

GEM: a Distributed Goal Evaluation Algorithm for Trust Management

DANIEL TRIVELLATO*, NICOLA ZANNONE*, SANDRO ETALLE*,⁺

^{*}*Eindhoven University of Technology, Eindhoven, The Netherlands*

⁺*University of Twente, Enschede, The Netherlands*

(*e-mail*: {d.trivellato,n.zannone,s.etalles}@tue.nl)

submitted 1 January 2011; revised 1 January 2011; accepted 1 January 2011

Abstract

Trust management is an approach to access control in distributed systems where access decisions are based on policy statements issued by multiple principals and stored in a distributed manner. In trust management, the policy statements of a principal can refer to other principals' statements; thus, the process of evaluating an access request (i.e., a goal) consists of finding a "chain" of policy statements that allows the access to the requested resource. Most existing goal evaluation algorithms for trust management either rely on a centralized evaluation strategy, which consists of collecting all the relevant policy statements in a single location (and therefore they do not guarantee the confidentiality of intensional policies), or do not detect the termination of the computation (i.e., when all the answers of a goal are computed). In this paper we present GEM, a distributed goal evaluation algorithm for trust management systems that relies on function-free logic programming for the specification of policy statements. GEM detects termination in a completely distributed way without the need of disclosing intensional policies, thereby preserving their confidentiality. We demonstrate that the algorithm terminates and is sound and complete with respect to the standard semantics for logic programs.

KEYWORDS: Trust management, distributed goal evaluation, policy confidentiality

1 Introduction

The widespread availability of the Internet has led to a significant increase in the number of collaborations, services and transactions carried out over networks spanning multiple administrative domains (e.g., web services). Such collaborations are frequently characterized by the interaction of users and institutions (hereafter indistinctly referred to as *principals*) who do not know each other beforehand. For this reason, in such distributed settings, attribute-based approaches to access control are mostly preferred to identity-based solutions (Ellison et al. 1999). Consider, for instance, an international medical research project *Alpha* involving several companies worldwide. Project *Alpha* is funded and coordinated by the multinational pharmaceutical company *mc* which, among its tasks, appoints the partners of the project consortium. In this scenario, it is likely that company *mc* does not know the project members of each partner company personally, i.e., does not know their identity. Therefore, rather than on the identity of the project members, the policies regulating

the access to project's documents will be based on their attributes (e.g., project membership, specialization) and their relationships with other principals (e.g., partner companies, departments within a company).

Trust management is an approach to access control in distributed systems where access decisions are based on the attributes of principals, which are attested by digitally signed certificates called *digital credentials* (Blaze et al. 1996). Digital credentials (or simply credentials) are the digital counterpart of paper credentials. Credentials are defined and derived by means of policy statements that specify the conditions upon which a credential is issued, where conditions are in turn represented by credentials. A distinguishing ingredient of trust management is that all the principals in a distributed system are free to define such policy statements and determine where to store them. The set of policy statements defined by a principal forms the *policy* of that principal. In the scenario above, for instance, the rules of company *mc* dictating the conditions for the membership of a user to project *Alpha* (e.g., a Master degree in chemistry) form the policy of *mc*. These statements can be stored by *mc* or at another principal's location (e.g., by each partner company).

In trust management languages, policy statements are often expressed as Horn clauses (Li and Mitchell 2003) where each atom represents a credential, and is possibly annotated with the storage location of the statements defining the credential. Depending on the language, the location can be expressed implicitly (Czenko and Etalle 2007; Li et al. 2003) or explicitly (Alves et al. 2006; Becker 2005). While typically principals do not have direct access to each other policies, the statements of a principal can refer to other principals' policies, thereby delegating authority to them. For instance, assume that a hospital *h* authorizes the members of project *Alpha* certified by the local pharmaceutical company *c1* to access the (anonymized) medical records of its patients suffering from genetic diseases. The policies governing this scenario can be represented by the following clauses:

1. $\text{mayAccessMedRec}(h, X) \leftarrow \text{memberOfAlpha}(c1, X)$.
2. $\text{memberOfAlpha}(c1, X) \leftarrow \text{projectPartner}(mc, Y), \text{memberOfAlpha}(Y, X)$.
3. $\text{projectPartner}(mc, c2)$.
4. $\text{projectPartner}(mc, c3)$.
5. $\text{projectPartner}(mc, c4)$.
6. $\text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c1, X)$.
7. $\text{memberOfAlpha}(c2, \text{alice})$.
8. $\text{memberOfAlpha}(c3, \text{bob})$.
9. $\text{memberOfAlpha}(c4, \text{charlie})$.

where the first parameter of each atom denotes the location where the credential that the atom represents is defined. Here, hospital *h* relies on the policy statements of *c1* to determine who is authorized to access the hospital's medical records (clause 1); in turn, *c1* relies on the policy statements of *mc* and the partner companies appointed by *mc* for the definition of project *Alpha*'s members (clause 2). Therefore, the process of evaluating a request to access the hospital's medical records (i.e., a *goal*) consists of deriving a "chain" of policy statements that delegates the authority from hospital *h* (i.e., the resource owner) to the members of project *Alpha* (i.e., the authorized principals). This process, usually referred to as *credential chain discovery* (Li et al. 2003), can be addressed using *goal evaluation algorithms*.

Since in trust management policies are stored at different locations, goal evaluation algorithms require principals to disclose their policy statements to other principals to enable credential chain discovery. In particular, for a successful computation, principals need to disclose at least (part of) their *extensional policy*, that is, the credentials that can be derived from their policy and are required for an access decision. For example, suppose that hospital h wants to determine who can access the medical records of its patients. To compute the answers of this goal, it is clear that $c1$ has to disclose to h the credentials certifying all the project members. Most of the existing goal evaluation algorithms (e.g., (Czenko and Etalle 2007; Li et al. 2003)), however, rely on a centralized evaluation strategy and require principals to disclose also (part of) their *intensional policy*, i.e., the policy statement used to derive those credentials (e.g., clause 2 in the example policy).

We argue that one of the advanced desiderata of goal evaluation algorithms for trust management is that the amount of information about intensional policies that principals reveal to each other should be minimized. In fact, intensional policies might contain sensitive information about the relationships among principals, whose disclosure would leak valuable business information that can be exploited by other principals in the domain (e.g., rival companies) (Yu and Winslett 2003). For example, if $c1$'s policy was disclosed to other principals for evaluation (e.g., to hospital h), the involvement of mc in project *Alpha* along with the list of all project partners would be exposed. As a consequence, some competitors of mc could start investing on similar projects, or could try to get at the project members to acquire sensitive information and project results. Furthermore, the loss of confidentiality of intensional policies can result in attempts by other principals to influence the policy evaluation process (Stine et al. 2008), and allows adversaries to know what credentials they need to forge to illegitimately get access to a resource (Frikken et al. 2006).

To protect the confidentiality of intensional policies, it is necessary to design a completely distributed goal evaluation algorithm that discloses as few information on intensional policies as possible. Since bottom-up approaches to goal evaluation (e.g., fixpoint semantics (Park 1969), magic templates (Ramakrishnan 1991) and magic sets (Chen 1997)) require knowledge of all the policy statements that depend on a given credential, they do not represent an applicable solution to our problem. Hence, a top-down approach to goal evaluation needs to be employed. The design of a distributed top-down algorithm, however, requires addressing two main problems: (a) loop detection, and (b) termination detection. In addition, to reduce network overhead, the goal evaluation algorithm should attempt to decrease the number of messages that principals exchange.

Loops are formed when the evaluation of a goal leads to a new request for the same goal. In our scenario, for example, to determine the set of project members without disclosing its intensional policy to hospital h , $c1$ should first request the list of project partners to mc , and then the list of their project members to $c2$, $c3$ and $c4$. Since $c2$ in turn relies on the policy statements of $c1$ to determine its project members, $c2$ would pose the same request back to $c1$, forming a loop. Figure 1 shows the “call graph” originating from this sequence of requests. Intuitively, $c1$ should detect the loop and refrain from evaluating $c2$'s request, as doing so could lead to a non-terminating chain of requests. Even though in the example scenario loops could be avoided, for instance, by requiring a single company (e.g., mc) to define the set of project members, this cannot be guaranteed in distributed systems characterized by the absence of a coordinating principal. Examples of such scenarios include self-

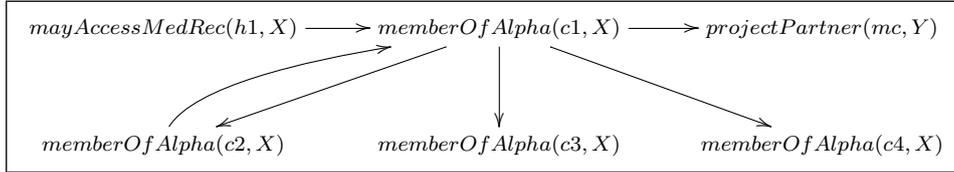


Fig. 1. Call Graph of the Evaluation of the Example Policies

organizing networks (Di Marzo Serugendo et al. 2004) and access control policies based on independent information sources (e.g., the Friend of a Friend - FOAF - project (FOA)).

Existing goal evaluation algorithms employ *tabling* techniques (Bry 1989; Chen and Warren 1996; Tamaki and Sato 1986; Vieille 1987) for the detection of loops. Although some of these algorithms resort to a distributed tabling strategy (e.g., (Damásio 2000; Hu 1997)), they rely on centralized data structures to detect termination - i.e., to detect when all the answers have been collected - thus leaking some policy information. In fact, the real challenge in designing a goal evaluation algorithm that does not disclose intensional policies lies in detecting termination distributedly. In the example, we have the following possible answer flow: *c3* returns *bob* as answer to *c1*, which forwards it to *h* and *c2*; *c4* returns *charlie* as answer to *c1*; *c1* sends *charlie* as additional answer to *h* and *c2*; *c2* returns *alice*, *bob*, and *charlie* as answer to *c1*, which sends *alice* as additional answer to *h* and *c2*. At this point, all the requests have been fully evaluated, but *c1* does not know whether *c2* will ever send additional answers. In other words, *c1* is waiting for *c2* to announce that its evaluation has terminated, and in turn *c2* is waiting for *c1* to announce that its evaluation has terminated. This situation is not acceptable in the context of access control, where a decision (positive or negative) always needs to be made. A few top-down goal evaluation algorithms are able to detect the termination of a computation distributedly (Alves et al. 2006; Zhang and Winslett 2008); however, they do not detect when the single goals within a computation are fully evaluated. In top-down goal evaluation, detecting when the evaluation of a goal has terminated is necessary to allow (a) for memory deallocation and (b) the use of negation, which is employed by some trust management systems to express non-monotonic constraints (e.g., separation of duty) (Czenko et al. 2006; Dong and Dulay 2010).

Finally, another non-trivial issue in designing a distributed goal evaluation algorithm is determining when a principal should send the answers to a request. The simplest solution is to force each principal to send an answer as soon as the principal has computed it, as done, for example, in (Alves et al. 2006). This is, however, suboptimal from the viewpoint of network overhead; in the example above, *c1* eventually sends three distinct messages to *h* and *c2*, one for each answer. A more network-efficient solution would be for *c1* to wait for the answers from *c3* and *c4* before sending its answers to the other principals. A naïve “wait” mechanism, on the other hand, might cause deadlocks. For instance, if *c1* also waits for *c2*’s answers, the computation deadlocks. In a trust management system, where network latency is likely to be a bigger bottleneck than computational power, it is preferable to have a mechanism that allows principals to wait until they collect the maximum possible set of answers before sending them to the requester, while avoiding deadlocks. Even though this solution may delay the identification of the answers of ground goals (i.e., goals

expecting a single answer), the “superfluous” computed answers might become relevant for the evaluation of other goals, reducing the delay in future computations.

In this paper we present GEM, a goal evaluation algorithm for trust management systems that addresses all the above mentioned problems. In GEM, policy statements are expressed as function-free logic programming clauses; each statement is stored by the principal defining it. GEM computes the answers of a goal in a completely distributed way without disclosing intensional policies of principals, thereby preserving their confidentiality. The algorithm deals with loops in three steps: (1) detection, (2) processing, and (3) termination. To enable loop detection, we employ a distributed tabling strategy and associate an identifier to each request for the evaluation of a goal. After its detection, a loop is processed by iteratively evaluating the goals in the loop until a fixpoint is reached, i.e., no more answers of the goals in the loop are computed, at which point their evaluation is terminated. This three-steps approach enables GEM to detect both when the whole computation has terminated, and when the single goals within a computation are fully evaluated, allowing for the use of non-monotonic constraints in policies. In addition, by exploiting the information stored in the table of a goal, principals are able to delay the response to a request until a “maximal” set of answers of the goal has been computed without running the risk of deadlocks. We demonstrate that GEM terminates and is sound and complete with respect to the standard semantics for logic programs.

The remainder of the paper is structured as follows. Section 2 presents preliminaries on logic programming and SLD resolution. Section 3 introduces GEM and its implementation. Section 4 demonstrates the soundness, completeness and termination of the algorithm, and discusses what information is disclosed by GEM during the evaluation of a goal. Section 5 presents the results of experiments conducted to evaluate the performance of GEM. A possible extension of GEM to deal with negation is presented in Section 6. Section 7 discusses related work. Finally, Section 8 concludes and gives directions for future work.

2 Preliminaries on Logic Programming

In this section we revisit the concepts of logic programming (Apt 1990) that are relevant to this paper. In particular, we review function-free logic programs.

An *atom* is an object of the form $p(t_1, \dots, t_n)$ where p is an n -ary predicate symbol and t_1, \dots, t_n are terms (i.e., variables or constants). An atom is *ground* if t_1, \dots, t_n are constants. A *clause* is an expression of the form $H \leftarrow B_1, \dots, B_n$ (with $n \geq 0$), where H is an atom called *head* and B_1, \dots, B_n (called *body*) are atoms. If $n = 0$, the clause is a *fact*. A *program* is a finite set of clauses. We say that an atom A is *defined in the program* P if and only if there is a clause in P that has an atom A' in its head such that A and A' are unifiable. Finally, a *goal* is a clause with no head atom, i.e., a clause of the form $\leftarrow B_1, \dots, B_n$. Without loss of generality, in this paper we restrict to goals with $0 \leq n \leq 1$, that is, consisting of at most one atom. The empty goal is denoted by \emptyset .

SLD resolution (Selective Linear Definite clause resolution) (Kowalski 1974) is the standard operational semantics for logic programs. In this paper, we refer to SLD resolution with leftmost selection rule (extending the algorithm to an arbitrary selection rule is trivial). Computations are constructed as sequences of “basic” steps. Consider a goal $G_0 = \leftarrow B_1, \dots, B_n$ and a clause c in a program P . Let $H \leftarrow B'_1, \dots, B'_m$ be a variant of c vari-

able disjoint from $\leftarrow B_1, \dots, B_n$. Let B_1 and H unify with *most general unifier* (mgu) θ . The goal $G_1 = \leftarrow (B'_1, \dots, B'_m, B_2, \dots, B_n)\theta$ is called a *resolvent of G_0 and c with selected atom B_1 and mgu θ* . An SLD *derivation step* is denoted by $G_0 \xrightarrow{\theta} G_1$. Clause $H \leftarrow B'_1, \dots, B'_m$ is called *input clause*, and atom B_1 is called the *selected atom* of G_0 .

An SLD derivation is obtained by iterating derivation steps. The sequence $\delta := G_0 \xrightarrow{\theta_1} G_1 \xrightarrow{\theta_2} \dots \xrightarrow{\theta_n} G_n \xrightarrow{\theta_{n+1}} \dots$ is called a *derivation of $P \cup \{G_0\}$* , where at every step the input clause employed is variable disjoint from the initial goal G_0 and from the substitutions and the input clauses used at earlier steps. Given a program P and a goal G_0 , SLD resolution builds a search tree for $P \cup \{G_0\}$, called (*derivation*) *tree* of G_0 , whose branches are SLD derivations of $P \cup \{G_0\}$. Any selected atom in the SLD resolution of $P \cup \{G_0\}$ is called a *subgoal*. SLD derivations can be finite or infinite. If $\delta := G_0 \xrightarrow{\theta_1} \dots \xrightarrow{\theta_n} G_n$ is a finite prefix of a derivation, we say that $\theta = \theta_1 \dots \theta_n$ is a *partial derivation* and θ is a *partial computed answer substitution* of $P \cup \{G_0\}$. If δ ends with the empty goal \emptyset , θ is called *computed answer substitution* (c.a.s.). Let $G_0 = \leftarrow B_1$. Then, we also call θ a *solution* of G_0 and $B_1\theta$ an *answer* of G_0 . The length of a (partial) derivation δ , denoted by $len(\delta)$, is the number of derivation steps in δ .

The most commonly employed technique to prevent infinite derivations is *tabling* (Bry 1989; Chen and Warren 1996; Guo and Gupta 2001; Shen et al. 2001; Tamaki and Sato 1986; Vieille 1987; Zhou and Sato 2003). Given a goal G_0 consisting of an atom defined in a program P , tabling-based goal evaluation algorithms create a table for each (sub)goal in the SLD resolution of $P \cup \{G_0\}$, to keep track of the previously evaluated goals and thus avoid the reevaluation of a subgoal. Tabling algorithms differ mainly in the data structures employed for the evaluation of goal G_0 . Linear tabling (Shen et al. 2001; Zhou and Sato 2003) and DRA (Guo and Gupta 2001), for instance, evaluate G_0 by building a single derivation tree of G_0 . In SLG resolution (Chen and Warren 1996), on the other hand, goal G_0 is evaluated by producing a forest of (partial) derivation trees, one for each subgoal in the resolution of $P \cup \{G_0\}$. In SLG, the evaluation of G_0 starts by ordinary resolution with the clauses in P ; as in SLD, a subgoal G_1 is selected in a resolvent of G_0 . If a tree for a variant of G_1 already exists, G_1 is added to the set of *consumers* of the corresponding table. Otherwise, a tree for G_1 is created. When a new answer of a subgoal is found, it is stored in the respective table and it is propagated to its consumer subgoals. The evaluation of a goal by means of a forest of derivation trees proposed by SLG resolution is at the basis of the distributed goal evaluation algorithm proposed in this paper.

3 The GEM Algorithm

In this section we first introduce some definitions and basic assumptions underlying our work. Then, we present GEM and discuss its implementation.

3.1 Definitions and Assumptions

Similarly to other works on trust management (e.g., (Li et al. 2003; Alves et al. 2006)), we consider policy statements expressed as function-free logic programming clauses. As in most trust management systems, policy statements are stored at different locations: each

location is controlled by a principal who is responsible for defining and evaluating the policy statements at that location. We assume a one-to-one correspondence between locations and principals; accordingly, we use a principal's identifier to refer to the location she controls. To represent the location where a policy statement is stored, we require every atom to have the form $p(loc, t_1, \dots, t_n)$, where loc is a mandatory term that represents the location where the atom is defined, and t_1, \dots, t_n are terms. For instance, $p(bob, \dots)$ refers to p as defined by Bob and thus stored at Bob's location.

Let a be a principal in the trust management system. We call the set of policy statements defined by a the *local trust management policy* (or simply the *policy*) of principal a . The set of clauses with non-empty body in a 's policy is the *intensional policy* of a , while the set of facts that can be derived from principal a 's policy forms the *extensional policy* of a . The set of all the local policies in the trust management system is called the *global trust management policy*.

Since we consider the confidentiality of intensional policies to be a main concern in trust management systems, we assume that principals do not have access to the policies at other principals' locations. As a consequence, the answers of a goal cannot be derived by building the derivation tree of the goal as done by SLD resolution, as this might involve input clauses defined by different principals. Similarly to SLG resolution, in GEM a principal computes the answers of a goal defined in her policy by building the *partial derivation tree* of the goal. Differently from a derivation tree, in the partial derivation tree of a goal G only the first derivation step is obtained by resolution with the clauses defining G ; all the subsequent steps are by substitution with the solutions of the subgoals of G .

Definition 1

Let $G \leftarrow A$ be a goal and P_A be the policy in which A is defined. A *partial derivation tree* of G is a derivation tree with the following properties:

- the root is the node $(A \leftarrow A)$;
- there is a derivation step $(A \leftarrow A) \xrightarrow{\theta} (A \leftarrow B_1, \dots, B_n)\theta$, where $(A \leftarrow A)$ is the root, iff there exists a clause $H \leftarrow B_1, \dots, B_n$ in P_A (renamed so that it is variable disjoint from A) s.t. A and H unify with $\theta = mgu(A, H)$;
- let $(A \leftarrow B_1, \dots, B_n)$ be a non-root node, and Ans be a set of answers of goal $\leftarrow B_1$; for each answer $B'_1 \in Ans$ (renamed so that it is variable disjoint from B_1) there is a derivation step $(A \leftarrow B_1, \dots, B_n) \xrightarrow{\theta} (A \leftarrow B_2, \dots, B_n)\theta$, where $\theta = mgu(B_1, B'_1)$;
- for each branch $(A \leftarrow A) \xrightarrow{\theta_0} (A \leftarrow B_1, \dots, B_n)\theta_0 \xrightarrow{\theta_1} \dots \xrightarrow{\theta_n} (A \leftarrow \emptyset)\theta_0\theta_1 \dots \theta_n$, we say that $A\theta$ (with $\theta = \theta_0\theta_1 \dots \theta_n$) is an answer of G using clause $H \leftarrow B_1, \dots, B_n$. \square

Notice that in order to enable the evaluation of an atom B_1 in the partial derivation tree of goal G , the location where B_1 is defined must be known by the principal evaluating G . A straightforward solution for guaranteeing that this requirement is satisfied would be to impose the location parameter of each atom in a policy to be ground at policy definition time. This, however, would limit the constraints that a principal can express. Consider, for instance, clause 2 on page 2. In the clause, the location parameter of the body atom *memberOfAlpha*(Y, X) is determined at runtime based on the answers of the first body

Level	Disclosed Information	Level	Disclosed Information
E0	None	I0	None
E1	Answers of a goal	I1	Part of the dependency graph
		I2	Full dependency graph
		I3	Clauses

(a) Extensional Policy Confidentiality Levels (b) Intensional Policy Confidentiality Levels

Fig. 2. Classification of Goal Evaluation Algorithms in Terms of Disclosed Policy Information

atom $projectPartner(mc, Y)$. Therefore, rather than relying on a “static” safety condition, we require the location parameter of an atom to be ground *when the atom is selected for evaluation*. If this is not the case, the computation *flounders*. A discussion on how to write flounder-free programs and queries is orthogonal to the scope of this paper. Here, we just mention that there exist well-established techniques based on *modes* (Apt and Marchiori 1994) which guarantee that certain parameters of an atom are ground when the atom is selected for evaluation.

Finally, we define a classification criteria for goal evaluation algorithms based on disclosed policy information. We will use such criteria to compare GEM with the existing algorithms (see Section 7). The classification criteria consists of two elements: an extensional policy confidentiality level and an intensional policy confidentiality level. Intuitively, the first characterizes algorithms in terms of how much information about extensional policies they disclose during goal evaluation, while the second refers to the disclosure of intensional policies. Confidentiality levels define an increasing scale used to characterized from the most conservative approaches where no policy information is disclosed, to the least confidentiality-preserving solutions which disclose respectively extensional and intensional policies in full. The extensional and intensional confidentiality levels are presented in Figure 2. In a goal evaluation algorithm classified as E1-I2, for example, principals disclose to a requester all the answers of the requested goals. In addition, all the dependencies among the goals in the global policy are disclosed to the principals responsible for goal evaluation.

3.2 Intuition

In this section we describe how GEM computes the answers of a goal. Given a goal G , GEM computes the answers of G by evaluating one branch of its partial derivation tree at a time; this may involve the generation of evaluation requests for subgoals that are processed by different principals at different locations. When all the answers from each branch of the tree of G have been computed, they are sent to the principal(s) that requested the evaluation of G . G is *completely evaluated* when no more answers of G can be computed.

