

# OPERATIONAL SEMANTICS FOR DKAL: APPLICATION AND ANALYSIS

YURI GUREVICH AND ARNAB ROY

ABSTRACT. DKAL is a new authorization language based on existential fixed-point logic and more expressive than existing authorization languages in the literature. We present some lessons learned during the first practical application of DKAL and some improvements that we made to DKAL as a result. We develop operational semantics for DKAL and present some complexity results related to the operational semantics.

## 1. INTRODUCTION

*Preamble.* DKAL is a new declarative authorization language for distributed systems [6]; it is based on existential fixed-point logic [3] and is considerably more expressive than existing authorization languages in the literature. The first practical application of DKAL was to the problem of automating source asset management (SAM) at Microsoft. We partook an active role in that endeavor. The automation problem involves more than DKAL and is still a work in progress. While DKAL proved itself sufficiently expressive as well as convenient and elegant, the application necessitated some improvements of DKAL that became a part of a larger revision of DKAL [8]. In this paper we present some lessons learned during the first practical application of DKAL and the improvements that we made to DKAL. In addition, we present some related theoretical results. We presume that the reader is familiar with [6]; otherwise this paper is self-contained.

*Source asset management (SAM).* SAM is an example of asset management. Large software companies have many partners, contractors and subcontractors who need to access the sources of the various software products produced by the company. The ever-growing number of such requests necessitates clear access control policies regulating who is entitled to use what sources, where, for how long, and so on. The company also needs processes in place to efficiently implement those policies. DKAL enables formulating rich and sophisticated policies that result in simpler audit and feasible automation.

*Lessons.* The most important lesson that we learned working on SAM was this: separate the access control policy from the workflow. The workflow includes the procedural aspects of access decisions. The workflow can be quite complex to ensure the flexibility of the the way access decisions are made. But the policies should be succinct and easy to understand. Furthermore, the policies should guarantee important safety properties of the access decision process for any workflow. Various compliance requirement and regulations typically can be formulated as safety

---

The work was done in the summer of 2008 when the second author was an intern at Microsoft Research.

properties. It is desirable that the policies alone guarantee the compliance with those requirements and regulations thus eliminating the need to understand large amounts of procedural code.

One other lesson was that in the case of automated access control regulated by clear access control policies it is crucial for the auditing purposes to properly log all human judgments. No set of policies and processes makes industrial access control deterministic. The inherent non-determinism is resolved by human judgments in making and delegating access decision. Logging these judgments is necessary for auditing purposes as it allows the auditor to understand who made what decisions and whether they had the right to do so (either originally or by delegation).

*Operational semantics and related complexity questions.* As a result of the lessons learned, we made some theoretical advances to DKAL. We separated policy and workflow and defined clear operational semantics. Logging the right information is a natural byproduct of the way operational semantics is defined. In addition we obtained some theoretical results characterizing the complexity of analyzing natural properties of policies. We explored two communication models for defining workflows and proved that deciding reachability questions is possible in polynomial time in certain cases whereas deciding invariants is coNP-complete.

## 2. RELATED WORK

Several languages have been developed in recent years to express authorization policies. The late genealogy of DKAL consists primarily of Binder, Delegation Logic and SecPAL. Binder [4] builds directly on Datalog, extending it with an import/export construct `says` that connects Datalog policies of different principals and makes the issuer of an imported assertion explicit. Delegation Logic [9, 10] does not have explicit issuers for all assertions but it has constructs specific to authorization including ones for delegation, representation, and thresholds.

SecPAL [1, 2, 5] has both explicit issuers and specific authorization constructs designed with distributed systems in mind. The number of constructs is deliberately kept low, but the language is expressive and captures numerous standard authorization scenarios. One important construct is *can say*: if A says B can say foo and if B says foo then A says foo. The can say construct can be nested. The semantics is defined by means of a few deduction rules but, for execution purposes, SecPAL reduces to safe Constraint Datalog. Nesting of the *can say* construct gives rise to relations whose arity depends on the nesting depth.

DKAL [6, 7] extends SecPAL in many directions as far as expressivity is concerned. But it is quite different and builds on Existential Fixed-Point Logic (EFPL) which is much richer than Constraint Datalog; EFPL is cast as Liberal Datalog in [3]. In particular DKAL allows one to freely use functions; as a result, principals can quote other principals. Contrary to SecPAL, communication in DKAL is targeted which is beneficial for security and liability protection. The additional expressivity is achieved within the same bounds of computational complexity as that of SecPAL.

The main novelty of this paper is that we take workflow into account and deal explicitly with the runtime evolution of policies. In this connection we extend DKAL with the *from* construct. In [6], the sender directs communication to particular receivers but the receivers are passive. The *from* construct allows the receivers to filter incoming information.

