

# Monitoring Compliance Policies over Incomplete and Disagreeing Logs<sup>\*</sup>

David Basin, Felix Klaedtke, Srdjan Marinovic, and Eugen Zălinescu

Institute of Information Security, ETH Zurich

**Abstract.** When monitoring system behavior to check compliance against a given policy, one is sometimes confronted with incomplete knowledge about system events. In IT systems, such incompleteness may arise from logging infrastructure failures and corrupted log files, or when the logs produced by different system components disagree on whether actions took place. In this paper, we present a policy language with a three-valued semantics that allows one to explicitly reason about incomplete knowledge and handle disagreements. Furthermore, we present a monitoring algorithm for an expressive fragment of our policy language. We illustrate through examples how our approach extends compliance monitoring to systems with logging failures and disagreements.

## 1 Introduction

Laws, inter-business contracts, security policies, and similar normative regulations define compliance requirements that IT systems need to enforce. For example, IT systems in US hospitals must enforce HIPAA [1], which regulates the dissemination of medical records and the subsequent obligations that medical staff are expected to fulfill. For banks, separation-of-duty constraints should reduce the risk of fraud [2]. Data-usage contracts between different businesses regulate how sensitive documents are exchanged and subsequently disposed. Checking whether implemented IT systems comply with a body of regulations or policies is a problem of growing importance, since non-compliant behavior can lead to serious security breaches, monetary penalties, and the erosion of stakeholder’s internal standards and commitments.

Runtime-verification techniques [4,5,19,22–24] offer a promising approach for automated compliance checking of IT systems. These techniques require logging mechanisms for recording policy-relevant system actions (represented as events), a suitable language for expressing policies and unambiguously defining permissible and prohibited system behavior, and a monitoring algorithm for determining and reporting policy violations.

In complex IT systems, which are usually composed of numerous interacting subsystems, the problem of incomplete knowledge about performed actions arises. In particular, logs may contain gaps due to corrupted files, logging-mechanism crashes, network failures, and so forth. Furthermore, when multiple

---

<sup>\*</sup> This work was partly supported by Google Inc.

logs are required to verify compliant behavior, they may disagree whether certain actions took place. For example, sharing a sensitive document between two parties may require the recipient to fulfill certain obligations. Thus, when analyzing the recipient’s and the sender’s logs against this policy, we need to treat all disagreements over the transfer of the document as incomplete knowledge, since favoring one log over the other may result in missed violations or false positives. Most runtime monitors, however, do not distinguish between a gap and a non-occurrence of an event. Thus applying them to incomplete logs can yield *wrong* results. For example, consider a policy like *a subject can access a document if the subject is not blacklisted*. If it is unknown whether a subject is blacklisted, then the subject is incorrectly reported as compliant.

In this paper, we present a policy language and an accompanying monitoring algorithm that accounts for possibly incomplete and disagreeing logs. At the core of our approach is a three-valued truth space [25]. In addition to the classical Boolean values  $t$  (true) and  $f$  (false), which respectively represent the occurrence and non-occurrence of an event, we represent a knowledge gap about an event’s occurrence by the third truth value  $\perp$ . Furthermore, when evaluating policies, their interpretation is as follows: the Boolean values  $t$  and  $f$  correspond to policy compliance and policy violation and  $\perp$  represents an inconclusive answer, which can be due to knowledge gaps of event occurrences or disagreeing events.

Our policy language is a variant of a first-order temporal logic [7, 17]. First-order temporal logics have been a good fit in various case studies for formally expressing and monitoring compliance policies, see, e.g., [5, 23]. Special care must be taken when defining the semantics of a logic with additional truth values besides the classical Boolean values. In particular, a vital requirement for monitoring incomplete and disagreeing logs is to ensure that reported violations cannot be retracted if or when the log is eventually completed, for example, by recovering lost files. Otherwise, these results are of no value. More precisely, formalized policies must be monotonic with respect to the underlying partial ordering on knowledge, i.e.,  $\perp$  is less than  $f$  and  $t$ , and  $f$  and  $t$  are incomparable [9, 10, 20]. Our policy language guarantees this monotonicity requirement. Furthermore, the third truth value  $\perp$  is a first-class citizen at the object-level of our policy language: the classical logical connectives are extended to the three-valued truth space and there are specific connectives that guarantee expressive-completeness with respect to the set of knowledge-monotonic operators. Such monotonic operators are needed in our application context to express at the logic’s object-level how disagreements between logged events should be resolved.

The monitoring algorithm presented in this paper for this three-valued setting is inspired by the one from [6, 7] for the standard Boolean setting. It iteratively scans the logged actions and soundly reports violations, i.e., whenever a violation is reported, it indeed is a policy violation. It also soundly reports potential violations, i.e., depending on how the knowledge gaps are filled, these might turn out to be real policy violations. However, our monitoring algorithm is not complete in the sense that some policy violations might not be reported. This limitation stems from the expressivity of our policy language over infinite

domains. Importantly, however, for an expressive fragment, which retains all the language’s connectives but limits the usage of free variables within a formula, we show that our monitoring algorithm guarantees completeness.

In summary, our main contribution is a solution to the problem of checking policy compliance in the presence of logging failures and disagreements between logged events. Our solution comprises a policy language and a monitoring algorithm. The policy language supports reasoning with incomplete knowledge. The monitoring algorithm may be used either off-line (for audit) or on-line (at runtime), and reports all policy violations and potential policy violations for an expressive fragment of our language. Although several features of our solution are present in related work—see Section 6 for a comparison—combining them to solve the stated problem is novel. In particular, our language is the first compliance language to consider three truth values at the object level, and our monitoring algorithm is the first algorithm to guarantee both soundness and completeness in a three-valued first-order setting.

The remainder of the paper is structured as follows. In Section 2, we describe our abstract logging setting. In Section 3, we introduce our policy language. In Section 4, we analyze our policy language with respect to monotonicity and expressiveness. In Section 5, we present our monitoring algorithm. Finally, in Sections 6 and 7, we discuss related work and draw conclusions. Technical details are omitted due to space limitations. These are given in the full version of the paper, which can be found on the authors’ web pages.

## 2 Logging Knowledge Base

We abstract from a particular *physical* log file structure, and view a logging infrastructure as producing a single logging knowledge base, which is evaluated against a compliance policy. A logging knowledge base uses the three-valued truth space  $\mathbf{3} := \{\mathbf{t}, \mathbf{f}, \perp\}$  to explicitly distinguish between what is known and unknown regarding event occurrences.

To formally define a logging knowledge base over  $\mathbf{3}$ , we introduce a *logging signature*  $S$ , which is a tuple  $(C, R, \iota)$ , where  $C$  is a finite set of constant symbols,  $R$  is a finite set of predicates disjoint from  $C$ , and the function  $\iota : R \rightarrow \mathbb{N}$  assigns each predicate  $r \in R$  an arity  $\iota(r)$ . Each predicate  $r$  denotes an action, and its arguments  $\bar{a}$  denote the action’s parameters,  $r(\bar{a})$  denoting an event. A *logging structure*  $\mathcal{D}$  over the signature  $S$  consists of a domain  $|\mathcal{D}| \neq \emptyset$  and interpretations  $c^{\mathcal{D}} \in |\mathcal{D}|$ , and  $r_{\mathbf{t}}^{\mathcal{D}} \subseteq |\mathcal{D}|^{\iota(r)}$  and  $r_{\mathbf{f}}^{\mathcal{D}} \subseteq |\mathcal{D}|^{\iota(r)}$ , for each  $c \in C$  and  $r \in R$ , such that  $r_{\mathbf{t}}^{\mathcal{D}}$  and  $r_{\mathbf{f}}^{\mathcal{D}}$  are disjoint. We let  $r_{\perp}^{\mathcal{D}} := |\mathcal{D}|^{\iota(r)} \setminus (r_{\mathbf{t}}^{\mathcal{D}} \cup r_{\mathbf{f}}^{\mathcal{D}})$ . We define a *logging knowledge base* over the signature  $S$  as a sequence  $\bar{\mathcal{D}} = (\mathcal{D}_0, \mathcal{D}_1, \dots)$  of logging structures over  $S$ , with the following properties:

1.  $\bar{\mathcal{D}}$  has constant domains, that is,  $|\mathcal{D}_i| = |\mathcal{D}_{i+1}|$ , for all  $i \geq 0$ . We denote the domain by  $|\bar{\mathcal{D}}|$ .
2. Each constant symbol  $c \in C$  has a rigid interpretation, that is,  $c^{\mathcal{D}_i} = c^{\mathcal{D}_{i+1}}$ , for all  $i \geq 0$ . We denote  $c$ ’s interpretation by  $c^{\bar{\mathcal{D}}}$ .