To illustrate how GEM works in detail, we consider the scenario presented in Section 1, where several pharmaceutical companies collaborate in the research project *Alpha*. However, we slightly modify the global (trust management) policy to better focus on the algorithm’s features. In particular, we assume that company $c1$ already knows which are the partner companies in project *Alpha*, without needing to request them to mc , and we reduce the partner companies to $c2$ and $c3$ only. Furthermore, we consider a research institute

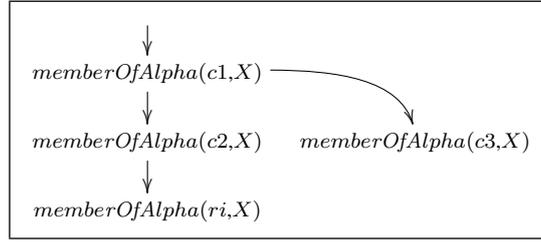


Fig. 3. Call Graph of the Example Global Policy

ri that works on project *Alpha* in partnership with company $c2$. As a result, we have the following global policy:

1. $memberOfAlpha(c1,X) \leftarrow memberOfAlpha(c2,X)$.
2. $memberOfAlpha(c1,X) \leftarrow memberOfAlpha(c3,X)$.
3. $memberOfAlpha(c2,X) \leftarrow memberOfAlpha(ri,X)$.
4. $memberOfAlpha(c2,alice)$.
5. $memberOfAlpha(c3,bob)$.

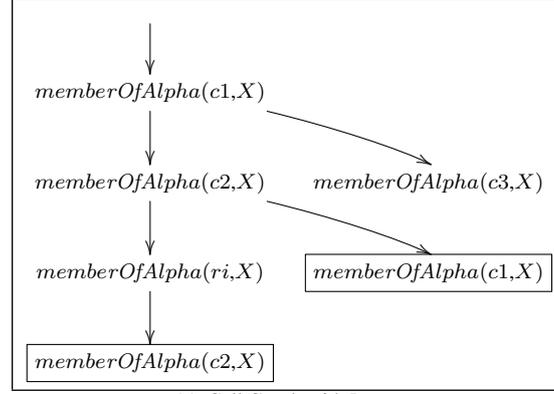
Recall that the first parameter of the atom in the head indicates the principal storing and evaluating a clause: clauses 1 and 2 are evaluated by $c1$, clauses 3 and 4 by $c2$, and clause 5 by $c3$. Suppose that hospital h sends to company $c1$ a request for (the evaluation of) goal $\leftarrow memberOfAlpha(c1,X)$ (for a matter of readability, from here on we omit the \leftarrow symbol when referring to a goal $\leftarrow A$, and we simply refer to it as A). Figure 3 shows the *call graph* of the evaluation of $memberOfAlpha(c1,X)$ with respect to the example global policy. A call graph is a directed graph where nodes represent goals and edges connect each goal to its subgoals (Leuschel et al. 1998). In other words, edges represent (evaluation) requests.

GEM performs a depth-first computation. When $c1$ receives the initial goal, it evaluates the first applicable clause in its policy (i.e., clause 1) and sends a request for $memberOfAlpha(c2,X)$ to $c2$. In turn, $c2$ sends a request for $memberOfAlpha(ri,X)$ to ri . ri does not have any clause applicable to $memberOfAlpha(ri,X)$ and returns an empty set of answers to $c2$. $c2$ evaluates the next applicable clause (i.e., $memberOfAlpha(c2,alice)$), which is a fact. Since $c2$ does not have any other clause to evaluate, it sends the computed answer to $c1$. $c1$ applies the next clause (clause 2) and sends a request for $memberOfAlpha(c3,X)$ to $c3$, that returns answer $memberOfAlpha(c3,bob)$ to $c1$ after applying clause 5. At this point, $memberOfAlpha(c1,X)$ is completely evaluated and $c1$ sends answers $memberOfAlpha(c1,alice)$ and $memberOfAlpha(c1,bob)$ to hospital h .

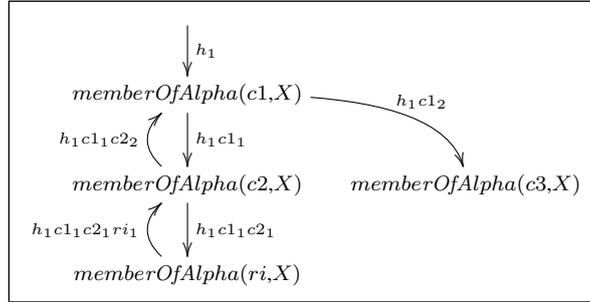
The evaluation of a subgoal of a goal G , however, may lead to new requests for G , forming a *loop*. In our scenario, this reflects the “sharing” of project members among partner companies. Consider, for instance, the global policy above with the following two additional clauses, stored by $c2$ and ri respectively:

6. $memberOfAlpha(c2,X) \leftarrow memberOfAlpha(c1,X)$.
7. $memberOfAlpha(ri,X) \leftarrow memberOfAlpha(c2,X)$.

The new call graph is shown in Figure 4(a). Now, when ri receives the request for *mem-*



(a) Call Graph with Loops



(b) Compact Call Graph with Request Identifiers and Loops

Fig. 4. Call Graph of the Example Global Policy with Loops

$berOfAlpha(ri,X)$ from $c2$, it applies clause 7 and sends a request for $memberOfAlpha(c2,X)$ back to $c2$, forming a loop. Similarly, the evaluation of clause 6 by $c2$ leads to another loop. The requests forming a loop are identified by boxed atoms in Figure 4(a). Formally, a loop is defined as follows.

Definition 2

Let C be the call graph of the evaluation of a goal G with respect to a global policy P . A *loop* is a maximal subgraph of C consisting of goals G_1, \dots, G_k such that for each $G_i \in \{G_1, \dots, G_k\}$ there exists a path that leads from G_1 to G_i and from G_i to (a variant of) G_1 . Then, we say that goals G_1, \dots, G_k are *involved* in the loop. \square

Intuitively, requests forming a loop should not be further evaluated. However, in the example above $c1$ and $c2$ cannot detect whether a request forms a loop, as in a distributed system several independent requests for the same goal can occur. In most of the existing goal evaluation algorithms (e.g., (Chen and Warren 1996; Damásio 2000; Li et al. 2003)), loop detection (and termination) is made possible by the system’s “global view” on the derivation process. For example, centralized goal evaluation algorithms such as SLG (Chen and Warren 1996) and RT (Li et al. 2003) identify loops by observing goal dependencies respectively in the call stack and in the call graph of the global policy. In a similar way, the distributed algorithm proposed by Damásio (Damásio 2000) requires the dependency graph of the global policy to be known to all principals in the trust management system.

Such global view, however, implies the loss of policy confidentiality. GEM detects loops and their termination in a completely distributed way without resorting to any centralized data structure. In GEM, loops are handled in three steps: (1) detection, (2) processing, and (3) termination.

Loop Detection. Loops are detected by dynamically identifying Strongly Connected Components (SCCs). A SCC is a set of mutually dependent subgoals. More precisely, a set of goals G_1, \dots, G_k is part of a SCC if for each $G_i \in \{G_1, \dots, G_k\}$ there exists a goal $G_j \in \{G_1, \dots, G_{i-1}, G_{i+1}, \dots, G_k\}$ such that G_i and G_j are involved in a common loop. To enable the identification of SCCs, we assign to each request a unique identifier from an *identifier domain*.

Definition 3

An *identifier domain* is a triple $\langle I, \sqsubset, \hookrightarrow \rangle$, where:

- I is a set of sequences of characters called *identifiers*;
- \sqsubset is a partial order on the identifiers in I . Given two identifiers $id_1, id_2 \in I$ s.t. $id_1 \sqsubset id_2$, we say that id_1 is *lower* than id_2 , and id_2 is *higher* than id_1 ;
- \hookrightarrow is a partial order on the identifiers in I . Given two identifiers $id_1, id_2 \in I$ s.t. $id_1 \hookrightarrow id_2$, we say that id_2 is *side* of id_1 .
- The following property holds: $\forall id_1, id_2, id_3, id_4 \in I$ if $id_1 \sqsubset id_2, id_3 \sqsubset id_4$, and $id_2 \hookrightarrow id_4$, then $id_1 \hookrightarrow id_3$. \square

Intuitively, \sqsubset defines a top-down ordering and \hookrightarrow defines a left-to-right ordering with respect to the call graph of the global policy. In other words, \sqsubset reflects the order in which the subgoals in a branch of the graph are evaluated, whereas \hookrightarrow reflects the order in which the branches are inspected. Several identifier domains can be employed whose identifiers respect these partial orders (e.g., based on alphanumeric ordering). For the sake of simplicity, in the sequel we consider identifiers obtained as follows: given a request for a goal G with identifier id_0 , the identifier of the request for a subgoal G_1 of G has the form id_0s_1 , denoting the concatenation of id_0 with a sequence of characters s_1 . Then, \sqsubset is a prefix relation, and we have that $id_0s_1 \sqsubset id_0$. Ordering \hookrightarrow is a partial order on the strings composing the identifiers. For example, consider another subgoal G_2 of G with identifier id_0s_2 , which is evaluated after G_1 . Then, we have that $id_0s_1 \hookrightarrow id_0s_2$. Even though identifiers from this domain leak some policy information (see Section 4.2 for more details), they allow for an easy visualization of the relationships among identifiers. When applying GEM in practice, more confidentiality-preserving identifier domains can be employed.

A loop is detected when a principal receives a request with identifier id_2 for a goal G such that a request id_1 for a variant of G has been previously received and $id_2 \sqsubset id_1$. Accordingly, we call request id_2 a *lower request* for G , while request id_1 is a *higher request* for G . We use the identifier of the higher request for G , id_1 , as the loop identifier. Goal G is called the *coordinator* of the loop. An SCC may contain several loops. Given two loops with identifiers id_1 and id_2 , we say that loop id_2 is *lower* than loop id_1 if $id_2 \sqsubset id_1$. The coordinator of the highest loop of the SCC (i.e., the loop with the highest identifier) is called the *leader* of the SCC.

Figure 4(b) represents a compact version of the call graph in Figure 4(a), where loop coordinators are depicted only once. In addition, in the call graph in Figure 4(b), edges are

labeled with the corresponding request identifier. In the remainder of the paper, we concatenate the identifier of a request for a goal evaluated by company $c1$ with meta-variables of the form $c1_i$. Thus, for instance, $c1_1$ and $c1_2$ are two distinct sequences of characters generated by $c1$. In the figure, identifiers h_1 , h_1c1_1 , h_1c1_2 , and $h_1c1_1c2_1$ identify higher requests for goals $memberOfAlpha(c1,X)$, $memberOfAlpha(c2,X)$, $memberOfAlpha(c3,X)$, and $memberOfAlpha(ri,X)$ respectively; identifiers $h_1c1_1c2_1ri_1$ and $h_1c1_1c2_2$ identify lower requests for goals $memberOfAlpha(c2,X)$ and $memberOfAlpha(c1,X)$ respectively. Goals inherit the ordering associated with the identifier of their higher request. Therefore, in Figure 4(b) goal $memberOfAlpha(c1,X)$ is higher than $memberOfAlpha(c2,X)$, $memberOfAlpha(c3,X)$, and $memberOfAlpha(ri,X)$. Goals $memberOfAlpha(c2,X)$ and $memberOfAlpha(c1,X)$ are the coordinators of loops h_1c1_1 and h_1 respectively. Loop h_1c1_1 is lower than loop h_1 , which is the highest loop of the SCC; therefore, goal $memberOfAlpha(c1,X)$ is the leader of the SCC. Looking at the identifier of the lower requests, $c1$ and $c2$ can determine the subgoals involved in the loop, which are $memberOfAlpha(c2,X)$ and $memberOfAlpha(ri,X)$ respectively.

Loop Processing. When a loop is detected, GEM sends the answers of the coordinator already computed to the requester of the lower request together with a notification about the loop. The loop is then processed iteratively as follows: in turn, each principal (a) processes the received answers, (b) “freezes” the evaluation of the subgoal involved in the loop and evaluates other branches of the partial derivation tree of the locally defined goal. Then, when all branches have been evaluated, (c) the new answers are sent to the requester of the higher request with a notification about the loop. We call the execution of actions (a), (b), and (c) a *loop iteration step*.

Definition 4

Let G be a goal, and G_1, \dots, G_k be the subgoals of G s.t. G, G_1, \dots, G_k are involved in a loop id . A *loop iteration step* for goal G is a three-phases process in which the principal evaluating G :

- (a) Receives a set of answers of the subgoals G_1, \dots, G_k of G .
- (b) Evaluates all the nodes in the partial derivation tree of G whose selected atom is not involved in a loop.
- (c) Sends the newly derived answers of G to the requester of G . □

The definition above applies to all the goals involved in a loop but the loop coordinator. The loop iteration step for a loop coordinator differs in the order in which the phases occur. In particular, for the coordinator phase (c) precedes (a) and (b), and the latter two are executed only after a loop iteration step for the other goals in the loop has been performed. In other words, the processing of the coordinator occurs only after all the other goals in the loop have been processed. This is because the coordinator, being the “highest goal” in the loop, is assigned the task of overseeing its processing. More precisely, it is in charge of starting (phase (c)) a new *loop iteration* whenever the answers of its subgoals lead to new answers of the coordinator (computed in phases (a) and (b)), i.e., until a fixpoint is reached. This difference is reflected in the definition below.

Definition 5

Let G_1, \dots, G_k be the goals involved in a loop id_1 s.t. goal G_1 is the loop coordinator. A *loop iteration* is a process where:

1. The answers of the coordinator G_1 that have not been previously sent are sent to the requesters of the lower requests for G_1 (phase (c) of the loop iteration step for the coordinator).
2. For each $G_i \in \{G_2, \dots, G_k\}$ a loop iteration step for G_i is performed, s.t. for each $G_j \in \{G_2, \dots, G_k\}$, if G_j is lower than G_i then the loop iteration step for G_j is executed before the loop iteration step for G_i .
3. The principal evaluating G_1 receives a set of answers of the subgoals of G_1 involved in loop id_1 (phase (a) of the loop iteration step for G_1). All the nodes in the partial derivation tree of G_1 whose selected atom is not involved in a loop are processed (phase (b) of the loop iteration step). \square

If the processing of the received answers leads to new answers of the coordinator, these new answers are sent to the requesters of lower requests, starting a new iteration. Otherwise, a fixpoint has been reached (i.e., all possible answers of the goals in the loop have been computed) and the answers of the coordinator are sent to the requester of the higher request. Notice that a goal in a higher loop may eventually provide new answers to a goal in a lower loop: the fixpoint for a loop must be recalculated every time new answers of its coordinator are computed.

In the example (Figure 4(b)), when $c2$ identifies loop h_1c1_1 , it informs ri that they are both involved in loop h_1c1_1 . Since ri has no more clauses to evaluate, it returns an empty set of answers to $c2$ notifying it that $memberOfAlpha(ri, X)$ is in loop h_1c1_1 . The further evaluation of $memberOfAlpha(c2, X)$ leads to the identification of loop h_1 and to a new answer $memberOfAlpha(c2, alice)$, which is sent first to ri in the context of loop h_1c1_1 . In turn, ri computes answer $memberOfAlpha(ri, alice)$ and sends it to $c2$. Now, a fixpoint for loop h_1c1_1 has been reached and $c2$ sends $memberOfAlpha(c2, alice)$ to $c1$ notifying it that $memberOfAlpha(c2, X)$ is in loop h_1 . Notice that $memberOfAlpha(c2, X)$ is also in loop h_1c1_1 , but since this loop does not involve $c1$, $c1$ is not notified of it. Next, $c1$ computes answers $memberOfAlpha(c1, alice)$ and $memberOfAlpha(c1, bob)$ (the latter being found through the evaluation of $memberOfAlpha(c3, X)$), and sends them to $c2$ in the context of loop h_1 . In turn, $c2$ computes $memberOfAlpha(c2, bob)$. Now, $c2$ has to find a fixpoint for loop h_1c1_1 given the new answer before proceeding with the evaluation of loop h_1 . It is worth noting that ri is not aware of loop h_1 . This is because loop notifications are only transmitted to higher requests (except for the lower request that has formed the loop).

Loop Termination. The termination of the evaluation of all the goals in an SCC is initiated by the principal handling the leader of the SCC when a fixpoint for the loop it coordinates has been reached. In the example, when the answers of $memberOfAlpha(c2, X)$ do not lead to new answers of $memberOfAlpha(c1, X)$, $c1$ informs $c2$ (which in turn informs ri) that the evaluation of $memberOfAlpha(c1, X)$ is terminated and sends answers $memberOfAlpha(c1, alice)$ and $memberOfAlpha(c1, bob)$ to h .

Side Requests. So far we have only considered “linear” loops, i.e., loops formed by lower requests. However, higher requests can also lead to a loop. Consider, for instance, the following additional clause stored by company $c3$:

$$8. \text{memberOfAlpha}(c3, X) \leftarrow \text{memberOfAlpha}(ri, X).$$

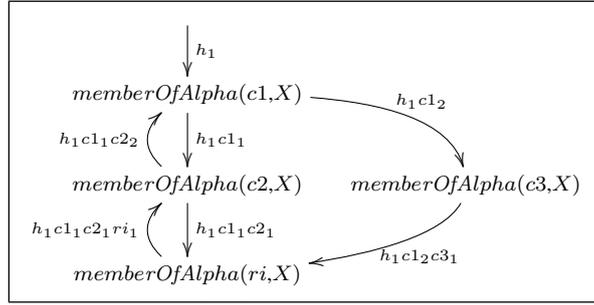
The new (compact) call graph is shown in Figure 5(a). Now, the evaluation of goal $\text{memberOfAlpha}(c3, X)$ by $c3$ leads to a request for goal $\text{memberOfAlpha}(ri, X)$, which is involved in loop h_1c1_1 . ri can identify that the request originates from the evaluation of a goal in the same SCC as $\text{memberOfAlpha}(ri, X)$, since the request identifier $h_1c1_2c3_1$ is side of the identifier $h_1c1_1c2_1$ of the initial request for $\text{memberOfAlpha}(ri, X)$ (i.e., $h_1c1_1c2_1 \hookrightarrow h_1c1_2c3_1$). However, it cannot identify the loop in which the goal evaluated by $c3$ is involved. This is because loop notifications are only transmitted to higher goals, and thus ri is not aware of loop h_1 . We refer to the request from $c3$ as *side request*, and we call $\text{memberOfAlpha}(c3, X)$ a *side goal*.

The main problem with side requests is that it is difficult to determine when they should be responded to. For example, if ri sends answers to $c3$ at every iteration of loop h_1c1_1 , $c1$ would not know when to stop waiting for answers from $c3$ (since $c1$ does not know on which goals $\text{memberOfAlpha}(c3, X)$ depends). On the other hand, ri cannot wait until a fixpoint is reached for loop h_1c1_1 , since only $c2$ (the principal handling the coordinator) is aware of that. To enable the detection of termination, however, a side request should be responded to only when a fixpoint is computed for all the loops lower than the loop in which the side goal is involved.

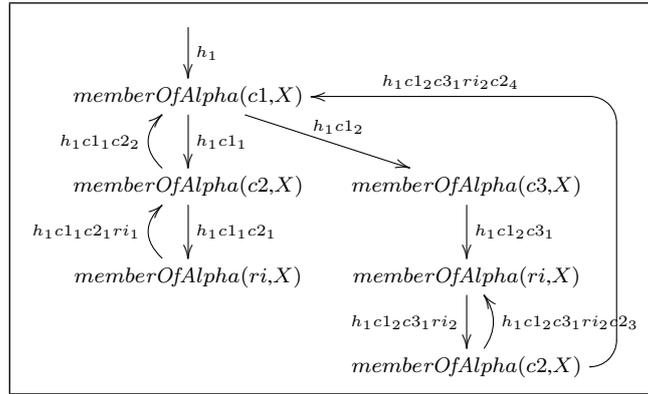
A simple yet effective solution to this problem is to treat a side request for a goal as a “new” request (i.e., a request for a goal that has not yet been evaluated) and to reevaluate the goal. Accordingly, when a side request is received, GEM creates a new partial derivation tree for the goal and proceeds with its evaluation. This corresponds to inspecting multiple times some branches of the call graph of the program (Figure 5(b)); however, it allows us to obtain a call graph formed only by linear loops which, as shown previously in this section, can be successfully evaluated by GEM. Notice that, despite in the unfolded graph in Figure 5(b) some nodes and edges are duplicated with respect to the folded graph in Figure 5(a), the flow of answers among goals is equivalent. This can be easily seen by the fact that edges connect the same nodes in both call graphs. In Section 4 (Theorem 3) we show that a call graph is always finite.

Even though possible in theory, we expect side requests not to occur frequently in the evaluation of a global policy. For instance, no side request was present in any of the example policies presented in the literature. Therefore, we believe the overhead imposed by the proposed solution to be negligible in practice. Nevertheless, we point out that the size of the unfolded graph is in the worst case exponential with respect to the number of nodes in the original graph. The reevaluation of side requests, in fact, resembles the approach adopted by SLD resolution, which reevaluates every goal encountered during a computation. However, thanks to our ability to detect loops, the number of goals evaluated by GEM during a computation is never higher than the number of goals that would be evaluated using SLD resolution.

Since we treat side requests in the same way as new requests, the partial order \hookrightarrow is



(a) Call Graph with Side Request



(b) Unfolded Call Graph

Fig. 5. Call Graphs for Side Requests

not exploited by the version of GEM presented here. An alternative solution that prevents the reevaluation of side requests and thus requires their identification could be achieved by transmitting loop notifications to the requesters of both higher and lower requests, so that all the principals evaluating a goal in an SCC would be aware of the identifiers of all the loops in the SCC. Furthermore, when a fixpoint for a loop is reached, the principal handling the loop coordinator should start a new loop iteration to inform all the goals in the loop that the fixpoint has been reached. These two modifications would enable principals that receive a side request to know (a) in which loop the side request is involved, and (b) when to reply to the side request, that is, when all the loops lower than the one in which the side goal is involved have reached a fixpoint. Given the complexity of the solution, in this paper we present only the “basic” version of GEM, which reevaluates side requests; the implementation of the described optimization is subject of future work.

To conclude, we point out that evaluating every higher request for a goal that is not yet completely evaluated is fundamental for enabling loop detection. Consider, for instance, a global policy consisting only of clauses 1 and 6 on page 9. Assume that hospital h issues at the same time a request for goals $memberOfAlpha(c1,X)$ and $memberOfAlpha(c2,X)$ with identifier id_1 and id_2 respectively. The evaluation of the request by principals $c1$ and $c2$ leads to a higher request for goals $memberOfAlpha(c2,X)$ and $memberOfAlpha(c1,X)$ respectively. When these second higher requests are received, a partial derivation tree for the requested goals already exists. If the requests were not further processed, the loop identi-

fication would not be possible as no lower request would be issued, and the computation would deadlock. Therefore, both initial requests must be processed independently. This “problem” is common to all the distributed goal evaluation algorithms whose termination detection exploits request identifiers (e.g., (Alves et al. 2006)). Even though relatively simple solutions to this problem can be found (e.g., using timestamped requests, where only the oldest is evaluated), in this paper we focus on a “basic” solution for distributed goal evaluation, and do not address efficiency-related issues.

3.3 Implementation

Here, we introduce the data structures and procedures used by GEM to evaluate a goal.

Data Structures. In GEM, principals communicate by exchanging *request* and *response* messages. We rely on blocking communication, that is, whenever a principal a sends (respectively receives) a request or response message, no other operation is performed by a until the sending (resp. receipt) process is completed. In addition, we assume that a message sent by principal a to a principal b is always received (once) by principal b .

Definition 6

A *request* is a triple $\langle id, req, G \rangle$, where:

- id is the request identifier;
- req is the principal issuing the request, called *requester*;
- G is a goal $\leftarrow p(loc, t_1, \dots, t_n)$, where loc is a constant. □

A *request* is an enquiry issued by principal req for the evaluation of goal G . Each *request* is uniquely identified by an identifier id and is sent to the principal defining G .