A limited policy evolution is possible in SecPAL: policy statements can be revoked but revocation is governed by static rules. There has been work on the analysis of dynamic evolution of authorization policies in the Trust Management framework [11]. The approach there has been to specify statically what kind of policy rules can be added or removed. They analyzed reachability and invariant checking. We do that too in our framework. The richness of DKAL compared to the simpler TM language makes these questions even more interesting.

### 3. OPERATIONAL SEMANTICS

We define operational semantics for communication in DKAL. This makes explicit the transaction of information that was implicit in [6]. Also we model situations where principals may listen only to information that they want to listen to by introducing a `from` construct. We deprecate the predicate `knows0` and the function `said0`; we also rename `tdon0` to `tdon`. We modify the TrustApplication rule to

$$x \leq (p \text{ said } x + p \text{ tdon } x)$$

**3.1. Terminology.** Every principal has an immutable, or at least stable, core policy and a dynamic policy where it can assert communication statements. We use the following notation:

$\Pi_p$	Core policy of p.
$\Delta_p$	Dynamic policy of p.
$\mathcal{K}_p$	$\{x : \Pi_p \cup \Delta_p \vdash x, \text{ no free variables in } x\}$

$\mathcal{K}_p$  is the *knowledge* of p.

**3.2. Communication Rules.** The main rule is COM. In simple form it says that if a principal  $A$  says an infon  $x$  to a principal  $B$  and if  $B$  is willing to accept  $x$  then  $x$  gets transmitted and  $B$  learns that  $A$  `said`  $x$ . In general, the rule is more complicated and involves a substitution. We restrict the infons being transmitted to ground infons.

$$\frac{\left\{ \begin{array}{l} (A \text{ to } q : x \leftarrow x_1, x_2, \dots, x_m, \delta) \in \Pi_A \cup \Delta_A \\ (B \text{ from } p : y) \in \Pi_B \cup \Delta_B \\ \{\eta x_1, \eta x_2, \dots, \eta x_m\} \subseteq \mathcal{K}_A, \text{Substrate}_A \models \eta \delta \\ \eta p = A, \eta q = B, \eta x = \eta y, \text{ no free variables in } \eta x \end{array} \right.}{\Delta_B += A \text{ said } \eta x} \text{ (COM)}$$

The last line means that  $\Delta_B$  is augmented with infon  $A$  `said`  $\eta x$ .

COM can be generalized to express common scenarios more succinctly. One such scenario is that if a principal  $B$  accepts an infon  $x$ , then it accepts infons of the form

$$z_1 \text{ tdon } z_2 \text{ tdon } \dots z_k \text{ tdon } x.$$

$$\frac{\left\{ \begin{array}{l} (A \text{ to } q : z_1 \text{ tdon } z_2 \text{ tdon } \dots z_k \text{ tdon } x \\ \quad \leftarrow x_1, x_2, \dots, x_m, \delta) \in \Pi_A \cup \Delta_A \\ (B \text{ from } p : y) \in \Pi_B \cup \Delta_B \\ \{\eta x_1, \eta x_2, \dots, \eta x_m\} \subseteq \mathcal{K}_A, \text{Substrate}_A \models \eta \delta \\ \eta p = A, \eta q = B, \eta x = \eta y, \text{ no free variables in } \eta x \end{array} \right.}{\Delta_B + = A \text{ said } \eta z_1 \text{ tdon } \eta z_2 \text{ tdon } \dots \eta z_k \text{ tdon } \eta x} \text{ (TCOM)}$$

#### 4. EXAMPLES

4.1. Alice wants to download an article from the Chux (an allusion to Chuck's) store. Best publishing house owns the copyright and delegates downloading permission to Chux. Chux has the policy that anybody meeting certain approval criteria can download the article. This gives rise to the following core policies.

$$\begin{aligned} \Pi_{best} : & \quad \text{best to } p : \text{chux tdon } p \text{ canDownload article} \\ \Pi_{chux} : & \quad \text{chux to } q : q \text{ canDownload article} \leftarrow \text{approve}(q, \text{article}) \\ \Pi_{alice} : & \quad \text{alice : best tdon}^* \text{ alice canDownload article} \end{aligned}$$

Alice starts the process by issuing a dynamic **from** assertion.

$$(1) \quad \text{alice asserts } \text{from } r : \text{alice canDownload article.}$$

Assuming that the relation  $\text{approve}(\text{alice}, \text{article})$  is true, we have the following evolution.

$$(2) \quad \Delta_{alice} + = \text{alice from } r : \text{alice canDownload article.} \quad \text{By (1)}$$

$$(3) \quad \Delta_{alice} + = \text{best said chux tdon} \\ \quad \text{alice canDownload article.} \quad \text{By (2), } \Pi_{best}, \text{TCOM}$$

$$(4) \quad \Delta_{alice} + = \text{chux said alice canDownload article.} \quad \text{By (2), COM}$$

$$(5) \quad \Pi_{alice} \cup \Delta_{alice} \vdash \text{alice knows} \\ \quad \text{alice canDownload article.} \quad \text{By (3, 4), } \Pi_{alice}$$

4.2. We add more details to the previous example. Chux lets a customer download an article if she authorizes the right amount of money for payment and has a perfect payrate. Chux trusts accounts.Chux on customer's payrate.