We call the indices of the elements in the sequence  $\bar{\mathcal{D}}$  *time points* and denote them with the Greek letter  $\tau$ . We interpret a logging knowledge base  $\bar{\mathcal{D}}$  as follows:

- If  $\bar{a} \in r_{\dagger}^{\mathcal{D}\tau}$ , then the event  $r(\bar{a})$  happened at the time point  $\tau$ .
- If  $\bar{a} \in r_{\ddagger}^{\mathcal{D}\tau}$ , then the event  $r(\bar{a})$  did not happen at the time point  $\tau$ .
- If  $\bar{a} \in r_{\perp}^{\mathcal{D}\tau}$ , then  $\bar{\mathcal{D}}$  contains a *knowledge gap* at the time point  $\tau$  with regard to whether the event  $r(\bar{a})$  happened at  $\tau$ . In practice, a gap is determined by additional information about logging failures.

Thus a logging knowledge base states explicitly whether logging information is complete at a time point  $\tau$ . In case of incomplete knowledge, we have  $r_{\perp}^{\mathcal{D}\tau} \neq \emptyset$ .

We extend the classical logging assumption, whereby there are only finitely many events happening at each time point, to a three-valued setting.

**Assumption 1.** *Let  $\bar{\mathcal{D}}$  be a logging knowledge base over the signature  $(C, R, \iota)$ . For each  $r \in R$  and  $\tau \in \mathbb{N}$ , either  $r_{\dagger}^{\mathcal{D}\tau}$  is finite and  $r_{\perp}^{\mathcal{D}\tau} = \emptyset$ , or  $r_{\perp}^{\mathcal{D}\tau} = |\bar{\mathcal{D}}|^{\iota(r)}$ .*

This assumption formalizes that as long as a particular logging process is running, it correctly records all events. If the process crashes, then nothing is recorded until the process is restarted. In line with our model of a logging knowledge base, this means that at each time point  $\tau$  and for each relation  $r$  either  $r_{\perp}^{\mathcal{D}\tau} = \emptyset$  or  $r_{\perp}^{\mathcal{D}\tau} = |\bar{\mathcal{D}}|^{\iota(r)}$ .

Note that a logging knowledge base does not differentiate between multiple instances of the same event happening at the same time point. To do so, one would have to ensure that either the time points' granularity is sufficient to render this scenario impossible, or to add unique artificial parameters (such as counters) for each such event instance.

### 3 Compliance Policy Language

In this section, we define our policy language  $\mathcal{L}_3$  and illustrate with examples how policies are formalized and evaluated in the presence of incomplete knowledge. We also show how disagreements can be handled with  $\mathcal{L}_3$ 's operators.

**Syntax and semantics.** In the following, let  $S = (C, R, \iota)$  be a signature and let  $V$  be a countably infinite set of variables, where  $V \cap (C \cup R) = \emptyset$ . Also, let  $\mathbb{I}$  be the set of nonempty intervals over  $\mathbb{N}$ . We often write an interval in  $\mathbb{I}$  as  $[b, b') := \{a \in \mathbb{N} \mid b \leq a < b'\}$ , where  $b \in \mathbb{N}$ ,  $b' \in \mathbb{N} \cup \{\infty\}$ , and  $b < b'$ .

**Definition 2.** *The  $\mathcal{L}_3$  formulas over the signature  $S$  are given by the grammar*

$$\varphi ::= \mathbf{f} \mid r(t_1, \dots, t_{\iota(r)}) \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \otimes \varphi \mid \forall x. \varphi \mid \varphi \mathbf{S}_I \varphi \mid \varphi \mathbf{U}_I \varphi,$$

where  $r$  ranges over the elements in  $R$ , the  $t_i$ s over the elements in  $C \cup V$ ,  $x$  over the elements in  $V$ , and  $I$  over the elements in  $\mathbb{I}$ .

$\neg$	$\wedge$	$\otimes$	$\vee$	$\rightarrow$
$\begin{array}{c c} \neg & \\ \hline \mathbf{t} & \mathbf{f} \\ \mathbf{f} & \mathbf{t} \\ \perp & \perp \end{array}$	$\begin{array}{c ccc} \wedge & \mathbf{t} & \mathbf{f} & \perp \\ \hline \mathbf{t} & \mathbf{t} & \mathbf{f} & \perp \\ \mathbf{f} & \mathbf{f} & \mathbf{f} & \perp \\ \perp & \perp & \mathbf{f} & \perp \end{array}$	$\begin{array}{c ccc} \otimes & \mathbf{t} & \mathbf{f} & \perp \\ \hline \mathbf{t} & \mathbf{t} & \perp & \perp \\ \mathbf{f} & \perp & \mathbf{f} & \perp \\ \perp & \perp & \perp & \perp \end{array}$	$\begin{array}{c ccc} \vee & \mathbf{t} & \mathbf{f} & \perp \\ \hline \mathbf{t} & \mathbf{t} & \mathbf{t} & \perp \\ \mathbf{f} & \mathbf{t} & \mathbf{f} & \perp \\ \perp & \mathbf{t} & \perp & \perp \end{array}$	$\begin{array}{c ccc} \rightarrow & \mathbf{t} & \mathbf{f} & \perp \\ \hline \mathbf{t} & \mathbf{t} & \mathbf{f} & \perp \\ \mathbf{f} & \mathbf{t} & \mathbf{t} & \perp \\ \perp & \mathbf{t} & \perp & \perp \end{array}$
(a) primitive operators			(b) derived operators	

**Fig. 1.** Truth tables for three-valued operators (strong Kleene logic [25]).

Before formally defining the evaluation semantics, Figure 1(a) shows  $\mathcal{L}_3$ 's interpretation of the logical connectives over  $\mathbf{3}$ . We mildly abuse notation and use same symbols to denote logical connectives and their corresponding three-valued operators. The classical connectives  $\neg$  and  $\wedge$  retain their interpretation when restricted to the Boolean values  $\mathbf{t}$  and  $\mathbf{f}$ . The  $\otimes$  connective does not have a classical counterpart. Intuitively, it represents a *consensus* on how much truth can be agreed upon and is useful for combining different sources of knowledge when neither  $\mathbf{t}$  nor  $\mathbf{f}$  should be preferred over the other.

In the following, a *valuation* is a mapping  $\theta : V \rightarrow |\bar{\mathcal{D}}|$ . For a valuation  $\theta$ , the variable vector  $\bar{x} = (x_1, \dots, x_n)$ , and  $\bar{d} = (d_1, \dots, d_n) \in |\bar{\mathcal{D}}|^n$ ,  $\theta[\bar{x} \mapsto \bar{d}]$  is the valuation mapping  $x_i$  to  $d_i$ , for  $i \in \{1, \dots, n\}$ , and the other variables' valuation is unaltered. We abuse notation by applying a valuation  $\theta$  also to constant symbols  $c \in C$ , with  $\theta(c) := c^{\bar{\mathcal{D}}}$ .