Definition 7

Let $r = \langle id, req, G \rangle$ be a request. A *response* to r is a tuple $\langle id, Ans, S_{ans}, Loops \rangle$, where:

- id is the response identifier;
- Ans is a (possibly empty) set of answers of G ;
- $S_{ans} \in \{active, loop(id_1), disposed\}$ is the status of the evaluation of G , where id_1 is a loop identifier;
- $Loops$ is a set of loop identifiers. □

A response has the same identifier of the request to which it refers. It contains a (possibly empty) set of answers of the requested goal G (Ans) together with the status of G 's evaluation (S_{ans}) and information about the loops in which it is involved ($Loops$). S_{ans} is *disposed* if G has been completely evaluated, *active* if additional answers of G may be computed, and *loop*(id_1) if the response is sent in the context of the evaluation of loop id_1 .

As discussed in Section 3.2, GEM computes the answers of a goal by a depth-first evaluation of its partial derivation tree (Definition 1), which may involve the generation of requests for subgoals evaluated at different locations. In the implementation, we represent a partial derivation tree as a data structure called *evaluation tree*. Compared to partial derivation trees, an evaluation tree keeps track of the identifier of the request and status of the evaluation of the selected atom of each *node* in the evaluation tree.

Definition 8

A *node* is a triple $\langle id, c, S \rangle$, where:

- id is the node identifier;
- c is a clause;
- $S \in \{new, active, loop(ID), answer, disposed\}$ is the status of the evaluation of the selected atom in c , where ID is a set of loop identifiers. \square

The status of a node is *new* if no atom from the body of c has yet been selected for evaluation. It is set to *active* when a body atom is selected, and to *disposed* when the selected atom is completely evaluated. The status is set to *loop(ID)* if the selected atom is involved in some loops, where ID is the set of identifiers of those loops, and to *answer* if c is a fact. As mentioned in Section 2, we employ the leftmost selection rule. Thus, the selected atom of c is always the first body atom.

Definition 9

The *evaluation tree* of a goal $G = \leftarrow A$ is a tree with the following properties:

- the root is node $\langle id_0, A \leftarrow A, S_0 \rangle$;
- there is an edge from the root to a node $\langle id_1, (A \leftarrow B_1, \dots, B_n)\theta, S_1 \rangle$, where $id_1 \sqsubset id_0$, iff there exists a derivation step $(A \leftarrow A) \xrightarrow{\theta} (A \leftarrow B_1, \dots, B_n)\theta$ in the partial derivation tree of G ;
- there is an edge from node $\langle id_2, A \leftarrow B_1, \dots, B_n, S_2 \rangle$ to node $\langle id_3, (A \leftarrow B_2, \dots, B_n)\theta, S_3 \rangle$, where $id_2 \sqsubset id_0$ and $id_3 \sqsubset id_0$, iff there exists a derivation step $(A \leftarrow B_1, \dots, B_n) \xrightarrow{\theta} (A \leftarrow B_2, \dots, B_n)\theta$ in the partial derivation tree of G ; \square

When a principal receives a higher request for a goal G , it creates a table for G . A table contains all the information about the evaluation of G .

Definition 10

The *table* of a goal G , denoted $Table(G)$, is a tuple $\langle HR, LR, ActiveGoals, AnsSet, Tree \rangle$, where:

- HR is a higher request for G ;
- LR is a set of lower requests for G ;
- $ActiveGoals$ is a set of pairs $\langle id, counter \rangle$ where id is a loop identifier and $counter$ is an integer value;
- $AnsSet$ is a set of pairs $\langle ans, ID \rangle$ where ans is an answer of G and ID is a set of request identifiers;
- $Tree$ is the evaluation tree of G . \square

The table of a goal G stores the higher request HR for which it has been created, the set of answers computed so far ($AnsSet$), and the evaluation tree of G ($Tree$). Possible lower requests for G are stored in LR . $ActiveGoals$ keeps a counter for each loop in which G is involved. The counter of a loop id indicates the number of subgoals of G which are involved in loop id , i.e., the number of nodes in $Tree$ with status *loop(ID)* such that $id \in ID$. The counter is decreased whenever an answer of one of these subgoals is received. The status of the root node of $Tree$ indicates the status of the evaluation of G . When G is completely evaluated, the fields of its table are erased, but the answers of G are maintained to speed up the evaluation of future requests for G .

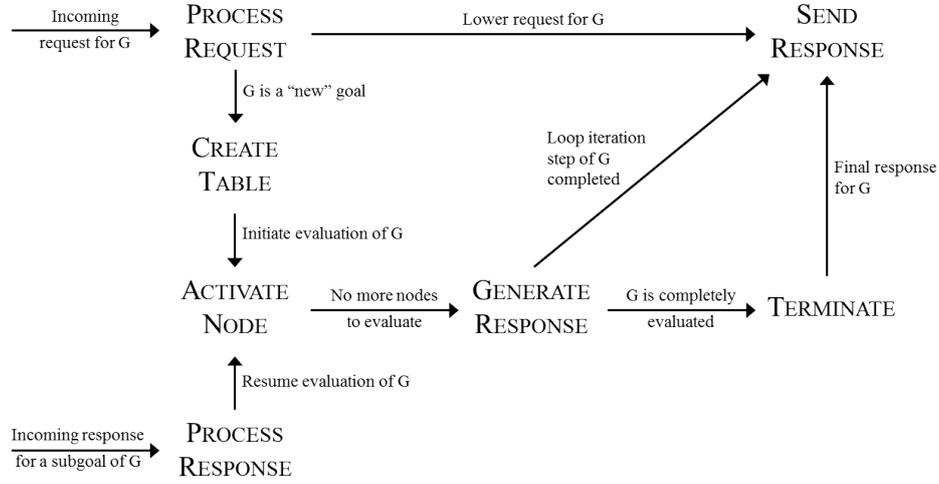


Fig. 6. Interaction among the Procedures for the Evaluation of a Goal G

Procedures. To initiate the evaluation of a goal G , a principal a generates a unique sequence of characters id_0 and sends a request $\langle id_0, a, G \rangle$ to the principal defining G . A response $\langle id_0, Ans, disposed, \{\} \rangle$ is returned to a when the evaluation of G terminates. GEM computes the answers of G (defined in the policy P_G) by means of the following procedures:

- **PROCESS REQUEST:** if the request is not a lower request, invokes **CREATE TABLE** to initiate the evaluation of G . Otherwise, it sends a loop notification to the requester;
- **CREATE TABLE:** creates a table for G and initializes its evaluation tree with the applicable clauses in P_G ;
- **ACTIVATE NODE:** activates a *new* node in the evaluation tree of G ;
- **PROCESS RESPONSE:** processes the answers received for a subgoal of G ;
- **GENERATE RESPONSE:** determines the requesters of G to whom a response must be sent. It is invoked when there are no more nodes in the evaluation tree of G to activate;
- **SEND RESPONSE:** sends the computed answers of G to the requesters of G ;
- **TERMINATE:** disposes the table of G . It is invoked when G is completely evaluated.

Each principal in the trust management system runs a listener that waits for incoming requests and responses. Whenever a request is received, the listener invokes **PROCESS REQUEST**. Similarly, **PROCESS RESPONSE** is invoked upon receiving a response to a previously issued request. The interactions and dependencies among the different procedures are shown in Figure 6.

PROCESS REQUEST (Algorithm 1) takes as input a request $\langle id_0, req, G \rangle$ and, if there exists no table for a variant of G , invokes procedure **CREATE TABLE** to create a table for G (lines 11-12). If another request for goal G (or a variant of G) has been previously received, three situations are possible:

1. *The request refers to a goal which has been completely evaluated* (lines 3-4). In this

Algorithm 1: PROCESS REQUEST

input: a request $\langle id_0, req, G \rangle$

- 1 **if** $\exists Table(G') = \langle \langle id_1, req', G' \rangle, LR, AG, AS, T \rangle$ s.t. G' is a variant of G **then**
- 2 let S_{root} be the status of the root node of T
- 3 **if** $S_{root} = disposed$ **then**
- 4 SEND RESPONSE($\langle id_0, req, G' \rangle, disposed, \{\}$)
- 5 **else if** $id_0 \sqsubset id_1$ **then**
- 6 $LR := LR \cup \{\langle id_0, req, G' \rangle\}$
- 7 SEND RESPONSE($\langle id_0, req, G' \rangle, active, \{id_1\}$)
- 8 **else**
- 9 let G'' be a variable renaming of G s.t. $\nexists Table(G'')$
- 10 CREATE TABLE($\langle id_0, req, G'' \rangle$)
- 11 **else**
- 12 CREATE TABLE($\langle id_0, req, G \rangle$)

Algorithm 2: CREATE TABLE

input: a request $\langle id_0, req, G = \leftarrow A \rangle$

- 1 create $Table(G)$
- 2 initialize $Table(G)$ to $\langle \langle id_0, req, G \rangle, \{\}, \{\}, \{\}, \langle id_0, A \leftarrow A, new \rangle \rangle$
- 3 **foreach** clause $H \leftarrow B_1, \dots, B_n$ applicable to G in the local policy **do**
- 4 let $H' \leftarrow B'_1, \dots, B'_n$ be a variable renaming of the clause s.t. it is variable disjoint from A , and $\theta = mgu(A, H')$
- 5 let s be a unique sequence of characters
- 6 add subnode $\langle id_0s, (H' \leftarrow B'_1, \dots, B'_n)\theta, new \rangle$ to the root
- 7 **end**
- 8 ACTIVATE NODE(G)

case, a response with all the answers of G is sent to the requester by invoking SEND RESPONSE.

2. *The request is a lower request for G* (lines 5-7). This corresponds to the detection of a loop id_1 , where id_1 is the identifier of HR . The request is added to the set of lower requests LR and the answers computed so far are sent to the requester together with a notification about loop id_1 , initiating the loop processing phase.
3. *The request is a side request or originates from a different initial request* (lines 8-10). We treat the request as a new request; accordingly, a new table for G is created by invoking CREATE TABLE.

CREATE TABLE (Algorithm 2) inputs a request $\langle id_0, req, G \rangle$ and creates a table for goal G with HR set to $\langle id_0, req, G \rangle$, and $Tree$ initialized with the clauses in the local policy applicable to G (renamed so that they share no variable with G) (lines 1-7). The identifiers of the subnodes of the root are obtained by concatenating id_0 with a unique sequence of characters s . When the initialization of the table of G is terminated, ACTIVATE NODE is invoked (line 8).

ACTIVATE NODE (Algorithm 3) activates a *new* node from the evaluation tree of a goal G . First, it sets the status of the root node of the evaluation tree T of goal G to *active* (lines 5-7). Then, a node with status *new* is selected from T (line 8). If the node's clause is a fact and represents a new answer, it is added to the set of answers AS (with an empty set of recipients), and ACTIVATE NODE is invoked again (lines 9-13). The answer subsumption check (line 11) is important to avoid sending the same answers of a goal more than once. If the clause is not a fact, the leftmost body atom B_1 of the node's clause is selected for

Algorithm 3: ACTIVATE NODE

input: a goal $G = \leftarrow A$

- 1 let $Table(G)$ be $\langle HR, LR, AG, AS, T \rangle$
- 2 **if** (\nexists a non-root node $t \in T$ with status *new*) **or** ($\langle A, ID \rangle \in AS$) **then**
- 3 GENERATE RESPONSE(G)
- 4 **else**
- 5 let S_{root} be the status of the root node of T
- 6 **if** $S_{root} = \textit{new}$ **then**
- 7 $S_{root} := \textit{active}$
- 8 select the leftmost non-root node $t = \langle id_1, H \leftarrow B_1, \dots, B_n, \textit{new} \rangle$ from T
- 9 **if** $n = 0$ **then**
- 10 set the status of t to *answer*
- 11 **if** H is not subsumed by any answer in AS **then**
- 12 $AS := AS \cup \{ \langle H, \{ \} \rangle \}$
- 13 ACTIVATE NODE(G)
- 14 **else**
- 15 **if** the location of B_1 is not ground **then**
- 16 halt with an error message /* floundering */
- 17 **else**
- 18 set the status of t to *active*
- 19 send request $\langle id_1, local, B_1 \rangle$ to the location of B_1

Algorithm 4: SEND RESPONSE

input: a request $\langle id_0, req, G \rangle$, a response status S_{ans} , a set of loop identifiers $Loops$

- 1 let $Table(G)$ be $\langle HR, LR, AG, AS, T \rangle$
- 2 $Ans := \{ \}$
- 3 **foreach** $\langle ans, ID \rangle \in AS$ s.t. $id_0 \notin ID$ **do**
- 4 $Ans := Ans \cup \{ ans \}$
- 5 $ID := ID \cup \{ id_0 \}$
- 6 **end**
- 7 send response $\langle id_0, Ans, S_{ans}, Loops \rangle$ to Req

evaluation. In case that the location parameter of B_1 is not ground, an error is raised and the computation is aborted by *floundering* (lines 15-16). Otherwise, a request for B_1 is sent to the corresponding location; the node identifier is used as request identifier (lines 17-19). If there are no more nodes with status *new*, or G is in the set of computed answers AS , GENERATE RESPONSE is invoked (lines 2-3).

SEND RESPONSE (Algorithm 4) inputs a request, a response status, and a set of loop identifiers and sends a response message to the requester, which includes the answers of G that have not been previously sent to that requester (lines 3-7).

Response messages are processed by PROCESS RESPONSE (Algorithm 5). The node t to which the response refers is identified by looking at the response identifier (line 1). If the status of the response is *disposed*, the selected atom B_1 of t is completely evaluated. Therefore, t is disposed and, if B_1 is in a loop, also all the other nodes in any loop of the SCC are disposed (lines 5-8). This is because the termination of a loop is ordered by the principal handling the leader of the SCC once all the goals (and consequently, all the loops) in the SCC are completely evaluated.

Otherwise, the status of t is updated depending on whether the response contains a loop notification, i.e., set $Loops$ contains some loop identifier (lines 10-13). In this case, an entry is added to the set of active goals AG for each new loop in $Loops$ (line 14). If the

Algorithm 5: PROCESS RESPONSE

input: a response $\langle id_0, Ans, S_{ans}, Loops \rangle$

- 1 let $t = \langle id_0, H \leftarrow B_1, \dots, B_n, S_t \rangle$ be the node in the evaluation tree of goal $G \leftarrow A$ to which the response refers
- 2 let $Table(G)$ be $\langle HR, LR, AG, AS, T \rangle$
- 3 let $\langle id_1, A \leftarrow A, S_{root} \rangle$ be the root node of T
- 4 **if** $S_{root} \neq disposed$ **then**
- 5 **if** $S_{ans} = disposed$ **then**
- 6 **if** $S_t = loop(ID)$ **then**
- 7 dispose all the nodes in T involved in any loop
- 8 $S_t := disposed$
- 9 **else**
- 10 **if** $S_t = loop(ID)$ **then**
- 11 $ID := ID \cup Loops$
- 12 **else if** $Loops \neq \{\}$ **then**
- 13 $S_t := loop(Loops)$
- 14 $AG := AG \cup \{\langle id_2, 0 \rangle \mid id_2 \in Loops \text{ and } \langle id_2, c \rangle \notin AG\}$
- 15 **if** $S_{ans} = loop(id_3)$ **then**
- 16 decrease the counter of id_3 in AG by 1
- 17 **if** $S_{root} = active$ **then**
- 18 $S_{root} := loop(\{id_3\})$
- 19 **foreach** $ans \in Ans$ **do**
- 20 let ans' be a variable renaming of ans s.t. it is variable disjoint from B_1 , and $\theta = mgu(B_1, ans')$
- 21 let s be a unique sequence of characters
- 22 add subnode $\langle id_1s, (H \leftarrow B_2, \dots, B_n)\theta, new \rangle$ of t
- 23 **end**
- 24 **if** $(S_{root} = active)$ **or** $(S_{root} = loop(ID))$ **and** $\forall id_4 \in ID, \langle id_4, 0 \rangle \in AG$ **then**
- 25 ACTIVATE NODE(G)

response has been sent in the context of the evaluation of a loop id_3 , the counter of id_3 in AG is decreased and the status of the table is changed to $loop(\{id_3\})$ (lines 15-18).

After updating the node and table status, the set of answers in the response is processed (lines 19-23). In particular, a new subnode of t is created for each answer. The clause of the new node is $(H \leftarrow B_2, \dots, B_n)\theta$, where θ is the mgu of B_1 and the answer, and its identifier is obtained by concatenating the identifier id_1 of the root node of T with a unique sequence of characters s . When all answers have been processed, if the principal is not waiting for a response for any subgoal in the evaluation tree of G , ACTIVATE NODE is invoked to proceed with the evaluation of G (line 25).

GENERATE RESPONSE (Algorithm 6) is invoked when all the clauses in the evaluation tree of a goal G (except for the ones in a loop) have been evaluated. If G is not part of a loop, TERMINATE is invoked (lines 2-3). Otherwise, we distinguish three cases:

1. If set LR is not empty, then goal G is the coordinator of a loop id_1 , where id_1 is the identifier of the higher request for G . If there are new answers of G that have not yet been sent to the lower requests in LR , a response with status $loop(id_1)$ is sent to each of them (lines 6-14). This corresponds to starting a new loop iteration for loop id_1 . The status of the root node of the evaluation tree T is updated to keep track of the loops that are currently being processed (lines 8-11) and the counter of id_1 in the set of active goals AG is set to the number of subgoals in T involved in loop id_1 (i.e., the number of nodes with status $loop(ID)$ such that $id_1 \in ID$, line 7). This number corresponds to the number of subgoals for which a response in the context of loop id_1 will be returned.

Algorithm 6: GENERATE RESPONSE

input: a goal $G = \leftarrow A$

```

1 let  $Table(G)$  be  $\langle HR, LR, AG, AS, T \rangle$ 
2 if  $(\exists \langle id_0, c, loop(ID) \rangle \in T)$  then
3   TERMINATE( $G$ )
4 else
5   let  $\langle id_1, A \leftarrow A, S_{root} \rangle$  be the root node of  $T$ 
6   if  $G$  is the coordinator of a loop  $id_1$  and  $\exists ans \in AS$  s.t.  $ans$  has not been sent to some request in  $LR$ 
   then
7     set the counter of  $id_1$  in  $AG$  to the number of subgoals in  $T$  involved in loop  $id_1$ 
8     if  $S_{root} = loop(ID_1)$  then
9        $S_{root} := loop(ID_1 \cup \{id_1\})$ 
10    else
11       $S_{root} := loop(\{id_1\})$ 
12    foreach  $\langle id_2, req, G \rangle \in LR$  do
13      SEND RESPONSE( $\langle id_2, req, G \rangle, loop(id_1), \{\}$ )
14    end
15  else if  $G$  is the leader of the SCC then
16    TERMINATE( $G$ )
17  else
18    let  $Loops$  be the set  $\{id_3 | \langle id_3, C \rangle \in AG \text{ and } id_1 \sqsubset id_3\}$ 
19    set the counter of each  $id_3 \in Loops$  to the number of subgoals in  $T$  in loop  $id_3$ 
20    if  $S_{root} = loop(ID_1)$  and  $\exists id_4 \in ID_1$  s.t.  $id_1 \sqsubset id_4$  then
21      SEND RESPONSE( $HR, loop(id_4), Loops$ )
22    else
23      SEND RESPONSE( $HR, active, Loops$ )
24     $S_{root} := active$ 

```

Algorithm 7: TERMINATE

input: a goal G

```

1 let  $Table(G)$  be  $\langle HR, LR, AG, AS, T \rangle$ 
2 dispose all non-answer nodes in  $T$ 
3 foreach  $\langle id_0, req, G \rangle \in \{HR\} \cup LR$  do
4   SEND RESPONSE( $\langle id_0, req, G \rangle, disposed, \{\}$ )
5 end
6  $HR := null$ 
7  $LR := AG := \{\}$ 

```

2. If G is the leader of the SCC and no new answers of G have been computed, the loop is terminated by invoking TERMINATE (lines 15-16). G is the leader of the SCC if the only loop identifier in set AG is the identifier of the higher request for G .
3. Otherwise, a response including the identifier of the loops in which G is involved is sent to the requester of HR (lines 18-24). The status of the response depends on whether HR is involved in one of the loops that are currently being processed (lines 20-23).

TERMINATE (Algorithm 7) is responsible of disposing a table once all the answers of its goal G have been computed. More precisely, all the table fields are erased except for the set AS of answers of G , which are kept in case of future requests for goal G . A response with status *disposed* is sent to the requesters of HR and LR (lines 3-5).

An example of the execution of GEM is given in Appendix A.

4 Properties of GEM

In this section we prove the soundness, completeness and termination of the GEM algorithm. Furthermore, we discuss what information is disclosed by GEM during the evaluation of a goal.

4.1 Soundness, Completeness and Termination

Here, we refer to an arbitrary but fixed set P_1, \dots, P_n of policies, and to the corresponding global policy $P = P_1 \cup \dots \cup P_n$. To prove its soundness and completeness, we demonstrate that GEM computes a solution if and only if such a solution can be derived via SLD resolution, which has been proved sound and complete (Apt 1990). As mentioned in Section 3.3, we assume that given a request or response message M sent by a principal a to a principal b , one and only one instance of message M is received by b . In other words, we assume no message duplication, and that messages are always received.

We introduce one last definition.

Definition 11

Let S be the set of tables resulting from running GEM on a goal G w.r.t. $P = P_1 \cup \dots \cup P_n$. Let G_1 be a goal whose table is in S . Let θ be a solution of G_1 using clause $H \leftarrow B_1, \dots, B_n$. Then, by construction $\exists \theta_0, \dots, \theta_n$ s.t. $\theta_0 = mgu(G_1, H)$ and θ_j is a solution of $B_j \theta_0 \dots \theta_{j-1}$ (with $j \in \{1, \dots, n\}$). The *ranking* of θ is defined inductively as follows:

- $rank(\theta) = 1$ if $n = 0$ (i.e., the clause is a fact),
- $rank(\theta) = 1 + \max(rank(\theta_1), \dots, rank(\theta_n))$ otherwise, where $rank(\theta_j)$ is the ranking of solution θ_j . □

The following theorem states that each solution computed by GEM can also be derived via SLD resolution using the global policy P , and is thus correct. Intuitively, this is due to the fact that the solutions generated by the algorithm are obtained using the clause resolution mechanism, which produces correct results.

Theorem 1 (Soundness)

Let G_1 be a goal. Let S be the set of tables resulting from running GEM on G_1 (w.r.t. $P = P_1 \cup \dots \cup P_n$). Let G_1, \dots, G_k be the goals for which there exists a table in S . For each goal $G_i \in \{G_1, \dots, G_k\}$ let $Sol_i = \{\theta_{i,1}, \dots, \theta_{i,k_i}\}$ be the (possibly empty) set of solutions of G_i generated by the algorithm. Then, for each $G_i \in \{G_1, \dots, G_k\}$ and for each $\theta_{i,j} \in Sol_i$ there exists an SLD derivation of $P \cup \{G_i\}$ with c.a.s. σ s.t. $G_i \theta_{i,j}$ is a renaming of $G_i \sigma$.

Proof. We proceed by contradiction and assume that there exists at least a “wrong” solution $\theta_{i,j}$ in Sol_i , i.e., a solution s.t. there is no corresponding SLD derivation of $P \cup \{G_i\}$ with c.a.s. σ where $G_i \theta_{i,j}$ is a renaming of $G_i \sigma$ (hypothesis).

Let us choose $\theta_{i,j}$ to be a “wrong” solution with minimal ranking (*). Let $G_i = \leftarrow A_i$. Since $\theta_{i,j}$ is a solution of G_i , there exists an evaluation tree of G_i in S created by CREATE TABLE (lines 2-7) with root $\langle id, A_i \leftarrow A_i, new \rangle$, a subnode with clause $c =$

$H \leftarrow B_1, \dots, B_n$ and substitutions $\theta_0, \dots, \theta_n$ s.t. $\theta_0 = mgu(A_i, H)$, and for each $l \in \{1, \dots, n\}$ there exists:

- A node in the evaluation tree of G_i with selected atom $B_l\theta_0 \cdots \theta_{l-1}$ (ACTIVATE NODE, lines 8, 18-19).
- An evaluation tree of $\leftarrow B_l\theta_0 \cdots \theta_{l-1}$ created by CREATE TABLE (lines 2-7) at the location of $B_l\theta_0 \cdots \theta_{l-1}$.
- A solution θ_l of $\leftarrow B_l\theta_0 \cdots \theta_{l-1}$; the answer $B_l\theta_0 \cdots \theta_l$ is sent to the requester of $\leftarrow B_l\theta_0 \cdots \theta_{l-1}$ by GENERATE RESPONSE (lines 12-14 or 20-23) if $\leftarrow B_l\theta_0 \cdots \theta_{l-1}$ is involved in a loop, or by TERMINATE (lines 3-5) otherwise.
- A node with clause $(H \leftarrow B_{l+1}, \dots, B_n)\theta_0 \cdots \theta_l$ added to the evaluation tree of G_i by PROCESS RESPONSE (lines 20-23).

