = \text{Deny}$, or $\hat{b}(q) = \text{NotApplicable}$ and $\hat{c}(q) = \text{Deny}$. If $\hat{b}(q) = \text{Deny}$, then by Lemma 5, rule (*first children*) must yield Deny. As above, if $\hat{c}(q) = \text{Deny}$, there must exist a rule that yields deny in (*rest children*) (see Judgement 17). Thus, under either case, $\hat{b}(q) \neq \text{Permit}$ and there exists a rule the yields Deny.

From the implications of $\hat{e}(q) = \text{Deny}$, no rule in (*rest children*) may not yield Permit. From implications of $\hat{d}(q) = \text{Deny}$, the rule (*first children*) may not yield permit and there must exist a rule in the policy that must yields Deny. Thus, Judgement 17 applies and $\langle p, q \rangle \models \text{Deny}$.

- (c) $\hat{f}(q) = \text{NotApplicable}$: For $\hat{f}(q)$ to equal NotApplicable, either $\hat{d}(q) = \text{NotApplicable}$, or $\hat{d}(q) = \text{Deny}$ and $\hat{e}(q) = \text{NotApplicable}$. However, $\hat{e}(q)$ cannot equal NotApplicable since \hat{e} has no NotApplicable terminal (it was replaced by a Deny terminal). Thus, $\hat{d}(q)$ must be NotApplicable. For $\hat{d}(q)$ to be NotApplicable, both $\hat{b}(q)$ and $\hat{c}(q)$ must be NotApplicable.

When $\hat{b}(q) = \text{NotApplicable}$, by Lemma 5, (*first children*) must yield **NotApplicable**.

When $\hat{c}(q) = \text{NotApplicable}$, by the inductive hypothesis, a policy with all the rules in (*rest children*) and on others, must yield **NotApplicable**. Since the only applicable judgement that yields **NotApplicable** is Judgement 19, the antecedents of that judgement must hold. Thus, all of the rules in (*rest children*) must yield **NotApplicable**.

Thus, every rule in p must yield **NotApplicable** and Judgement 19 is applicable. Thus, $\langle p, q \rangle \models \text{NotApplicable}$.

Completeness:

- (a) $\langle p, q \rangle \models \text{Permit}$: Only one judgement can be applied to yield **Permit**: Judgement 15. According to the antecedent of Judgement 15, there must exist a rule in the policy that yields **Permit**. Thus, either rule $n + 1$ yields **Permit** or at least on rule in (*rest children*) yields **Permit**.

In the first case where rule $n + 1$ yields **Permit**, by Lemma 5 \hat{b} will yield **Permit**. Since the creation of \hat{f} from \hat{b} does not change the **Permit** terminal or any of the internal nodes, \hat{f} will yield **Permit** for the same requests as \hat{b} . Thus, \hat{f} yields **Permit**.

In the second case where rule $n + 1$ does not yield **Permit** but another rule in (*rest children*) does, by the inductive hypothesis the MTBDD \hat{c} will yield **Permit**. Either rule $n + 1$ yields **NotApplicable** or **Deny**. If it yields **NotApplicable**, then since in \hat{f} it has its **NotApplicable** terminal replaced by \hat{c} , it will now produce what \hat{c} produces: **Permit**. If it yields **Deny**, then since in \hat{f} it has its **Deny** terminal replaced by \hat{c} , it will yield what this MTBDD \hat{c} yields. Since \hat{c} is just \hat{c} with its **NotApplicable** terminal replaced by **Deny** and \hat{c} yields **Permit**, \hat{c} will yield **Permit**. Thus, \hat{c} will yield **Permit**.

- (b) $\langle p, q \rangle \models \text{Deny}$: Only one judgement could can be applied to yield **deny**: Judgement 17. According to the antecedent of Judgement 17, there must exist a rule in the policy that yields **Deny** and no rules that yield **Permit** on q . Thus, either rule $n + 1$ yields **Deny** or at least on rule in (*rest children*) yields **Deny**.

In the first case where rule $n + 1$ yields **Deny**, by Lemma 5 \hat{b} yields **Deny**. The creation of \hat{f} from \hat{b} changes the **Deny** terminal to be \hat{c} . Thus, \hat{f} yields what \hat{c} yields. Since no rules may yield **Permit**, by the inductive hypothesis, \hat{c} does not yield **Permit**. Thus, \hat{c} either yields **Deny** or **NotApplicable**. \hat{e} is \hat{c} with its **NotApplicable** terminal replaced with a **Deny** terminal. Thus, \hat{e} always yields **Deny**. Thus, \hat{f} yields **Deny** as required.

