

The Role of Abduction in  
Declarative Authorization Policies

Moritz Y. Becker  
*Microsoft Research, Cambridge, UK*  
moritzb@microsoft.com

Sebastian Nanz  
*Informatics and Mathematical Modelling*  
*Technical University of Denmark*  
nanz@imm.dtu.dk

August 2007

Technical Report

Microsoft Research  
Roger Needham Building  
7 J.J. Thomson Avenue  
Cambridge, CB3 0FB  
United Kingdom

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Adding Abduction to Policy Evaluation</b>	<b>3</b>
2.1	An Extensible Scheme for Policy Evaluation . . . . .	3
2.2	Deductive Policy Evaluation . . . . .	6
2.3	Constructing Proof Graphs . . . . .	7
2.4	Abductive Policy Evaluation . . . . .	7
<b>3</b>	<b>Termination Conditions</b>	<b>9</b>
3.1	Subsumption Weakening . . . . .	10
3.2	Static Termination Analysis . . . . .	10
<b>4</b>	<b>Application Scenarios</b>	<b>12</b>
4.1	Explaining Access Denials to Users . . . . .	12
4.2	Administration of Authorization Policies . . . . .	12
4.3	Automated Delegation . . . . .	13
<b>5</b>	<b>Discussion</b>	<b>14</b>
<b>A</b>	<b>Proofs</b>	<b>18</b>

# The Role of Abduction in Declarative Authorization Policies

Moritz Y. Becker

*Microsoft Research, Cambridge, UK*

`moritzb@microsoft.com`

Sebastian Nanz

*Informatics and Mathematical Modelling*

*Technical University of Denmark*

`nanz@imm.dtu.dk`

August 2007

## Abstract

Declarative authorization languages promise to simplify the administration of access control systems by allowing the authorization policy to be factored out of the implementation of the resource guard. However, writing a correct policy is an error-prone task by itself, and little attention has been given to tools and techniques facilitating the analysis of complex policies, especially in the context of access denials. We propose the use of abduction for policy analysis, for explaining access denials and for automated delegation. We show how a deductive policy evaluation algorithm can be conservatively extended to perform abduction on Datalog-based authorization policies, and present soundness, completeness and termination results.

**Keywords:** access control, abduction, authorization language, Datalog

## 1 Introduction

Authorization is the task of granting or denying access to a system's resources according to a policy. Traditionally, authorization policies have been implemented by access control lists (ACL) or capabilities provided by the operating system, sometimes augmented by groups or roles. However, there are many applications for which these mechanisms are too inflexible, not sufficiently expressive and provide the wrong level of abstraction. For example, access to electronic health records is regulated by a huge number of laws that are both complex and prone to change. Decentralized applications such as grid systems require support for delegation of authority and attribute-based constraints.

Such requirements have led to the development of trust management systems and declarative authorization languages for flexible, expressive application-level access control (e.g. [15, 10, 14, 6, 3]). An authorization policy is then written as a set of rules that are both human readable and machine enforceable. This

approach aims to increase the usability and scalability of access control systems: policies written in such languages are more concise and have a lower viscosity than ACLs, and provide a much higher level of abstraction, thus facilitating a closer representation of the intended policy.

However, comprehending and predicting the consequences of a policy is difficult, as policies can be complex and contain hundreds of rules, and each access grant is based on the construction of a logical proof of compliance with respect to this policy. Thus, writing a correct policy is still a highly error-prone task. We conjecture that the current lack of tools for analyzing policies remains a major obstacle to a wider adoption of authorization languages.

In this paper, we develop algorithms for analyzing the consequences of declarative authorization policies. Many existing authorization languages are based on negation-free Datalog or are translated into Datalog for evaluating access requests. (A Datalog clause is a first-order definite Horn clause without function symbols.) Datalog is sufficiently expressive for a wide range of policies, including delegation, which requires recursion. Furthermore, Datalog is decidable and can be evaluated efficiently. Hence, to maximize generality, our algorithms work on policies specified in Datalog.

In particular, we focus on the tasks of *explaining* access grants and access denials. In the former case, the question we are trying to answer is “*why is a given request granted?*”. It is easy to see that the proof graph contains exactly the necessary information for constructing the (possibly textual) explanation. The basic evaluation algorithm that is used for deciding access requests can be easily extended to construct a copy of the proof graph during evaluation.

In the case of access denial, the question is “*which authorization facts or credentials were missing that would have led to an access grant?*”. This turns out to be a harder question, as the failed partial proof does not contain enough information to answer it. Moreover, there are in general infinitely many different answers, but often only finitely many “interesting” or “useful” ones. We propose to apply *abductive* techniques for finding the set of meaningful answers. Abduction [16] is a reasoning paradigm that has been used for planning, fault diagnosis and other areas in AI, but has not previously been considered for analyzing authorization policies. Many algorithms have been developed for various variants of abduction (see [11] for an extensive survey); in this paper, we show that a deductive evaluation algorithm that is used for deciding access requests can be conservatively extended to perform abduction of authorization facts and credentials. Thus we show that existing implementations of Datalog-based authorization engines can be leveraged and extended with little effort to facilitate this kind of analysis. Moreover, we show that this algorithm can be used for multiple purposes: (1) as the basis of a tool helping security administrators to write and to debug policies, (2) for providing users with an answer in the case of an access denial that is more helpful than a mere “no”, and (3) to compute sets of missing credentials in automated distributed delegation scenarios.

The remainder of the paper is structured as follows: Section 2 first presents a non-deterministic terminating algorithm for evaluating policies. The algorithm is extended to construct proof graphs. A second extension is developed that computes sets of missing facts that, if added, would lead to a positive access decision. Section 3 presents techniques for guaranteeing that the algorithm terminates. Section 4 discusses three application scenarios to illustrate the different ways in which the algorithm could be used. In Section 5 we discuss

related work and conclude.

## 2 Adding Abduction to Policy Evaluation

The basic authorization problem consists of deciding whether an access request complies with an authorization policy and a set of credentials. For Datalog-based policies, this amounts to deductive query evaluation. Tabling resolution algorithms are proposed in [5] and [3] for evaluating queries against Datalog-based authorization policies. These algorithms are easy to implement and are guaranteed to terminate due to tabling [17, 8].

Here we provide a generalization of these algorithms, presented as state transition systems. The non-deterministic presentation lends itself to parallel implementations; moreover, it also leads to simpler soundness, completeness and termination proofs. Section 2.2 instantiates this generalized tabling scheme to a purely deductive policy evaluation algorithm. The evaluation algorithms in [5] and [3] can be seen as straightforward deterministic implementations of this general scheme. Section 2.3 illustrates a simple extension of the first one, which facilitates the construction of proof graphs, e.g. for explaining positive access decisions. Finally, Section 2.4 extends it further to perform *abduction*, which computes the *dual* of the basic authorization problem, namely the sets of facts or credentials which, according to the policy, would grant an authorization request. As shown in Section 4, this algorithm can be applied to explain access denials, to analyze policies and to provide automated distributed delegation.

### 2.1 An Extensible Scheme for Policy Evaluation

**Preliminaries.** We assume a denumerable set of variables  $\mathcal{X}$  and a first-order signature with a countable (possibly infinite) set of constants  $\mathcal{C}$  and a finite set of predicate names (but no function symbols). An *atom*  $P$  consists of a predicate name applied to an ordered list of terms, each of which is either a variable or a constant. *Clauses* are of the form  $P_0 \leftarrow \vec{P}$ . The atom  $P_0$  is referred to as the *head*, and the (possibly empty) finite sequence of atoms  $\vec{P}$  as the *body* of the clause. A clause with an empty body is also called a *fact*. A *policy*  $\mathcal{P}$  is a finite set of clauses.

The semantics of a policy  $\mathcal{P}$  is given by the least fixed point of the *immediate consequence operator*  $T_{\mathcal{P}}$  [1]:

$$T_{\mathcal{P}}(I) = \{P_0\theta : (P_0 \leftarrow P_1, \dots, P_n) \in \mathcal{P}, P_i\theta \in I \text{ for each } i, P_0\theta \text{ ground}\}$$

We denote the least fixed point of  $T_{\mathcal{P}}$  by  $T_{\mathcal{P}}^{\omega}(\emptyset)$ . Intuitively, it contains all ground atoms that are deducible from the policy.