Then, $\theta_{i,j} = \theta_0 \cdots \theta_n$. If the body of c is empty, then there is a trivial 1-step SLD derivation of $P \cup \{G_i\}$ with c.a.s. σ_i (namely the mgu of G_i and c), therefore contradicting the hypothesis. So, let us now assume that $n > 0$; by construction, for each $l \in \{1, \dots, n\}$, $rank(\theta_l) < rank(\theta_{i,j})$. So, by the minimality argument (*), for each $l \in \{1, \dots, n\}$ there exists an SLD derivation of $P \cup \{\leftarrow B_l\theta_0 \cdots \theta_{l-1}\}$ with c.a.s. σ_l s.t. $B_l\sigma_0 \cdots \sigma_{l-1}\sigma_l = B_l\theta_0 \cdots \theta_{l-1}\theta_l$. But then, by standard logic programming results (given the presence of clause c), there exists a successful SLD derivation of $P \cup \{G_i\}$ with c.a.s. σ s.t. $G_i\sigma = G_i\theta_{i,j}$, contradicting the hypothesis. \square

Since GEM employs a “wait” mechanism to determine when the answers of a goal should be sent to the requester, both the completeness and termination properties of the algorithm depend on the correctness of this mechanism. Therefore, before demonstrating that GEM is complete and always terminates, we prove that the “wait” mechanism is correctly implemented, i.e., that the answers of a goal are eventually sent to the requester. This is particularly challenging in the presence of loops.

In the implementation of GEM proposed in Section 3.3, the “wait” mechanism for goals involved in a loop consists of loop counters: at each iteration of a loop id , the answers of a goal G are only sent when the counter of loop id in set *ActiveGoals* is 0 (procedure PROCESS RESPONSE, line 24). Since the counter is set to the number k of subgoals of G which are involved in loop id (GENERATE RESPONSE, lines 7 and 19), at each iteration of loop id the principal evaluating G should thus receive k response messages. In order to prove this, we first show that GEM correctly keeps track of the loops in which the subgoals of G are involved.

Proposition 1

Let G_1, \dots, G_m be the goals involved in a loop id_1 . Let $G_i, G_j \in \{G_1, \dots, G_m\}$ be two goals s.t. G_j is a subgoal of G_i . Then, the node in the evaluation tree of G_i with selected atom G_j has status *loop(ID)*, where $id_1 \in ID$.

Proof. Let G_1 be the coordinator of loop id_1 . Let G_1, \dots, G_k be a subset of G_1, \dots, G_m s.t. for each $i \in \{2, \dots, k\}$ goal G_i is a subgoal of G_{i-1} , and G_1 is a subgoal of G_k . The node in the evaluation tree of G_1, \dots, G_k with selected atom G_2, \dots, G_k, G_1 respectively has status *loop(ID)*, where $id_1 \in ID$, because of the following observations:

- The identifiers id_2, \dots, id_k of the requests for goals G_1, \dots, G_k and the identifier id_{k+1} of the request for goal G_1 are constructed by procedures CREATE TABLE (lines 5-6) and PROCESS RESPONSE (lines 21-22) in such a way that $id_j \sqsubset id_1$, for each $j \in \{2, \dots, k+1\}$, and thus the lower request id_{k+1} for G_1 can be identified.
- Upon receiving the lower request id_{k+1} , the principal evaluating G_1 returns a response $\langle id_{k+1}, Ans_{k+1}, S_{k+1}, \{id_1\} \rangle$ to the principal evaluating G_k (procedure PROCESS REQUEST, lines 5-7).
- The status of the node in the evaluation tree of G_k with selected atom G_1 is set to $loop(\{id_1\})$ (PROCESS RESPONSE, lines 12-13).
- A counter for loop id_1 is added to set *ActiveGoals* in the table of goal G_k (PROCESS RESPONSE, line 14).
- For each $i \in \{2, \dots, k\}$, the principal evaluating goal G_i sends to the principal evaluating goal G_{i-1} a response of the form $\langle id_i, Ans_i, S_i, ID \rangle$, where ID is the set of all loops in *ActiveGoals* whose identifier is higher than id_i (GENERATE RESPONSE, lines 18, 21, and 23), and thus $id_1 \in ID$.
- The status of the node in the evaluation tree of goal G_{i-1} with selected atom G_i is set to $loop(ID)$, where $id_1 \in ID$ (PROCESS RESPONSE, lines 12-13). \square

Corollary 1

Let G be a goal involved in a loop id . Let k be the number of nodes in the evaluation tree of G with status $loop(ID)$ s.t. $id \in ID$. When a response is sent to the requester of the higher request for G (or lower request, if G is the loop coordinator), the counter of loop id in set *ActiveGoals* in the table of G is set to k .

At each loop iteration, the counters of the loops in which a goal G is involved are set to the number of subgoals of G involved in those loops by procedure GENERATE RESPONSE, lines 7 and 19. Hence, we now need to show that at each iteration of a loop id the number of response messages with status $loop(id)$ received by the principal evaluating G is equal to the number of subgoals of G involved in loop id , i.e., that counters correctly keep track of the number of response messages received by the principal evaluating G at each iteration of loop id .

Informally, the correctness of counters stems from the fact that at each loop iteration step for a goal G there is *only one* choice of loop identifier to include in the response to the requester of a higher request for G . This is because of the following considerations:

1. Let G be a goal involved in one or more loops. In the loop processing phase, the loop identifier included by the principal evaluating G in the response sent to the requester of a higher request for G is taken from the status of the root node of the evaluation tree of G (procedure GENERATE RESPONSE, lines 20-21).
2. After a response for G is sent by GENERATE RESPONSE (lines 20-23), the status of the root node of the evaluation tree of G is set to *active* (line 24).
3. If G is a non-coordinator goal, then there can be *at most one* loop identifier per time in the status of the root node of its evaluation tree. Therefore, when sending a response for G , procedure GENERATE RESPONSE has only one choice of loop identifier to include in the response status. The reason why a non-coordinator goal can have at most one loop identifier in the status of the root of its evaluation tree is

the following. The only point where the status of the root node of a non-coordinator goal G is modified to take into account the loop being processed is on line 18 of procedure PROCESS RESPONSE, and the check on line 17 updates the status only in case it is currently set to *active*. We point out that when the response for a subgoal of G is processed by PROCESS RESPONSE the status of the root of the evaluation tree of G is always *active*, due to point (2) above and the fact that GEM only processes one goal at a time (which is due to condition on line 24 of PROCESS RESPONSE), and thus no response will be received by the principal evaluating G in the context of a loop unless a response for G was previously sent.

4. if G is the coordinator of a loop id_l , then there can be *at most two* loop identifiers per time in the status of the root node of its evaluation tree: one for loop id_l , and *at most one* for a higher loop id_h . Remember that as loop identifiers we use the identifier of the higher request for the coordinator; hence, in this case id_l is the identifier of the higher request for G . Given the condition on line 20 of GENERATE RESPONSE, only id_h can be included in the status of a response for G sent to the requester of a higher request. In fact, id_h (denoted id_4 in the procedure) is the only identifier in the status of the root node of the evaluation tree of G that is higher than id_l (denoted id_1 in the procedure), i.e., higher than the identifier of the higher request for G . Therefore, when sending a response for G , GENERATE RESPONSE has only one choice of loop identifier to include in the response status, namely id_h .

Technically, a coordinator can have at most two loop identifiers in the status of the root node of its evaluation tree because of the following. Similarly to non-coordinators, due to condition on line 17 of PROCESS RESPONSE only *one* loop identifier can be added to the root's status on line 18 of PROCESS RESPONSE. This occurs when G receives a response from one of its subgoals in the context of a higher loop id_h . A second loop identifier (the identifier of loop id_l) can be added to the root status on lines 8-9 of GENERATE RESPONSE if the response received by the principal evaluating G in the context of loop id_h leads to new answers of G , which need to be sent to the goals involved in loop id_l . No more than two loops at a time will be processed by the principal evaluating G (i.e., id_l and at most one higher loop id_h) because of the following reasons. Upon receiving a response in the context of a higher loop id_h :

- a response for G in the context of loop id_h will not be sent to a higher goal until a fixpoint for the loop id_l of which G is the coordinator is reached, during which time the status of the root node of the evaluation tree of G is $loop(\{id_h, id_l\})$, and
- due to the condition on line 24 of PROCESS RESPONSE no responses for higher goals can be received by the principal evaluating G until a response for G in the context of loop id_h is sent upwards. In fact, the counter of loop id_h in the table of higher goals cannot be 0, because no response for G in the context of loop id_h was sent upwards yet. When a response for G is sent upwards, the status of the root node of its evaluation tree becomes *active* again (see point (2)).

Formally, the correctness of counters is demonstrated by the following Proposition.

Proposition 2

Let G be a goal and G_1, \dots, G_k be the subgoals of G s.t. G, G_1, \dots, G_k are involved in a loop id_l . At each iteration of loop id_l , the principal evaluating G receives k response messages, one for each subgoal $G_i \in \{G_1, \dots, G_k\}$.

Proof. Let $G, G_1, \dots, G_k, \dots, G_m$ be all the goals involved in loop id_l , where $m \geq k$. Let id, id_1, \dots, id_m be the identifiers of the requests for goals G, G_1, \dots, G_m respectively. The proof is by induction on the number ℓ of goals $G_j \in \{G_1, \dots, G_m\}$ s.t. $id_j \sqsubset id$, that is, the number of goals whose request identifier is lower than the identifier id of the request for G .

Base case: $\ell = 1$. Then, also $k = 1$. Let $G_j \in \{G_1, \dots, G_m\}$ be the only goal s.t. $id_j \sqsubset id$. It is straightforward to see that goal G_j is (a variant of) the coordinator of loop id_l , and id_j denotes the lower request for G_j . When there are no more nodes with status *new* in the evaluation tree of G_j , i.e., when all the branches of the evaluation tree of G_j have been evaluated, procedure GENERATE RESPONSE is invoked by ACTIVATE NODE (lines 2-3). By procedure GENERATE RESPONSE (lines 12-14), at each loop iteration one and only one response to the request for the coordinator G_j is sent by GEM to the principal evaluating G . Thus, at each iteration of loop id_l the principal evaluating G receives $k = 1$ response messages. Q.e.d.

Inductive case: Now, assume that G has ℓ such goals $G_j \in \{G_1, \dots, G_m\}$ s.t. $id_j \sqsubset id$, where $\ell > 1$. In this case, each subgoal $G_i \in \{G_1, \dots, G_k\}$ of G is either the coordinator of loop id_l or a goal with at most $\ell - k$ subgoals $G_p \in \{G_1, \dots, G_m\}$ s.t. $id_p \sqsubset id_i$. If G_i is the coordinator of loop id_l , by the same reasoning done in the base case, one and only one response to the request for G_i is sent by GEM to the principal evaluating G at each loop iteration.

On the other hand, if G_i is not the loop coordinator, there exist at most $\ell - k$ goals $G_p \in \{G_1, \dots, G_m\}$ s.t. $id_p \sqsubset id_i$. Let t be the number of subgoals of G_i involved in loop id_l . Since $\ell - k \leq \ell - 1$, by the inductive hypothesis (*) the principal evaluating G_i receives t response messages at each iteration of loop id_l . By Corollary 1, at each loop iteration the counter of loop id_l in the table of goal G_i is set to t , and is decreased by 1 every time a response to the requests for its subgoals involved in the loop is received (procedure PROCESS RESPONSE, lines 15-16). Therefore, after t response messages, the counter of loop id_l in the table of G_i is 0, and procedure PROCESS RESPONSE (lines 24-25) resumes the evaluation of goal G_i . When there are no more nodes with status *new* in the evaluation tree of G_i , procedure ACTIVATE NODE (lines 2-3) invokes GENERATE RESPONSE. By procedure GENERATE RESPONSE (lines 20-21), one and only one response to the request for G_i is sent by GEM to the principal evaluating G at each iteration of loop id_l . Therefore, at each iteration of loop id_l the principal evaluating goal G receives k response messages. \square

Finally, we show that procedure TERMINATE is eventually invoked for any goal in a computation.

Proposition 3

Let G_1 be a goal. Procedure TERMINATE is eventually called for G_1 .

Proof. The proof is divided into two parts. First, we show that TERMINATE is eventually called for a goal G_1 that is not involved in a loop. Then, we show that it is always invoked also if G_1 is involved in one or more loops.

The first part of the proof is straightforward, and is given by the fact that the number of answers of goal G_1 is finite. This is because of the following observations:

1. The global policy P is finite, and the terms in P that are not variables are constants defined in P ; thus, the Herbrand model of P is finite.
2. Let $P_{G_1} \in P$ be the policy where goal G_1 is defined. The answers of G_1 are computed by GEM through the clauses in P_{G_1} applicable to G_1 (procedure CREATE TABLE, lines 3-7). Each clause can be either a fact or have the form $H \leftarrow B_1, \dots, B_m$, such that B_1, \dots, B_m are defined in a policy in P . By (1), both the number of facts in P_{G_1} and the number of answers of subgoals B_1, \dots, B_n are finite.

Thus, the number of answers of goal G_1 is finite. When all the answers of G_1 have been computed and all the nodes in the partial tree of G_1 have been evaluated, procedure ACTIVATE NODE (lines 2-3) invokes GENERATE RESPONSE, which in turn (lines 2-3) invokes TERMINATE.

Consider now the case in which G_1 is part of an SCC consisting of loops id_1, \dots, id_k , s.t. $id_k \sqsubset \dots \sqsubset id_1$. Let G_1, \dots, G_m be all the goals involved in loops id_1, \dots, id_k (where $m \geq k$), and goal $G_{c_i} \in \{G_1, \dots, G_m\}$ be the coordinator of loop $id_i \in \{id_1, \dots, id_k\}$. Because the number of answers of each goal G_1, \dots, G_m is finite, we have that:

- At each iteration of loop id_i , if new answers of the loop coordinator G_{c_i} are derived, they are sent to the requesters of the lower requests for G_{c_i} , starting a new iteration of loop id_i (procedure GENERATE RESPONSE, lines 6-14). On the contrary, if no answer of G_{c_i} is computed, the answers of G_{c_i} are sent to the requester of the higher request for G_{c_i} (GENERATE RESPONSE, lines 20-23). The loops higher than id_i in the SCC are then processed.
- At each iteration of loop id_1 , if new answers of the leader G_{c_1} are derived, they are sent to the requesters of the lower requests for G_{c_1} , starting a new iteration of loop id_1 (procedure GENERATE RESPONSE, lines 6-14). Notice that this might cause a fixpoint for the loops lower than id_1 in the SCC to be recomputed. On the contrary, if no answer of G_{c_1} is computed, the answers of G_{c_1} are sent to the requester of the higher request for G_{c_1} , and a response with status *disposed* is sent to the requesters of the lower requests for G_{c_1} (GENERATE RESPONSE, lines 15-16 and TERMINATE, lines 3-5).
- For each goal $G_j \in \{G_1, \dots, G_m\}$, all the nodes in the evaluation tree of G_j are disposed (PROCESS RESPONSE, lines 5-8); then, procedure ACTIVATE NODE is invoked, which immediately invokes GENERATE RESPONSE (lines 2-3).
- The principal evaluating goal G_j sends a response with status *disposed* to the requester of the higher request for G_j . If G_j is a loop coordinator, the principal evaluating G_j also sends a response with status *disposed* to the requesters of the lower requests for G_j (GENERATE RESPONSE, lines 2-3 and TERMINATE, lines 3-5).

Therefore, procedure TERMINATE is always invoked for goal G_1 . □

Proposition 3 implies that the table of a goal involved in a computation is always disposed. In fact, the disposal of the table of a goal is carried out by procedure TERMINATE (lines 2, 6, and 7). Consider, for instance, the following variation of the global policy introduced in Section 3.2, where the research insitute ri refers to goal $memberOfAlpha(c1,X)$ instead of $memberOfAlpha(c2,X)$:

$$\begin{aligned} memberOfAlpha(c1,X) &\leftarrow memberOfAlpha(c2,X). \\ memberOfAlpha(c2,X) &\leftarrow memberOfAlpha(ri,X). \\ memberOfAlpha(ri,X) &\leftarrow memberOfAlpha(c1,X). \end{aligned}$$

First of all, let us recall that the termination of the evaluation of the goals involved in a loop is commanded by the leader of the SCC (goal $memberOfAlpha(c1,X)$ in the example policy). When no new answer of the leader is computed by $c1$ during a loop iteration, procedure TERMINATE is invoked (lines 15-16 of GENERATE RESPONSE), which disposes the table of the goal and sends a response with status *disposed* both to the requesters of the higher and lower requests for $memberOfAlpha(c1,X)$ (lines 3-5). When ri receives the response, it disposes all the nodes in the evaluation tree of $memberOfAlpha(ri,X)$ involved in a loop (lines 5-8 of PROCESS RESPONSE), which in this case corresponds to disposing all the non-root nodes. At this point, the status of the root node of the evaluation tree of $memberOfAlpha(ri,X)$ is *active* (see point 2 of the discussion preceding Proposition 2). Therefore, the condition on line 24 of PROCESS RESPONSE is satisfied, and procedure ACTIVATE NODE is invoked for $memberOfAlpha(c1,X)$. Since all the non-root nodes in the evaluation tree of $memberOfAlpha(c1,X)$ have status *disposed*, GENERATE RESPONSE is invoked (lines 2-3 of ACTIVATE NODE), which in turn (lines 2-3) invokes procedure TERMINATE. TERMINATE disposes the table of goal $memberOfAlpha(ri,X)$ and sends a response with status *disposed* to $c2$. Similarly to $memberOfAlpha(ri,X)$, PROCESS RESPONSE disposes all the nodes in the evaluation tree of goal $memberOfAlpha(c2,X)$, and a response with status *disposed* is sent by $c2$ to $c1$ by procedure TERMINATE. Since the root of the evaluation tree of $memberOfAlpha(c1,X)$ had already been disposed, in this case the response message is ignored by $c1$ (line 4 of PROCESS RESPONSE).

We now state and prove the completeness result.

Theorem 2 (Completeness)

Let G_1 be a goal. Let S be the set of tables resulting by running GEM on G_1 (w.r.t. $P = P_1 \cup \dots \cup P_n$). Assume that running GEM on G_1 (w.r.t. $P = P_1 \cup \dots \cup P_n$) did not result in floundering. If there exists an SLD derivation of $P \cup \{G_1\}$ with c.a.s. θ , then there exists a solution σ of G_1 in S s.t. $G_1\theta$ is a renaming of $G_1\sigma$.

Proof. We proceed by contradiction, and assume that S is missing a solution of G_1 . That is, there exists a successful SLD derivation of $P \cup \{G_1\}$ with c.a.s. θ and there is no solution σ of G_1 generated by the algorithm s.t. $G_1\theta = G_1\sigma$ (hypothesis).

This implies that there exist a (maximal) set of goals G_1, \dots, G_k in S s.t. for each $i \in \{1, \dots, k\}$ there is a non-empty maximal set of substitutions $\{\theta_{i,1}, \dots, \theta_{i,m_i}\}$ s.t.:

- (a) G_i is a goal in S .
- (b) $\theta_{i,1}, \dots, \theta_{i,m_i}$ are correct solutions of G_i according to SLD resolution: for each $\theta_{i,j}$ there exists a successful SLD derivation of $P \cup \{G_i\}$ with c.a.s. $\theta_{i,j}$ (up to renaming).

(c) The algorithm does not generate the answers $G_i\theta_{i,1}, \dots, G_i\theta_{i,m_i}$ (up to renaming).

The set G_1, \dots, G_k is not empty as it contains at least G_1 (the finiteness of the construction is demonstrated in the proof of Proposition 3).

For each i, j , let $der_{i,j}$ be the SLD derivation of $P \cup \{G_i\}$ with c.a.s. $\theta_{i,j}$ of minimal length. Let us choose integers p, q in such a way that $der_{p,q}$ has minimal length among the derivations in the set $\{der_{i,j}\}$. The fact that $der_{p,q}$ has minimal length implies that for any goal G' in S , the following holds: if there exists an SLD derivation of $P \cup \{G'\}$ of length smaller than $len(der_{p,q})$ with c.a.s. θ' , then the algorithm generates a solution ϑ' for which $G'\theta'$ is a renaming of $G'\vartheta'$ (*).

Let c be the clause used in the first step of the derivation $der_{p,q}$. If c is a fact, we immediately have a contradiction: since G_p is a goal in S , this means that there exists an evaluation tree of $G_p \leftarrow A_p$ created by CREATE TABLE (lines 2-7) with root node $\langle id, A_p \leftarrow A_p, new \rangle$ and a node with clause c as subnode of the root node. Therefore, the algorithm will compute a c.a.s. equivalent to $\theta_{p,q}$ (ACTIVATE NODE), contradicting the hypothesis.

If c is a rule $H \leftarrow B_1, \dots, B_n$, and $\sigma_0 = mgu(G_p, H)$, then by hypothesis there exist SLD derivations $der_{B_1}, \dots, der_{B_n}$, and substitutions $\sigma_1, \dots, \sigma_n$ s.t. $H\sigma_0 \dots \sigma_n = G_p\theta_{p,q}$, and for each $i \in \{1, \dots, n\}$:

- der_{B_i} is an SLD derivation of $P \cup \{\leftarrow B_i\sigma_0 \dots \sigma_{i-1}\}$.
- The c.a.s. of der_{B_i} is σ_i , and $len(der_{B_i}) < len(der_{p,q})$. (**)

Since G_p is a goal in S , there exists an evaluation tree of G_p created by CREATE TABLE (lines 2-7) with root node $\langle id, A_p \leftarrow A_p, new \rangle$ and a node with clause c as subnode of the root node. Then, it is easy to see that for each $i \in \{1, \dots, n\}$:

- There exists a node in the evaluation tree of G_p with selected atom $B_i\sigma_0 \dots \sigma_{i-1}$ (ACTIVATE NODE, lines 8, 18-19).
- There exists an evaluation tree of $\leftarrow B_i\sigma_0 \dots \sigma_{i-1}$ created by CREATE TABLE (lines 2-7) at the location of $B_i\sigma_0 \dots \sigma_{i-1}$.
- Since $len(der_{B_i}) < len(der_{p,q})$, by (*) and (**) the algorithm computes a solution equivalent to σ_i of the goal $\leftarrow B_i\sigma_0 \dots \sigma_{i-1}$.
- By Propositions 2 and 3, the answer $B_i\sigma_0 \dots \sigma_i$ is sent to the requester of $\leftarrow B_i\sigma_0 \dots \sigma_{i-1}$ by GENERATE RESPONSE (lines 12-14 or 20-23) if $\leftarrow B_i\sigma_0 \dots \sigma_{i-1}$ is involved in a loop, or by TERMINATE (lines 3-5) otherwise.
- There exists a node with clause $(H \leftarrow B_{i+1}, \dots, B_n)\sigma_0 \dots \sigma_i$ added to the evaluation tree of G_p by PROCESS RESPONSE (lines 20-23).

Therefore, $\sigma_1 \dots \sigma_n$ is (equivalent to) a solution of the evaluation tree of G_p , contradicting (a), (b), and (c). \square

Finally, we prove that GEM always terminates.

Theorem 3 (Termination)

Given a goal G evaluated w.r.t. a finite global policy P , GEM terminates.