$\Pi_{chux}$  :

chux **to**  $a$  :  $a$  *canDownload*  $s \leftarrow$   
 $a$  *authorized*  $\$k$  to chux for  $s$ ,  $a$  *hasPayRate* Perfect,  
 $price(s) = k$ .

chux : accounts.Chux **tdon**  $a$  *hasPayRate*  $e$ .

chux :  $a$  **tdon**  $a$  *authorized*  $\$k$  to chux for  $s$ .

chux **from**  $a$  :  $a$  *authorized*  $\$k$  to chux for  $s$ .

$\Pi_{accounts.Chux}$  :

accounts.Chux **to** chux :  $a$  *hasPayRate*  $e \leftarrow a$  *hasPayRate*  $e$ .

The last assertion is not a tautology. If accounts.chux knows that  $a$  has PayRate  $e$  then it says that to Chux.

Alice starts by authorizing payment to Chux for the article and then asks for the article. Chux then asks whether Alice has perfect payrate. The dynamic assertions are as follows:

alice asserts **to** chux : alice *authorized*  $\$40$  to chux for article.  
 alice asserts **from**  $r$  : alice *canDownload* article.  
 chux asserts **from**  $a$  : alice *hasPayRate*  $e$ .

If the price of the article is  $\$40$  and Alice has a perfect payrate, we can again infer using the operational semantics that Alice knows she can download the article.

4.3. In this example, Chux gives discounts to employees of Fabricam. Chux trusts Fabricam to state whether a customer is an employee. Chux also trusts Crypto, a certifying intermediary, to confirm that a company claims that someone is its employee.

$\Pi_{chux}$  :

chux : crypto **tdon**  $r$  **said**  $q$  *is an employee of*  $r$ .

chux : fabricam **tdon**  $q$  *is an employee of* fabricam.

chux :  $q$  *can take discount* 5X4302  $\leftarrow q$  *is an employee of* fabricam.

chux **from**  $a$  :  $r$  **said**  $q$  *is an employee of*  $r$ .

chux **from**  $a$  :  $q$  *is an employee of* fabricam.

Consider a scenario where

crypto asserts **to** chux : fabricam **said** chris *is an employee of* fabricam.

Using the operational semantics, we can infer that Chux knows Chris can take discount 5X4302.

## 5. MODELING SOURCE ASSET MANAGEMENT

In this section we explore the application of DKAL to a practical problem in large software companies, namely Source Asset Management (SAM). This encompasses formulating a policy and processes for making code access decisions.

A Project Manager in a large software company  $A$  wants a codeset to be accessed by staff at another company. In  $A$ , the codeset is owned by an Information Asset Owner (IAO) who usually delegates such access decisions to a Source Rep. The access decision takes into account the codeset, the company to access the codeset, the type of permission — read/write or read-only.

In our DKAL model, authorizations are managed by Authorization Managers (AMs) of the respective companies. AMs are automated tools modeled as principals. As far as code authorization is concerned, the AM of the company is (or represents) the company itself. The Project Manager, IAO and Source Rep are also modeled as principals. The policy of the authorization manager  $A$ -AM of  $A$  may be as follows:

$\Pi_{A\text{-AM}}$  :

$p \text{ tdon } q \text{ canGet}(\text{codeset}, \text{parameters}) \leftarrow \text{assetOwner}(p, \text{codeset}).$   
 $p \text{ tdon } q \text{ tdon } r \text{ canGet}(\text{codeset}, \text{parameters}) \leftarrow$   
 $\quad p \text{ tdon } r \text{ canGet}(\text{codeset}, \text{parameters}), p \text{ is the manager of } q.$   
 $p \text{ tdon } q \text{ canAccess}(\text{code}, \text{parameters}) \leftarrow$   
 $\quad p \text{ is a PM}, r \text{ canGet}(\text{codeset}, \text{parameters}),$   
 $\quad r \text{ is the AM of } q, \text{code is in codeset.}$   
 $p \text{ tdon } q \text{ tdon } r \text{ canAccess}(\text{code}, \text{parameters}) \leftarrow$   
 $\quad p \text{ tdon } q \text{ canAccess}(\text{code}, \text{parameters}).$

Consider a scenario where  $A$  wants to allow a vendor  $B$  an access to a codeset, and where  $B$ -AM has the policy

$\Pi_{B\text{-AM}}$  :

$A\text{-AM tdon } B\text{-AM canGet}(\text{codeset}, \text{parameters}).$   
 $A\text{-AM tdon } q \text{ canAccess}(\text{code}, \text{parameters}).$

Suppose that Alan, Alfred, Andrew, Anthony and Alice work for  $A$ , and let Alan  $\triangleright$  Alfred  $\triangleright$  Andrew  $\triangleright$  Anthony where  $p \triangleright q$  means  $p$  is the manager of  $q$ . The role of Alice in the hierarchy will not matter. Suppose further that Alan is the IAO of *drivercodes*, a codeset that is to be given access to by  $B$ . A simplified workflow for the access decision may be as follows:

Access decision starts; Anthony goes to  $A$ -AM to request that  $B$  be given access to *drivercodes*.