The operator  $T_{\mathcal{P}}$  is monotonic and continuous, and its least fixed point  $T_{\mathcal{P}}^{\omega}(\emptyset)$  contains all ground atoms deducible from  $\mathcal{P}$ . In a *bottom-up* approach to query evaluation, we could compute this fixed point once and for all, as it is frequently done in the context of deductive databases. This has the advantage that it is a complete, terminating procedure, and query evaluation is fast once the fixed-point has been constructed. However, bottom-up evaluation is not usually suitable for authorization queries, as the set of credentials at hand is not constant. Computing the model for every request is thus not efficient as

it results in the evaluation of irrelevant goals and a *top-down*, goal-directed approach seems more appropriate.

The most widely-known top-down evaluation algorithm is SLD resolution, as used for Prolog. Unfortunately, SLD resolution can run into infinite loops even for safe Datalog programs if some of the predicates have recursive definitions. The problem remains even if instead of a depth-first a breadth-first search strategy is employed: the looping occurs because the SLD search tree is infinite. *Tabling*, or *memoing*, is an efficient approach for guaranteeing termination by incorporating some bottom-up techniques into a top-down resolution strategy. The basic idea is to prune infinite search trees by keeping a table of encountered subgoals and their answers, and to compute a subgoal only if it is not already in the table [18]. In the following, we present a non-deterministic algorithm for Datalog evaluation based on tabling. We have to fix a number of general terms first.

We also need the following general definitions. Let  $P$  and  $Q$  range over atoms. A *variable renaming* for  $P$  is an injective substitution whose range is a subset of  $\mathcal{X}$ . A *fresh renaming* of  $P$  is a variable renaming for  $P$  such that the variables in its range have not appeared anywhere else. A substitution  $\theta$  is *more general* than  $\theta'$  iff there exists a substitution  $\rho$  such that  $\theta' = \theta\rho$ . A substitution  $\theta$  is a *unifier* of  $P$  and  $Q$  iff  $P\theta = Q\theta$ . A substitution  $\theta$  is a *most general unifier* of  $P$  and  $Q$  iff it is more general than any other unifier of  $P$  and  $Q$ . When  $P$  and  $Q$  are unifiable, they also have a most general unifier that is unique up to variable renaming; if the most general unifier exists, we may denote it by  $\text{mgu}(P, Q)$ . We say that  $P$  is *subsumed by*  $Q$  (also written  $P \preceq Q$ ) iff  $P = Q\theta$  for some substitution  $\theta$ .

In our examples, we write variables in *italics*, constants in **typewriter** font, and predicate names in **sans serif**.

An authorization policy defines authorization-relevant predicates such as `canRead`, `canWrite` etc. Upon an access request, the resource guard issues a query (e.g. “`canRead(Alice, Foo)?`”) to be evaluated against the policy. Usually, the policy is composed of the locally stored policy plus a set of facts obtained from (user-submitted or fetched) *credentials* that may support the request. We will later show examples of policies and their uses.

**Description of the Scheme.** The algorithms described in the following subsections are instantiations of a state transition system that processes nodes of the following form:

**Definition 2.1 (Nodes)** A *node* is either a *root node*  $\langle P \rangle$  where  $P$  is an atom, or a *tuple* (with at least 3 fields) of the form  $\langle P; \vec{Q}; S; \dots \rangle$ , where the atom  $P$  is called the *index*, the (possibly empty) sequence of atoms  $\vec{Q}$  the *subgoals*, and the atom  $S$  the *partial answer*. If the list of subgoals  $\vec{Q}$  is empty, a node is called an *answer node* with *answer*  $S$ . Otherwise it is called a *goal node*, and the first atom in  $\vec{Q}$  is its *current subgoal*.

Intuitively, the list of subgoals  $\vec{Q}$  contains the atoms that still have to be solved for the goal  $P$ . The subgoals are solved from left to right, hence the head of the list is the current subgoal. The current subgoal can be resolved against another answer node, which may entail instantiations of variables which will

narrow down the partial answer  $S$ . The partial answer may eventually become a proper answer if all subgoals have been solved.

Furthermore, the algorithm makes use of two tables:

**Definition 2.2 (Answer and Wait Tables)** An *answer table* is a partial function from atoms to sets of answer nodes. A *wait table* is a partial function from atoms to sets of goal nodes.

We denote the answer and wait tables in the algorithm by  $Ans$  and  $Wait$ , respectively. The set  $Ans(P)$  contains all answer nodes pertaining to the goal  $\langle P \rangle$  found so far. The set  $Wait(P)$  contains all those nodes whose current subgoal is waiting for answers from  $\langle P \rangle$ . Whenever a new answer for  $\langle P \rangle$  is produced, the computation of these waiting nodes is resumed.

$$\begin{array}{l}
(\text{root}) \quad (\{\langle P \rangle\} \uplus \mathcal{N}, Ans, Wait) \rightarrow (\mathcal{N} \cup \mathcal{N}', Ans, Wait) \\
\quad \text{if } \mathcal{N}' = \mathbf{generate}_{\mathcal{P}}(P) \\
\\
(\text{ans}) \quad (\{n\} \uplus \mathcal{N}, Ans, Wait) \rightarrow (\mathcal{N} \cup \mathcal{N}', Ans[P \mapsto Ans(P) \cup \{n\}], Wait) \\
\quad \text{if } n \text{ is an answer node with index } P \\
\quad \quad \nexists n' \in Ans(P) : n \preceq n' \\
\quad \quad \mathcal{N}' = \bigcup_{n'' \in Wait(P)} \mathbf{resolve}(n'', n) \\
\\
(\text{goal}_1) \quad (\{n\} \uplus \mathcal{N}, Ans, Wait) \rightarrow (\mathcal{N} \cup \mathcal{N}', Ans, Wait[Q' \mapsto Wait(Q') \cup \{n\}]) \\
\quad \text{if } n \text{ is a goal node with current subgoal } Q \\
\quad \quad \exists Q' \in dom(Ans) : Q \preceq Q' \\
\quad \quad \mathcal{N}' = \bigcup_{n' \in Ans(Q')} \mathbf{resolve}(n, n') \\
\\
(\text{goal}_2) \quad (\{n\} \uplus \mathcal{N}, Ans, Wait) \rightarrow (\mathcal{N} \cup \{\langle Q \rangle\}, Ans[Q \mapsto \emptyset], Wait[Q \mapsto \{n\}]) \\
\quad \text{if } n \text{ is a goal node with current subgoal } Q \\
\quad \quad \forall Q' \in dom(Ans) : Q \not\preceq Q'
\end{array}$$

Table 1: Generic tabling algorithm

The algorithm is given in Table 1 as a transition system defined by a relation  $\rightarrow$  on states of the following form:

**Definition 2.3 (States)** A *state* is a triple  $(\mathcal{N}, Ans, Wait)$  where  $\mathcal{N}$  is a set of nodes,  $Ans$  is an answer table, and  $Wait$  is a wait table.

A state of the form  $(\{\langle P \rangle\}, \{P \mapsto \emptyset\}, \{P \mapsto \emptyset\})$  is an *initial state*. A state  $\mathcal{S}$  is a *final state* iff there is no state  $\mathcal{S}'$  and such that  $\mathcal{S} \rightarrow \mathcal{S}'$ .

We have left the description of the algorithm generic with respect to the following choices:

1. the structure of tuples (beyond the first three fields)
2. the subsumption relation  $\preceq$  on answer nodes
3. the procedure  $\mathbf{resolve}(n, n')$
4. the procedure  $\mathbf{generate}_{\mathcal{P}}(P)$

Intuitively, if  $n \preceq n'$  ( $n$  is subsumed by  $n'$ ) holds, then the answer node  $n$  provides no more information than  $n'$ ; in the algorithm, we can thus discard  $n$  and potentially ensure that the answer set is kept finite. The procedure **resolve** $(n, n')$  is intended to take a goal node  $n$  and an answer node  $n'$  and combine the current subgoal of  $n$  with the answer provided by  $n'$  to get a new node with a simpler subgoal. The procedure **generate** $_P(P)$  is intended to generate a set of tuples for a given query  $\langle P \rangle$  by resolving  $P$  against the rules of program  $\mathcal{P}$ .

Starting in an initial state, rule (root) generates answer and goal nodes for a query  $\langle P \rangle$ . Answer nodes are processed by (ans) which inserts them into  $Ans(P)$  if they are not subsumed by the answers already present there; likewise they are resolved against all nodes currently waiting for an answer to  $P$ . Goal nodes are either handled by (goal<sub>1</sub>) or (goal<sub>2</sub>), depending on whether the current subgoal is subsumed by an atom in the domain of the answer table. If it is, the already existing answers to that atom can be reused for the current subgoal, and the goal node is added to the wait table in (goal<sub>1</sub>). Otherwise, (goal<sub>2</sub>) spawns a new root node, and initializes the answer and wait tables.