**Definition 3.** Let  $\bar{\mathcal{D}} = (\mathcal{D}_0, \mathcal{D}_1, \dots)$  be a temporal structure over the signature  $S$ ,  $\theta$  a valuation, and  $\tau \in \mathbb{N}$  a time stamp. We inductively define the mapping  $\llbracket \cdot \rrbracket^{\bar{\mathcal{D}}, \theta, \tau}$  from formulas over  $S$  to values in  $\mathbf{3}$  as follows:

$$\begin{aligned}
\llbracket \mathbf{f} \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \mathbf{f} \\
\llbracket r(t_1, \dots, t_{u(r)}) \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= v \text{ if } (\theta(t_1), \dots, \theta(t_{u(r)})) \in r_v^{\bar{\mathcal{D}}\tau}, \text{ where } v \in \mathbf{3} \\
\llbracket \neg\varphi \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \neg \llbracket \varphi \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} \\
\llbracket \varphi_1 \wedge \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \llbracket \varphi_1 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} \wedge \llbracket \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} \\
\llbracket \varphi_1 \otimes \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \llbracket \varphi_1 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} \otimes \llbracket \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} \\
\llbracket \forall x. \varphi \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \bigwedge_{d \in |\bar{\mathcal{D}}|} \llbracket \varphi \rrbracket^{\bar{\mathcal{D}}, \theta[x \mapsto d], \tau} \\
\llbracket \varphi_1 \mathbf{S}_I \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \bigvee_{\tau - \tau' \in I} (\llbracket \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau'} \wedge \bigwedge_{\tau'' \in (\tau', \tau]} \llbracket \varphi_1 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau''}) \\
\llbracket \varphi_1 \mathbf{U}_I \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} &:= \bigvee_{\tau' - \tau \in I} (\llbracket \varphi_2 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau'} \wedge \bigwedge_{\tau'' \in [\tau, \tau')} \llbracket \varphi_1 \rrbracket^{\bar{\mathcal{D}}, \theta, \tau''})
\end{aligned}$$

In this definition,  $\bigwedge$  and  $\bigvee$  are respectively the (possibly infinitary) meet and join over the ordering  $\mathbf{f} \leq \perp \leq \mathbf{t}$ . Note that they match the corresponding operators in Figure 1. The temporal connectives are accompanied by intervals and a formula of the form  $\varphi \mathbf{S}_I \psi$  or  $\varphi \mathbf{U}_I \psi$  is only satisfied in  $\bar{\mathcal{D}}$  at the time point  $\tau$  if it is satisfied within the bounds given by the interval  $I$  of the respective temporal operator. We may omit the interval  $I$  if it is  $[0, \infty)$ .

We introduce the following additional syntactic sugar. We write  $\mathbf{t}\neg\varphi$ ,  $\varphi \vee \psi$  for  $\neg(\neg\varphi \wedge \neg\psi)$ ,  $\varphi \rightarrow \psi$  for  $\neg\varphi \vee \psi$ , and  $\exists x. \varphi$  for  $\neg\forall x. \neg\varphi$ . For a vector of variables  $\bar{x} = (x_1, x_2, \dots, x_n)$ , with  $n \geq 0$ , we write  $\forall \bar{x}. \varphi$  for  $\forall x_1. \forall x_2 \dots \forall x_n. \varphi$ . Moreover, we define the temporal connectives  $\blacklozenge_I \psi$  and  $\blacktriangleright_I \psi$  as  $\mathbf{tS}_I \psi$  and  $\mathbf{tU}_I \psi$ ,

respectively. Intuitively,  $\blacklozenge_{[b,b']} \psi$  is **t** at  $\tau$ , if  $\psi$  is **t** at least at one *past* time point in the time interval  $[\max(0, \tau - b' - 1), \tau - b]$ . If  $\psi$  is **f** at all these time points, then  $\blacklozenge_{[b,b']} \psi$  is **f** at  $\tau$ . The presence of at least one  $\perp$  and no **t** results in the truth value  $\perp$  for  $\blacklozenge_{[b,b']} \psi$  at  $\tau$ , since depending on how the incompleteness is resolved either outcome (**t** or **f**) is possible. The interpretation of  $\blacklozenge_{[b,b']} \psi$  is similar for *future* time points. The dual temporal connectives are  $\square_I \psi := \neg \blacklozenge_I \neg \psi$  and  $\blacksquare_I \psi := \neg \blacklozenge_I \neg \psi$ . We use standard conventions concerning the binding strength of connectives to omit parentheses. For instance, temporal connectives bind weaker than the other connectives. Furthermore,  $\rightarrow$  binds weaker than  $\vee$ , which in turn binds weaker than  $\wedge$  and  $\otimes$ .

Finally, we introduce some additional notation. Given a formula  $\varphi$ , we denote by  $fv(\varphi)$  and  $\bar{fv}(\varphi)$  the set and respectively the vector of free variables of  $\varphi$ . We call a formula  $\varphi$  *closed* if  $fv(\varphi) = \emptyset$ . For a formula  $\varphi$  with  $\bar{fv}(\varphi) = \bar{x} = (x_1, \dots, x_n)$ , we define the set of elements of  $|\bar{\mathcal{D}}|^n$  for which  $\varphi$  evaluates to  $v \in \mathbf{3}$  at a time point  $\tau \in \mathbb{N}$  as

$$\llbracket \varphi \rrbracket_v^{\bar{\mathcal{D}}, \tau} := \{ \bar{d} \in |\bar{\mathcal{D}}|^n \mid \llbracket \varphi \rrbracket^{\bar{\mathcal{D}}, \theta[\bar{x} \mapsto \bar{d}], \tau} = v, \text{ for some valuation } \theta \}.$$

**Compliance policies.** Regardless of the policy language, compliance policies are usually given as a set of regulative *normative* statements (norms), which express what an agent is obliged to do given some actions it has performed, or which conditions need to hold (or to have held) for an agent to be permitted to execute some actions. Norms are meant to be applied at all times within a system, and it has also been argued [11, 12] that deadlines are of essential importance in regulating temporal norms. Following these notions, compliance policies in  $\mathcal{L}_3$  are formalized as follows:

**Definition 4.** *A compliance policy represented in  $\mathcal{L}_3$  is a closed formula of the form  $\square \forall \bar{x}. \psi$ , where each future temporal connective in  $\psi$  is bounded.*

The outermost unbounded  $\square$  connective specifies that a policy must be fulfilled at each time point. Bounded inner future temporal connectives guarantee that each obligation has a deadline.

We map the truth values onto policy evaluations as follows: **t/f** denotes that a policy is satisfied/violated, and  $\perp$  denotes that it is unknown whether a policy is satisfied or violated. Furthermore, for a compliance policy  $\square \forall \bar{x}. \psi$ , it is often useful to report additional information regarding its violations, which is given by the aforementioned sets  $\llbracket \psi \rrbracket_f^{\bar{\mathcal{D}}, \tau}$ ,  $\llbracket \psi \rrbracket_{\perp}^{\bar{\mathcal{D}}, \tau}$ , and  $\llbracket \psi \rrbracket_t^{\bar{\mathcal{D}}, \tau}$ , for a time point  $\tau$ . Their interpretation is as follows:

- The elements in  $\llbracket \psi \rrbracket_f^{\bar{\mathcal{D}}, \tau}$  witness a policy violation at time point  $\tau$ .
- For elements in  $\llbracket \psi \rrbracket_{\perp}^{\bar{\mathcal{D}}, \tau}$ , it is unknown whether they violate the policy at time point  $\tau$ . They are potential violations.
- The elements in  $\llbracket \psi \rrbracket_t^{\bar{\mathcal{D}}, \tau}$  satisfy the policy at time point  $\tau$ .

In Section 4, we show that all reported violations and satisfactions at  $\tau$  persist regardless of how incompleteness is resolved.