$A\text{-AM asserts from } p: B\text{-AM canGet}(\text{drivercodes}, \text{params1}).$

Decision is escalated to Andrew, and then to Alfred.

Alfred *asserts to*  $A\text{-AM}: B\text{-AM canGet}(\text{drivercodes}, \text{params1}).$

Alan *asserts to*  $A\text{-AM}: \text{Alfred tdon } B\text{-AM canGet}(\text{drivercodes}, \text{params1}).$

$A\text{-AM}$  communicates decision to  $B\text{-AM}$ .

$A\text{-AM asserts to } B\text{-AM}: B\text{-AM canGet}(\text{drivercodes}, \text{params1}).$

A developer Bruce in  $B$  needs a portion of drivercodes called gfx. Policy of Bruce is as follows:

$$\Pi_{bruce} : \\ B\text{-AM } \mathbf{tdon} \text{ Bruce } \mathit{canAccess}(code, parameters).$$

A typical workflow to acquire this code would be as follows:

Bruce asks for access to the code.

Bruce *asserts* **from**  $p$  : Bruce *canAccess*(gfx, params1).

In the workflow, this goes to  $B$ -AM.  $B$ -AM forwards the query along.

$B$ -AM *asserts* **from**  $p$  : Bruce *canAccess*(gfx, params1).

In the workflow, query goes to  $A$ -AM, which again forwards the query.

$A$ -AM *asserts* **from**  $p$  : Bruce *canAccess*(gfx, params1).

The workflow “chooses” who to ask about the query. It happens to be Alice. She decides to grant the code access to Bruce.

Alice *asserts* **to**  $A$ -AM : Bruce *canAccess*(gfx, params1).

Anthony confirms that Alice has the authority to grant access.

Anthony *asserts* **to**  $A$ -AM : Alice **tdon** Bruce *canAccess*(gfx, params1).

$A$ -AM infers that Bruce can access the code and informs  $B$ -AM accordingly.

$A$ -AM *asserts* **to**  $B$ -AM : Bruce *canAccess*(gfx, params1).

$B$ -AM trusts  $A$ -AM on such statements and concludes that Bruce can access the code.  $B$ -AM communicates the decision to Bruce.

$B$ -AM *asserts* **to** Bruce : Bruce *canAccess*(gfx, params1).

Bruce concludes that he can access the code.

As you can imagine, the workflow can proceed in myriad possible ways. Yet, the core policies of  $A$ -AM,  $B$ -AM and Bruce lets us state the following theorem, which holds independent of the workflow.

**Theorem 1** (Safety).

- (1) If  $B$ -AM *knows*  $B$ -AM *canGet*(codeset, parameters)  
then  $B$ -AM *knows*  $A$ -AM *said*  $B$ -AM *canGet*(codeset, parameters)
- (2) If Bruce *knows* Bruce *canAccess*(code, parameters)  
then Bruce *knows*  $B$ -AM *said* Bruce *canAccess*(code, parameters)

## 6. COMPLEXITY ANALYSIS OF SAFETY PROPERTIES

**6.1. Unrestricted Communication in Restricted DKAL.** In this subsection we consider a restricted, yet useful fragment  $DKAL^-$  of DKAL which is DKAL without variables, without  $+$ ,  $\mathit{canActAs}$ ,  $\mathit{canSpeakAs}$ ,  $\mathit{exists}$ ,  $\mathit{tdon}^*$  and without substrate constraints in the rules. Note that attributes allow us to deal with

roles even in the absence of `canActAs`, `canSpeakAs`. As opposed to `canActAs` and `canSpeakAs`, attributes take parameters. In the restricted fragment, we have three forms of infons. Firstly, there are *primitive infons* of the form `Attribute(a1, ..., ar)` where elements  $a_i$  are regular. Secondly, we have *said infons* of the form `p said x`. Thirdly, we have *tdon infons* of the form `p tdon x`.

The communication is unrestricted in the following sense. We are interested in the consequences of the core policy of a given principal, say  $A$ , without restricting information that  $A$  can learn from other principals.

To simplify exposition, we let  $\Sigma_A$  be the set of said infons derived by the COM rule from  $\Delta_A$  and communications to  $A$  by other principal. For example, if  $(B \text{ to } A : x) \in \Delta_A$  and  $(A \text{ from } B : x) \in \Delta_B$  then  $(B \text{ said } x) \in \Sigma_A$ . From the perspective of knowledge, it is sufficient to work with  $\Sigma_A$ , since the `from` and `to` statements by themselves do not add any knowledge for  $A$ ; it is only when we combine corresponding `from` and `to` statements by the COM rule that additional knowledge is derived.