Proof. We assume that nodes (i.e., goals) in the call graph of P inherit the identifier (and the associated ordering) of the request for which they are created. Termination follows from

two observations: (i) the call graph of P is finite, and (ii) the number of response messages exchanged by the principals involved in the evaluation of G is finite.

The call graph of P is finite (i) for the following reasons:

1. The set of goals over predicates in P (up to renaming) is finite. This is because terms that are not variables are constants in P .
2. There is no infinite path in the call graph of P composed of nodes id_1, \dots, id_n s.t. $id_n \sqsubset \dots \sqsubset id_1$. This is because of (1) and because the algorithm never creates a new node with identifier id_i for a goal if a node with identifier id_j already exists for a variant of that goal and $id_i \sqsubset id_j$.
3. The outdegree of each node in the call graph of P is finite. This is because the number of atoms in the body of each clause in P is finite.

The number of response messages is finite (ii) because:

1. The number of answers of each goal defined in P is finite (see the proof of Proposition 3).
2. The (possibly empty) set of answers of a goal are transmitted only when a table for the goal is first created (and a node representing the goal is added to the call graph of P) or new answers of its subgoals are received.
3. For any nodes id_1 and id_2 , a set of answers that flows from id_2 to id_1 in response to a request id_2 never contains answers previously communicated in response to request id_2 (SEND RESPONSE, lines 3-4).
4. An empty set of answers may flow from id_2 to id_1 only if $id_2 \sqsubset id_1$ (GENERATE RESPONSE, lines 20-23, and TERMINATE, lines 3-5), or id_1 identifies a lower request and a loop id_2 has just been identified (PROCESS REQUEST, lines 5-7).
5. There is no infinite path composed of nodes id_n, \dots, id_1 in the call graph of P through which the answers flow s.t. $id_n \sqsubset \dots \sqsubset id_1$.
6. By Proposition 3, procedure TERMINATE is eventually invoked for any goal. \square

4.2 Disclosed Information

A primary objective of GEM is to preserve the confidentiality of intensional policies. Here, we discuss what policy information principals are able to collect during the evaluation of a goal, and classify GEM according to the confidentiality levels defined in Section 2.

First, let us define the following notation. Let P be a global policy, and G_a and G_b be two goals in P defined by principals a and b respectively. We say that goal G_a *depends* on goal G_b if there is a path from G_a to G_b in the call graph of the evaluation of G_a with respect to P . Since each edge in the call graph represents a request in GEM, and in trust management each request corresponds to a delegation of authority, if G_a depends on G_b then we say that there is a *chain of trust* from principal a to principal b .

We also introduce some notation on request identifiers. As mentioned in Section 3.2, the identifiers in an identifier domain can be defined in several ways (e.g., applying a hash function to the identifier of a higher request). In this paper, we have considered an identifier domain where identifiers are obtained by concatenating the identifier of a higher request with a sequence of characters. Here, we thus employ this identifier domain to discuss what information is disclosed by GEM during goal evaluation. We classify identifiers obtained

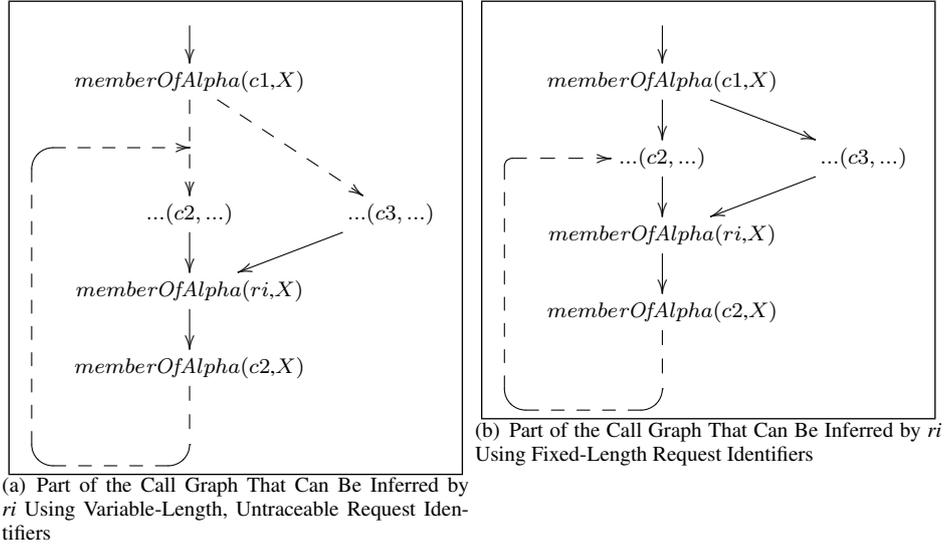
by concatenation according to two dimensions: *traceability* and *length*. Given two request identifiers id_1 and id_2 for goals G_1 and G_2 respectively, such that $id_2 \sqsubset id_1$, the traceability dimension refers to the ability of a principal to infer which principals are involved in the evaluation of the goals in the path from G_1 to G_2 . On the other hand, the length dimension defines the ability to determine the number of goals in the path from G_1 to G_2 . Let $id_0s_1 \cdots s_n$ be a request identifier, where id_0 is the identifier of the initial request in a computation and each s_i (for $i \in \{1, \dots, n\}$) is a sequence of characters added by a principal p_i to the identifier $id_0s_1 \cdots s_{i-1}$ of a higher request. For what concerns the traceability dimension, we say that $id_0s_1 \cdots s_n$ is a *traceable* identifier if each string s_i uniquely identifies (the location of) principal p_i (as done, for instance, by the identifiers in the example in Section 3.2); otherwise, we say that identifier $id_0s_1 \cdots s_n$ is *untraceable*. The length dimension is defined based on the number of characters concatenated by each principal p_i to the request identifier $id_0s_1 \cdots s_{i-1}$. Let $len(s_i)$ denote the number of characters in the string s_i . If $len(s_1) = \dots = len(s_n)$, then we say that $id_0s_1 \cdots s_n$ is a *fixed-length* identifier; otherwise, we say that $id_0s_1 \cdots s_n$ is a *variable-length* identifier (we assume that cryptographic techniques are in place to avoid collision of identifiers (Hoch and Shamir 2008)). Notice that a traceable identifier does not necessarily disclose information about the number of goals in the path between two goals; this is because a goal defined by a principal a can have several subgoals defined in a 's policy. Consider, for instance, the traceable request identifier $h : 12c1 : 345$, obtained by concatenating a request identifier with a principal's identifier and a variable-length sequence of digits for each goal evaluated by the principal. Even though identifier $h : 12c1 : 345$ shows that the principals involved in the computation are hospital h and company $c1$, it does not confer information about the number of goals evaluated by those principals. Company $c1$, for example, might have evaluated two locally defined goals, concatenating the higher request $h : 12$ received from hospital h first with its identifier $c1$ and digit "3" (separated by a semicolon), and then with the sequence of digits "45".

We are now ready to present what information a principal b is able to learn about the local policy of a principal a where a goal G_a is defined. First of all, by requesting the evaluation of G_a , b learns the set of answers to the request, i.e., the extensional policy relative to G_a ; this is necessary for any goal evaluation algorithm. As mentioned in Section 1, the confidentiality of extensional policies can be protected, for instance, by relying on hidden credentials (Bradshaw et al. 2004; Frikken et al. 2006) or trust negotiation algorithms (Winsborough et al. 2000; Winslett 2003). Here, we are more interested in what b can learn about the intensional policy defining G_a .

By sending a request for G_a (say, with identifier id_1), b can learn whether G_a depends on some goal G_b defined in her policy. Indeed, if b receives a request for G_b with identifier id_2 such that $id_2 \sqsubset id_1$, then b knows that G_a depends on G_b .

If G_a depends on a number of goals defined by b , then by requesting G_a b learns:

- what are the goals G_{b_1}, \dots, G_{b_n} defined in her policy on which G_a depends;
- for each $G_{b_i} \in \{G_{b_1}, \dots, G_{b_n}\}$, b knows who is the principal p_i that requested G_{b_i} ; therefore, b learns that G_a depends on a goal defined by p_i , i.e., that there is a chain of trust from a to p_i (and from p_i to b). Principal b , however, does not necessarily learn which is the goal defined by p_i on which G_a depends;

Fig. 7. Part of the Call Graph That Can Be Inferred by ri

- depending on how the identifiers are constructed, b might be able to learn additional information about the path from G_a to G_{b_i} . In particular, if the identifier of the request for G_{b_i} is fixed-length, b is able to infer the number of goals in the path from G_a to G_{b_i} . Additionally, if the identifier is traceable, b also learns who are the principals defining those goals.

Thus, GEM can be classified as E1-I1 according to the classification criteria in Section 2. In fact, principals learn all the answers of a goal along with some dependencies among the goals involved in an evaluation.

We now illustrate the concepts presented above with an example, using the global policy introduced in Section 3.2 and the call graph shown in Figure 5(a) on page 14 (ignore, for now, the request identifiers depicted in the figure). Assume that the research institute ri requests goal $memberOfAlpha(c1,X)$ to $c1$. If the identifiers used in the computation were variable-length and untraceable, ri would be able to learn that:

- $memberOfAlpha(c1,X)$ depends on goal $memberOfAlpha(ri,X)$ defined in its policy;
- $memberOfAlpha(c1,X)$ depends on *some* goal G_{C2} defined in $c2$'s policy and on *some* goal G_{C3} defined by $c3$; however, it does not learn which goals they are. Furthermore, due to the loop notification received from $c2$ following the evaluation of goal $memberOfAlpha(c2,X)$, ri learns that $memberOfAlpha(c2,X)$ depends on *some* goal in the path from $memberOfAlpha(c1,X)$ to G_{C2} .

Figure 7(a) represents ri 's "view" of Figure 5(a), that is, the part of the call graph that ri can infer. In the graph, we denote with dots ("...") the predicate symbols and terms that ri does not learn; a dashed edge from a goal G_1 to a goal G_2 indicates that ri is able to infer that G_1 depends on G_2 , but not the (number of) goals and principals in the path from G_1 to G_2 . Since ri does not learn that G_{C2} is actually goal $memberOfAlpha(c2,X)$, with respect to the call graph in Figure 5(a), in Figure 7(a) goal $memberOfAlpha(c2,X)$ is "duplicated". The

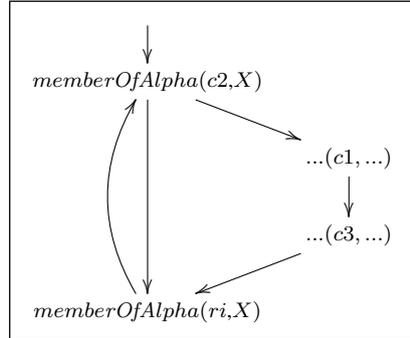
reason why ri is not able to infer that G_{C2} is goal $memberOfAlpha(c2,X)$, and even more, does not learn whether $memberOfAlpha(c2,X)$ is in the path from $memberOfAlpha(c1,X)$ to $memberOfAlpha(ri,X)$ or is a lower goal, is that the only information that ri receives in response to the request for $memberOfAlpha(c2,X)$ (besides the answers of the goal) is a notification that $memberOfAlpha(c2,X)$ is in a loop, say with identifier id_l . ri can observe that id_l corresponds to a request higher than the request for $memberOfAlpha(ri,X)$, i.e., that the loop coordinator is higher than $memberOfAlpha(ri,X)$. However, ri cannot infer whether the loop was formed by its own request (in which case ri would learn that G_{C2} is $memberOfAlpha(c2,X)$), or by a request issued by $c2$ when evaluating goal $memberOfAlpha(c2,X)$, or even by the evaluation of a goal on which $memberOfAlpha(c2,X)$ depends. In other words, because of the variable-length and untraceable nature of identifiers, ri does not know the number of goals in the path from $memberOfAlpha(ri,X)$ to the loop coordinator.

In addition to the information above, if the identifiers used in the computations were fixed-length, ri would also be able to learn that:

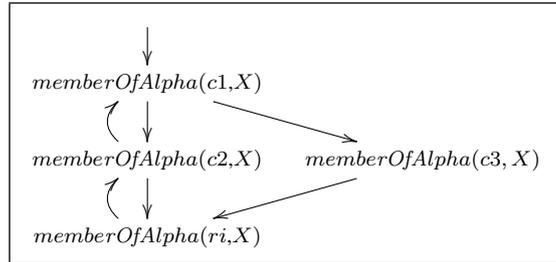
- one of the paths from $memberOfAlpha(c1,X)$ to $memberOfAlpha(ri,X)$ consists of *three* goals: $memberOfAlpha(c1,X)$, G_{C2} , and $memberOfAlpha(ri,X)$. Furthermore, ri can infer that G_{C2} is the coordinator of loop id_l . However, ri still does not learn whether G_{C2} is $memberOfAlpha(c2,X)$;
- the other path from $memberOfAlpha(c1,X)$ to $memberOfAlpha(ri,X)$ consists of *three* goals: $memberOfAlpha(c1,X)$, G_{C3} , and $memberOfAlpha(ri,X)$.

The part of the call graph that ri can infer in a computation with fixed-length identifiers is shown in Figure 7(b). Note that, in this example, the information that ri would be able to infer is the same independently from the traceability of the identifiers. This is because ri already knows all the principals in the paths from $memberOfAlpha(c1,X)$ to $memberOfAlpha(ri,X)$: $c1$ is the principal to whom ri sent the initial request, and $c2$ and $c3$ are the principals from whom ri received the request for $memberOfAlpha(ri,X)$. Even with traceable identifiers, ri would not be able to infer more information about the path from $memberOfAlpha(c2,X)$ to G_{C2} , as the only information received by ri from $c2$ is the loop identifier.

A principal might attempt to infer a bigger portion of the call graph by issuing several requests for the goals defined by different principals in the trust management system. For instance, ri can infer more information by issuing a request for each goal defined by companies $c1$, $c2$, and $c3$. In particular, by issuing a request for goal $memberOfAlpha(c2,X)$, in a computation with fixed-length and traceable identifiers ri would learn that $memberOfAlpha(c2,X)$ is the goal defined by $c2$ that depends on $memberOfAlpha(ri,X)$ (Figure 8(a)). By also issuing a request for goal $memberOfAlpha(c3,X)$, ri could infer the whole call graph (Figure 8(b)). Notice, however, that the global policy considered here is a relatively simple policy with few goals and few principals defining them. A more complex policy would complicate and sometimes prevent the inference of goal dependencies. Furthermore, some information about the example global policy would not be deducible by ri when using variable-length and untraceable identifiers. All the edges in the call graph in Figure 8(b), for instance, would be dashed edges if variable-length untraceable identifiers were used. Finally, it is worth noting that even though ri might be able to learn the whole call graph,



(a) Part of the Call Graph That Can Be Inferred by ri By Requesting Goal $memberOfAlpha(c2, X)$



(b) Part of the Call Graph That Can Be Inferred by ri by Issuing a Request for the Goals Defined by Each Principal in the Trust Management System

Fig. 8. Part of the Call Graph That Can Be Inferred by ri by Issuing Multiple Requests Using Fixed-Length Identifiers

that call graph might correspond to different intensional policies (Costantini 2001). For example, ri is not able to learn whether $memberOfAlpha(c2, X)$ and $memberOfAlpha(c3, X)$ are connected by disjunction or conjunction in $c1$'s policy.

To conclude, we argue that when using an appropriate identifier domain the knowledge about goal dependencies disclosed by GEM is not sufficient for a principal b to infer the intensional policy relative to a goal G_a defined by a principal a . Principal b always learns whether goal G_a depends on goals defined in her policy, but most likely not all the goals in the global policy on which G_a depends. Consider, for instance, clause 2 on page 2. A principal other than $c1$ and mc cannot learn that $memberOfAlpha(c1, X)$ depends on $projectPartner(mc, Y)$. We believe that the information that principal b can infer is consistent with the concept of trust management. In fact, if G_a depends on a goal defined by b , then there is a chain of trust from principal a to principal b ; it seems legitimate that the existence of such a chain of trust may not remain secret to b .

5 Practical Evaluation

We implemented the algorithms presented in Section 3.3 in Java and conducted several experiments to evaluate the performance of GEM. In particular, we first tested the implementation with the example policies defined in Section 3.2 and Appendix A. Then, we modified those policies to assess the scalability of GEM with respect to an increase in the

number of principals in the trust management system and the number of clauses in the global policy.

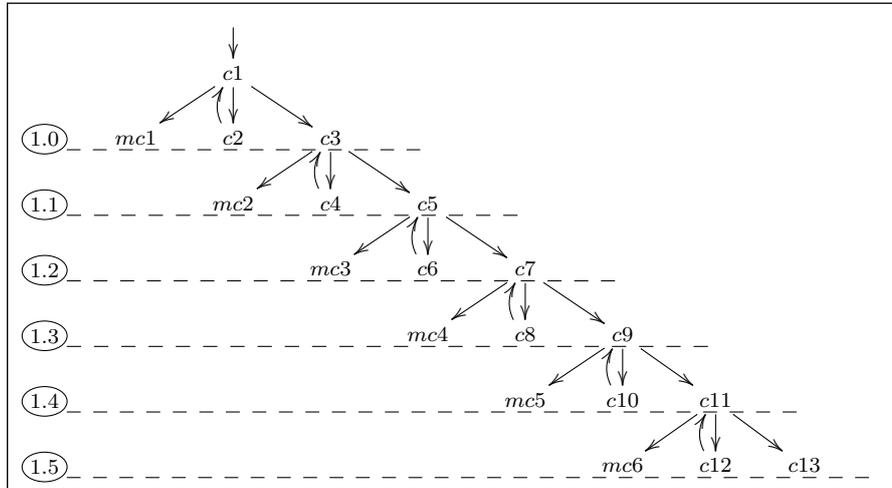
We carried out the experiments by running GEM on four machines located in different area networks. More precisely, we employed two machines located within the Eindhoven University of Technology (TU/e) network, and two located at the University of Twente (UT). The two TU/e machines mount an Intel Core 2 Quad 2.4 GHz processor with 3 GB of RAM and a 32 bit Windows operating system. The UT machines are 32 bit Ubuntu machines with the same processor but 2 GB of RAM. In each experiment, we have assigned approximately one fourth of the principals (i.e., one fourth of the local policies) in the global policy to each machine. The exchange of messages between principals on different machines is via HTTP (`javax.servlet.http` Servlet API).

To present the results of the experiments, we group them into two sets. The first set of experiments (Section 5.1) studies the performance of GEM for an increasing number of principals, clauses, and loops in the global policy. The second set (Section 5.2) shows, for some of the global policies in the first set, the effects of increasing the size of the extensional policy (i.e., the number of facts in the global policy). For each experiment we report the following results:

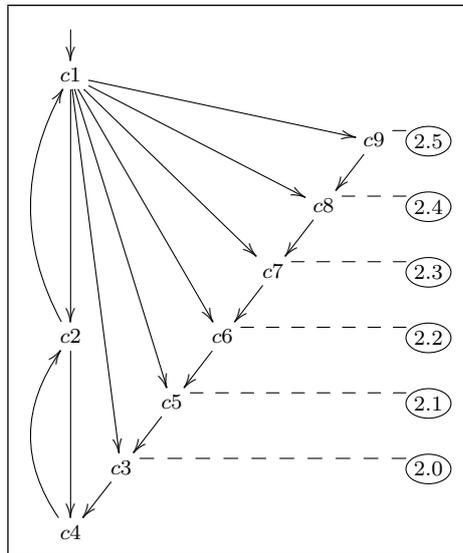
- the number of principals involved in the computation (denoted by *Princ*);
- the number of tables created by GEM during the computation (*Tab*);
- the number of clauses evaluated (*Clauses*);
- the sum of the computation times on the four machines, expressed in milliseconds (*CTime*);
- the total time (*TTime*), expressed in milliseconds, given by the *CTime* plus network communication time;
- the total memory occupied by GEM on the four machines, expressed in kilobytes (*TMem*);
- the maximum memory occupied by the tables of the goals created by GEM on the four machines, expressed in kilobytes (*TabMem*);
- the memory occupied by the tables of the goals created by GEM on the four machines after the tables' disposal, expressed in kilobytes (*EndTabMem*);
- the number of requests issued during the computation between principals on different machines (*Req*);
- the number of loops identified during the computation (*Loops*);
- the number of response messages exchanged between principals on different machines (*Resp*);
- the number of non-empty response messages (i.e., response messages containing at least one answer) exchanged between principals on different machines (*Resp&Ans*);
- the total number of answers computed by GEM during the evaluation of the policy (*Ans*).

5.1 Experiments Set 1: Increasing the Number of Principals, Clauses, and Loops

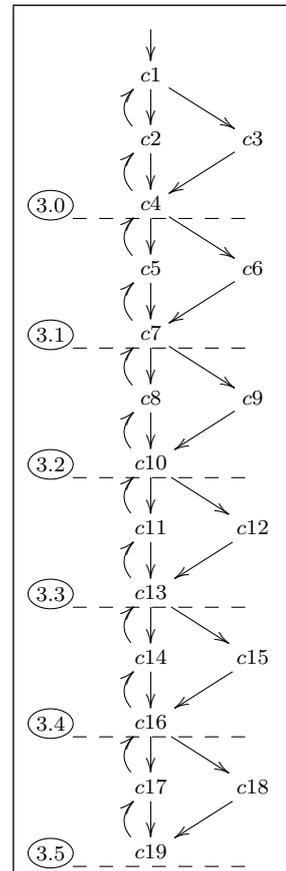
In the first set of experiments we conducted three groups of (sub)experiments to evaluate the performance of GEM in response to an increase in: (1) the number of principals and



(a) Call Graph of the Global Policies 1.0 to 1.5



(b) Call Graph of the Global Policies 2.0 to 2.5



(c) Call Graph of the Global Policies 3.0 to 3.5

Fig. 9. Call Graph of the Global Policies Used in the Experiments Set 1

ID	Princ	Tab	Clauses	TTime (CTime) in ms	TMem in kB	TabMem (EndTabMem) in kB	Req	Loops	Resp (Resp&Ans)	Ans
1.0	4	4	6	286,1 (6,3)	32	18 (15)	5	1	9 (6)	9
1.1	7	7	12	305,9 (8,5)	53	39 (27)	9	2	17 (11)	26
1.2	10	10	18	315,2 (11,2)	78	66 (44)	13	3	25 (17)	49
1.3	13	13	24	322,8 (14,7)	108	97 (63)	17	4	33 (22)	78
1.4	16	16	30	327,0 (17,0)	140	134 (85)	21	5	41 (27)	113
1.5	19	19	36	331,7 (19,8)	180	175 (108)	25	6	49 (33)	154
2.0	4	6	8	830,3 (10,7)	47	37 (32)	10	4	31 (17)	20
2.1	5	10	10	1037,2 (13,0)	68	61 (50)	16	6	48 (19)	32
2.2	6	15	12	1283,8 (16,5)	92	91 (74)	23	8	72 (30)	46
2.3	7	21	14	1459,9 (19,9)	124	122 (100)	31	10	95 (34)	62
2.4	8	28	16	1673,6 (23,4)	163	153 (127)	40	12	125 (47)	80
2.5	9	36	18	1840,0 (27,2)	206	189 (159)	50	14	154 (53)	100
3.0	4	6	8	830,3 (10,7)	47	37 (32)	10	4	31 (17)	20
3.1	7	16	16	1653,6 (23,9)	151	149 (110)	28	12	113 (71)	112
3.2	10	36	24	3029,0 (45,2)	454	427 (297)	64	28	309 (199)	384
3.3	13	76	32	5672,4 (78,8)	1210	1133 (752)	136	60	773 (535)	1088
3.4	16	156	40	10479,7 (134,4)	3024	2827 (1824)	280	124	1789 (1276)	2800
3.5	19	316	48	21939,5 (310,2)	7285	6862 (4267)	568	252	4569 (3477)	6816

Table 1. Performance Evaluation Results: Experiments Set 1

clauses, (2) the number of loops, and (3) both the number of principals, clauses and loops in a global policy. To this end, we modified the global policies defined in Section 3.2 and Appendix A as shown in Figure 9. In particular, we created six variants of the global policy in Appendix A to evaluate GEM in response to an increase in the number of principals and clauses (experiments group 1). The call graph of the six variants of the global policy is shown in Figure 9(a); in the figure, we use identifiers from 1.0 to 1.5 to denote the variants, where variant 1.0 represents the original policy in Appendix A. To keep the figure as simple yet informative as possible, we label the nodes in the graph with the identifier of the principal evaluating the goal they represent rather than with the goal itself, as for the purpose of this section the number of principals involved in a computation is more relevant than the goals they evaluate.