**Examples.** We begin with the following security policy requiring that *if a request is serviced at a web-server then it must not have been denied by a firewall*. In practice, this policy would be a part of a larger specification. However, this excerpt is enough to illustrate how  $\mathcal{L}_3$ 's semantics deal with logging failures. We formalize this policy as  $\Box \forall r. \psi_1$ , where

$$\psi_1 := \text{service}(r) \rightarrow \neg \blacklozenge_{[0,4]} \text{deny}(r).$$

When there are no failures, then any serviced request that has previously been denied violates the policy, and is contained in  $\llbracket \psi_1 \rrbracket_f^{\bar{D}, \tau}$ . If the web-server's logger crashes at a time point  $\tau$ , i.e.  $\text{service}_{\perp}^{\bar{D}, \tau} = |\bar{D}|$ , then all requests that had been denied at the previous four time points by the firewall potentially violate the policy, i.e.  $\llbracket \psi_1 \rrbracket_{\perp}^{\bar{D}, \tau} = \bigcup_{\tau'} \text{deny}_{\tau'}^{\bar{D}, \tau}$ , where  $\tau - 4 < \tau' \leq \tau$ . If, however, there are no denied requests in the designated interval, the set  $\llbracket \psi_1 \rrbracket_{\perp}^{\bar{D}, \tau}$  is empty and the policy is therefore satisfied. This shows that not all logging failures must result in potential violations. We note that if all unknown events are treated as not to have happened, then the policy would be *wrongly* reported as satisfied.

For our second example, we focus on formalizing inter-business contracts. These contracts often specify obligations that the signing parties must enforce regarding the treatment of sensitive documents used during the collaborations. To ensure that each party complies with its obligations, a policy must specify how events are combined from different logs belonging to different stakeholders. For example, when two companies exchange sensitive information, the contract might say that *all received documents must be paid for within 5 days*. A straightforward, but naive, formalization of this policy is  $\Box \forall d. \psi_2$ , where

$$\psi_2 := \text{receive}(d) \rightarrow \blacklozenge_{[0,6]} \text{pay}(d).$$

The *receive* event is taken from the receiving stakeholder's log. This specification assumes that the receiving stakeholder is honest, since if its IT system does not log a received document, the stakeholder's behavior is trivially compliant according to the given specification. We can attempt to expand the formalization to include the sender's *send* event (from the sender's log) as follows

$$\psi'_2 := \text{send}(d) \vee \text{receive}(d) \rightarrow \blacklozenge_{[0,6]} \text{pay}(d).$$

In this case, the receiver is obliged to pay if either it receives a document, or the sender says that it has sent the document. However, this is also unsatisfactory, as the sender can cheat and insert fictitious *send* events causing policy violations. In  $\mathcal{L}_3$  we can combine the logs with the  $\otimes$  operator and obtain<sup>1</sup>

$$\psi''_2 := \text{send}(d) \otimes \text{receive}(d) \rightarrow \blacklozenge_{[0,6]} \text{pay}(d).$$

<sup>1</sup> We assume that the time granularity is coarse enough to allow *receive* and *send* happen at the same time point. If a *receive* can happen with a delay of, e.g., at most one time unit after a *send*, a more elaborate formalization is required:

$$\begin{aligned} \Box \forall d. & (\text{send}(d) \wedge (\text{send}(d) \otimes \blacklozenge_{[0,2]} \text{receive}(d)) \rightarrow \blacklozenge_{[0,6]} \text{pay}(d)) \wedge \\ & (\text{receive}(d) \wedge (\text{receive}(d) \otimes \blacklozenge_{[0,2]} \text{send}(d)) \rightarrow \blacklozenge_{[0,5]} \text{pay}(d)). \end{aligned}$$

In this case, all disagreements at some  $\tau$  about payments are in  $\llbracket \psi'_2 \rrbracket_{\perp}^{\bar{\mathcal{D}}, \tau}$ , since  $\perp \rightarrow \mathbf{f}$  is  $\perp$ . The specification no longer favors one stakeholder over the other. This has the benefit of not requiring additional pre-processing of logs, which would need its own language and semantics. We remark that the given specification cannot be directly expressed in existing compliance policy languages because  $\perp$  does not exist at the object level in those languages.

For our third example, we consider a form of separation-of-duty constraint [2]: *a subject  $s$  may access an object  $o$  if it has not previously accessed some object  $o'$ , where  $o'$ 's dataset conflicts with  $o$ 's*. One possible formalization of this requirement is

$$\Box \forall s. \forall o. \forall d. \forall o'. \forall d'. \text{access}(s, o, d) \wedge (\blacklozenge \text{access}(s, o', d')) \rightarrow \neg \text{conflict}(d, d').$$

In this example,  $\text{access}(s, o, d)$  records that  $s$  accessed  $o$  in a dataset  $d$ . The predicate  $\text{conflict}$  does not correspond to an event; it describes a property of a system state. When having the events  $\text{conflict}_s$  and  $\text{conflict}_f$  at hand, which mark the start point and the end point of two datasets being conflicting, the formula  $\neg \text{conflict}_f(d, d') \text{S} \text{conflict}_s(d, d')$  can be used to describe this state property. For the sake of brevity, we assume that an object belongs to at most one dataset. In case  $s$  accessed an  $o$ , and it is unknown whether  $s$  had any other accesses, then if there exists  $d'$  in conflict with  $d$ , such an access is a potential violation.

Notice that the above formalization only considers whether the data items are in conflict at the time point when  $o$  is accessed. This means that even if the datasets are in conflict just before the access, the policy is not violated. With respect to the separation-of-duty requirement, one may say that this behavior is in a compliance *gray* area. In  $\mathcal{L}_3$ , we define the following temporal connective  $\mathbf{C}_I$  that treats such *gray* areas as  $\perp$ , signaling that it is unclear whether the policy is satisfied or violated:

$$\mathbf{C}_I \psi := (\blacklozenge_I \psi) \otimes (\blacksquare_I \psi).$$

Intuitively,  $\mathbf{C}_I \psi$  insists that the truth value of  $\psi$  does not change in the given past interval  $I$ . Any change results in  $\perp$ , and otherwise the truth value is not changed. We can define a similar temporal connective using  $\Box$  and  $\blacklozenge$  to mark a future gray zone. We make use of  $\mathbf{C}_I$  by changing the original formalization to

$$\Box \forall s. \forall o. \forall d. \forall o'. \forall d'. \text{access}(s, o, d) \wedge (\blacklozenge \text{access}(s, o', d')) \rightarrow \mathbf{C}_{[0,2)} \neg \text{conflict}(d, d'),$$

where  $[0, 2)$  is a two-day *gray* zone interval.

## 4 Monotonicity and Compositional Expressiveness

A logging knowledge base may grow in knowledge by resolving missing information about the occurrences and non-occurrences of events, i.e., moving elements from  $r_{\perp}^{\mathcal{D}\tau}$  to the relations  $r_{\mathbf{t}}^{\mathcal{D}\tau}$  or  $r_{\mathbf{f}}^{\mathcal{D}\tau}$ .

**Definition 5.** An extension of a logging knowledge base  $\bar{\mathcal{D}} = (\mathcal{D}_0, \mathcal{D}_1, \dots)$  over  $S = (C, R, \iota)$  is a logging knowledge base  $\bar{\mathcal{D}}^* = (\mathcal{D}_0^*, \mathcal{D}_1^*, \dots)$  over  $S$  with  $|\bar{\mathcal{D}}^*| = |\bar{\mathcal{D}}|$ ,  $c^{\bar{\mathcal{D}}} = c^{\bar{\mathcal{D}}^*}$  for all  $c \in C$ , and  $r_b^{\mathcal{D}\tau} \subseteq r_b^{\mathcal{D}^*\tau}$  for all  $b \in \{\mathbf{t}, \mathbf{f}\}$ ,  $\tau \in \mathbb{N}$ , and  $r \in R$ .



Under Assumption 1, an extension either does not alter a relation  $r_{\perp}^{\mathcal{D}\tau}$  or empties  $r_{\perp}^{\mathcal{D}\tau}$  by moving finitely many elements to  $r_{\mathbf{t}}^{\mathcal{D}\tau}$  and the remaining elements to  $r_{\mathbf{f}}^{\mathcal{D}\tau}$ .