One interesting property is the reachability relation which is defined as follows: an infon  $I$  is reachable from a core policy  $\Pi_A$  if there exists a set  $\Sigma_A$  such that  $\Pi_A \cup \Sigma_A \vdash A \text{ knows } I$ .

**Theorem 2** (Reachability). *The reachability relation for  $DKAL^-$  is polynomial time.*

*Proof.* Without loss of generality, we may assume that initially  $\Sigma_A$  is empty. We transform the total policy  $\Pi_A \cup \Sigma_A$  of  $A$  in ways that preserve the reachability status of the given infon  $I$ .

*Remove some said infons from core policy.*

- (1) Remove any rule  $(p \text{ said } x) \leftarrow x_1, \dots, x_k$  from  $\Pi_A$  and add the assertion  $(p \text{ said } x)$  to  $\Sigma_A$ .
- (2) If an infon  $p \text{ said } x$  occurs in the body of a rule  $R$  in  $\Pi_A$ , remove the infon from the body of  $R$  and add the assertion  $(p \text{ said } x)$  to  $\Sigma_A$ .

The new policy  $\Pi_A^1 \cup \Sigma_A^1$  has the following property. In any core rule  $y_0 \leftarrow y_1, \dots, y_m$ , every infon  $y_i$  is either primitive or a tdon infon. Obviously,  $\Pi_A \cup \Sigma_A^1 \vdash A \text{ knows } I$  if and only if  $\Pi_A^1 \cup \Sigma_A^1 \vdash A \text{ knows } I$ .

*Augment tdon rules.* For any core rule  $p_1 \text{ tdon } p_2 \text{ tdon } \dots \text{ tdon } p_j \text{ tdon } x \leftarrow x_1, \dots, x_k$  where  $x$  is primitive or a said infon, do the following. Add core rules

$$\begin{aligned}
& p_2 \text{ tdon } p_3 \text{ tdon } \dots p_j \text{ tdon } x \leftarrow p_1 \text{ tdon } p_2 \text{ tdon } \dots p_j \text{ tdon } x \\
& p_3 \text{ tdon } \dots p_j \text{ tdon } x \leftarrow p_2 \text{ tdon } p_3 \text{ tdon } \dots p_j \text{ tdon } x \\
& \dots \\
& p_j \text{ tdon } x \leftarrow p_{j-1} \text{ tdon } p_j \text{ tdon } x \\
& x \leftarrow p_j \text{ tdon } x
\end{aligned}$$

Also add the following assertions to  $\Sigma_A$ :

$$\begin{aligned}
 & p_1 \text{ said } p_2 \text{ tdon } \dots p_j \text{ tdon } x \\
 & p_2 \text{ said } p_3 \text{ tdon } \dots p_j \text{ tdon } x \\
 & \dots \\
 & p_{j-1} \text{ said } p_j \text{ tdon } x \\
 & p_j \text{ said } x
 \end{aligned}$$

This way we obtain  $\Pi_A^2 \cup \Sigma_A^2$ . Observe that  $\Pi_A \cup \Sigma_A^2 \vdash A \text{ knows } I$  if and only if  $\Pi_A^2 \cup \Sigma_A^2 \vdash A \text{ knows } I$ .

Any said infon  $p \text{ said } x$  is reachable from any  $\Pi_A$  with a witness  $\Sigma_A$  asserting ( $A \text{ from } p : x$ ) and  $\Sigma_p$  asserting ( $p \text{ to } A : x$ ). So we may assume that  $I$  is primitive or a tdon infon. Recall that the DKAL derivability relation is polynomial time. Thus it suffice to prove that  $I$  is reachable from  $\Pi_A$  if and only if  $\Pi_A^2 \cup \Sigma_A^2 \vdash A \text{ knows } I$ . The if part is obvious. It remains to prove the only if part.

We claim that  $A$  cannot derive a new primitive or tdon infon as a result of learning any additional information, that is as a result of extending  $\Sigma_A^2$  with any additional assertion  $R$  of said infon  $p \text{ said } x$  (after using the COM rule with corresponding **from** and **to** assertions). Indeed, there are no said infons in the bodies of the rules in  $\Pi_A^2 \cup \Sigma_A^2$ . So adding  $R$  triggers no rule. Further, if  $\Pi_A^2 \cup \Sigma_A^2 \vdash p \text{ tdon } x$  ( $x$  can be an infon of any type) then, by the second transformation above,  $\Pi_A^2 \cup \Sigma_A^2 \vdash x$ . So adding  $R$  does not help either. But these are the only ways that a said assertion can be helpful.  $\square$

We say that  $(A \text{ knows } I_1) \longrightarrow (A \text{ knows } I_2)$  is an invariant for a core policy  $\Pi_A$  if  $\Pi_A \cup \Sigma_A \vdash A \text{ knows } I_1$  implies  $\Pi_A \cup \Sigma_A \vdash A \text{ knows } I_2$  for any set  $\Sigma_A$ . This gives rise to a binary invariant relation  $\text{Inv}(I_1, I_2)$  for  $DKAL^-$ .