In the second group of experiments we created six variant of the global policy defined in Section 3.2 to increase the number of loops in the global policy. Figure 9(b) shows the call graph of the six variants of the policy, which are denoted by identifiers 2.0 to 2.5; variant 2.0 corresponds to the original global policy. Similarly, other six variants of the same policy were created for the experiments in the third group, in order to increase the number of both principals, clauses, and loops; the call graph of the six variants, denoted by identifiers 3.0 to 3.5 (where 3.0 corresponds to the original global policy defined in Section 3.2), is shown in Figure 9(c). Notice that, for the sake of compactness, Figures 9(b) and 9(c) show the folded graph of the global policies, i.e., they do not represent the reevaluation of goals due to side requests. Since in GEM the computation is based on the unfolded versions of the graph (e.g., the graph in Figure 5(b) on page 14 for variants 2.0 and 3.0), the number of lower requests occurring in the actual computation is higher than the one displayed by the figures. For instance, the number of lower requests for the leader of the SCC goes up to seven in experiments 2.5 and 3.5.

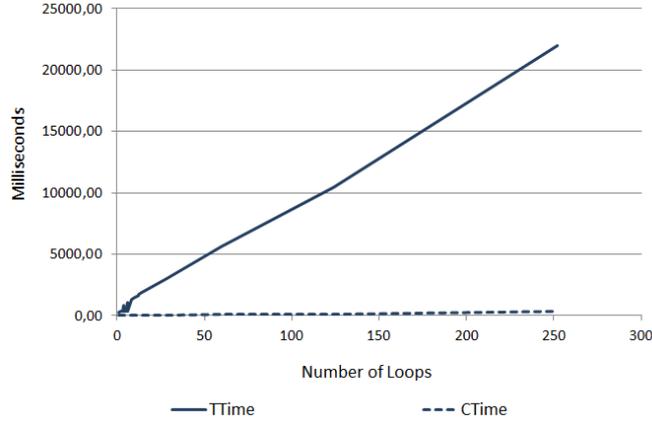
Table 1 presents the results of the three groups of experiments. Each row in the table

shows the results for the variant of the global policy with identifier indicated in column *ID*. The total time (*TTime*) and the computation time (*CTime*) are the average times for 100 runs of each experiment; the values in all the other columns are constant for every run, as they depend on the structure of the global policy.

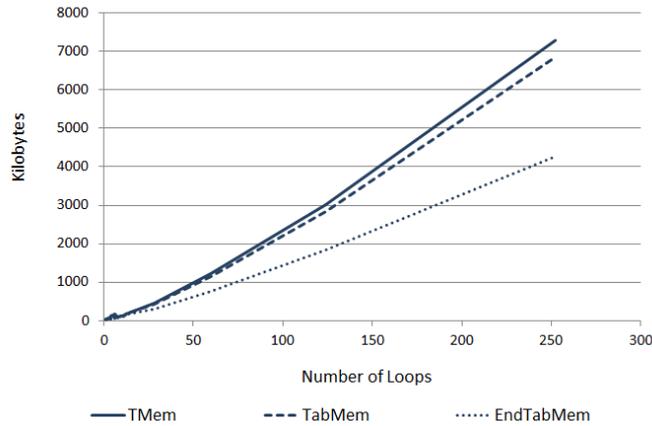
Before interpreting the results of the experiments, let us provide some general comments on the relationship between the values in different columns of Table 1. First of all, in every experiment the number of requests (*Req*) is equal to the number of tables generated (*Tab*) plus the number of loops identified by GEM (*Loops*); this is because in GEM a request is either a higher request, which leads to the creation of a table, or a lower request, and thus forms a loop. Second, the number of response messages (*Resp*) is always at least as large as the number of requests, since to every request is given a response, even if with an empty set of answers. The number of empty response messages (i.e., response messages with an empty set of answers) can be observed by subtracting the number of response messages containing at least one answer (*Resp&Ans*) from the total number of response messages *Resp*. Finally, for some of the experiments (namely in the computation of variants 2.2 to 2.5 and 3.2 to 3.5), the number of tables generated by GEM is higher than the number of clauses (*Clauses*) and thus the number of goals defined in the global policy. This is because while in column *Tab* we report the *total* number of tables generated during a computation, column *Clauses* shows the number of *different* clauses being evaluated. In other words, we do not count twice the clauses used for the evaluation of a goal that is reevaluated because of a side request. The reason behind this choice is that, on the one hand, we are interested in showing the impact of the number of tables generated and answers computed (*Ans*) during the evaluation of a policy on the memory usage (*TMem* and *TabMem*); on the other hand, considering the number of different clauses gives a better insight on the size of a policy. Therefore, we point out that the number of clauses actually evaluated during a computation is higher than the value reported in column *Clauses*, since some of those clauses are used in more than one table.

A first interesting outcome of the experiments is that the time and memory results in Table 1 increase approximately linearly with the number of loops in the global policy (Figure 10). When the number of loops is low (experiments in groups 1 and 2) the time and memory usage are negligible, but as the number of loops increases considerably (experiments group 3), the *TTime*, *TMem*, and *TabMem* get substantially higher. This is because an increase in the number of loops (especially if nested, as in the global policy in Figure 9(c)), leads to an increase in the number of response messages, answers, and (since GEM reevaluates side requests) tables generated in a computation.

Another interesting result is represented by the difference between total time and computation time (Figure 10(b)). In fact, the *TTime* ranges from 16.7 times the *CTime* for policy variant 1.5, up to 80.1 times the *CTime* for variant 2.1. This implies that most of the *TTime* is devoted to network communication. In the graph in Figure 10(b), the number of messages represented on the x-axis is given by the sum of the request (*Req*) and response (*Resp*) messages in Table 1; therefore, we can conclude that the larger the number of requests and response messages and the number of answers per response message, the higher the proportion between *TTime* and *CTime*. Furthermore, we point out that in a real distributed system where to each principal corresponds a different machine (possibly



(a) Total and Computation Time for an Increasing Number of Loops in the Computation



(b) Total and Tables Memory for an Increasing Number of Loops in the Computation

Fig. 10. Time and Memory Results for Experiments Set 1

in a different area network) the $TTime$ would be much higher than the results in Table 1, especially when the number of principals increases.

Finally, Table 1 shows that the number of answers in a computation is always larger than the number of response messages containing at least one answer; in other words, each message in $Resp\&Ans$ carries more than one answer on average. This information, combined with the observation on the difference between $TTime$ and $CTime$ displayed in Figure 10(b), suggests that GEM can consistently reduce network overhead with respect to other goal evaluation algorithms which send one message for each computed answer (e.g., (Alves et al. 2006)), especially when the number of principals and facts in the global policy increases.

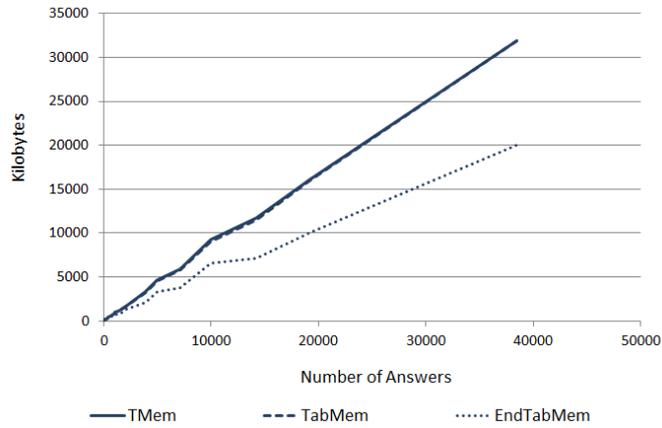
ID	Princ	Tab	Clauses	TTime (CTime) in ms	TMem in kB	TabMem (EndTabMem) in kB	Req	Loops	Resp (Resp&Ans)	Ans
1.0	4	4	6	286,1 (6,3)	32	18 (15)	5	1	9 (6)	9
1.0a	4	4	24	293,7 (13,5)	81	69 (58)	5	1	9 (6)	72
1.0b	4	4	104	334,8 (33,9)	308	300 (255)	5	1	9 (6)	352
1.0c	4	4	204	519,8 (78,8)	590	586 (498)	5	1	9 (6)	702
1.5	19	19	36	331,7 (19,8)	180	175 (108)	25	6	49 (33)	154
1.5a	19	19	144	391,6 (80,3)	1215	1210 (790)	25	6	49 (33)	1432
1.5b	19	19	624	1589,4 (817,7)	5857	5764 (3786)	25	6	49 (33)	7112
1.5c	19	19	1224	4356,5 (2986,9)	11711	11541 (7098)	25	6	49 (33)	14212
2.0	4	6	8	830,3 (10,7)	47	37 (32)	10	4	31 (17)	20
2.0a	4	6	26	843,7 (23,9)	192	184 (144)	10	4	31 (17)	200
2.0b	4	6	106	948,4 (68,6)	846	844 (651)	10	4	31 (17)	1000
2.0c	4	6	206	1441,2 (161,3)	1672	1658 (1279)	10	4	31 (17)	2000
2.5	9	36	18	1840,0 (27,2)	206	189 (159)	50	14	154 (53)	100
2.5a	9	36	36	1871,8 (58,1)	1017	976 (726)	50	14	154 (53)	1000
2.5b	9	36	116	2490,1 (259,1)	4665	4534 (3315)	50	14	154 (53)	5000
2.5c	9	36	216	3875,9 (755,2)	9225	8984 (6538)	50	14	154 (53)	10000
3.2	10	36	24	3029,0 (45,2)	454	427 (297)	64	28	309 (196)	384
3.2a	10	36	78	3142,3 (158,1)	3306	3259 (2081)	64	28	309 (196)	3840
3.2b	10	36	318	6277,8 (1789,6)	16103	15954 (10032)	64	28	309 (196)	19200
3.2c	10	36	618	13906,3 (6346,1)	31938	31883 (19968)	64	28	309 (196)	38400

Table 2. Performance Evaluation Results: Experiments Set 2

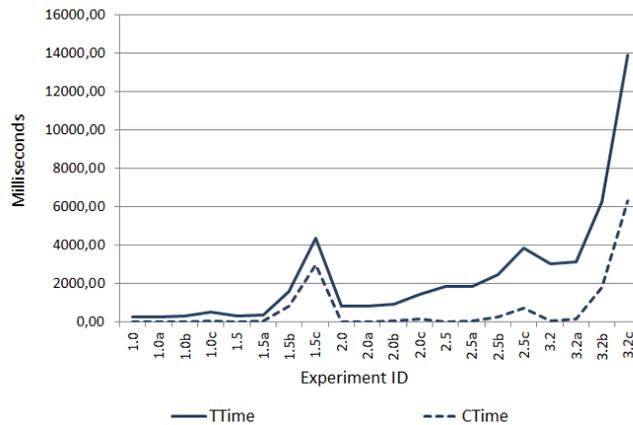
5.2 Experiments Set 2: Increasing the Size of Extensional Policies

In a real-world policy, we expect the size of the extensional policy (i.e., the number of facts that can be derived from the policy) to substantially exceed the size of the intensional policy (i.e., the number of clauses used to derive new facts). Consider, for example, the students of a university: while there are only a few rules that define the procedure for becoming a student, the number of students usually goes beyond several thousands. The goal of the experiments presented in this section is to evaluate the impact on the performance of GEM of an increase in the number of facts in a global policy. To this end, we considered some of the global policies introduced in Section 5.1 and increased the size of their extensional policies by a factor of 10, 50, and 100; in particular, we modified variants 1.0, 1.5, 2.0, 2.5, and 3.2. We could not perform experiments on variants 3.3 to 3.5 due to the limited memory available on some of the machines used in the experiments.

Table 2 shows the results of the second set of experiments. In the table, suffix “a” on a variant’s identifier indicates an increase by a factor of 10 of the number of facts in that variant of the global policy, suffix “b” indicates an increase by a factor of 50, and suffix “c” indicates an increase by a factor of 100. Notice that the number of answers (*Ans*) computed in the experiments on variants 1.0 and 1.5 grows less than a factor of 10, 50, and 100 because, in order not to modify the structure of the global policy, the number of facts in the (local) policies of principals *mc1* to *mc6* in Figure 9(a) was not increased; more precisely, the policy of principals *mc1* to *mc6* always consists of only two facts. Similarly to the previous experiments, the total time (*TTime*) and the computation time (*CTime*) are the average times for 100 runs of each experiment.



(a) Memory Results with Respect to the Number of Answers Derived in the Computation



(b) Time Results with Respect to the Number of Messages Exchanged in the Computation

Fig. 11. Time and Memory Results for Experiments Set 2

The results in Table 2 show that memory usage (*TMem* and *TabMem*) and computation time grows faster than the other values for policies with a very large number of computed answers (i.e., variants 1.5c, 2.5c, and 3.2c). For what concerns memory usage (Figure 11(a)), the extra overhead is due to the accompanying increase of the information that needs to be stored in tables (i.e., clauses, loops, answers). After the disposal of the tables employed in the computation, there is a decrease of up to 38% of the memory usage (*EndTabMem*). This suggests that it is very important to delete as much information as possible from the table of a goal when the goal is completely evaluated, as this leads to a substantial reduction of memory usage. In this respect, GEM has the advantage of enabling principals to detect when the evaluation of the single goals involved in a computation is completed, and immediately clean up the table of those goals.

Further to the increase in the number of exchanged messages, we believe the extra over-

head in $TTime$ and $CTime$ in computations with a large number of answers (the peaks for variants 1.5c, 2.5c, and 3.2c in Figure 11(b)) to be due to the growth of the data structures used by GEM to search, for example, for new answers and clauses to be activated. In addition, contrarily to experiments set 1, in this set of experiments the $TTime$ is not always dominated by network communication time. In particular, for variant 1.5c the $CTime$ is twice the network time, and for variant 1.5b the $CTime$ is slightly larger than the network time. In the remaining experiments the $TTime$ ranges from 2.2 times the $CTime$ for policy variant 3.2c, up to 77.7 times the $CTime$ for variant 2.0. Furthermore, we can notice that the difference between $TTime$ and $CTime$ always decreases as the number of facts in the global policy grows. This is due to the fact that the number of messages exchanged between principals remains constant while the size of the extensional policy is increased.

To conclude, we highlight again how the “wait” mechanism that GEM employs to collect a maximum set of answers before sending a response can contribute to reduce the network overhead, especially for global policies with a large extensional policy. For example, in the experiment on variant 3.2c, GEM sends “only” 196 response messages, while other distributed goal evaluation algorithms (e.g., (Alves et al. 2006)) would send 38400 messages, one for each computed answer. Intuitively, the latter approach would lead to a network communication time much higher than the 7.5 seconds spent by GEM.

6 Dealing with Negation

GEM is devised to work with *definite* logic programs, i.e., programs without negation. Negation is used by some trust management systems (e.g., (Czenko et al. 2006; Dong and Dulay 2010)) to express non-monotonic constraints in the body of a clause, such as separation of duty or “distrust” in principals with certain attributes (e.g., employees of a rival company). Here, we discuss how GEM can be extended to support the use of *negation as failure*.

Negation as failure is an inference rule that derives the truth of a negated body atom $not(B)$ by the failure to derive B . The problem when allowing the use of negation (as failure) is that in the presence of *loops through negation* (i.e., loops involving negated atoms) a program may have several minimal models (Gelfond and Lifschitz 1988). For instance, program $p \leftarrow not(q), q \leftarrow not(p)$ has two minimal models: $\{p\}$ and $\{q\}$. In addition, these two models are not well-founded (Van Gelder et al. 1991), as there is no clause in the program demonstrating that p respectively q are true. Another undesired consequence of loops through negation is that they may introduce “inconsistencies” in a program, as shown in the example at the end of this section. Actually, there are additional consequences of loops through negation (Apt and Bol 1994; Van Gelder et al. 1991), which we do not discuss further as they go beyond the scope of this paper. In fact, our goal here is not to have a full-fledge handling of negation, but to allow the use of negation in policies though preventing loops through negation.

Loops through negation are a very well-studied issue in the logic programming literature. There are three standard solutions to the problems they raise: (a) forbidding the presence of loops through negation, as done by the weakly perfect model semantics (Przymusinska and Przymusinski 1990), which is defined only for *weakly stratified programs* (that include *locally stratified* (Przymusinski 1988) and *stratified* programs (Apt et al. 1988)); (b) using

a three-valued semantics (Przymusiński 1990), including the truth value *undefined* next to *true* and *false*, as done, e.g., by the well-founded semantics (Van Gelder et al. 1991) and the Fitting’s semantics (Fitting 1985), where the semantics of p in the program $p \leftarrow \text{not}(p)$ is *undefined*; (c) following a multi-model approach, as in the case of stable models (Gelfond and Lifschitz 1988).

Our solution follows an approach similar to (a), since solutions (b) and (c) are not suitable for trust management systems. In fact, relying on a three-valued semantics (b) requires additional mechanisms to determine whether the truth value of the goals involved in a loop through negation is true, false, or *undefined* (e.g., delaying in SLG resolution (Chen and Warren 1996)). In trust management, however, loops through negation are inherently wrong, as they indicate conflicting policy statements issued by principals among which there is a mutual trust relationship, and thus should not be processed. Similarly, solution (c) would imply that an access request should be either granted or revoked depending on which (truth) value we “choose” to assign to the goals in a computation, which is clearly not a safe approach. Solution (a) can also not be applied straightforwardly in our context because the definition of weakly stratified program relies on a “global ordering” among *all* the (ground) atoms in a global policy. This would require principals to agree beforehand on the allowed dependencies among (ground) goals; however, in a trust management system principals often do not know each other until their first interaction. Therefore, rather than forbidding the presence of loops through negation, we prevent their evaluation. In this respect, the added value of GEM is its ability to detect loops at runtime. We exploit this feature by introducing an additional runtime check to the algorithm, which causes the computation to flounder if a loop involving a negated goal is detected. The check is safe in that if the computation does not flounder, then it always returns a correct answer.

In summary, GEM can be extended to allow the use of negation in policy statements as follows. Given a clause with a literal $\text{not}(B)$ selected for evaluation:

1. if B is not ground, an error is raised and the computation flounders;
2. if the evaluation of atom B succeeds with an answer, then $\text{not}(B)$ fails and the clause is disposed;
3. if B is completely evaluated and has no answers, then $\text{not}(B)$ succeeds and a new node is added to the evaluation tree of the goal, removing $\text{not}(B)$ from the body;
4. if a loop notification for atom B is received, an error is raised and the computation flounders.

Conditions (1), (2) and (3) are standard when defining negation as failure: (1) is necessary to guarantee correctness (Apt 1990), while (2) and (3) define the semantics of negation. Notice that condition (3) captures also the case of *infinite failure*, as done, for instance, by the well-founded semantics (Van Gelder et al. 1991). For example, given a policy composed of clauses $q \leftarrow \text{not}(p)$ and $p \leftarrow p$ and a goal q , GEM first completely evaluates clause $p \leftarrow p$ by detecting the loop and deducing that no answer of p can be derived; consequently, it concludes that q is true. Condition (4) states that the algorithm flounders if it detects a loop through negation. The “floundering message” is propagated through all the goals involved in the loop similarly to a response message, so that their evaluation is aborted.

Note that floundering due to condition (1) can be avoided by restricting to *well-moded*

programs (Apt and Marchiori 1994). Differently from weakly stratified programs, which require an ordering among all the goals in a global policy, the definition of well-moded program requires each clause *independently* to be well-moded. Therefore, by requiring local policies to be well-moded, this type of floundering can be prevented.

It is straightforward to demonstrate that the proposed extension of GEM:

- always terminates (for arbitrary global policies and requests), because the only difference with the standard GEM algorithm (which terminates) is the presence of an additional termination condition, and
- for non-floundering computations, it is sound and complete with respect to the stable models and well-founded semantics.

We now show how the extended algorithm deals with negation by means of an example. We consider a scenario inspired by the one introduced in Section 1, where the pharmaceutical company $c1$ needs to determine the set of principals participating to project *Alpha*. Project *Alpha* is a multidisciplinary project which requires the collaboration of experts from several fields: physicians, biologists, chemists, etc. Company $c1$ already formed a team of qualified chemists to work on the project, and delegates to the partner company $c2$ the authority of determining the remaining project members. To avoid interference with the work of its trusted chemists, however, $c1$ wants to prevent chemists of $c2$ to take part to the project. In its definition of project members, $c2$ includes also the members of project *Alpha* at $c1$. This scenario can be represented by the following policy statements:

1. $\text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c2, X), \text{not}(\text{chemist}(c2, X)).$
2. $\text{memberOfAlpha}(c1, \text{david}).$
3. $\text{chemist}(c1, \text{david}).$
4. $\text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c1, X).$
5. $\text{memberOfAlpha}(c2, \text{alice}).$
6. $\text{chemist}(c2, \text{alice}).$
7. $\text{memberOfAlpha}(c2, \text{eric}).$

To compute the answers of goal $\text{memberOfAlpha}(c1, X)$, GEM proceeds as follows. First, clause 1 is evaluated by $c1$, leading to a request for goal $\text{memberOfAlpha}(c2, X)$ to $c2$. The evaluation of the first applicable clause in $c2$'s policy (clause 4) forms a loop, identified by $c1$. The loop processing phase continues at $c2$, which identifies the first two answers of goal $\text{memberOfAlpha}(c2, X)$ (i.e., $\text{memberOfAlpha}(c2, \text{alice})$ and $\text{memberOfAlpha}(c2, \text{eric})$, clauses 5 and 7 respectively) and sends them to $c1$. For each of these answers, $c1$ requests to $c2$ whether the project member is a chemist. The evaluation of goal $\text{chemist}(c2, \text{alice})$ succeeds at $c2$ (clause 6), while goal $\text{chemist}(c2, \text{eric})$ fails; therefore, their negated counterpart in clause 1 fails and succeeds respectively, leading to a new answer at $c1$, namely $\text{memberOfAlpha}(c1, \text{eric})$. This answer, together with the answer derived by evaluating clause 2 (i.e., $\text{memberOfAlpha}(c1, \text{david})$), is sent by $c1$ to $c2$, starting the second iteration of the loop. In this iteration, $c2$ finds one new answer, $\text{memberOfAlpha}(c2, \text{david})$, which is immediately returned to $c1$. Now, $c1$ evaluates clause 1 based on the new answer received from $c2$. Since David is not a chemist at $c2$, $c1$ derives again answer $\text{memberOfAlpha}(c1, \text{david})$, which had already been computed in the previous iteration. As a result, the evaluation

of goal $memberOfAlpha(c1,X)$ terminates with two answers: $memberOfAlpha(c1,eric)$ and $memberOfAlpha(c1,david)$.

The example above shows that GEM can easily support policies including both loops and negation. We now show an example of how GEM operates in the presence of loops through negation. Consider, for instance, the following policy statements complementing the global policy above:

8. $chemist(c2,X) \leftarrow memberOfAlpha(c1,X), chemist(c3,X)$.
9. $chemist(c3,eric)$.

Clause 8 states that all the members of project *Alpha* at $c1$ that work as chemists at the other partner company $c3$ are also chemists at $c2$. Notice that this clause is “inconsistent” with clause 1. In fact, clause 1 defines as members of project *Alpha* principals that *are not* chemists at company $c2$; at the same time, clause 8 states that chemists at $c2$ are members of project *Alpha* at $c1$. For this reason, when $c1$ ascertains that goals $memberOfAlpha(c1,eric)$ and $chemist(c2,eric)$ are in a loop, it raises an error and the computation flounders. In fact, in the example computation above, the evaluation of goal $chemist(c2,eric)$ by $c2$ would lead to a contradiction: if Eric were not a chemist at $c2$, he would be a member of project *Alpha*; however, if Eric were a member of project *Alpha*, he would be a chemist at $c2$.