## 2.2 Deductive Policy Evaluation

In its simplest instantiation, the algorithm performs ordinary deduction, i.e. starting in an initial state with root node  $\langle P \rangle$  it will terminate in a state where  $Ans(P)$  represents all instantiations of  $P$  which are deducible from the policy  $\mathcal{P}$ . We obtain this instantiation by defining:

1. Tuples are of the form  $\langle P; \vec{Q}; S \rangle$ .
2.  $\langle P; []; S \rangle \preceq_D \langle P; []; S' \rangle$  iff  $S \preceq S'$ .
3. Let  $n = \langle \cdot; []; Q' \rangle$  be an answer node, and  $Q''$  a fresh renaming of  $Q'$ .  

$$\mathbf{resolve}^D(\langle P; [Q, \vec{Q}]; S, n) = \begin{cases} \{ \langle P; \vec{Q}\theta; S\theta \rangle \} & \text{if } \theta = \text{mgu}(Q, Q'') \text{ exists,} \\ \emptyset & \text{otherwise} \end{cases}$$
4.  $\mathbf{generate}_P^D(P) = \bigcup_{(Q \leftarrow \vec{Q}) \in \mathcal{P}} \mathbf{resolve}^D(\langle P; [Q, \vec{Q}]; Q, \langle P; []; P' \rangle)$   
 where  $P'$  is a fresh renaming of  $P$ .

The subsumption relation  $\preceq_D$  causes all answer nodes to be discarded whose partial solutions are “more instantiated” and therefore less general than already existing answers.

**Example 2.4** The following policy allows the file `Foo` to be read by `Bob` and every employee who is associated with any work group (in particular, `Alice`):

```

canRead(x, Foo) ← isEmployee(x), inWorkgroup(x, y).
canRead(Bob, Foo).
isEmployee(Alice).
inWorkgroup(Alice, WG23).

```

Suppose the algorithm is started in an initial state with query  $\langle \text{canRead}(z, \text{Foo}) \rangle$ . The only possible start transition is (root), thus  $\mathbf{generate}_P^D(\text{canRead}(z, \text{Foo}))$  is called and produces a goal node

$$n_0 \equiv \langle \text{canRead}(z, \text{Foo}); [\text{isEmployee}(x), \text{inWorkgroup}(x, y)]; \text{canRead}(x, \text{Foo}) \rangle$$

and an answer node  $\langle \text{canRead}(z, \text{Foo}); []; \text{canRead}(\text{Bob}, \text{Foo}) \rangle$ . Eventually, the goal node will be resolved against the last two facts in the policy to yield a second answer  $\langle \text{canRead}(z, \text{Foo}); []; \text{canRead}(\text{Alice}, \text{Foo}) \rangle$ . The algorithm terminates with no further answers.

### 2.3 Constructing Proof Graphs

A simple extension of the above instantiation reconstructs the proof graphs for every answer in  $\text{Ans}(P)$ :

1. Tuples are of the form  $\langle P; \vec{Q}; S; \vec{n}; Cl \rangle$ , where  $Cl$  is a clause in  $\mathcal{P}$  and  $\vec{n}$  is a sequence of answer nodes called *child nodes*.
2.  $\langle P; []; S; \cdot; \cdot \rangle \preceq_{\mathbf{G}} \langle P; []; S'; \cdot; \cdot \rangle$  iff  $S \preceq S'$ .
3. Let  $n = \langle \cdot; []; Q'; \cdot; \cdot \rangle$  be an answer node, and  $Q''$  a fresh renaming of  $Q'$ .
 
$$\text{resolve}^{\mathbf{G}}(\langle P; [Q, \vec{Q}]; S; \vec{n}; Cl \rangle, n) = \begin{cases} \{ \langle P; \vec{Q}\theta; S\theta; [\vec{n}, n]; Cl \rangle & \text{if } \theta = \text{mgu}(Q, Q'') \text{ exists,} \\ \emptyset & \text{otherwise} \end{cases}$$
4. Let  $P'$  be a fresh renaming of  $P$ .
 
$$\text{generate}_{\mathcal{P}}^{\mathbf{G}}(P) = \bigcup_{(Q \leftarrow \vec{Q}) \in \mathcal{P}} \text{resolve}^{\mathbf{G}}(\langle P; [Q, \vec{Q}]; Q; []; Cl \rangle, \langle P; []; P'; []; Cl \rangle)$$

When resolving a goal against an answer node, the answer node is inserted as a new child node, as justification for the resolution step. In order to reconstruct the proof graph, an answer node is interpreted to have edges pointing to each of its child nodes and an edge pointing to the rule  $Rl$  which has been used to derive that particular answer. Figure 1 shows a proof graph for the derivation of  $\text{canRead}(\text{Alice}, \text{Foo})$  in Example 2.4.

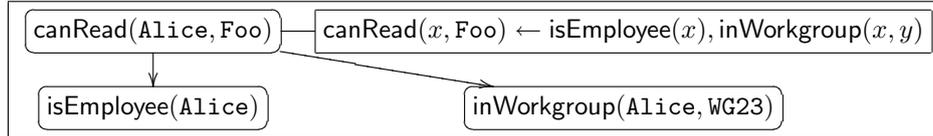


Figure 1: Proof graph related to Example 2.4

Proof graphs are useful for auditing and explaining positive access decisions. If the predicates are associated with meta-information on how they can be translated into natural language, the proof graph could also be represented as a sentence such as “Alice can read Foo because Alice is an employee and Alice is in workgroup WG23”.

### 2.4 Abductive Policy Evaluation

In our setting, the term abduction relates to the following problem. Given an atom  $P$  and a policy  $\mathcal{P}$ , find all sets  $A$  of atoms such that  $P$  is deducible from  $\mathcal{P}$  augmented by  $A$ . The set  $A$  is called an *abductive solution* for query  $P$ , and we require that the predicate names occurring in  $A$  are from a given set of *abducible* predicate names. The choice of the abducibles usually depends on the application domain and the kind of analysis we want to perform. In many cases, we are interested in all possible abductive solutions, so we specify all

predicate names in  $\mathcal{P}$  to be abducible. We define  $\mathcal{A}$  to be the set of all ground instantiations of abducible predicates.

In the context of decentralized authorization, the parameters of the abductive solution may be unknown and may thus have to be left uninstantiated. For example, the solution could specify a delegation chain where the identities of the intermediate delegators cannot be fixed a priori. Therefore, the abductive solutions we are interested in may contain variables that can be arbitrarily instantiated; this is sometimes referred to as *floundering* abduction. This way, each solution can represent an infinite number of ground solutions. We can provide a relatively simple algorithm to solve the floundering abduction problem (compared to e.g. [2, 9]) mainly because our policies are monotonic.

The generic tabling scheme is instantiated as follows:

1. Tuples are of the form  $\langle P; \vec{Q}; S; \Delta \rangle$ , where  $\Delta$  is a set of atoms called the *residue*.
2.  $\langle P; []; S; \Delta \rangle \preceq_A \langle P; []; S'; \Delta' \rangle$  iff  $|\Delta| \geq |\Delta'|$  and there exists a substitution  $\theta$  such that  $S = S'\theta$  and  $\Delta \supseteq \Delta'\theta$ .
3. Let  $n = \langle -; []; Q'; \Delta' \rangle$  be an answer node, and  $Q'', \Delta''$  fresh renamings of  $Q', \Delta'$ .

$$\mathbf{resolve}^D(\langle P; [Q, \vec{Q}]; S; \Delta \rangle, n) = \begin{cases} \{ \langle P; \vec{Q}\theta; S\theta; \Delta\theta \cup \Delta''\theta \rangle \} & \text{if } \theta = \text{mgu}(Q, Q'') \text{ exists,} \\ \emptyset & \text{otherwise} \end{cases}$$

4.  $\mathbf{generate}_{\mathcal{P}, \mathcal{A}}^A(P) = \bigcup_{(Q, \vec{Q}) \in \mathcal{P}} \mathbf{resolve}^A(\langle P; [Q, \vec{Q}]; Q; \emptyset \rangle, \langle -; []; P; \emptyset \rangle) \cup \{ \langle P; []; P; \{P\} \rangle : P \text{ is abducible} \}$

The main idea is thus to extend tuples with a *residue*  $\Delta$ , containing atoms which are just *assumed* to hold in the process of the algorithm. Such atoms are initially inserted into the residue by  $\mathbf{generate}_{\mathcal{P}, \mathcal{A}}^A$  whenever an abducible goal  $\langle P \rangle$  is encountered. They are then propagated using  $\mathbf{resolve}^A$  such that for each hypothetical answer  $S$  obtained,  $\Delta$  expresses which atoms must be added to  $\mathcal{P}$  in order to be able to deduce it.