**Theorem 3** (Invariant Checking - Lower Bound). *The invariant relation for  $DKAL^-$  is coNP-hard.*

*Proof.* Reduction from 3-SAT. Given a 3-SAT instance  $\phi \equiv (x_{11} \vee x_{12} \vee x_{13}) \wedge \dots \wedge (x_{n1} \vee x_{n2} \vee x_{n3})$ , where each  $x_{ij}$  is either one of the literals  $x_1, \dots, x_k$  or one of them negated, we construct a policy  $\Pi_A$ :

$$\begin{aligned}
 & c_1 \leftarrow B \text{ said } x_{11} \quad c_1 \leftarrow B \text{ said } x_{12} \quad c_1 \leftarrow B \text{ said } x_{13} \\
 & \dots \\
 & c_n \leftarrow B \text{ said } x_{n1} \quad c_n \leftarrow B \text{ said } x_{n2} \quad c_n \leftarrow B \text{ said } x_{n3}
 \end{aligned}$$

$$\begin{aligned}
 & F \leftarrow B \text{ said } x_1, B \text{ said } x'_1 \\
 & \dots \\
 & F \leftarrow B \text{ said } x_k, B \text{ said } x'_k
 \end{aligned}$$

$$\phi \leftarrow c_1, \dots, c_n$$

Here  $c_1, \dots, c_n, F, \phi$  are infons. Now we set  $I_1 = \phi$  and  $I_2 = F$ . Observe that  $\phi$  has a satisfying assignment iff there exists a  $\Sigma_A$  such that  $\Pi_A \cup \Sigma_A \vdash A \text{ knows } \phi$ , but  $\Pi_A \cup \Sigma_A \not\vdash A \text{ knows } F$ .  $\square$

**Theorem 4** (Invariant Checking - Upper Bound). *The invariant relation for  $DKAL^-$  is in  $coNP$ .*

*Proof. Case 1 :*  $I_1$  is of the form  $p \text{ said } x$ . In this case we set  $\Sigma_A = \{p \text{ said } x\}$  and check whether  $\Pi_A \cup \Sigma_A \vdash A \text{ knows } I_2$ . If not then  $\text{Inv}(I_1, I_2)$  fails. If yes then  $\text{Inv}(I_1, I_2)$  holds; this is because both  $I_1$  and  $I_2$  will continue to hold if more assertions are added.

**Case 2 :**  $I_1$  is not a said infon. We claim this: if  $\text{Inv}(I_1, I_2)$  fails then there is a subset  $\Sigma'_A$  of  $\Sigma_A^2$ , where  $\Sigma_A^2$  is the dynamic policy constructed in the proof of theorem 2, such that  $\Pi_A \cup \Sigma'_A \vdash A \text{ knows } I_1$  but  $\Pi_A \cup \Sigma'_A \not\vdash A \text{ knows } I_2$ . Recall that no additional (to  $\Sigma_A^2$ ) communication allows  $A$  to derive any new primitive or  $\text{tdon}$  infon. Thus if  $\Pi_A \cup \Sigma_A^2 \not\vdash A \text{ knows } I_1$ , the invariant trivially holds. Suppose therefore,  $\Pi_A \cup \Sigma_A^2 \vdash A \text{ knows } I_1$ . Now if  $\Pi_A \cup \Sigma_A^2 \not\vdash A \text{ knows } I_2$ , then  $\Sigma_A^2$  is itself a certificate of non-membership. So we additionally suppose  $\Pi_A \cup \Sigma_A^2 \vdash A \text{ knows } I_2$ .

Let  $\Sigma''$  be the set of minimum cardinality such that  $\Pi_A \cup \Sigma'' \vdash A \text{ knows } I_1$  and  $\Pi_A \cup \Sigma'' \not\vdash A \text{ knows } I_2$ . Let  $\Sigma' = \Sigma'' \cap \Sigma_A^2$  and let  $\Sigma_\delta = \Sigma'' - \Sigma'$ . It must be that  $\Pi_A \cup \Sigma' \not\vdash A \text{ knows } I_1$  and  $\Pi_A \cup \Sigma_\delta \not\vdash A \text{ knows } I_2$ , since otherwise  $\Sigma'$  will be of cardinality at most equal to  $\Sigma''$  and it will be a subset of  $\Sigma_A^2$ . Now, if  $p \text{ said } x \in \Sigma_\delta$ , then  $p \text{ said } x$  cannot be in the body of some rule, nor can any rule having a head of the form  $q_1 \text{ tdon } q_2 \text{ tdon } \dots \text{ tdon } p \text{ tdon } x$  be in  $\Pi_A$ , since these are already in  $\Sigma_A^2$ . Therefore, starting from  $\Sigma'$  we cannot add any *useful* assertions of the form  $p \text{ said } x$  to make  $A \text{ knows } I_1$  true as  $I_1$  does not begin with a **said**. Hence  $\Sigma_\delta$  must be empty, making  $\Sigma''$  a subset of  $\Sigma_A^2$ . Since  $\Sigma_A^2$  is of poly-size in the input, we have a poly-sized certificate of non-membership.  $\square$