We say that a policy specification is *monotonic* if the  $\mathbf{t}$  and  $\mathbf{f}$  evaluations, over a given logging knowledge base, can never be retracted for any of its extensions. In other words, regardless of how the logging base's incompleteness is resolved, the policy violations and satisfactions persist. Monotonicity is a vital requirement for a compliance policy, because monotonic specifications prevent a non-compliant behavior from being turned into a compliant behavior by holding back information. In the following, we establish that for  $\mathcal{L}_3$  all policy specifications are monotonic by construction. To formalize monotonicity, we first order the truth values with a partial ordering  $\leq_k$  as follows:  $\perp \leq_k \mathbf{f}$ ,  $\perp \leq_k \mathbf{t}$ , and  $\mathbf{f}$  and  $\mathbf{t}$  are incomparable. In short,  $\mathbf{f}$  and  $\mathbf{t}$  contain more knowledge than  $\perp$ . The following theorem states that the evaluations of  $\mathcal{L}_3$ 's formulas do not reduce the amount of knowledge, when incompleteness is resolved in a logging knowledge base's extension.

**Theorem 6.** *Given an  $\mathcal{L}_3$  formula  $\psi$ , a valuation  $\theta$ , and a logging knowledge base  $\bar{\mathcal{D}}$ , then  $\llbracket \psi \rrbracket^{\bar{\mathcal{D}}, \theta, \tau} \leq_k \llbracket \psi \rrbracket^{\bar{\mathcal{D}}^*, \theta, \tau}$ , for all extensions  $\bar{\mathcal{D}}^*$  of  $\bar{\mathcal{D}}$  and all  $\tau \in \mathbb{N}$ .*

*Proof.* From the definition of a logging knowledge base's extension, and by structural induction using the fact that all of  $\mathcal{L}_3$  connectives' corresponding operators are  $\leq_k$ -monotonic, including the infinitary operators for temporal connectives.

As a corollary, given a compliance policy  $\square \forall \bar{x}. \psi$ , a logging knowledge base  $\bar{\mathcal{D}}$ , a valuation  $\theta$ , and a time point  $\tau$ , if  $\llbracket \psi \rrbracket^{\bar{\mathcal{D}}, \theta, \tau}$  is  $\mathbf{t}$  or  $\mathbf{f}$ , then this evaluation persists at  $\tau$ , for all extensions  $\bar{\mathcal{D}}^*$ . Moreover, we have  $\llbracket \psi \rrbracket_{\mathbf{f}}^{\bar{\mathcal{D}}^*, \tau} \supseteq \llbracket \psi \rrbracket_{\mathbf{f}}^{\bar{\mathcal{D}}, \tau}$  and  $\llbracket \psi \rrbracket_{\mathbf{t}}^{\bar{\mathcal{D}}^*, \tau} \supseteq \llbracket \psi \rrbracket_{\mathbf{t}}^{\bar{\mathcal{D}}, \tau}$ , for all extensions  $\bar{\mathcal{D}}^*$  and  $\tau \in \mathbb{N}$ . Therefore, even with incomplete knowledge it is sound to report the elements in  $\llbracket \psi \rrbracket_{\mathbf{f}}^{\bar{\mathcal{D}}, \tau}$  as policy violations when monitoring  $\bar{\mathcal{D}}$ .

Given that all  $\mathcal{L}_3$  policies are monotonic, an important question is: *Can all monotonic compositional operators for combining events from different logs be defined as syntactic sugar in  $\mathcal{L}_3$ ?* If the answer is positive, then  $\mathcal{L}_3$  does not need to be further extended. An  $n$ -ary three-valued operator  $O : \mathbf{3}^n \rightarrow \mathbf{3}$  is *representable* using a set  $\mathcal{C}$  of operators if  $O$  can be written as the functional composition of operators in  $\mathcal{C}$ . We utilize the following theorem to show that any monotonic operator can be expressed in  $\mathcal{L}_3$ .

**Theorem 7 (Blamey [10]).** *For any  $n \in \mathbb{N}$ , every  $\leq_k$ -monotonic  $n$ -ary operator over the  $\mathbf{3}$  truth space is representable using the set  $\{\mathbf{f}, \neg, \wedge, \otimes\}$  of operators.*

Blamey's proof is constructive and yields a function that given a monotonic operator produces an expression showing how to compose the operators  $\mathbf{f}$ ,  $\neg$ ,  $\wedge$ , and  $\otimes$ . As  $\mathcal{L}_3$  has all the corresponding connectives, such an expression can be seen as a formula in  $\mathcal{L}_3$ . Hence  $\mathcal{L}_3$  can express any  $n$ -ary three-valued  $\leq_k$ -monotonic operator, including those for combining different logs.

## 5 Monitoring Algorithm

The input of our algorithm consists of a compliance policy  $\Box \forall \bar{x}. \varphi$  and a logging knowledge base  $\bar{\mathcal{D}}$  over a signature  $S = (C, R, \iota)$ . The algorithm iteratively processes the logging structures  $\mathcal{D}_\tau$ , for each  $\tau \in \mathbb{N}$ . To process a structure  $\mathcal{D}_\tau$  for formulas with bounded future operators, the algorithm might need to process structures  $\mathcal{D}_{\tau'}$  with  $\tau' > \tau$  as well. When run in the *on-line* mode, the algorithm waits until such structures become available. For the rest of this section, we fix the signature  $S$ , the logging knowledge base  $\bar{\mathcal{D}}$ , and the policy  $\Box \forall \bar{x}. \varphi$ . Furthermore, we assume that the domain  $|\bar{\mathcal{D}}|$  is infinite.

At each iteration  $\tau$ , the algorithm outputs a triple  $(S_t^\tau, S_f^\tau, S_\perp^\tau)$ , where for each  $v \in \mathbf{3}$ , the element  $S_v^\tau$  is either  $\text{Fin } V$ ,  $\text{CoFin}$ , or  $\text{None}$ , where  $\text{Fin}$ ,  $\text{CoFin}$ , and  $\text{None}$  are labels standing respectively for “finite set”, “cofinite set”, and “inconclusive”, and  $V$  is a finite set.

Our algorithm is sound, i.e. if  $S_v^\tau = \text{Fin } V$  then  $V = \llbracket \varphi \rrbracket_v^{\bar{\mathcal{D}}, \tau}$ , for all  $v \in \mathbf{3}$  and  $\tau \in \mathbb{N}$ . However, our algorithm is not *complete*, where completeness means that the algorithm always returns a value from which one can deduce all compliant tuples  $(\llbracket \varphi \rrbracket_t^{\bar{\mathcal{D}}, \tau})$ , all violations  $(\llbracket \varphi \rrbracket_f^{\bar{\mathcal{D}}, \tau})$ , and all potential violations  $(\llbracket \varphi \rrbracket_\perp^{\bar{\mathcal{D}}, \tau})$ . Note that when  $\varphi$  has free variables, all these sets cannot be explicitly output, as at least one is infinite. However, if two sets are finite, then the third one is cofinite, and it is thus implicitly determined. Therefore our algorithm is complete when at least two of the elements of the returned triples are of the form  $\text{Fin } V$ . When  $\varphi$  is closed, completeness means that at each iteration a truth value is returned, as the triples  $(\text{Fin } \{\emptyset\}, \text{Fin } \emptyset, \text{Fin } \emptyset)$ ,  $(\text{Fin } \emptyset, \text{Fin } \{\emptyset\}, \text{Fin } \emptyset)$ , and  $(\text{Fin } \emptyset, \text{Fin } \emptyset, \text{Fin } \{\emptyset\})$  correspond respectively with the truth values  $\mathbf{t}$ ,  $\mathbf{f}$ , and  $\perp$ .