In general, there are infinitely many abductive solutions: by monotonicity, any extension of an abductive solution is trivially also an abductive solution. Clearly, the abductive algorithm should only consider solutions that are not simple supersets of already existing ones. Similarly, we are not interested in a new solution that is an instantiation of (and thus less general than) an already existing one. The subsumption relation  $\preceq_A$  makes sure that such “uninteresting” answers are not considered.

The correctness of the algorithm is formalized by the following theorems.

**Theorem 2.5 (Soundness)** *If  $(\mathcal{N}, \text{Ans}, \text{Wait})$  is reachable from an initial state  $\mathcal{S}_0$  then for all  $P \in \text{dom}(\text{Ans})$ :  $\langle P'; []; S; \Delta \rangle \in \text{Ans}(P)$  implies that for all substitutions  $\vartheta$  such that all elements of  $\Delta\vartheta$  are ground it holds that  $P = P'$ ,  $S\vartheta \preceq P$ , and  $S\vartheta \in T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset)$ .*

**Theorem 2.6 (Completeness)** *If  $\mathcal{S}_f \equiv (\mathcal{N}, \text{Ans}, \text{Wait})$  is a final state reachable from an initial state  $\mathcal{S}_0$  then for all  $P \in \text{dom}(\text{Ans})$ :  $S \in T_{\mathcal{P} \cup \mathcal{A}}^\omega(\emptyset)$  and  $S \preceq P$  implies that there exists a substitution  $\vartheta$ , an atom  $S'$ , and a residue  $\Delta$  such that  $S'\vartheta = S$ ,  $\Delta\vartheta \subseteq \mathcal{A}$ , and  $\langle P; []; S'; \Delta \rangle \in \text{Ans}(P)$ .*

**Example 2.7** Consider again the program of Example 2.4, but assume that it no longer contains the atom `inWorkgroup(Alice, WG23)`. Furthermore, suppose that both `isEmployee` and `inWorkgroup` are abducible predicate names.

For query  $\langle \text{canRead}(z, \text{Foo}) \rangle$  the procedure  $\text{generate}_{\mathcal{P}, \mathcal{A}}^{\mathcal{A}}$  executes as the one described in Example 2.4, and produces in particular the goal node  $n_0$ . Using  $(\text{goal}_2)$ , a new root  $\langle \text{isEmployee}(x) \rangle$  will be inserted. The call to the function  $\text{generate}_{\mathcal{P}, \mathcal{A}}^{\mathcal{A}}(\text{isEmployee}(x))$  produces answer nodes

$$\begin{aligned} &\langle \text{isEmployee}(x); []; \text{isEmployee}(\text{Alice}); \emptyset \rangle \\ &\langle \text{isEmployee}(x); []; \text{isEmployee}(x); \{\text{isEmployee}(x)\} \rangle \end{aligned}$$

Upon termination of the algorithm,  $\text{Ans}(\text{canRead}(z, \text{Foo}))$  contains answer node which exhibit the following answer/residue pairs:

$$\begin{aligned} &(\text{canRead}(\text{Bob}, \text{Foo}), \quad \emptyset) \\ &(\text{canRead}(\text{Alice}, \text{Foo}), \quad \{\text{inWorkgroup}(\text{Alice}, y)\}) \\ &(\text{canRead}(x, \text{Foo}), \quad \{\text{isEmployee}(x), \text{inWorkgroup}(x, y)\}) \end{aligned}$$

The first one does not require any assumptions, and the other two give sets of hypothetical assumptions in order for answers of a particular shape to hold. For example, to grant read access for `Alice`, she would have to show that she is member of some work group  $y$ . All other possible abductive solutions are subsumed by these three answers.

### 3 Termination Conditions

The abduction algorithm from Section 2.4 is guaranteed to terminate if there is a finite set of answers such that every valid answer would be subsumed by some element in the set. However, there are cases in which every complete set of answers is infinite and the algorithm does not terminate.

**Example 3.1** Consider the policy

$$\text{canRead}(user, file) \leftarrow \text{deleg}(delegator, user, file), \text{canRead}(delegator, file).$$

In this example, `canRead`( $x, y$ ) indicates that principal  $x$  has read access to resource  $y$ , and `deleg`( $x, y, z$ ) indicates that principal  $x$  delegates read access for resource  $z$  to principal  $y$ . This policy implements a simple variant of discretionary access control: users can delegate read access if they have read access themselves. The abductive query `canRead(Alice, Foo)` has an infinite set of answers with growing residues:

$$\begin{aligned} &\{\text{canRead}(\text{Alice}, \text{Foo})\} \\ &\{\text{deleg}(x_1, \text{Alice}, \text{Foo}), \text{canRead}(x_1, \text{Foo})\} \\ &\{\text{deleg}(x_1, \text{Alice}, \text{Foo}), \text{deleg}(x_2, x_1, \text{Foo}), \text{canRead}(x_2, \text{Foo})\} \\ &\{\text{deleg}(x_1, \text{Alice}, \text{Foo}), \text{deleg}(x_2, x_1, \text{Foo}), \text{deleg}(x_3, x_2, \text{Foo}), \text{canRead}(x_3, \text{Foo})\} \\ &\vdots \end{aligned}$$

The answers do not subsume each other: being able to provide the missing facts corresponding to one of these answer does not imply being able to provide

the facts corresponding to any other answer. Clearly, the algorithm does not terminate in this case.

There are different ways to approach this problem. If the algorithm is used for debugging policies or for explaining access denials to users, non-termination may not be a serious problem, if the answers can be returned one by one. Ideally, the answers would be returned in some meaningful order, e.g. sorted by simplicity. This can be achieved by serializing the non-deterministic algorithm into a deterministic one with a fixed order of transitions.

Sometimes, however, it is important to ensure termination and to compute the complete set of answers. For example, the algorithm could be used to verify that there is *no* abductive answer of a certain form; this would require a complete set of abductive solutions. This section discusses various strategies of ensuring termination.

### 3.1 Subsumption Weakening

One such strategy is to replace the subsumption relation  $\preceq_A$  by a weaker relation. Intuitively, a weaker subsumption relation has the effect of filtering out more answers in the (ans)-transition; in other words, fewer answers are deemed “relevant” or “interesting”. As long as this alternative subsumption relation enjoys a sort of compactness property (essentially that it is not possible to indefinitely keep adding new answers that are not subsumed) then the algorithm is guaranteed to terminate.

**Theorem 3.2** *Let  $\sqsubseteq$  be a partial order on nodes such that in every infinite sequence of nodes  $n_1, n_2, \dots$  containing only a finite number of distinct constants, there are nodes  $n_i$  and  $n_j$  with  $i < j$  and  $n_j \sqsubseteq n_i$ . If the subsumption relation  $\preceq_A$  on nodes in the algorithm is replaced by  $\sqsubseteq$  then all transition paths starting from an initial state are of finite length.*

The following definition specifies two examples of subsumption relations that satisfy the condition of the theorem.

**Definition 3.3** Let  $n = \langle P; []; S; \Delta \rangle$  and  $n' = \langle P; []; S'; \Delta' \rangle$ . Then  $n \sqsubseteq_0 n'$  iff  $S \preceq S'$  and the predicate names occurring in  $\Delta'$  are a subset of the predicate names occurring in  $\Delta$ . Let  $M$  be a positive integer. Then  $n \sqsubseteq_M n'$  iff  $|\Delta| > M$  or  $n \preceq_A n'$ .

The first relation,  $\sqsubseteq_0$ , is useful if one is only interested in the predicate names of the missing facts, not their parameters. For example, a security administrator may be interested in the question “*is it possible for Alice to gain read access to Foo if someone, no matter who, is granted write access?*”. Here, the administrator is only interested in whether an abductive answer exists that contains the predicate name `canWrite`.

The second relation,  $\sqsubseteq_M$ , is parameterized on a constant  $M$  and filters out answers with more than  $M$  missing facts. This method could be used in the non-termination example above to ensure termination by cutting off the delegation chain at a certain maximum length.