**Corollary 1** (Invariant Checking). *The invariant relation for  $DKAL^-$  is  $coNP$ -complete.*

Remarks.

- If we constrain the principals in theorem 2 to assert at most one of more than one possible choices, then reachability is  $NP$ -hard. The proof is simple: just get rid of the ‘F’ rules in a reduction proof similar to theorem 3.

**6.2. Restricted Communication in Unrestricted DKAL.** In this subsection we analyze the complexity of deciding classes of safety properties for unrestricted DKAL. The policy does not have any **from** rules. The workflow interacts with the DKAL inference machine through **from** assertions only. The content of **from** assertions is unrestricted within the bounds of syntactic well-formedness.

Since the **to** assertions can only be derived from the core policy, we need to look at the union of the policies of all the principals  $\bigcup_\alpha \Pi_\alpha$  in order to analyze reachability and invariant checking. We assume that the number of variables in each rule in each policy is fixed a-priori and that  $\alpha$  ranges over all principals in the system. Also since the set of dynamic assertions  $\Delta_A$  for a principal  $A$  is only going to contain **from** assertions,  $\Sigma_A$  as described in the previous section would be irrelevant here. Hence we work with  $\Delta$ 's in this section.

**Theorem 5** (Reachability Checking). *The reachability relation for DKAL is in polynomial time.*

*Proof.* We issue **from** statements for all the **to** rules in the policies. Let us call the union of these assertions  $\Delta'$ . We then check whether  $\bigcup_\alpha \Pi_\alpha \cup \Delta' \vdash A \text{ knows } I$ .

This is doable in polytime since  $\Delta'$  is polysized in the policies. It is easy to see that  $\Delta'$  is the required  $\Delta$ , since if  $A$  knows  $I$  is not derivable using  $\Delta'$  with the policies then no further **from** assertion will help derive it.  $\square$

**Theorem 6** (Invariant Checking - Lower Bound). *The invariant relation for DKAL is coNP-hard.*

*Proof.* The proof is similar to theorem 3. We will reduce 3-SAT to  $\overline{\text{INVARIANT}}$ . Given a 3-SAT instance  $\phi \equiv (x_{11} \vee x_{12} \vee x_{13}) \wedge \cdots \wedge (x_{n1} \vee x_{n2} \vee x_{n3})$ , where each  $x_{ij}$  is either one of the literals  $x_1, \dots, x_k$  or one of them negated, we construct the following policy:

$\Pi_A :$

$$\begin{array}{l} c_1 \leftarrow B \text{ said } x_{11} \quad c_1 \leftarrow B \text{ said } x_{12} \quad c_1 \leftarrow B \text{ said } x_{13} \\ \dots \\ c_n \leftarrow B \text{ said } x_{n1} \quad c_n \leftarrow B \text{ said } x_{n2} \quad c_n \leftarrow B \text{ said } x_{n3} \\ \\ F \leftarrow B \text{ said } x_1, B \text{ said } x'_1 \\ \dots \\ F \leftarrow B \text{ said } x_k, B \text{ said } x'_k \\ \\ \phi \leftarrow c_1, \dots, c_n \end{array}$$

$\Pi_B :$

$$\begin{array}{l} B \text{ to } A : x_1 \\ B \text{ to } A : x'_1 \\ \dots \\ B \text{ to } A : x_n \\ B \text{ to } A : x'_n \end{array}$$

Now we set  $I_1 = \phi$  and  $I_2 = F$ . Observe that  $\phi$  has a satisfying assignment iff there exists a  $\Delta_A$  such that  $\Pi_A \cup \Delta_A \vdash A$  knows  $\phi$ , but  $\Pi_A \cup \Delta_A \not\vdash A$  knows  $F$ . Hence the proof.  $\square$

**Theorem 7** (Invariant Checking - Upper Bound). *The invariant relation for DKAL is in coNP.*

*Proof.* Let  $\Delta = \{P \text{ from } Q : X(a, b, \dots, k) \mid P, Q \in \alpha\}$ , where

- $Q \text{ to } P : X(A, B, \dots, K) \leftarrow \text{body}$  is a rule in the policy of  $Q$  and  $(a, b, \dots, k)$  enumerate all possible instantiations of the variables  $(A, B, \dots, K)$ .
- $Q \text{ to } r : X(A, B, \dots, K) \leftarrow \text{body}$  is a rule in the policy of  $Q$  and  $(a, b, \dots, k)$  enumerate all possible instantiations of the variables in  $(A, B, \dots, K)$ .