Incompleteness of our algorithm is rooted in the standard issues that arise when dealing with infinite domains [3], which  $\mathcal{L}_3$  inherits from first-order queries in the Boolean setting. Consider for instance the formula  $\psi = p(x) \vee q(y)$  with  $x \neq y$  and assume that  $p_t^{\bar{\mathcal{D}}, \tau}$  and  $q_t^{\bar{\mathcal{D}}, \tau}$  are finite and non-empty, and  $p_\perp^{\bar{\mathcal{D}}, \tau} = q_\perp^{\bar{\mathcal{D}}, \tau} = \emptyset$ , for some  $\tau \in \mathbb{N}$ . Then  $\llbracket \psi \rrbracket_t^{\bar{\mathcal{D}}, \tau}$  and  $\llbracket \psi \rrbracket_f^{\bar{\mathcal{D}}, \tau}$  are neither finite nor cofinite, hence our algorithm cannot deal with it: at  $\tau$ , it returns  $(\text{None}, \text{None}, \text{Fin } \emptyset)$ . Formulas such as  $\psi$  are problematic in the Boolean setting, since their evaluation results are domain-dependent [3]. In the three-valued setting, there are similar issues, even for formulas that are non-problematic in the Boolean setting. Consider the formula  $\psi' = p(x) \wedge q(y)$  with  $p_t^{\bar{\mathcal{D}}, \tau}$  finite and non-empty and  $q_\perp^{\bar{\mathcal{D}}, \tau} = |\bar{\mathcal{D}}|$ , for some  $\tau \in \mathbb{N}$ . Then both  $\llbracket \psi' \rrbracket_t^{\bar{\mathcal{D}}, \tau}$  and  $\llbracket \psi' \rrbracket_\perp^{\bar{\mathcal{D}}, \tau}$  are infinite and domain-dependent.

Even though the algorithm is incomplete on  $\mathcal{L}_3$ , we obtain completeness for a fragment of  $\mathcal{L}_3$ , presented at the end of this section.

**Algorithmic overview.** We briefly describe the main ideas underlying the algorithm. Due to space constraints, a detailed presentation is deferred to the full paper.

The algorithm’s core is the procedure `eval`, whose arguments are a formula  $\psi$ , a finite set  $\Gamma = \{(r, E_r) \mid r \in R\}$  representing the relations of the logging structure  $\mathcal{D}_\tau$ , and a time point  $\tau$ . The values  $E_r$ , i.e., the second component of

```

proc init( $\varphi$ )
  for each  $\psi \in \text{sf}(\varphi)$  with  $\psi = \psi S_I \psi'$  do
     $L_\psi \leftarrow \langle \rangle$ 

proc eval( $\varphi, \Gamma, \tau$ )
  case  $\varphi = \mathbf{f}$ 
    return (Fin  $\emptyset$ , Fin  $\{\langle \rangle\}$ , Fin  $\emptyset$ )
  case  $\varphi = r(\bar{t})$ 
     $E_r \leftarrow \text{get\_value}(r, \Gamma)$ 
    return eval_predicate( $\varphi, E_r$ )
  case  $\varphi = \neg\psi$ 
    return eval_neg(eval( $\psi, \Gamma, \tau$ ))
  case  $\varphi = \psi \wedge \psi'$ 
     $E_\psi \leftarrow (\psi, \text{eval}(\psi, \Gamma, \tau))$ 
     $E_{\psi'} \leftarrow (\psi', \text{eval}(\psi', \Gamma, \tau))$ 
    return eval_and( $E_\psi, E_{\psi'}$ )
  case  $\varphi = \psi \otimes \psi'$ 
     $E_\psi \leftarrow (\psi, \text{eval}(\psi, \Gamma, \tau))$ 
     $E_{\psi'} \leftarrow (\psi', \text{eval}(\psi', \Gamma, \tau))$ 
    return eval_times( $E_\psi, E_{\psi'}$ )
  case  $\varphi = \forall \bar{x}. \psi$ 
     $E_\psi \leftarrow \text{eval}(\psi, \Gamma, \tau)$ 
    return eval_forall( $\bar{x}, \psi, E_\psi$ )
  case  $\varphi = \psi S_I \psi'$ 
     $E_\psi \leftarrow \text{eval}(\psi, \Gamma, \tau)$ 
     $E_{\psi'} \leftarrow \text{eval}(\psi', \Gamma, \tau)$ 
    return eval_since( $\varphi, \tau, E_\psi, E_{\psi'}$ )

```

**Fig. 2.** The init and eval procedures.

elements in  $\Gamma$ , as well as the return value of the eval procedure, are triples of the form  $(S_t, S_f, S_\perp)$ , where each  $S_v$  with  $v \in \mathbf{3}$  is either Fin  $V$ , CoFin, or None. Such values satisfy (either by Assumption 1 or by construction) the following invariant with regard to some formula  $\gamma$  and time point  $\tau$ : if  $S_v = \text{Fin } V$ , then  $\llbracket \gamma \rrbracket_v^{\mathcal{D}, \tau}$  is a finite subset of  $|\mathcal{D}|^{|\text{fv}(\gamma)|}$  and  $V = \llbracket \gamma \rrbracket_v^{\mathcal{D}, \tau}$ ; if  $S_v = \text{CoFin}$ , then  $\llbracket \gamma \rrbracket_v^{\mathcal{D}, \tau}$  is a cofinite subset of  $|\mathcal{D}|^{|\text{fv}(\gamma)|}$  and the other two elements of the triple are of the form  $S_{v'} = \text{Fin } V'$ , for  $v' \in \mathbf{3} \setminus \{v\}$ . This invariant is denoted as  $\text{Inv}(\gamma, \tau, E)$ , where  $E = (S_t, S_f, S_\perp)$ . By Assumption 1, the values  $E_r$  from the set  $\Gamma$  satisfy the invariant  $\text{Inv}(r(\bar{x}), \tau, E_r)$ , where  $\bar{x}$  is a sequence of distinct variables of length  $\iota(r)$ . We prove in Theorem 9 that the return value  $E$  of  $\text{eval}(\varphi, \Gamma, \tau)$  satisfies the invariant  $\text{Inv}(\varphi, \tau, E)$ , thus establishing the correctness of our algorithm.

The eval procedure, given in Figure 2, is called recursively over  $\psi$ 's sub-formulas. The procedure performs a case distinction on all possible top-level connectives. Some of the sub-procedures used by eval are in Figure 3, while the remaining the pseudo-code is given in the full paper.

Next, we sketch each case of the eval procedure. The simplest case is when  $\psi$  is the truth value  $\mathbf{f}$ . In this case we simply return the triple  $(\text{Fin } \emptyset, \text{Fin } \{\langle \rangle\}, \text{Fin } \emptyset)$ . When  $\psi$  is of the form  $r(\bar{t})$  for some predicate  $r$ , we first retrieve the value  $E_r$  associated with  $r$  from the set  $\Gamma$  of pairs. We then retrieve the sets  $\llbracket r(\bar{t}) \rrbracket_v^{\mathcal{D}, \tau}$  from  $r_v^{\mathcal{D}, \tau}$ , for each  $v \in \mathbf{3}$ , by filtering the relations  $r_v^{\mathcal{D}, \tau}$  according to the implicit constraints present in the sequence  $\bar{t}$  of constants and variables.

To evaluate formulas  $\psi$  whose top-most connective is a non-temporal connective, we first evaluate the direct sub-formulas of  $\psi$  and then compute, whenever possible, the sets  $\llbracket \psi \rrbracket_v^{\mathcal{D}, \tau}$  for  $v \in \mathbf{3}$ , using the equalities given in Lemma 8 below. These equalities extend the standard equalities that express the relationship between first-order logic and relational algebra, from the Boolean to the three-valued setting. They use the relational algebra operators *projection* and *join* [3]. We refer to the full paper for their formal definitions, and here we proceed with their intuitive description. As the temporal aspect is not relevant in this case of

eval, we also fix the time point  $\tau$  and drop the superscript in  $\llbracket \psi \rrbracket_v^{\bar{\mathcal{D}}, \tau}$ , i.e., we just write  $\llbracket \psi \rrbracket_v$ , for  $v \in \mathbf{3}$  and a formula  $\psi$ .