### 3.2 Static Termination Analysis

The advantage of weakening the subsumption relation is a strong termination guarantee for all possible policies, queries and sets of abducibles. The downside

of this approach is a correspondingly weaker completeness result: completeness only holds with respect to the subsumption relation.

We now develop an alternative approach that guarantees termination under the original subsumption relation, as long as the policy satisfies a certain property. To gain an intuition for this property, consider the necessary conditions for non-termination. The algorithm does not terminate only if there is an infinite sequence of abductive answers, each of which is not subsumed by any previous answer. This is only possible if the residues  $\Delta$  are growing unboundedly. Since the answers can only contain the (finitely many) predicate names and constants occurring in the policy and the query, it must be the case that there is a predicate name  $p$  such that for all integers  $N$  there is always a residue in the sequence in which there are more than  $N$  occurrences of  $p$ . But this is only possible if the policy is recursive.

Verbaeten [19] provides a sufficient termination condition for general abductive logic programs that requires non-recursivity. But recursion plays an important role in authorization policies, for example for specifying delegation. Fortunately, due to the subsumption check, recursion does not always lead to non-termination. In order for the sequence to pass the subsumption condition, the  $p$ -atoms in the residues must form increasingly bigger structures that are connected via an unbounded number of shared variables.

In Example 3.1, we can see that non-termination stems from the recursive `canRead` condition, and furthermore the sharing of the variable `delegator` with a second body predicate which is not in the head of the clause: this essentially causes the creation of an increasing linked structure of `deleg` atoms with newly created, shared variables.

This leads to a necessary condition for non-termination, the negation of which is then a sufficient condition for termination. The following definition is used in the formalization of the condition.

**Definition 3.4** A clause  $R \leftarrow \vec{R}$  is an *unfolding* of a clause  $P \leftarrow P_1, \dots, P_i, \dots, P_n$  if there exists a clause  $Q \leftarrow \vec{Q} \in \mathcal{P}$  such that  $P_i$  and  $Q$  unify with mgu  $\theta$ , and  $R = P\theta$  and  $\vec{R} = (P_1, \dots, P_{i-1}, \vec{Q}, P_{i+1}, \dots, P_n)\theta$ . A clause  $C$  can be *unfolded* to yield a clause  $C'$  if  $C'$  is obtained from  $C$  by 0 or more unfolding transformations.

**Theorem 3.5** *Let  $\mathcal{P}$  be a set of clauses such that no clause can be unfolded to yield a clause with the following property:*

- *The head predicate occurs in a body atom  $P$ .*
- *It has a second body atom  $Q$  that is abducible and shares a variable with  $P$  which does not occur in the head.*

*Then all transition paths starting from an initial state are of finite length.*

The condition in Theorem 3.5 is decidable; in fact, a static analyzer for checking it can be written in Prolog in less than a hundred lines. The static analyzer could then be run before an abductive query, and if the condition is not satisfied, the user could be warned that evaluation of the query may not terminate. Preliminary experiments have shown that the condition gives a relatively accurate approximation for non-termination. In particular, many recursive policies can be shown to guarantee termination.

## 4 Application Scenarios

This section illustrates three possible applications of the abduction algorithm in the area of policy-based access control.

### 4.1 Explaining Access Denials to Users

In current access control systems, the user is often left without guidance if access is denied. Consider the following policy of some company:

$$\begin{aligned} \text{canRead}(x, \text{/workgroup23/}) &\leftarrow \text{isEmployee}(x), \text{inWorkgroup}(x, \text{WG23}). \\ \text{canRead}(x, \text{/workgroup23/}) &\leftarrow \text{isManager}(x). \end{aligned}$$

Suppose employee `Alice` submits the credential `isEmployee(Alice)` upon log-on to the authorization system, but forgets to submit the credential that she is member of workgroup `WG23`. If she then tries to access the folder `/workgroup23/`, she will just get the answer “access denied”. Using the abduction technique, the produced residues `{inWorkgroup(Alice, WG23)}` and `{isManager(Alice)}` could be used to construct the more helpful message “access would have been granted if you had shown that you are a member of work group `WG23` or that you are a manager”.

If parts of the policy itself are considered confidential, the abductive solutions can be filtered by a *disclosure* meta-policy. Disclosure policies have been studied extensively in the area of automated trust negotiation (e.g. [21, 20, 13]). A simpler (but slightly less fine-grained) approach would be to tag particular atoms in clause bodies that must not be abduced.

### 4.2 Administration of Authorization Policies

The next example illustrates the use of the abduction algorithm in a policy analysis and debugging tool. Consider the following example rules that are part of a policy of an electronic health record (EHR) service:

$$\begin{aligned} \text{treatingClinician}(cli, pat) &\leftarrow & (1) \\ &\text{roleMember}(pat, \text{Patient}), \text{roleMember}(cli, \text{Clinician}), \text{consent}(pat, cli). \end{aligned}$$
$$\begin{aligned} \text{canReadEHR}(cli, pat, subj) &\leftarrow & (2) \\ &\text{treatingClinician}(cli, pat), \text{nonSensitive}(subj). \end{aligned}$$
$$\begin{aligned} \text{canReadEHR}(cli, pat, \text{Psych}) &\leftarrow & (3) \\ &\text{treatingClinician}(cli, pat), \text{isCertifiedPsychiatrist}(cli). \end{aligned}$$
$$\begin{aligned} \text{canReadEHR}(pat, pat, subj) &\leftarrow & (4) \\ &\text{roleMember}(pat, \text{Patient}), \text{nonSensitive}(subj). \end{aligned}$$

Rule (1) specifies that a clinician `cli` is a *treating clinician* of patient `pat` if the patient has given consent to treatment. The predicate `canReadEHR(x, pat, subj)` is used to check if principal `x` is permitted to read patient `pat`'s record items on subject matter `subj`, where subject matters range over categories such as psychiatry, cardiology, radiology etc. Some subject matters, such as `Psych`, are deemed sensitive and have stricter access requirements. The predicate `nonSensitive(subj)`

defines the range of subjects that are deemed non-sensitive. Rule (2) allows clinicians to read their patients' record items on non-sensitive subjects. Rule (3) specifies that only psychiatrists are permitted to access their patients' psychiatric data. Finally, Rule (4) permits patients to view their own data, but only the items regarding non-sensitive subject matters.

The rationale behind Rule (4) is that patients should not be allowed to access data that could potentially distress them if read without professional guidance. In particular, they should not be able to autonomously access their psychiatric data. In order to check if the policy really implements the intended behavior, the security administrator can issue the abductive query  $\text{canReadEHR}(pat, pat, \text{Psych})$  to see if there is a way for patients to access their own psychiatric data.

Assuming that all predicates apart from  $\text{canReadEHR}$  and  $\text{treatingClinician}$  are abducible, we obtain an answer with residues

$$\{\text{roleMember}(pat, \text{Patient}), \text{nonSensitive}(\text{Psych})\}.$$

This is easily dismissed as unproblematic if the administrator can verify that there is no way of inserting a fact  $\text{nonSensitive}(\text{Psych})$ . But we also get a second answer with residue

$$\{\text{roleMember}(pat, \text{Patient}), \text{roleMember}(pat, \text{Clinician}), \\ \text{isCertifiedPsychiatrist}(pat), \text{consent}(pat, pat)\}.$$

This answer is more troublesome: a patient can read what her psychiatrist has written if she happens to be a certified psychiatrist too and has given consent to treat herself. This may or may not be regarded as a bug in the policy; but in any case, as there are no further abductive answers, and by completeness and termination of the algorithm, it is guaranteed that there are no further loopholes in the policy.

### 4.3 Automated Delegation

The following scenario takes place in a multi-domain grid computing environment. Alice is a user who wishes to submit a job to be computed on a grid compute cluster. She knows that during the execution of her job, a node from the compute cluster will have to access her file `alice.dat` stored on a file server in a different domain. Therefore, at some point, an authorization query of the form  $\text{canRead}(\text{Node}, \text{alice.dat})$  will be evaluated on the file server. Suppose the file server's policy contains the rule from Example 3.1 and the fact  $\text{canRead}(\text{Alice}, \text{alice.dat})$ .

As Alice's job may take many days to complete, she wants to know in advance which delegation credentials she has to submit to the file server, so she sends the abductive query  $\text{canRead}(node, \text{alice.dat})$  to the server. Her query contains a variable *node* in place of the node that will eventually access her file, because she cannot know its identity in advance.