7 Related Work

Research on goal evaluation has been carried out both in the field of logic programming and trust management. In this section we compare our work with existing frameworks focusing on the information disclosed during the evaluation process, i.e., based on the classification criteria defined in Section 3.1. Additionally, we indicate whether the analyzed systems employ a centralized or distributed goal evaluation strategy and discuss the termination detection mechanism they adopt. Within termination detection, we distinguish between termination of the whole computation initiated by a particular request and termination of the single goals involved in the computation (i.e., detecting when a goal is completely evaluated). Table 3 summarizes the results of this analysis. In the table, LP denotes the algorithms proposed in the logic programming domain, while TM denotes trust management systems.

SLG resolution (Chen and Warren 1996), TP-resolution (Shen et al. 2001), DRA (Guo and Gupta 2001), OPTYap (Rocha et al. 2005), and the work by Hulin (1989) are centralized tabling systems in which the complete program (i.e., the global policy) is available during the evaluation. Therefore, these five systems are classified as E1-I3 according to the classification criteria defined in Section 3.1, that is, they do not preserve the confidentiality of neither extensional nor intensional policies. SLG identifies loops by observing goal dependencies in the “call stack” of the program; termination is detected when no more operations can be applied to the goals in the stack. SLG resolution is employed in a number of Prolog systems such as, for instance, XSB (Swift and Warren 2012). The evaluation strategy employed by GEM is similar to the XSB scheduling strategy called *local evaluation*, which completely evaluates a SCC before returning the answers of the leader to a goal

	Frameworks	Evaluation	Computation Termination	Goal Termination	Classification
LP	SLG (Chen and Warren 1996)	centralized	centralized	centralized	E1-I3
	TP resolution (Shen et al. 2001)	centralized	centralized	centralized	E1-I3
	DRA (Guo and Gupta 2001)	centralized	centralized	centralized	E1-I3
	OPTYap (Rocha et al. 2005)	centralized	centralized	centralized	E1-I3
	Hulin (1989)	centralized	centralized	centralized	E1-I3
	Damasio (2000)	distributed	distributed	distributed	E1-I2
TM	Hu (1997)	distributed	distributed	distributed	E1-I2
	RT (Li et al. 2003)	centralized	centralized	centralized	E1-I3
	Tulip (Czenko and Etalle 2007)	centralized	centralized	centralized	E1-I3
	SecPAL (Becker et al. 2010)	centralized	centralized	centralized	E1-I3
	SD3 (Jim and Suciu 2001)	distributed	N/A	N/A	E1-I3
	Becker et al. (2009)	distributed	N/A	N/A	E1-I3
	Cassandra (Becker 2005)	distributed	no	no	E1-I0
	PeerTrust (Alves et al. 2006)	distributed	distributed	no	E1-I1
		distributed	distributed	distributed	E1-I2
	MTN (Zhang and Winslett 2008)	distributed	distributed	no	E1-I1
GEM	distributed	distributed	distributed	E1-I1	

Table 3. Comparison Between Goal Evaluation Algorithms

outside the SCC. Similarly to SLD resolution (Kowalski 1974), in TP-resolution and DRA a goal is evaluated by building a single derivation tree for the goal. Loops are detected when a subgoal appears more than once in a branch of the tree, and the evaluation of a goal terminates when there are no more nodes in the derivation tree to be evaluated. OPTYap and Hulin propose a parallel tabled execution strategy to improve the efficiency of goal evaluation. OPTYap resorts to centralized data structures to identify loops and detect termination. In (Hulin 1989), each process communicates its termination to a global variable, whose access is limited to one process at a time by means of a deadlock mechanism.

Distributed goal evaluation frameworks are presented in (Hu 1997) and (Damásio 2000). To detect termination, the work by Hu (1997) assumes the presence of global data structures and requires goal dependencies to be propagated among the different principals. In (Damásio 2000), termination detection resorts to a static dependency graph known to all principals and determined at compile time. Consequently, the confidentiality of (part of) the intensional policies is not preserved, and both algorithms are classified as E1-I2.

In trust management, distributed goal evaluation is a main issue since policies are distributed among principals. The trust management frameworks RT (Li et al. 2003) and Tulip (Czenko and Etalle 2007) rely on a centralized goal evaluation strategy, where all the clauses necessary for the evaluation of a goal are collected in a single location. Similarly, SecPAL (Becker et al. 2010) assumes all the clauses in a global policy to be available to the principal responsible for the evaluation of a goal. In SD3 (Jim and Suciu 2001), when queried for a goal, a principal returns to the requester the clauses defining the goal, with the locally defined (body) atoms already evaluated. Becker et al. (2009) present an algorithm in which the body atoms of the clauses defining a goal are sent in turn to the principals defining them; each principal evaluates the atom(s) defined in her policy and sends its answers and the remaining atoms to the next principal, until the evaluation fails or all the atoms are evaluated. As a result, policy confidentiality is not preserved by any of

these algorithms, which are thus classified as E1-I3. Furthermore, neither Jim and Suciu (2001) nor Becker et al. (2009) discuss how termination is detected. Cassandra (Becker 2005) employs a distributed evaluation strategy in which no information about intensional policies is disclosed. However, it does not detect neither the complete evaluation of single goals, nor the termination of the whole computation.

PeerTrust (Alves et al. 2006) and MTN (Zhang and Winslett 2008) detect termination of the computation started by a particular request in a fully distributed way; this is achieved by “observing” when no more messages are exchanged among principals and all the goals are quiescent. In (Alves et al. 2006), the authors present two solutions: the first, based on the work in (Damásio 2000), is also able to detect the completion of single goals, but requires the dependency graph of the global policy to be known to all principals beforehand. The second solution, which is also adopted in (Zhang and Winslett 2008), detects termination of the computation without disclosing information about intensional policies. However, since all request and response messages are tagged with the identifier of the initial request, some information about goal dependencies can be inferred (hence the E1-I1 classification); more precisely, a principal can learn whether a given goal depends on a goal defined in her policy. In addition, neither PeerTrust nor MTN features a loop identification mechanism. Consequently, they are not able to detect termination of individual goals, which is required to free the resources used during the computation and to allow the use of negation. Furthermore, when using negation, the detection of loops through negation allows to preserve the soundness and completeness of the computation with respect to the standard semantics for logic programs. We enable the identification of loops and the detection of goal termination at the cost of possibly revealing more information about goal dependencies. In fact, in GEM all the principals involved in a loop are notified about the loop: on the one hand, this enables the principal(s) handling negated goals to terminate the computation with floundering. On the other hand, this implies that GEM discloses information about the presence of mutual dependencies among goals to more principals than PeerTrust and MTN. In the example in Section 4.2, for instance, with PeerTrust and MTN the research institute ri would not receive any loop notification from company $c2$; therefore, ri would not learn that there exists a mutual dependency between goals $memberOfAlpha(ri, X)$ and $memberOfAlpha(c2, X)$.

Besides the protection of intensional policies, preserving the confidentiality of extensional policies is also an important requirement of trust management systems, as the answers of a goal might contain sensitive information (e.g., the list of patients of a mental hospital). Even though none of the existing goal evaluation algorithms satisfies this requirement (see Table 3), GEM can be easily adapted to protect the confidentiality of extensional policies. In particular, by enabling the distributed evaluation of policies, GEM allows principals to discriminate between goals that may be accessed by other principals and goals that may only be used for internal computations, because of their sensitivity. This distinction is not possible when using an algorithm that relies on a centralized evaluation strategy. A finer-grained protection of extensional policies can be achieved by integrating GEM with trust negotiation algorithms (Winsborough et al. 2000; Winslett 2003). Trust negotiation algorithms protect the disclosure of extensional policies (i.e., possibly sensitive credentials) by means of *disclosure policies* that specify which credentials a requester must provide to get access to the requested credentials. Some trust negotiation algorithms

deal also with the protection of disclosure policies (e.g., (Seamons et al. 2001)); however, they assume that all the credentials of the principals in a trust management system have been already derived when a transaction takes place (Winsborough and Li 2002). GEM, on the other hand, provides a way of deriving those credentials. Thus, GEM and trust negotiation algorithms can be combined in such a way that a GEM request is evaluated only if the requester satisfies the disclosure policy of that goal, i.e., if she is trustworthy enough to see the answers to the request. The resulting integrated algorithm enables distributed goal evaluation while preserving the confidentiality of both intensional and extensional policies. A similar approach is presented in (Koshutanski and Massacci 2008; Lee et al. 2009). However, in (Koshutanski and Massacci 2008) the authors do not discuss how to deal with recursive policy statements, while the algorithm presented in (Lee et al. 2009) raises an error in the case that cyclic dependencies are detected, and for this reason is not complete. MTN (Zhang and Winslett 2008) also applies trust negotiation strategies to distributed goal evaluation, but as discussed in Section 7 this algorithm is not able to detect termination of individual goals within a computation. In (Minami et al. 2011), the authors present a framework to analyze and compare distributed goal evaluation algorithms based on the information about extensional policies that they disclose during a computation.

To conclude, we point out that contrarily to other works on goal evaluation (e.g., (Lee et al. 2010)), the distributed evaluation strategy of GEM does not allow to build the complete “proof” of a goal. Building such a proof is in fact similar to constructing the derivation tree of a goal having as root the initial goal, as branches the clauses used to evaluate the goal, and as leaf nodes the answers of its subgoals. Even though cryptographic techniques can be employed to prevent the disclosure of the facts used in the derivation process (Lee et al. 2009), the construction of such a proof cannot be obtained without disclosing the intensional policies of the principals involved in the evaluation, which is what GEM aims at avoiding. We argue that the approach followed by GEM is consistent with the concept of trust management. In trust management, in fact, if the policy of a principal a refers to the policy statements of a principal b , then principal a trusts b for the definition and evaluation of those statements. In circumstances where the proof of a goal is required, the confidentiality requirement should be put aside in favor of a goal evaluation strategy that allows the construction of such a proof (e.g., RT (Li et al. 2003)).

8 Conclusions

In this paper we have presented GEM, a distributed goal evaluation algorithm for trust management systems. Differently from many of the existing algorithms, GEM detects the termination of a computation in a completely distributed way without disclosing intensional policies, thereby preserving their confidentiality. In addition, GEM is able to detect when the single goals within a computation are fully evaluated, by enabling the identification of strongly connected components. Even though this may lead to the disclosure of some additional information about goal dependencies, it also enables the use of negation (as failure) in policies. In Section 4.2 we show that the information disclosed by GEM is not sufficient to infer the intensional policy of a principal; thus, we believe that the benefits of our solution overcome the drawbacks. GEM always terminates and is sound and com-

plete with respect to the standard semantics for logic programs. As future work, we plan to extend GEM to support constraint rules (Li and Mitchell 2003) and subsumptive tabling.

Although efficiency is not a primary objective of this paper, GEM can contribute to keep network traffic low. In fact, in most distributed goal evaluation systems (e.g., (Alves et al. 2006)) answers are sent as soon as they are computed. On the contrary, GEM delays the communication of the answers of a goal until all possible answers have been computed, i.e., until all the branches of the partial derivation tree of the goal have been inspected. This strategy may delay the identification of the answers of ground goals. However, it simplifies the termination detection mechanism and we believe reduces the number of messages exchanged by principals during a computation. The experiments presented in Section 5 suggest that since the computation time is dominated by network communication, a reduction in the number of messages exchanged between principals leads to a consistently lower computation time. In addition, since the answers of a goal can be reused for future computations, the proposed solution may reduce the computation time of later evaluations.

Based on the results of the experiments presented in Section 5, we can conclude that GEM performs well both in terms of computation time and memory occupation even for very large global policies. To confirm this conviction, we have employed GEM in some prototype of real-world distributed systems in the maritime safety and security (Trivellato et al. 2011) and employability (Böhm et al. 2010) domains. In addition, we are currently designing an advanced version of GEM that implements an “early loop detection” strategy to avoid the reevaluation of side requests. Finally, since the policy language proposed in this paper can be used to represent the semantics of several existing trust management languages (e.g., RT (Li et al. 2003) and PeerTrust (Alves et al. 2006)), we point out that GEM can be used to evaluate goals over policies expressed in any of those languages.

Acknowledgements

This work has been carried out as part of the POSEIDON project under the responsibility of the Embedded Systems Institute (ESI). This project is partially supported by the Dutch Ministry of Economic Affairs under the BSIK03021 program. This work has been also funded by the EU-IST-IP-216287 TAS³ project.

References

- The Friend of a Friend (FOAF) project. <http://www.foaf-project.org/>.
- ALVES, M., DAMASIO, C. V., NEJDL, W., AND OLMEDILLA, D. 2006. A Distributed Tabling Algorithm for Rule Based Policy Systems. In *Proceedings of the 7th IEEE International Workshop on Policies for Distributed Systems and Networks, POLICY'06*. IEEE Computer Society, Washington, DC, USA, 123–132.
- APT, K. R. 1990. Logic programming. In *Handbook of theoretical computer science (vol. B): formal models and semantics*. MIT Press, Cambridge, MA, USA, 493–574.
- APT, K. R., BLAIR, H. A., AND WALKER, A. 1988. *Towards a theory of declarative knowledge*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 89–148.
- APT, K. R. AND BOL, R. N. 1994. Logic programming and negation: A survey. *Journal of Logic Programming* 19/20, 9–71.
- APT, K. R. AND MARCHIORI, E. 1994. Reasoning about Prolog programs: from Modes through Types to Assertions. *Formal Aspects of Computing* 6, 6A, 743–765.

- BECKER, M. Y. 2005. Cassandra: flexible trust management and its application to electronic health records. Ph.D. thesis, Computer Laboratory, University of Cambridge, UK.
- BECKER, M. Y., FOURNET, C., AND GORDON, A. D. 2010. SecPAL: Design and semantics of a decentralized authorization language. *Journal of Computer Security* 18, 619–665.
- BECKER, M. Y., MACKAY, J. F., AND DILLAWAY, B. 2009. Abductive Authorization Credential Gathering. In *Proceedings of the 10th IEEE international conference on Policies for distributed systems and networks, POLICY'09*. IEEE Press, Piscataway, NJ, USA, 1–8.
- BLAZE, M., FEIGENBAUM, J., AND LACY, J. 1996. Decentralized Trust Management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy, SP'96*. IEEE Computer Society, 164–173.
- BÖHM, K., ETALLE, S., DEN HARTOG, J., HÜTTER, C., TRABELSI, S., TRIVELLATO, D., AND ZANNONE, N. 2010. Flexible Architecture for Privacy-Aware Trust Management. *Journal of Theoretical and Applied Electronic Commerce Research* 5, 77–96.
- BRADSHAW, R. W., HOLT, J. E., AND SEAMONS, K. E. 2004. Concealing complex policies with hidden credentials. In *Proceedings of the 11th ACM conference on Computer and Communications Security, CCS'04*. ACM, New York, NY, USA, 146–157.
- BRY, F. 1989. Query Evaluation in Recursive Databases: Bottom-up and Top-down Reconciled. In *Proceedings of the First International Conference on Deductive and Object-Oriented Databases, DOOD'89*. North-Holland/Elsevier Science Publishers, 25–44.
- CHEN, W. AND WARREN, D. S. 1996. Tabled Evaluation With Delaying for General Logic Programs. *Journal of the ACM* 43, 1 (January), 20–74.
- CHEN, Y. 1997. Magic sets and stratified databases. *International Journal of Intelligent Systems* 12, 3, 203–231.
- COSTANTINI, S. 2001. Comparing different graph representations of logic programs under the Answer Set semantics. In *Proceedings of the 1st International Workshop on Answer Set Programming, ASP'01*. AAAI Press.
- CZENKO, M. AND ETALLE, S. 2007. Core TuLiP logic programming for trust management. In *Proceedings of the 23rd international conference on Logic programming, ICLP'07*. LNCS, vol. 4670. Springer-Verlag, Berlin, Heidelberg, 380–394.
- CZENKO, M., TRAN, H., DOUMEN, J., ETALLE, S., HARTEL, P., AND DEN HARTOG, J. 2006. Nonmonotonic Trust Management for P2P Applications. *Electronic Notes in Theoretical Computer Science* 157, 3 (May), 113–130.
- DAMÁSIO, C. V. 2000. A Distributed Tabling System. In *Proceedings of the 2nd Conference on Tabulation in Parsing and Deduction, TAPD'00*. 65–75.
- DI MARZO SERUGENDO, G., FOUKIA, N., HASSAS, S., KARAGEORGOS, A., MOSTÉFAOUI, S. K., RANA, O. F., ULIERU, M., VALCKENAERS, P., AND VAN AART, C. 2004. Self-Organisation: Paradigms and Applications. *Engineering Self-Organising Systems* 2977, 1–19.
- DONG, C. AND DULAY, N. 2010. Shinren: Non-monotonic Trust Management for Distributed Systems. In *Proceedings of the 4th IFIP WG 11.11 International Conference on Trust Management, IFIPTM'10*. IFIP, vol. 321. Springer Boston, 125–140.
- ELLISON, C., FRANTZ, B., LAMPSON, B., RIVEST, R., THOMAS, B., AND YLONEN, T. 1999. SPKI Certificate Theory.
- FITTING, M. 1985. A Kripke-Kleene-semantics for general logic programs. *Journal of Logic Programming* 2, 4, 295–312.
- FRIKKEN, K., ATALLAH, M., AND LI, J. 2006. Attribute-Based Access Control with Hidden Policies and Hidden Credentials. *IEEE Transactions on Computers* 55, 1259–1270.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Proceedings of the 5th International Conference and Symposium on Logic Programming, ICLP'88*. MIT Press, 1070–1080.

- GUO, H.-F. AND GUPTA, G. 2001. A Simple Scheme for Implementing Tabled Logic Programming Systems Based on Dynamic Reordering of Alternatives. In *Proceedings of the 17th International Conference on Logic Programming, ICLP'01*. Springer-Verlag, London, UK, UK, 181–196.
- HOCH, J. AND SHAMIR, A. 2008. On the Strength of the Concatenated Hash Combiner When All the Hash Functions Are Weak. In *Automata, Languages and Programming*. LNCS, vol. 5126. Springer-Verlag, Berlin, Heidelberg, 616–630.
- HU, R. 1997. Efficient tabled evaluation of normal logic programs in a distributed environment. Ph.D. thesis, State University of New York at Stony Brook.
- HULIN, G. 1989. Parallel processing of recursive queries in distributed architectures. In *Proceedings of the 15th international conference on Very large data bases, VLDB '89*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 87–96.
- JIM, T. AND SUCIU, D. 2001. Dynamically distributed query evaluation. In *Proceedings of the 20th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of database systems, PODS '01*. ACM, New York, NY, USA, 28–39.
- KOSHUTANSKI, H. AND MASSACCI, F. 2008. Interactive access control for autonomic systems: From theory to implementation. *ACM Transactions on Autonomous and Adaptive Systems* 3, 3 (August), 1–31.
- KOWALSKI, R. 1974. Predicate logic as a programming language. *Information Processing* 74, 556–574.
- LEE, A. J., MINAMI, K., AND BORISOV, N. 2009. Confidentiality-preserving distributed proofs of conjunctive queries. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security, ASIACCS'09*. ACM, New York, NY, USA, 287–297.
- LEE, A. J., MINAMI, K., AND WINSLETT, M. 2010. On the Consistency of Distributed Proofs with Hidden Subtrees. *ACM Transactions on Information and System Security* 13, 3 (July), 1–32.
- LEUSCHEL, M., MARTENS, B., AND SAGONAS, K. 1998. Preserving Termination of Tabled Logic Programs While Unfolding. In *Proceedings of the 7th International Workshop on Logic Program Synthesis and Transformation, LOPSTR'98*. LNCS, vol. 1463. Springer-Verlag, Berlin, 189–205.
- LI, N. AND MITCHELL, J. C. 2003. Datalog with Constraints: A Foundation for Trust Management Languages. In *Proceedings of the 5th International Symposium on Practical Aspects of Declarative Languages, PADL'03*. LNCS, vol. 2562. Springer-Verlag, London, UK, UK, 58–73.
- LI, N., WINSBOROUGH, W. H., AND MITCHELL, J. C. 2003. Distributed credential chain discovery in trust management. *Journal of Computer Security* 11, 1 (February), 35–86.
- MINAMI, K., BORISOV, N., WINSLETT, M., AND LEE, A. J. 2011. Confidentiality-preserving proof theories for distributed proof systems. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security, ASIACCS'11*. ACM, New York, NY, USA, 145–154.
- PARK, D. 1969. Fixpoint induction and proofs of program properties. *Machine Intelligence* 5, 59–78.
- PRZYMUSINSKA, H. AND PRZYMUNSINSKI, T. C. 1990. Weakly stratified logic programs. *Fundamenta Informaticae* 13, 1, 51–65.
- PRZYMUNSINSKI, T. 1990. The Well-Founded Semantics Coincides with Three-Valued Stable Semantics. *Fundamenta Informaticae* 13, 4, 445–463.
- PRZYMUNSINSKI, T. C. 1988. *On the declarative semantics of deductive databases and logic programs*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 193–216.
- RAMAKRISHNAN, R. 1991. Magic templates: a spellbinding approach to logic programs. *Journal of Logic Programming* 11, 189–216.
- ROCHA, R., SILVA, F., AND COSTA, V. S. 2005. On applying or-parallelism and tabling to logic programs. *Theory and Practice of Logic Programming* 5, 1-2, 161–205.
- SEAMONS, K. E., WINSLETT, M., AND YU, T. 2001. Limiting the Disclosure of Access Control Policies during Automated Trust Negotiation. In *Proceedings of the Network and Distributed System Security Symposium, NDSS'01*. The Internet Society.

- SHEN, Y.-D., YUAN, L.-Y., YOU, J.-H., AND ZHOU, N.-F. 2001. Linear tabulated resolution based on Prolog control strategy. *Theory and Practice of Logic Programming 1*, 71–103.
- STINE, K., KISSEL, R., BARKER, W. C., LEE, A., AND FAHLSING, J. 2008. Guide for Mapping Types of Information and Information Systems to Security Categories. Special Publication SP 800-60 Rev. 1, National Institute of Standards and Technology (NIST).
- SWIFT, T. AND WARREN, D. S. 2012. XSB: Extending Prolog with Tabled Logic Programming. *Theory and Practice of Logic Programming 12*, Special Issue 1-2, 157–187.
- TAMAKI, H. AND SATO, T. 1986. OLD Resolution with Tabulation. In *Proceedings of the Third International Conference on Logic Programming, ICLP'86*. Springer-Verlag, London, UK, UK, 84–98.
- TRIVELLATO, D., ZANNONE, N., AND ETALLE, S. 2011. A Security Framework for Systems of Systems. In *Proceedings of the 12th IEEE international conference on Policies for distributed systems and networks, POLICY'11*. IEEE Computer Society, Piscataway, NJ, USA.
- VAN GELDER, A., ROSS, K. A., AND SCHLIPF, J. S. 1991. The well-founded semantics for general logic programs. *Journal of the ACM 38*, 619–649.
- VIEILLE, L. 1987. A Database-Complete Proof Procedure Based on SLD-Resolution. In *Proceedings of the Fourth International Conference on Logic Programming, ICLP'87*. MIT Press, 74–103.
- WINSBOROUGH, W. H. AND LI, N. 2002. Protecting sensitive attributes in automated trust negotiation. In *Proceedings of the 2002 ACM workshop on Privacy in the Electronic Society, WPES'02*. ACM, New York, NY, USA, 41–51.
- WINSBOROUGH, W. H., SEAMONS, K. E., AND JONES, V. E. 2000. Automated Trust Negotiation. In *Proceedings of the DARPA Information Survivability Conference and Exposition, DISCEX'00*. Vol. 1. IEEE Computer Society, Los Alamitos, CA, USA, 88–102.
- WINSLETT, M. 2003. An introduction to trust negotiation. In *Proceedings of the 1st international conference on Trust management, iTrust'03*. LNCS, vol. 2692. Springer-Verlag, Berlin, Heidelberg, 275–283.
- YU, T. AND WINSLETT, M. 2003. A Unified Scheme for Resource Protection in Automated Trust Negotiation. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy, SP'03*. IEEE Computer Society, Washington, DC, USA, 110–122.
- ZHANG, C. C. AND WINSLETT, M. 2008. Distributed Authorization by Multiparty Trust Negotiation. In *Proceedings of the 13th European Symposium on Research in Computer Security, ESORICS '08*. LNCS, vol. 5283. Springer-Verlag, Berlin, Heidelberg, 282–299.
- ZHOU, N.-F. AND SATO, T. 2003. Efficient fixpoint computation in linear tabling. In *Proceedings of the 5th ACM SIGPLAN International conference on Principles and Practice of Declarative Programming, PPDP'03*. ACM, New York, NY, USA, 275–283.