In the second case where rule $n + 1$ does not yield **Deny** but another rule in (*rest children*) does, by the inductive hypothesis MTBDD \hat{c} will yield **Deny**. Since rule $n + 1$ may yield neither **Permit** nor **Deny**, it must yield **NotApplicable**. Since \hat{f} is \hat{b} with its **NotApplicable** terminal changed to \hat{c} , \hat{f} will yield what \hat{c} yields. Thus, \hat{f} will yield **Deny**.

Thus, in either case, \hat{f} under q will yield **Deny** if $\langle p, q \rangle \models \text{Deny}$.

- (c) $\langle p, q \rangle \models \text{NotApplicable}$: Given that the target matches, only one judgement can be applied to yield **deny**: Judgement 19. According to the antecedent of Judgement 19, every rule with in the policy must yields **NotApplicable**. Thus, by Lemma 5, $\hat{b}(q) = \text{NotApplicable}$. By the inductive hypothesis, $\hat{c}(q) = \text{NotApplicable}$. Since \hat{f} is \hat{b} with its **NotApplicable** terminal replaced by \hat{c} , $\hat{f}(q) = \hat{c}(q) = \text{NotApplicable}$.

Thus, in every case, $\langle p, q \rangle \models d \implies f(q) = d$

3. Deny-Overrides: Almost the same as Permit-Overrides.

□

5.4 PolicySets

Theorem 7 (Equivalence).

$$\forall p \in \mathcal{P}, \forall q \in \mathcal{Q}, [(\text{MAR}(p))(q) = d \iff \langle p, q \rangle \models d]. \quad (64)$$

Proof. Let the *height* of a PolicySet be defined as 1 for any PolicySet that has only Policies as children (no PolicySets) and as one more than the maximum height of any of children for PolicySets that do have other PolicySets as children. Theorem will be proved by induction over the height of the PolicySet.

Base Case: height 1, that is, a PolicySet with all its children being Policies. Given a PolicySet of height 1, it can have any number of children. By Lemma 6, we know that each of these children are both sound and complete. Now we must show that Margrave combines them correctly to yield a MTBDD that is sound and complete. Margrave uses the same function to combine Policies into PolicySets as it does to combine Rules into Policies. Inspecting the judgements for PolicySets, one can see that every rule for a PolicySet has a corresponding rule for Policy (listed right after it). The only differences in each pair of corresponding judgements are the syntactic category of children and the use of \models instead of \models_r . The proof above for Lemma 6 only makes use of these two facts in using Lemma 5; it does not make use of any internal structure of rules or \models_r , only that Margrave is sound and complete with respect to rule and the relation \models_r . Lemma 6 shows that Margrave is sound and complete with respect to Policies and the relation \models when used with Policies. Thus, Lemma 6 can be substituted for Lemma 5 everywhere in the proof of Lemma 6 to prove the base case of this theorem.

Inductive Case: Assume the inductive hypothesis that for all PolicySets of height n or less, $\forall p \in \mathcal{P}, \forall q \in \mathcal{Q}, [(\text{MAR}(p))(q) = d \iff \langle p, q \rangle \models d]$. Note that all the children of a PolicySet of height $n + 1$ has a height of n or less. Thus, the inductive hypothesis proves Margrave to be equivalent on all the children. Thus, as Lemma 6 stood in place of Lemma 5 in proof of the base case, the inductive hypothesis can stand in place of Lemma 5 in proof of the inductive case.

Thus, it has been proven that for a PolicySet of any finite height. Since only PolicySets of finite height are permissible by the grammar, the theorem is proven. □

6 Conclusion

This paper has presented a subset of XACML, provided a natural semantics for that subset, formalized an algorithm evaluating policies written in this subset, and proven the soundness and completeness of this algorithm.

The reader should bear in mind that full XACML is a much more rich set of operators. For example, full XACML results in not simply a decision being produced by each policy-request pair, but also a

list of obligations, actions that must be fulfilled upon the rendering of the decision (*e.g.*, notifying an auditor). How these lists of obligations are produced from the policy introduces non-determinism into the language. Future work shall have to provide a semantics for full XACML.

The algorithm presented is incapable of running on full XACML. Furthermore, due to the limitations of MTBDD construction, no way to extend the algorithm to full XACML presents itself. Future work should produce a tool capable of the analysis of XACML policies written in full XACML.

Acknowledgements

This paper was produced in the fulfillment of the requirements of the course CS 296 section 01 in Department of Computer Science, Brown University. The authors wish to thank the instructor Shriram Krishnamurthi for offering this course.

References

- [1] FISLER, K., KRISHNAMURTHI, S., MEYEROVICH, L. A., AND TSCHANTZ, M. C. Verification and change impact analysis of access-control policies. In *27th International Conference on Software Engineering ICSE '05* (St. Louis, Missouri, May 2005). To appear.
- [2] MOSES, T. eXtensible Access Control Markup Language (XACML) version 1.0. Tech. rep., OASIS, Feb. 2003.