The first returned answer is the trivial answer where *node* is instantiated to `Alice` and the residue is empty. The second answer is uninstantiated and has the singleton residue  $\{\text{deleg}(\text{Alice}, node, \text{alice.dat})\}$ . This would require direct delegation to the node, which is not convenient as its identity is not known to Alice. The third answer has residue

$\{\text{deleg}(\text{Alice}, x, \text{alice.dat}), \text{deleg}(x, \text{node}, \text{alice.dat})\}$ . This answer represents a delegation chain of depth 2 and is the most useful in this situation, because Alice knows the identity of the compute cluster’s scheduling service. Thus she can submit a delegation credential  $\text{deleg}(\text{Alice}, \text{Scheduler}, \text{alice.dat})$  to the scheduling service along with her job and a partially instantiated missing-credential “template”  $\{\text{deleg}(\text{Scheduler}, \text{node}, \text{alice.dat})\}$ . The service will then execute the job on some node, e.g. `Node42`, passing along Alice’s delegation credential as well as a newly created (or cached) credential instantiated from the template, namely  $\text{deleg}(\text{Scheduler}, \text{Node42}, \text{alice.dat})$ . When the node eventually requests access to Alice’s file on the file server, it submits both Alice’s and the scheduler’s delegation credentials to support the request. Access is then guaranteed to be granted as long as the file server’s policy has not been changed in the meantime.

## 5 Discussion

**Related work** There has been very little research on improving the usability of authorization systems in the case of access denial. The *Know* system [12] can provide helpful feedback to the user in the form of a list of conditions under which the policy allows access. A separate disclosure policy restricts the information revealed by the feedback. However, the authors only consider policies of rather limited expressiveness, namely those that can be written as propositional boolean formulas; hence the feedback can be computed using Ordered Binary Decision Diagrams (OBDDs).

Bonatti et al. [7] have developed a framework for explaining both positive and negative access decisions in the context of a Datalog-based authorization language. For explaining access denials, they essentially compute a tabled failed proof graph (called explanation graph) for the query. Users can navigate through the graph, which is represented in controlled natural language, to see where the proof failed. To keep the overhead as low as possible, they do not attempt to search for the missing facts that would complete the failed proof. We have shown in this paper that we can compute the sets of missing facts while maintaining a low implementation overhead.

Koshutanski and Massacci [13] employ abduction for interactive credential retrieval: if the credentials presented by the user are not sufficient to allow access, the service computes a set of missing credentials (filtered by a disclosure policy) and returns it to the user who can either supply the required credentials or decline, in which case the process iterates. This process terminates because it is assumed that the set of constants that could be used as credential parameters is finite and known to the service in advance. The policy can then be reduced to propositional (variable-free) formulas. However, we believe that this is an unreasonable assumption, particularly in decentralized applications where authorization is based on attributes as the identities of principals in delegation chains are not known a priori.

Becker and Nanz [4] have developed an algorithm for analyzing authorization policies in which facts (such as current role activations) can be added and removed dynamically by commands (such as activating or deactivating a role). The paper presupposes a function for computing sufficient preconditions for executing such commands, but does not explain how it can be implemented. The

abductive procedure presented in this paper could be used to implement the required function.

**Implementation** Prototypes of the abductive algorithm and the static termination analysis have been implemented in OCaml and in XSB Prolog, respectively. The SecPAL system [3] supports proof graph generation based on the algorithm presented here. Current and future work on the SecPAL system will include integration of tools based on our abduction algorithm. As SecPAL is compiled into Datalog with constraints, the implementation will also have to handle constraints. Extending the abduction algorithm with constraints is relatively simple, assuming the existence of satisfiability and subsumption checking operations. In essence, tuples are extended to include a constraint on variables occurring in the predicate and the residue. In the resolution step, the conjunction of the constraints from the two input nodes is computed and checked for satisfiability. Furthermore, the subsumption check is extended to also check if the constraint from one node is subsumed by the constraint from the other node.

**Conclusion** Declarative authorization languages can increase the flexibility, scalability, and manageability of access control systems. However, they are not without their own usability problems, as many authorization policies are intrinsically complex. To alleviate this complexity, tools are needed for facilitating access auditing and review, meaningful user feedback, policy diagnosis and debugging, and automated credential retrieval. In this paper, we have shown how a tabled resolution algorithm for policy evaluation can be extended to create proof graphs and to perform abduction, the computation of sets of missing proof facts. For the abductive algorithm, we have explored methods for guaranteeing termination and complete answers. The algorithms are general enough to be applicable to a wide range of existing systems. We believe that tools based on these algorithms will help the declarative policy approach gain wider adoption in the industry.

## References

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] J. J. Alferes, L. M. Pereira, and T. Swift. Abduction in well-founded semantics and generalized stable models via tabled dual programs. *Theory and Practice of Logic Programming*, 4(4):383–428, 2004.
- [3] M. Y. Becker, C. Fournet, and A. D. Gordon. Design and semantics of a decentralized authorization language. In *IEEE Computer Security Foundations Symposium*, 2007.
- [4] M. Y. Becker and S. Nanz. A logic for state-modifying authorization policies. In *European Symposium on Research in Computer Security*, 2007.
- [5] M. Y. Becker and P. Sewell. Cassandra: distributed access control policies with tunable expressiveness. In *IEEE International Workshop on Policies for Distributed Systems and Networks*, pages 159–168, 2004.

- [6] M. Y. Becker and P. Sewell. Cassandra: Flexible trust management, applied to electronic health records. In *IEEE Computer Security Foundations Workshop*, 2004.
- [7] P. A. Bonatti, D. Olmedilla, and J. Peer. Advanced policy explanations on the web. In *European Conference on Artificial Intelligence*, pages 200–204, 2006.
- [8] W. Chen and D. S. Warren. Tabled evaluation with delaying for general logic programs. *Journal of the ACM*, 43(1):20–74, 1996.
- [9] M. Denecker and D. D. Schreye. SLDNFA: An abductive procedure for abductive logic programs. *Journal of Logic Programming*, 34(2):111–167, 1998.
- [10] J. DeTreville. Binder, a logic-based security language. In *IEEE Symposium on Security and Privacy*, pages 105–113, 2002.
- [11] A. C. Kakas, R. A. Kowalski, and F. Toni. The role of abduction in logic programming. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 5, pages 235–324, 1998.
- [12] A. Kapadia, G. Sampemane, and R. H. Campbell. Know why your access was denied: regulating feedback for usable security. In *ACM conference on computer and communications security*, pages 52–61, 2004.
- [13] H. Koshutanski and F. Massacci. Interactive access control for web services. In *International Information Security Conference*, pages 151–166, 2004.
- [14] N. Li and J. C. Mitchell. Datalog with constraints: A foundation for trust management languages. In *Practical Aspects of Declarative Languages*, pages 58–73, 2003.
- [15] N. Li, J. C. Mitchell, and W. H. Winsborough. Design of a role-based trust management framework. In *Symposium on Security and Privacy*, pages 114–130, 2002.
- [16] C. Peirce. *Philosophical Writings of Peirce*, chapter Abduction and Induction. 1955.
- [17] H. Tamaki and T. Sato. OLD resolution with tabulation. In *International Conference on Logic Programming*, pages 84–98. Springer-Verlag, 1986.
- [18] D. Toman. Memoing evaluation for constraint extensions of Datalog. *Constraints*, 2(3/4):337–359, 1997.
- [19] S. Verbaeten. Termination analysis for abductive general logic programs. In *International Conference on Logic Programming*, pages 365–379, 1999.
- [20] W. H. Winsborough and N. Li. Towards practical automated trust negotiation. In *IEEE International Workshop on Policies for Distributed Systems and Networks*, 2002.

- [21] W. H. Winsborough, K. E. Seamons, and V. E. Jones. Automated trust negotiation. In *DARPA Information Survivability Conference and Exposition*, volume 1, pages 88–102, 2000.