Appendix A Example

In this section we show how GEM computes the answers of a goal using the procedures presented in Section 3.3. As an example global policy, we use a fragment of the policy introduced in Section 1. In particular, we consider the following policy statements:

1. $\text{memberOfAlpha}(c1, X) \leftarrow \text{projectPartner}(mc, Y), \text{memberOfAlpha}(Y, X)$.
2. $\text{projectPartner}(mc, c2)$.
3. $\text{projectPartner}(mc, c3)$.
4. $\text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c1, X)$.
5. $\text{memberOfAlpha}(c2, \text{alice})$.
6. $\text{memberOfAlpha}(c3, \text{bob})$.

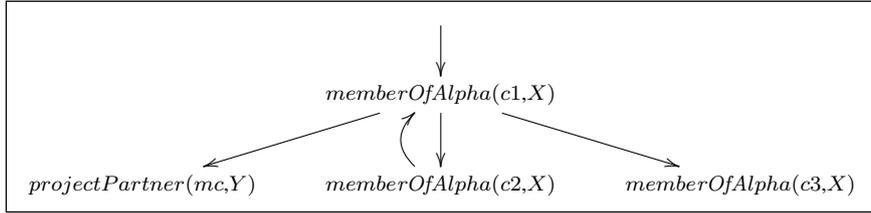


Fig. A 1. Call Graph of the Evaluation of $memberOfAlpha(c1,X)$ with Respect to the Example Global Policy

The call graph of the global policy is shown in Figure A 1. We illustrate the computation for an initial request $(h_1, h, memberOfAlpha(c1, X))$ from hospital h to company $c1$. Table A 1 shows the list of all procedure calls made by GEM to produce the response to the initial request. The first column of the table indicates the order in which the calls are made; the second column denotes the principal and location where each procedure is evaluated. GEM computes the answers of goal $memberOfAlpha(c1, X)$ by making 53 procedure calls; the number of messages exchanged between different principals, however, is only 14, consisting of 5 request messages and 9 response messages (including the initial request and its response). Next, we present and discuss some “screenshots” showing the status of the computation at various stages.

When principal $c1$ receives the request for goal $memberOfAlpha(c1, X)$ from h , it calls procedure PROCESS REQUEST (Algorithm 1 in Section 3.3) that initializes the table of the goal. Table A 2 shows the table of $memberOfAlpha(c1, X)$ resulting from the execution of PROCESS REQUEST on the initial request. The table field HR (higher request) is set to the initial request, and the evaluation tree of the goal, $Tree$, is initialized by adding to the root node a subnode representing the only clause in $c1$ ’s local policy applicable to the goal, i.e., clause 1. The node status is set to *new*, and the node identifier is obtained by concatenating the request identifier h_1 with string $c1_1$. To keep the representation more compact, in Table A 2 and in the other tables presented in this section the evaluation tree of a goal is represented as a list of nodes rather than as the structure defined in Section 3.3.

In order to compute the list of project members, $c1$ needs to first retrieve from mc the list of partner companies in the project, and then for each of these companies the list of its project members. Table A 3 shows the status of the computation after goal $projectPartner(mc, Y)$ has been completely evaluated by mc (procedure calls 2 to 7 in Table A 1), i.e., after the set of project partners has been computed. The request for goal $projectPartner(mc, Y)$ from $c1$ to mc is generated by the activation of node h_1c1_1 in $c1$ ’s table (procedure call 2, ACTIVATE NODE), which results in a change of status from *new* to *active* of both the root node and the node itself. Similarly to $c1$, when mc receives the request it creates the table of the goal (call 3 in Table A 1), setting HR to the higher request and initializing $Tree$ with clauses 2 and 3 of the global policy presented above. Two calls to procedure ACTIVATE NODE (calls 4 and 5) lead to the identification of two answers of the goal, namely $projectPartner(mc, c2)$ and $projectPartner(mc, c3)$, which are added to $AnsSet$ with an empty list of request identifiers. At the next call to ACTIVATE NODE (call 6), the evaluation tree of goal $projectPartner(mc, Y)$ has no more nodes to activate (i.e., all the branches of the evaluation tree have been inspected) and procedure GENER-

Call | Principal | Procedure

Call	Principal	Procedure
1	c1	PROCESS REQUEST($(h_1, h, \text{memberOfAlpha}(c1, X))$)
2	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
3	mc	PROCESS REQUEST($(h_1 c1_1, c1, \text{projectPartner}(mc, Y))$)
4	mc	ACTIVATE NODE($\text{projectPartner}(mc, Y)$)
5	mc	ACTIVATE NODE($\text{projectPartner}(mc, Y)$)
6	mc	ACTIVATE NODE($\text{projectPartner}(mc, Y)$)
7	mc	GENERATE RESPONSE($\text{projectPartner}(mc, Y)$)
8	mc	TERMINATE($\text{projectPartner}(mc, Y)$)
9	mc	SEND RESPONSE($(h_1 c1_1, c1, \text{projectPartner}(mc, Y), \text{disposed}, \{\})$)
10	c1	PROCESS RESPONSE($(h_1 c1_1, \{\text{projectPartner}(mc, c2), \text{projectPartner}(mc, c3)\}, \text{disposed}, \{\})$)
11	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
12	c2	PROCESS REQUEST($(h_1 c1_2, c1, \text{memberOfAlpha}(c2, X))$)
13	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
14	c1	PROCESS REQUEST($(h_1 c1_2 c2_1, c2, \text{memberOfAlpha}(c1, X))$)
15	c1	SEND RESPONSE($(h_1 c1_2 c2_1, c2, \text{memberOfAlpha}(c1, X), \text{active}, \{h_1\})$)
16	c2	PROCESS RESPONSE($(h_1 c1_2 c2_1, \{\}, \text{active}, \{h_1\})$)
17	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
18	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
19	c2	GENERATE RESPONSE($\text{memberOfAlpha}(c2, X)$)
20	c2	SEND RESPONSE($(h_1 c1_2, c1, \text{memberOfAlpha}(c2, X), \text{active}, \{h_1\})$)
21	c1	PROCESS RESPONSE($(h_1 c1_2, \{\text{memberOfAlpha}(c2, \text{alice})\}, \text{active}, \{h_1\})$)
22	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
23	c3	PROCESS REQUEST($(h_1 c1_3, c1, \text{memberOfAlpha}(c3, X))$)
24	c3	ACTIVATE NODE($\text{memberOfAlpha}(c3, X)$)
25	c3	ACTIVATE NODE($\text{memberOfAlpha}(c3, X)$)
26	c3	GENERATE RESPONSE($\text{memberOfAlpha}(c3, X)$)
27	c3	TERMINATE($\text{memberOfAlpha}(c3, X)$)
28	c3	SEND RESPONSE($(h_1 c1_3, c1, \text{memberOfAlpha}(c3, X), \text{disposed}, \{\})$)
29	c1	PROCESS RESPONSE($(h_1 c1_3, \{\text{memberOfAlpha}(c3, \text{bob})\}, \text{disposed}, \{\})$)
30	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
31	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
32	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
33	c1	GENERATE RESPONSE($\text{memberOfAlpha}(c1, X)$)
34	c1	SEND RESPONSE($(h_1 c1_2 c2_1, c2, \text{memberOfAlpha}(c1, X), \text{loop}(h_1), \{\})$)
35	c2	PROCESS RESPONSE($(h_1 c1_2 c2_1, \{\text{memberOfAlpha}(c1, \text{alice}), \text{memberOfAlpha}(c1, \text{bob})\}, \text{loop}(h_1), \{\})$)
36	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
37	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
38	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
39	c2	GENERATE RESPONSE($\text{memberOfAlpha}(c2, X)$)
40	c2	SEND RESPONSE($(h_1 c1_2, c1, \text{memberOfAlpha}(c2, X), \text{loop}(h_1), \{h_1\})$)
41	c1	PROCESS RESPONSE($(h_1 c1_2, \{\text{memberOfAlpha}(c2, \text{bob})\}, \text{loop}(h_1), \{h_1\})$)
42	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
43	c1	ACTIVATE NODE($\text{memberOfAlpha}(c1, X)$)
44	c1	GENERATE RESPONSE($\text{memberOfAlpha}(c1, X)$)
45	c1	TERMINATE($\text{memberOfAlpha}(c1, X)$)
46	c1	SEND RESPONSE($(h_1, h, \text{memberOfAlpha}(c1, X), \text{disposed}, \{\})$)
47	c1	SEND RESPONSE($(h_1 c1_2 c2_1, c2, \text{memberOfAlpha}(c1, X), \text{disposed}, \{\})$)
48	c2	PROCESS RESPONSE($(h_1 c1_2 c2_1, \{\}, \text{disposed}, \{\})$)
49	c2	ACTIVATE NODE($\text{memberOfAlpha}(c2, X)$)
50	c2	GENERATE RESPONSE($\text{memberOfAlpha}(c2, X)$)
51	c2	TERMINATE($\text{memberOfAlpha}(c2, X)$)
52	c2	SEND RESPONSE($(h_1 c1_2, c1, \text{memberOfAlpha}(c2, X), \text{disposed}, \{\})$)
53	c1	PROCESS RESPONSE($(h_1 c1_2, \{\}, \text{disposed}, \{\})$)

Table A 1. Procedure Call Stack For the Example Global Policy

ATE RESPONSE is invoked (call 7). Since the goal is not involved in any loop, its evaluation

Principal c1

HR	$(h_1, h, \leftarrow \text{memberOfAlpha}(c1, X))$
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{\}$
Tree	$(h_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c1, X), \text{new})$ $(h_1 c1_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{projectPartner}(mc, Y), \text{memberOfAlpha}(Y, X), \text{new})$

Table A 2. *Status of the Computation After Procedure Call 1 in Table A 1*

Principal c1

HR	$(h_1, h, \leftarrow \text{memberOfAlpha}(c1, X))$
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{\}$
Tree	$(h_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c1, X), \text{active})$ $(h_1 c1_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{projectPartner}(mc, Y), \text{memberOfAlpha}(Y, X), \text{active})$

Principal mc

HR	$(h_1 c1_1, c1, \leftarrow \text{projectPartner}(mc, Y))$
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{(\text{projectPartner}(mc, c2), \{\}), (\text{projectPartner}(mc, c3), \{\})\}$
Tree	$(h_1 c1_1, \text{projectPartner}(mc, Y) \leftarrow \text{projectPartner}(mc, Y), \text{new})$ $(h_1 c1_1 mc_1, \text{projectPartner}(mc, c2), \text{answer})$ $(h_1 c1_1 mc_2, \text{projectPartner}(mc, c3), \text{answer})$

Table A 3. *Status of the Computation After Procedure Call 7 in Table A 1*

is completed and procedure TERMINATE is executed next (line 3 of Algorithm 6 in Section 3.3).

As a result of the execution of procedure TERMINATE, the root node of the evaluation tree of goal $\text{projectPartner}(mc, Y)$ is disposed and the answers identified are sent to $c1$ through procedure SEND RESPONSE (procedure call 9, results shown in Table A 4). The response message received by $c1$ is processed by procedure PROCESS RESPONSE; the message contains the two answers ($\text{projectPartner}(mc, c2)$ and $\text{projectPartner}(mc, c3)$) and an empty set of loop identifiers, and has status *disposed*, indicating that no more answers of goal $\text{projectPartner}(mc, Y)$ will be received. The evaluation tree of goal $\text{memberOfAlpha}(c1, X)$ is updated by adding two subnodes to node $h_1 c1_1$, one for each project partner (see $c1$'s table in Table A 4).

The activation of node $h_1 c1_2$ by $c1$ leads to the request for goal $\text{memberOfAlpha}(c2, X)$ to $c2$. Accordingly, $c2$ creates a table for the goal; the evaluation tree of the goal consists of three nodes: the root node and two subnodes, representing clauses 4 and 5 of the global pol-

Principal c1

HR	$(h_1, h, \leftarrow \text{memberOfAlpha}(c1, X))$
LR	$\{(h_1 c1_2 c2_1, c2, \leftarrow \text{memberOfAlpha}(c1, X))\}$
ActiveGoals	$\{\}$
AnsSet	$\{\}$
Tree	$(h_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c1, X), \text{active})$ $(h_1 c1_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{projectPartner}(mc, Y), \text{memberOfAlpha}(Y, X), \text{disposed})$ $(h_1 c1_2, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c2, X), \text{active})$ $(h_1 c1_3, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c3, X), \text{new})$

Principal mc

HR	null
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{(\text{projectPartner}(mc, c2), \{h_1 c1_1\}), (\text{projectPartner}(mc, c3), \{h_1 c1_1\})\}$
Tree	$(h_1 c1_1, \text{projectPartner}(mc, Y) \leftarrow \text{projectPartner}(mc, Y), \text{disposed})$ $(h_1 c1_1 mc_1, \text{projectPartner}(mc, c2), \text{answer})$ $(h_1 c1_1 mc_2, \text{projectPartner}(mc, c3), \text{answer})$

Principal c2

HR	$(h_1 c1_2, c1, \leftarrow \text{memberOfAlpha}(c2, X))$
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{\}$
Tree	$(h_1 c1_2, \text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c2, X), \text{active})$ $(h_1 c1_2 c2_1, \text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c1, X), \text{active})$ $(h_1 c1_2 c2_2, \text{memberOfAlpha}(c2, \text{alice}), \text{new})$

Table A 4. Status of the Computation After Procedure Call 15 in Table A 1

icy, with identifiers $h_1 c1_2 c2_1$ and $h_1 c1_2 c2_2$ respectively. The activation of node $h_1 c1_2 c2_1$ by $c2$, in turn, leads to a request for goal $\text{memberOfAlpha}(c1, X)$ to $c1$, forming a loop. The loop is identified by $c1$ in procedure PROCESS REQUEST (call 14 in Table A 1): in fact, the identifier of the higher request for $\text{memberOfAlpha}(c1, X)$ (h_1) is a prefix of the identifier of $c2$'s request ($h_1 c1_2 c2_1$). Therefore, the lower request is added by $c1$ to set LR , and a response is sent from $c1$ to $c2$ with a notification of loop h_1 (call 15).

The loop notification sent from $c1$ to $c2$ starts the loop processing phase, which involves procedure calls from 16 to 34 in Table A 1. The results of the loop processing phase are shown in Table A 5. Upon receiving the loop notification, $c2$ sets the status of the node whose evaluation formed the loop to $\text{loop}(\{h_1\})$ and "freezes" its evaluation; then, it proceeds with the evaluation of the other nodes of the evaluation tree. The activation of node $h_1 c1_2 c2_2$ (procedure call 17), in particular, leads to the first answer of the goal, i.e., $\text{memberOfAlpha}(c2, \text{alice})$. Since at this point there are no more nodes to be activated, the computed answer can be sent to $c1$ with a notification about the loop.

Principal c1	
HR	$(h_1, h, \leftarrow \text{memberOfAlpha}(c1, X))$
LR	$\{(h_1 c1_2 c2_1, c2, \leftarrow \text{memberOfAlpha}(c1, X))\}$
ActiveGoals	$\{(h_1, 1)\}$
AnsSet	$\{(\text{memberOfAlpha}(c1, \text{alice}), \{h_1 c1_2 c2_1\}), (\text{memberOfAlpha}(c1, \text{bob}), \{h_1 c1_2 c2_1\})\}$
Tree	$(h_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c1, X), \text{loop}(\{h_1\}))$ $(h_1 c1_1, \text{memberOfAlpha}(c1, X) \leftarrow \text{projectPartner}(mc, Y), \text{memberOfAlpha}(Y, X), \text{disposed})$ $(h_1 c1_2, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c2, X), \text{loop}(\{h_1\}))$ $(h_1 c1_3, \text{memberOfAlpha}(c1, X) \leftarrow \text{memberOfAlpha}(c3, X), \text{disposed})$ $(h_1 c1_4, \text{memberOfAlpha}(c1, \text{alice}), \text{answer})$ $(h_1 c1_5, \text{memberOfAlpha}(c1, \text{bob}), \text{answer})$
Principal mc	
HR	null
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{(\text{projectPartner}(mc, c2), \{h_1 c1_1\}), (\text{projectPartner}(mc, c3), \{h_1 c1_1\})\}$
Tree	$(h_1 c1_1, \text{projectPartner}(mc, Y) \leftarrow \text{projectPartner}(mc, Y), \text{disposed})$ $(h_1 c1_1 mc_1, \text{projectPartner}(mc, c2), \text{answer})$ $(h_1 c1_1 mc_2, \text{projectPartner}(mc, c3), \text{answer})$
Principal c2	
HR	$(h_1 c1_2, c1, \leftarrow \text{memberOfAlpha}(c2, X))$
LR	$\{\}$
ActiveGoals	$\{(h_1, 1)\}$
AnsSet	$\{(\text{memberOfAlpha}(c2, \text{alice}), \{h_1 c1_2\})\}$
Tree	$(h_1 c1_2, \text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c2, X), \text{active})$ $(h_1 c1_2 c2_1, \text{memberOfAlpha}(c2, X) \leftarrow \text{memberOfAlpha}(c1, X), \text{loop}(\{h_1\}))$ $(h_1 c1_2 c2_2, \text{memberOfAlpha}(c2, \text{alice}), \text{answer})$
Principal c3	
HR	null
LR	$\{\}$
ActiveGoals	$\{\}$
AnsSet	$\{(\text{memberOfAlpha}(c3, \text{bob}), \{h_1 c1_3\})\}$
Tree	$(h_1 c1_3, \text{memberOfAlpha}(c3, X) \leftarrow \text{memberOfAlpha}(c3, X), \text{disposed})$ $(h_1 c1_3 c3_1, \text{memberOfAlpha}(c3, \text{bob}), \text{answer})$

Table A 5. Status of the Computation After Procedure Call 34 in Table A 1

Before sending the answer, $c2$ sets the counter in *ActiveGoals* to 1 (procedure GENERATE RESPONSE, call 19) and adds the identifier of HR to the set of recipients of answer $\text{memberOfAlpha}(c2, \text{alice})$ in *AnsSet* (procedure SEND RESPONSE, call 20).

The loop is now processed at $c1$. After adding a subnode to the evaluation tree of goal $\text{memberOfAlpha}(c1, X)$ for the answer received from $c2$, $c1$ freezes node $h_1 c1_2$ and

starts the evaluation of node h_1c1_3 (procedure call 22). This results in a request from $c1$ to $c3$ for the evaluation of goal $memberOfAlpha(c3,X)$. The only clause applicable to $memberOfAlpha(c3,X)$ (clause 6 of the global policy) is a fact; therefore, the goal is completely evaluated after one call to procedure ACTIVATE NODE (call 24). The answer of the goal, $memberOfAlpha(c3,bob)$, is returned to $c1$ (procedure call 28). Since the status of the response message is *disposed*, $c1$ disposes node h_1c1_3 and adds subnode h_1c1_5 to it reflecting the answer received from $c3$ (procedure PROCESS RESPONSE, call 29). The next two executions of procedure ACTIVATE NODE at $c1$ lead to the identification of two answers of goal $memberOfAlpha(c1,X)$, namely $memberOfAlpha(c1,alice)$ and $memberOfAlpha(c1,bob)$. Before returning these answers to the requester of HR (i.e., h), however, all the loops need to be fully processed. For this reason, $c1$ sends the two answers to $c2$ in response to LR first; the status of the response message is $loop(h_1)$, and the status of the root node of the evaluation tree in $c1$'s table is changed accordingly (procedure calls 33 and 34 in Table A 1).

Now, the second iteration of the loop processing phase starts (procedure calls 35-44). In this second iteration, $c2$ identifies a new answer of its goal, i.e., $memberOfAlpha(c2,bob)$, which is sent back to $c1$. This answer, however, does not lead to new answers at $c1$. Since h_1 is the only loop in the SCC (and hence $memberOfAlpha(c1,X)$ is the leader of the SCC), and no new answers of $memberOfAlpha(c1,X)$ have been computed, the loop termination phase can start (line 15 of Algorithm 6 in Section 3.3). In this phase, $c1$ sends a response message with status *disposed* to both $c2$ (the other principal in the loop) and h (to which also the answers are sent). Upon receiving this message, $c2$ disposes all the nodes in the evaluation tree of $memberOfAlpha(c2,X)$ that are involved in some loop (procedure PROCESS RESPONSE), and forwards the message back to $c1$ (calls 51 and 52). $c1$ simply ignores the message, as the status of the root node of the evaluation tree of $memberOfAlpha(c1,X)$ is already *disposed* (line 4 of Algorithm 5 in Section 3.3), and the computation terminates. Table A 6 shows the status of the tables of all the goals at the end of the computation.

Principal c1	
HR	null
LR	{}
ActiveGoals	{}
AnsSet	{(memberOfAlpha(c1,alice),{h ₁ c ₁ c ₂ c ₁ ,h ₁ }), (memberOfAlpha(c1,bob),{h ₁ c ₁ c ₂ c ₁ ,h ₁ })}
Tree	(h ₁ ,memberOfAlpha(c1,X)← memberOfAlpha(c1,X),disposed) (h ₁ c ₁ ,memberOfAlpha(c1,X)← projectPartner(mc,Y), memberOfAlpha(Y,X),disposed) (h ₁ c ₁ c ₂ ,memberOfAlpha(c1,X)← memberOfAlpha(c2,X),disposed) (h ₁ c ₁ c ₃ ,memberOfAlpha(c1,X)← memberOfAlpha(c3,X),disposed) (h ₁ c ₁ c ₄ ,memberOfAlpha(c1,alice),answer) (h ₁ c ₁ c ₅ ,memberOfAlpha(c1,bob),answer) (h ₁ c ₁ c ₆ ,memberOfAlpha(c1,bob),answer)
Principal mc	
HR	null
LR	{}
ActiveGoals	{}
AnsSet	{(projectPartner(mc,c2),{h ₁ c ₁ }), (projectPartner(mc,c3),{h ₁ c ₁ })}
Tree	(h ₁ c ₁ ,projectPartner(mc,Y)← projectPartner(mc,Y),disposed) (h ₁ c ₁ mc ₁ ,projectPartner(mc,c2),answer) (h ₁ c ₁ mc ₂ ,projectPartner(mc,c3),answer)
Principal c2	
HR	null
LR	{}
ActiveGoals	{}
AnsSet	{(memberOfAlpha(c2,alice),{h ₁ c ₁ c ₂ }), (memberOfAlpha(c2,bob),{h ₁ c ₁ c ₂ })}
Tree	(h ₁ c ₁ c ₂ ,memberOfAlpha(c2,X)← memberOfAlpha(c2,X),disposed) (h ₁ c ₁ c ₂ c ₁ ,memberOfAlpha(c2,X)← memberOfAlpha(c1,X),disposed) (h ₁ c ₁ c ₂ c ₂ ,memberOfAlpha(c2,alice),answer) (h ₁ c ₁ c ₂ c ₃ ,memberOfAlpha(c2,alice),answer) (h ₁ c ₁ c ₂ c ₄ ,memberOfAlpha(c2,bob),answer)
Principal c3	
HR	null
LR	{}
ActiveGoals	{}
AnsSet	{(memberOfAlpha(c3,bob),{h ₁ c ₁ c ₃ })}
Tree	(h ₁ c ₁ c ₃ ,memberOfAlpha(c3,X)← memberOfAlpha(c3,X),disposed) (h ₁ c ₁ c ₃ c ₁ ,memberOfAlpha(c3,bob),answer)

Table A 6. *Final Status of the Computation for the Example Global Policy*