We note that since the number of variables in a rule is fixed and all possible instantiations is upper bounded by the policy size,  $\Delta$  is poly-sized in the input.

We claim that a certificate of non-membership is a subset of  $\Delta$ . We prove this by arguing that given any set of **from** statements  $F$ , there is a set  $F' \in \Delta$  such that  $\bigcup_{\alpha} \Pi_{\alpha} \cup F \vdash I$  iff  $\bigcup_{\alpha} \Pi_{\alpha} \cup F' \vdash I$ . This is because

- Any **from** statement, that does not unify with any of the **to** rule heads, will never lead to additional knowledge.
- A **from** statement, that unifies with the head of a rule  $Q$  **to**  $P : X(A, B, \dots, K)$ , say  $P$  **from**  $Q : X(A', B', \dots, K')$ , leads to the same set of additional knowledge as does a set of statements in  $\Delta$ :  $P$  **from**  $Q : X(a, b, \dots, k)$  which enumerates all possible instantiations of the variables in  $(A', B', \dots, K')$ .
- A **from** statement, that unifies with the head of a rule  $Q$  **to**  $r : X(A, B, \dots, K)$ , say  $P$  **from**  $Q : X(A', B', \dots, K')$ , leads to the same set of additional knowledge as does a set of statements in  $\Delta$ :  $P$  **from**  $Q : X(a, b, \dots, k)$  which enumerates all possible instantiations of the variables in  $(A', B', \dots, K')$ .

Finally, the certificate of non-membership is a set  $\Delta' \in \Delta$ , such that  $\bigcup_{\alpha} \Pi_{\alpha} \cup \Delta' \vdash A$  **knows**  $I_1$ , but  $\bigcup_{\alpha} \Pi_{\alpha} \cup \Delta' \not\vdash A$  **knows**  $I_2$ .  $\square$

**Corollary 2** (Invariant Checking). *The invariant relation for DKAL is coNP-complete.*

## REFERENCES

- [1] Moritz Y. Becker, Cédric Fournet and Andrew D. Gordon, “SecPAL: Design and Semantics of a Decentralized Authorization Language”, Technical Report MSR-TR-2006-120, Microsoft Research, September 2006.
- [2] Moritz Y. Becker, Cédric Fournet and Andrew D. Gordon, “SecPAL: Design and Semantics of a Decentralized Authorization Language”, 20th IEEE Computer Security Foundations Symposium (CSF), 3–15, 2007.
- [3] Andreas Blass and Yuri Gurevich, “Two Forms of One Useful Logic: Existential Fixed Point Logic and Liberal Datalog”, Bulletin of the European Association for Theoretical Computer Science, Number 95 (June 2008), 164–182.
- [4] John DeTreville, “Binder, a Logic-Based Security Language”, in IEEE Symposium on Security and Privacy, 105–113, 2002.
- [5] Blair Dillaway, “A Unified Approach to Trust, Delegation, and Authorization in Large-Scale Grids”, Technical Paper, Microsoft Corporation, September 2006, <http://research.microsoft.com/en-us/projects/SecPAL/>, seen on Dec 24, 2008.
- [6] Yuri Gurevich and Itay Neeman, “DKAL: Distributed-Knowledge Authorization Language,” 21st IEEE Computer Security Foundations Symposium (CSF 2008), 149–162.
- [7] Yuri Gurevich and Itay Neeman, “DKAL: Distributed-Knowledge Authorization Language,” Microsoft Research Tech Report MSR-TR-2008-09, January 2008.
- [8] Yuri Gurevich and Itay Neeman, “DKAL 2.0,” (tentative title)
- [9] Ninghui Li, “Delegation Logic: A Logic-Based Approach to Distributed Authorization”, Ph.D. thesis, New York University, September 2000.
- [10] Ninghui Li, Benjamin N. Grosf and Joan Feigenbaum, “Delegation Logic: A Logic-Based Approach to Distributed Authorization”, ACM Trans. on Information and System Security (TISSEC) 6:1 (February 2003), 128–171.
- [11] Ninghui Li, William H. Winsborough, and John C. Mitchell, “Beyond Proof-of-Compliance: Safety and Availability Analysis in Trust Management”, in Proceedings of 2003 IEEE Symposium on Security and Privacy, 123–139, May 2003.

MICROSOFT RESEARCH REDMOND  
*E-mail address:* [gurevich@microsoft.com](mailto:gurevich@microsoft.com)

DEPARTMENT OF COMPUTER SCIENCE, STANFORD UNIVERSITY  
*E-mail address:* [arnab@cs.stanford.edu](mailto:arnab@cs.stanford.edu)