## A Proofs

**Lemma A.1** *If  $\mathcal{S} \equiv (\mathcal{N}, \text{Ans}, \text{Wait})$  is reachable from an initial state, and  $\mathcal{S}' \equiv (\mathcal{N}', \text{Ans}', \text{Wait}')$  is reachable from  $\mathcal{S}$ , then  $\text{dom}(\text{Ans}) \subseteq \text{dom}(\text{Ans}')$ ,  $\text{dom}(\text{Wait}) \subseteq \text{dom}(\text{Wait}')$ , and  $\text{dom}(\text{Wait}) \subseteq \text{dom}(\text{Ans})$ . For all  $P \in \text{dom}(\text{Ans})$ :  $\text{Ans}(P) \subseteq \text{Ans}'(P)$ . For all  $P \in \text{dom}(\text{Wait})$ :  $\text{Wait}(P) \subseteq \text{Wait}'(P)$ .*

*Proof.* By induction on the transition rules. The statements follow from the observation that (ans) and (goal<sub>1</sub>) only increase  $\text{Ans}(P)$  and  $\text{Wait}(P)$ , respectively; (goal<sub>2</sub>) always increases the domains of  $\text{Ans}$  and  $\text{Wait}$ ; and (root) leaves  $\text{Wait}$  unchanged and either increases the domain of  $\text{Ans}$  (that can only happen if in the very first transition from the initial state) or leaves  $\text{Ans}$  unchanged.  $\square$

**Lemma A.2 (Node Invariant)** *If  $(\mathcal{N}, \text{Ans}, \text{Wait})$  is reachable from some initial state and  $\langle P; \vec{Q}; S; \Delta \rangle \in \mathcal{N}$  with  $Rl \equiv R \leftarrow \vec{R}$ , then:*

1.  $S \preceq P$ ;
2.  $Rl \in \mathcal{P} \cup \Delta$ ;
3. each node in  $\vec{n}$  is a fresh renaming of some node in  $\bigcup \text{Ans}$ ;
4.  $\langle \cdot; \cdot; \cdot; \Delta_i \rangle \in \vec{n}$  implies  $\Delta_i \subseteq \Delta$ ; and,
5. there is some  $\theta$  such that  $R\theta = S$ , and  $\vec{R}\theta = \vec{Q}' @ \vec{Q}$  (where  $\vec{Q}'$  are the answers in  $\vec{n}$ ).

*Proof.* By induction on the transition rules. The statements follow directly from the the definition of the transition rules.  $\square$

**Theorem 2.5.** *If  $(\mathcal{N}, \text{Ans}, \text{Wait})$  is reachable from an initial state  $\mathcal{S}_0$  then for all  $P \in \text{dom}(\text{Ans})$ :  $\langle P'; [ ]; S; \Delta \rangle \in \text{Ans}(P)$  implies that for all substitutions  $\vartheta$  such that all elements of  $\Delta\vartheta$  are ground it holds that  $P = P'$ ,  $S\vartheta \preceq P$ , and  $S\vartheta \in T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset)$ .*

*Proof.* By induction on transition rules. The statement holds by definition for  $\mathcal{S}_0$ . Now assume the state is not an initial state, and let  $\text{Ans}'$  be the answer table of the preceding state.

For (ans), we only have to consider the new answer  $n \equiv \langle P; [ ]; S; \Delta \rangle$ . Assume we have a substitution  $\vartheta$  such that all elements of  $\Delta\vartheta$  are ground.

By Lemma A.2,  $S \preceq P$ , therefore also  $S\vartheta \preceq P$ ; furthermore,  $Rl \equiv R \leftarrow \vec{R} \in \mathcal{P} \cup \Delta$ . If  $Rl \in \Delta$ , then  $n$  is of the form  $\langle P; [ ]; S; \{S\} \rangle$  and  $S\vartheta \in \Delta\vartheta \subseteq T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset)$ .

In the case of  $Rl \in \mathcal{P}$ , we have again by Lemma A.2 that there exists a substitution  $\theta$  such that  $R\theta = S$  and  $\vec{R}\theta$  is equal to the set of answers in  $\vec{n}$ , each of the nodes of which is a fresh renaming of some node in  $\bigcup \text{Ans}'$ . With  $\theta' = \theta\vartheta$ , we have  $R\theta' = S\vartheta$ . By definition of  $T_{\mathcal{P} \cup \Delta\vartheta}^\omega$ , it remains to be shown that  $\vec{R}\theta' \subseteq T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset)$ .

Consider  $\langle \cdot; [ ]; R_i\theta; \Delta_i \rangle \in \vec{n}$ , i.e. a child node of  $n$ . From Lemma A.2, together with the fact that all elements of  $\Delta\vartheta$  are ground, we have that all elements of  $\Delta_i\vartheta$  are ground. By the inductive hypothesis we have therefore  $R_i\theta\vartheta = R_i\theta' \in T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset)$ , and, since  $\Delta_i \subseteq \Delta$ , also  $R_i\theta' \in T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset)$ . Therefore, by definition of  $T_{\mathcal{P}}^\omega$ ,  $S\vartheta \in T_{\mathcal{P} \cup \Delta\vartheta}^\omega(T_{\mathcal{P} \cup \Delta\vartheta}^\omega(\emptyset))$ , as required.

For the other transition rules the statement trivially holds by the inductive hypothesis.  $\square$

**Theorem 2.6.** *If  $\mathcal{S}_f \equiv (\mathcal{N}, \text{Ans}, \text{Wait})$  is a final state reachable from an initial state  $\mathcal{S}_0$  then for all  $P \in \text{dom}(\text{Ans})$ :  $S \in T_{\mathcal{P} \cup \mathcal{A}}^\omega(\emptyset)$  and  $S \preceq P$  implies that there exists a substitution  $\vartheta$ , an atom  $S'$ , and a residue  $\Delta$  such that  $S'\vartheta = S$ ,  $\Delta\vartheta \subseteq \mathcal{A}$ , and  $\langle P; []; S'; \Delta \rangle \in \text{Ans}(P)$ .*

*Proof.* By induction on  $n$  in  $T_{\mathcal{P} \cup \mathcal{A}}^n(\emptyset)$ . The statement vacuously holds for  $n = 0$ . For  $n > 0$ , consider any  $S \in T_{\mathcal{P} \cup \mathcal{A}}^n(\emptyset)$ ,  $P \in \text{dom}(\text{Ans})$  and  $S \preceq P$ . If  $P$  is already in the domain of  $\mathcal{S}_0$ 's answer table the statement holds by definition of initial state and by monotonicity of the transition rules with respect to the answer table (Lemma A.1). So now assume that  $P$  was added to the domain as a result of a (goal<sub>2</sub>) transition.

By definition of  $T_{\mathcal{P} \cup \mathcal{A}}$ , there exists a rule  $Rl \equiv R_0 \leftarrow R_1, \dots, R_n \in \mathcal{P} \cup \mathcal{A}$  and a substitution  $\theta$  such that  $R_0\theta = S$  and  $R_i\theta \in T_{\mathcal{P} \cup \mathcal{A}}^{n-1}(\emptyset)$ . We choose  $\theta$  such that the variables in its domain are in  $Rl$ .

We distinguish two cases,  $Rl \in \mathcal{P}$  and  $Rl \in \mathcal{A}$ . Assume  $Rl \in \mathcal{A}$  first. Since the elements of  $\mathcal{A}$  are ground,  $S \in \mathcal{A}$ . Together with  $S \preceq P$  this gives *abducible*( $P$ ). Thus there will eventually be a (root) transition producing a new node  $\langle P; []; P; \{P\} \rangle$ . Take  $\vartheta = \text{mgu}(S, P)$ . Then  $P\vartheta = S$ , as required.

Assume now  $Rl \in \mathcal{P}$ . By the inductive hypothesis, for  $i \in \{1, \dots, n\}$  and all  $R'_i \in \text{dom}(\text{Ans})$  such that  $R_i\theta \preceq R'_i$  there exists a substitution  $\vartheta_i$ , an atom  $S'_i$ , and a residue  $\Delta_i$  such that  $S'_i\vartheta_i = R_i\theta$ ,  $\Delta_i\vartheta_i \subseteq \mathcal{A}$ , and  $n'_i \equiv \langle R'_i; []; S'_i; \Delta_i \rangle$  is a fresh renaming of an answer node in  $\text{Ans}(R'_i)$ . We choose each  $\vartheta_i$  such that all variables in its domain are in the freshly renamed answer node, thus their domains are pairwise mutually disjoint.

For the remainder of the proof, let  $P'$  be a fresh renaming of  $P$  such that  $P'\vartheta_0 = R_0\theta = S$ , for some  $\vartheta_0$  (which exists because of  $S \preceq P$ ).

Let  $\theta_0 = \text{mgu}(R_0, P')$ , and for  $i \in \{1, \dots, k\}$ , let  $\theta_i = \text{mgu}(R_i\theta_0 \dots \theta_{i-1}, S'_i)$ .

Let  $\Delta'_0 = \emptyset$ , and for  $i \in \{1, \dots, k\}$ , let  $\Delta'_i = \Delta'_{i-1}\theta_i \cup \Delta_i\theta_i$ .