Given a formula  $\psi$  and a truth value  $v \in \mathbf{3}$ , we can see the set  $\llbracket \psi \rrbracket_v$  as a *named relation*, where columns in  $\llbracket \psi \rrbracket_v$  are named by the free variables in  $\bar{fv}(\psi)$ . Given a free variable  $x$  of  $\psi$ , the *projection* of the tuples in the relation  $\llbracket \psi \rrbracket_v$  on the columns corresponding to other free variables is denoted  $\pi_x(\llbracket \psi \rrbracket_v)$ . For instance, if  $\llbracket p(x, y) \rrbracket_t = \{(0, 2), (1, 2), (1, 3)\}$ , then  $\pi_x(\llbracket p(x, y) \rrbracket_t) = \{(2), (3)\}$ . For  $v, v' \in \mathbf{3}$ , the *natural join* of the sets  $\llbracket \psi \rrbracket_v$  and  $\llbracket \psi' \rrbracket_{v'}$ , denoted  $\llbracket \psi \rrbracket_v \bowtie \llbracket \psi' \rrbracket_{v'}$ , is the set of tuples for which the projections on the columns, corresponding to  $\psi$ 's and  $\psi'$ 's free variables, are in  $\llbracket \psi \rrbracket_v$  and respectively in  $\llbracket \psi' \rrbracket_{v'}$ , and the fields of which match on the common free variables. For instance, if  $\llbracket q(y, z) \rrbracket_t = \{(2, 4)\}$ , then  $\llbracket p(x, y) \rrbracket_t \bowtie \llbracket q(y, z) \rrbracket_t = \{(0, 2, 4), (1, 2, 4)\}$ . We adopt the convention that  $\bowtie$  binds stronger than  $\cup$ .

**Lemma 8.** *Let  $\bar{\mathcal{D}}$  be a logging knowledge base,  $\tau$  be a time point, and  $\psi$  and  $\psi'$  be  $\mathcal{L}_3$  formulas. The following equalities hold:*

$$\begin{aligned}
\llbracket \neg \psi \rrbracket_v &= \llbracket \psi \rrbracket_{-v}, \text{ if } v \in \mathbf{3} \\
\llbracket \psi \wedge \psi' \rrbracket_t &= \llbracket \psi \rrbracket_t \bowtie \llbracket \psi' \rrbracket_t \\
\llbracket \psi \wedge \psi' \rrbracket_f &= \llbracket \psi \rrbracket_f \cup \llbracket \psi' \rrbracket_f, \text{ if } fv(\psi) = fv(\psi') \\
\llbracket \psi \wedge \psi' \rrbracket_{\perp} &= \llbracket \psi \rrbracket_t \bowtie \llbracket \psi' \rrbracket_{\perp} \cup \llbracket \psi \rrbracket_{\perp} \bowtie \llbracket \psi' \rrbracket_t \cup \llbracket \psi \rrbracket_{\perp} \bowtie \llbracket \psi' \rrbracket_{\perp} \\
\llbracket \psi \otimes \psi' \rrbracket_b &= \llbracket \psi \rrbracket_b \bowtie \llbracket \psi' \rrbracket_b, \text{ if } b \in \{\mathbf{t}, \mathbf{f}\} \\
\llbracket \psi \otimes \psi' \rrbracket_{\perp} &= \llbracket \psi \rrbracket_{\perp} \bowtie \llbracket \psi' \rrbracket_{\perp} \cup \llbracket \psi \rrbracket_t \bowtie \llbracket \psi' \rrbracket_f \cup \llbracket \psi \rrbracket_f \bowtie \llbracket \psi' \rrbracket_t \\
\llbracket \forall x. \psi \rrbracket_t &= \emptyset, \text{ if } \llbracket \psi \rrbracket_t \text{ is finite and } x \in fv(\psi) \\
\llbracket \forall x. \psi \rrbracket_f &= \pi_x(\llbracket \psi \rrbracket_f), \text{ if } x \in fv(\psi) \\
\llbracket \forall x. \psi \rrbracket_{\perp} &= \pi_x(\llbracket \psi \rrbracket_{\perp}) \setminus \pi_x(\llbracket \psi \rrbracket_f), \text{ if } x \in fv(\psi)
\end{aligned}$$

These equalities provide a method to compute, under the stated conditions, the relations  $\llbracket \psi \rrbracket_v$  from the corresponding relations for  $\psi$ 's direct sub-formulas. For instance, if  $\psi = \psi_1 \wedge \psi_2$  and  $\llbracket \psi_1 \rrbracket_t, \llbracket \psi_2 \rrbracket_t$  are finite relations, then  $\llbracket \psi \rrbracket_t$  is a finite relation given by the join of the other two relations. Furthermore, when  $\llbracket \psi_1 \rrbracket_t$  is finite,  $\llbracket \psi_2 \rrbracket_t$  is cofinite, and  $fv(\psi_2) \subseteq fv(\psi_1)$ , then  $\llbracket \psi \rrbracket_t$  is a finite relation that we can compute as  $\llbracket \psi_1 \rrbracket_t \bowtie \llbracket \psi_2 \rrbracket_t = \llbracket \psi_1 \rrbracket_t \bowtie (|\bar{\mathcal{D}}|^{fv(\psi_2)} \setminus (\llbracket \psi_2 \rrbracket_f \cup \llbracket \psi_2 \rrbracket_{\perp}))$ . Note that the condition  $fv(\psi_2) \subseteq fv(\psi_1)$  is essential, as otherwise  $\llbracket \psi_1 \rrbracket_t \bowtie \llbracket \psi_2 \rrbracket_t$  may be infinite. For example, if  $\bar{fv}(\psi_1) = (x)$  and  $\bar{fv}(\psi_2) = (x, y)$  with  $\llbracket \psi_1 \rrbracket_t = \{(1)\}$ ,  $\llbracket \psi_2 \rrbracket_f = \{(1, 2)\}$ , and  $\llbracket \psi_2 \rrbracket_{\perp} = \{(3, 4)\}$ , then  $\llbracket \psi \rrbracket_t = \{1\} \times (|\bar{\mathcal{D}}| \setminus \{2, 4\})$ . The same method is applied to each of the other sub-cases of the binary connectives.