Finally, let  $n_i \equiv \langle P; [R_{i+1}, \dots, R_k]\theta_0 \dots \theta_i; R_0\theta_0 \dots \theta_i; \Delta'_i \rangle$  for  $i \in \{0, \dots, k\}$ .

We will now show that nodes  $n_0, \dots, n_k$  eventually get produced along any path leading from  $\mathcal{S}_0$  to  $\mathcal{S}_f$  and that  $n_k$  is the answer node we are looking for.

After  $P$  is added to the domain of the answer table in the (goal<sub>2</sub>) transition, there will eventually be a (root) transition producing a new set of nodes that contains  $n_0$ , because  $\langle P; [R_0, R_1, \dots, R_k]; R_0; \emptyset \rangle$  and  $P$  (with  $P'$  as fresh renaming of  $P$ ) are resolvable with resolution unifier  $\theta_0$ .

Suppose for some  $i \in \{1, \dots, k-1\}$  that  $n_i$  gets produced as a node along a path leading from  $\mathcal{S}_0$  to  $\mathcal{S}_f$ . Then there must be a later (goal<sub>1</sub>) or a (goal<sub>2</sub>) transition where  $n_i$  is added to the wait table for some  $R'_{i+1}$  where  $R_{i+1}\theta_i \preceq R'_{i+1}$ . Since  $\theta_i$  is more general than  $\theta$ , we also have  $R_{i+1}\theta \preceq R'_{i+1}$ , so  $S'_{i+1}$  is the answer of some answer node in  $\text{Ans}(R'_{i+1})$ , by the inductive hypothesis. Therefore, this answer node is resolved with  $n_i$  either in a (ans) or a (goal<sub>1</sub>) transition, and the set of nodes produced by this transition contains  $n_{i+1}$ .

We have now established that along all paths leading from  $\mathcal{S}_0$  to  $\mathcal{S}_f$  the nodes  $n_0, \dots, n_k$  are produced. It follows that  $n_k \equiv \langle P; []; R_0\theta_0 \dots \theta_k; \Delta'_k \rangle$  is eventually added to the answer table for  $P$  in a (ans) transition, and hence it is in  $\text{Ans}(P)$  (by Lemma A.1). It remains to prove that there exists  $\vartheta$  such that  $R_0\theta_0 \dots \theta_k\vartheta = S$  and  $\Delta'_k\vartheta \subseteq \mathcal{A}$ .

Consider the substitution  $\theta_0 \dots \theta_k$ . It is by definition equal to

$$\text{mgu}(R_0, P')\text{mgu}(R_1\theta_0, S'_1)\dots\text{mgu}(R_k\theta_0 \dots \theta_{k-1}, S'_k).$$

Note that  $\theta_0 \dots \theta_{i-1}$  does not affect  $S'_i$  (since all  $S'_i$  only contain fresh variables), for  $i \in \{1, \dots, k\}$ , so the above is also equivalent to

$$\text{mgu}(R_0, P') \text{mgu}(R_1 \theta_0, S'_1 \theta_0) \dots \text{mgu}(R_k \theta_0 \dots \theta_{k-1}, S'_k \theta_0 \dots \theta_{k-1}).$$

This, in turn, can be written as  $\text{mgu}((R_0, R_1, \dots, R_k), (P', S'_1, \dots, S'_k))$ . Recall that  $(R_0, R_1, \dots, R_k) \theta = (P' \vartheta_0, S'_1 \vartheta_1, \dots, S'_k \vartheta_k)$ . Since the domains of  $\vartheta_i$  are pairwise mutually disjoint, we also have  $(R_0, R_1, \dots, R_k) \theta = (P', S'_1, \dots, S'_k) \vartheta_0 \dots \vartheta_k$ . Moreover, as the  $R_i \theta$  are ground, and none of the variables in  $P'$  or  $S'_i$  are in the domain of  $\theta$ , we have  $(R_0, R_1, \dots, R_k) \theta \vartheta_0 \dots \vartheta_k = (P', S'_1, \dots, S'_k) \theta \vartheta_0 \dots \vartheta_k$ . We have thus established that  $\theta \vartheta_0 \dots \vartheta_k$  is a unifier of the two tuples. But  $\theta_0 \dots \theta_k$  is their most general unifier, hence  $\theta \vartheta_0 \dots \vartheta_k = \theta_0 \dots \theta_k \vartheta$  for some  $\vartheta$ . Therefore,  $S = R_0 \theta = R_0 \theta \vartheta_0 \dots \vartheta_k = R_0 \theta_0 \dots \theta_k \vartheta$ , as required.

Now observe that  $\Delta_k$  is a union of residues of the form  $\Delta_i \theta_i \dots \theta_k$ , and we have that  $\Delta_i \vartheta_i \subseteq \mathcal{A}$ . Since all  $\Delta_i$  are taken from freshly renamed nodes, none of the variables in  $\Delta_i$  can be in the domain of  $\theta$  or any  $\theta_j$  or  $\vartheta_j$  except for  $i = j$ . Therefore,  $\Delta_i \vartheta_i = \Delta_i \theta \vartheta_0 \dots \vartheta_i \subseteq \mathcal{A}$ , and as this residue is ground, we have in fact  $\Delta_i \vartheta_i = \Delta_i \theta \vartheta_0 \dots \vartheta_k \subseteq \mathcal{A}$ . The substitutions  $\theta_0$  up to  $\theta_{i-1}$  do not affect  $\Delta_i$ , thus  $\Delta_i \theta_0 \dots \theta_k \vartheta = \Delta_i \theta_i \dots \theta_k \vartheta$ , and hence  $\Delta_k \vartheta \subseteq \mathcal{A}$ .  $\square$

**Theorem 3.2.** *Let  $\sqsubseteq$  be a partial order on nodes such that in every infinite sequence of nodes  $n_1, n_2, \dots$  containing only a finite number of distinct constants there are nodes  $n_i$  and  $n_j$  with  $i < j$  and  $n_j \sqsubseteq n_i$ . If the subsumption relation  $\preceq_{\mathcal{A}}$  on nodes in the algorithm is replaced by  $\sqsubseteq$  then all transition paths starting from an initial state are of finite length.*

*Proof.* Similar to the proof for Theorem 3.5 below.  $\square$

**Theorem 3.5.** *Let  $\mathcal{P}$  be a set of clauses such that no clause can be unfolded to yield a clause with the following property:*

- *The head predicate occurs in a body atom  $P$ .*
- *It has a second body atom  $Q$  that is abducible and shares a variable with  $P$  which does not occur in the head.*

*Then all transition paths starting from an initial state are of finite length.*

*Proof.* When a root node  $\langle P \rangle$  is produced in a  $(\text{goal}_2)$  transition,  $P$  is permanently added to the domain of the answer table. Due to the side conditions of  $(\text{goal}_2)$ , such a node is only produced if there is no  $P'$  in the domain of the answer table for which  $P \preceq P'$ . Therefore, the number of  $(\text{goal}_2)$  transitions is finite, and thus the number of (root) transitions is also finite. Each (root) transition produces only a finite number of nodes, since  $\mathcal{P}$  is finite. Assuming that the size of  $\text{Ans}$  is bounded within any path (to be proved below), there are only finitely many (ans) transitions. Since  $\text{Wait}$  is finite, there are only finitely many nodes produced by (ans) transitions along the path.  $(\text{goal}_1)$  replaces a node with a number of nodes whose subgoal lists are strictly shorter until an answer node is produced. Assuming that the size of  $\text{Ans}$  is bounded within any path, each  $(\text{goal}_1)$  transition produces only finitely many nodes. Therefore, there can also only be finitely many  $(\text{goal}_1)$  transitions.

We now prove that the size of  $\text{Ans}$  is bounded. It is sufficient to prove that there is no infinite sequence of answer nodes produced along the path such that

each node is not subsumed by any previous node in the sequence. Suppose there is such an infinite sequence. Then the size of the residues is unbounded. Therefore there exists an abducible predicate  $p$  such that for every integer  $N$  there is a answer node in the sequence with a residue that contains more than  $N$  occurrences of  $p$ . Furthermore, the  $p$ -atoms within a residue are linked by an unbounded number of shared variables. This is only possible if these atoms are produced by a clause with some head predicate  $p_0$  whose body can be unfolded to contain two distinct body atoms with predicates  $p_0$  and  $p$ . Furthermore, these two body atoms share a variable that is not in the head of the clause. This contradicts the assumed condition.  $\square$