The described approach is implemented through the procedures `eval_neg`, `eval_and`, `eval_times`, and `eval_forall`, given in Figure 3. Each procedure returns a triple  $(R_t, R_f, R_{\perp})$ , where  $R_v$  is a value computed based on the identities in Lemma 8 using the procedures `join` and `union`, which are given in the full paper. The `join` procedure takes as arguments tuples  $(\psi, E)$  and  $(\psi', E')$ , and truth values  $v$  and  $v'$ . Provided that the invariants  $Inv(\psi, \tau, E)$  and  $Inv(\psi', \tau, E')$  are satisfied, the return value is either `Fin` ( $\llbracket \psi \rrbracket_v \bowtie \llbracket \psi' \rrbracket_{v'}$ ) or `None`, depending on whether a finite relation can be computed. The `union` procedure has similar arguments and return values. The auxiliary procedures `update_cofin` and `update` from Figure 3 handle the following corner case: If two elements of the newly

```

proc eval_and( $H_\psi, H_{\psi'}$ )
   $R_t \leftarrow \text{join}(H_\psi, H_{\psi'}, t, t)$ 
   $R_f \leftarrow \text{union}(H_\psi, H_{\psi'}, f, f)$ 
   $R_\perp \leftarrow \text{eval\_and}_\perp(H_\psi, H_{\psi'})$ 
  return update_cofin( $\psi \wedge \psi', R_t, R_f, R_\perp$ )

proc eval_and_⊥( $H_\psi, H_{\psi'}$ )
   $R_1 \leftarrow \text{join}(H_\psi, H_{\psi'}, t, \perp)$ 
   $R_2 \leftarrow \text{join}(H_\psi, H_{\psi'}, \perp, t)$ 
   $R_3 \leftarrow \text{join}(H_\psi, H_{\psi'}, \perp, \perp)$ 
  case  $R_1, R_2, R_3 = \text{Fin } V_1, \text{Fin } V_2, \text{Fin } V_3$ 
    return  $\text{Fin } (V_1 \cup V_2 \cup V_3)$ 
  otherwise
    return None

proc eval_neg( $S_t, S_f, S_\perp$ )
  return ( $S_f, S_t, S_\perp$ )

proc eval_forall( $\bar{x}, \psi, (S_t, S_f, S_\perp)$ )
  ( $R_t, R_f, R_\perp$ )  $\leftarrow$  (None, None, None)
  case  $S_t = \text{Fin } T$ 
     $R_t \leftarrow \text{Fin } \emptyset$ 
    case  $S_\perp = \text{Fin } U$ 
       $R_\perp \leftarrow \text{Fin } \emptyset$ 
    case  $S_f = \text{Fin } F$ 
       $\bar{s} \leftarrow \text{get\_positions}(\bar{x}, \psi)$ 
       $R_f \leftarrow \text{Fin } (\pi_{\bar{s}}(F))$ 
      case  $S_\perp = \text{Fin } U$ 
         $R_\perp \leftarrow \text{Fin } (\pi_{\bar{s}}(U) \setminus \pi_{\bar{s}}(F))$ 
  return update_cofin( $\forall \bar{x}. \psi, R_t, R_f, R_\perp$ )

proc eval_times( $H_\psi, H_{\psi'}$ )
   $R_t \leftarrow \text{join}(H_\psi, H_{\psi'}, t, t)$ 
   $R_f \leftarrow \text{join}(H_\psi, H_{\psi'}, f, f)$ 
   $R_\perp \leftarrow \text{eval\_times}_\perp(H_\psi, H_{\psi'})$ 
  return update_cofin( $\psi \otimes \psi', R_t, R_f, R_\perp$ )

proc eval_times_⊥( $H_\psi, H_{\psi'}$ )
   $R_1 \leftarrow \text{union}(H_\psi, H_{\psi'}, \perp, \perp)$ 
   $R_2 \leftarrow \text{join}(H_\psi, H_{\psi'}, t, f)$ 
   $R_3 \leftarrow \text{join}(H_\psi, H_{\psi'}, f, t)$ 
  case  $R_1, R_2, R_3 = \text{Fin } V_1, \text{Fin } V_2, \text{Fin } V_3$ 
    return  $\text{Fin } (V_1 \cup V_2 \cup V_3)$ 
  otherwise
    return None

proc update_cofin( $\psi, R_t, R_f, R_\perp$ )
   $R_t \leftarrow \text{update}(\psi, R_t, R_f, R_\perp)$ 
   $R_f \leftarrow \text{update}(\psi, R_f, R_t, R_\perp)$ 
   $R_\perp \leftarrow \text{update}(\psi, R_\perp, R_t, R_f)$ 
  return ( $R_t, R_f, R_\perp$ )

proc update( $\psi, R_1, R_2, R_3$ )
  case  $R_2 = \text{Fin } \_$  and  $R_3 = \text{Fin } \_$ 
    if  $fv(\psi) \neq \emptyset$  then return CoFin
    else if  $R_2 = \text{Fin } \emptyset$  and  $R_3 = \text{Fin } \emptyset$  then
      return  $\text{Fin } \{\{\}\}$ 
    else
      return  $\text{Fin } \emptyset$ 
  otherwise
    return  $R_1$ 

```

**Fig. 3.** The eval\_neg, eval\_and, eval\_times, and eval\_forall procedures.

formed triple  $(R_t, R_f, R_\perp)$  are of the form  $\text{Fin } V$  and the remaining element is None, then  $\text{update\_cofin}(\psi, R_t, R_f, R_\perp)$  changes None to either CoFin if  $fv(\varphi) \neq \emptyset$ , or otherwise (when  $fv(\varphi) = \emptyset$ ) to  $\text{Fin } \{\{\}\}$  or  $\text{Fin } \emptyset$  depending on the truth value that should be returned. This ensures that the invariant  $Inv$  is preserved by the return value of the eval\_and, eval\_times, and eval\_forall procedures.

Finally, we consider the temporal operators. Let  $\psi = \alpha S_I \beta$ . For efficiency, eval maintains between iterations a sequence  $L_\psi$ , which is initialized by the init procedure with the empty sequence. The sequence  $L_\psi$  contains values  $E_{\tau'}$  that satisfy the invariant  $Inv(\alpha S_{[\delta, \delta]} \beta, \tau, E_{\tau'})$ , where  $\delta = \tau - \tau'$  and  $\tau'$  is such that  $0 \leq \tau - \tau' < b$ , with  $I = [a, b)$ . In this way, the sub-formulas  $\alpha$  and  $\beta$  are not re-evaluated at previous time points  $\tau'$ . Instead, the result of their evaluation is stored in  $L_\psi$ . The return value is computed by iteratively calling eval\_or on the elements  $E_{\tau'}$  of  $L_\psi$  for which  $(\tau - \tau') \in I$ . This last step reflects the equivalence between  $\alpha S_I \beta$  and  $\bigvee_{\delta \in I} \alpha S_{[\delta, \delta]} \beta$ . Given two formulas  $\psi_1$  and  $\psi_2$  and two values  $E_1$  and  $E_2$  satisfying respectively the invariants  $Inv(\psi_1, \tau, E_1)$  and  $Inv(\psi_2, \tau, E_2)$ , the procedure eval\_or returns a value  $E$  that satisfies  $Inv(\psi_1 \vee \psi_2, \tau, E)$ .

The case for Until is analogous to Since. The only significant difference is that the procedure must delay its answer until all relevant events have occurred. Various optimizations, which we mention in the full paper, can further improve the efficiency of handling temporal operators.

The following theorem establishes termination and soundness of our algorithm. To state it formally, we first explicitly define the relationship between the arguments  $\Gamma_\tau$  of the `eval` procedure, and the logging structures  $\mathcal{D}_\tau$  of  $\bar{\mathcal{D}}$ . We let

$$\text{triples}(\mathcal{D}_\tau) := \{(r, (\text{val}(r_{\bar{t}}^{\mathcal{D}_\tau}), \text{val}(r_{\bar{t}}^{\mathcal{D}_\tau}), \text{val}(r_{\perp}^{\mathcal{D}_\tau}))) \mid r \in R\},$$

where  $\text{val}(V)$  is `Fin`  $V$  if  $V$  is finite, and is `CoFin` otherwise. Thus  $\Gamma_\tau = \text{triples}(\mathcal{D}_\tau)$ .

**Theorem 9.** *Let  $\bar{\mathcal{D}}$  be a logging knowledge base,  $\varphi$  a formula in  $\mathcal{L}_3$ , and  $\tau \in \mathbb{N}$  a time point. The procedure `eval`( $\varphi, \Gamma_\tau, \tau$ ) returns a value  $E$  that satisfies the invariant  $\text{Inv}(\varphi, \tau, E)$ , whenever `init`( $\varphi$ ), `eval`( $\varphi, \Gamma_0, 0$ ),  $\dots$ , `eval`( $\varphi, \Gamma_{\tau-1}, \tau-1$ ) were called previously in this order, where  $\Gamma_{\tau'} = \text{triples}(\mathcal{D}_{\tau'})$ , for  $\tau' \leq \tau$ .*

**A complete fragment.** In general, our algorithm is incomplete. However, by limiting the usage of free variables, we obtain the fragment  $\mathcal{L}_3^c$  for which we guarantee completeness.

**Definition 10.** *The set  $\mathcal{L}_3^c$  of formulas is inductively defined:*

- $f \in \mathcal{L}_3^c$  and  $r(t_1, \dots, t_{i(r)}) \in \mathcal{L}_3^c$ ,
- if  $\varphi \in \mathcal{L}_3^c$ , then  $\neg\varphi \in \mathcal{L}_3^c$  and  $\forall x. \varphi \in \mathcal{L}_3^c$ ,
- if  $\varphi, \psi \in \mathcal{L}_3^c$  and either  $\text{fv}(\varphi) = \text{fv}(\psi)$ ,  $\text{fv}(\varphi) = \emptyset$ , or  $\text{fv}(\psi) = \emptyset$ , then  $\varphi \wedge \psi \in \mathcal{L}_3^c$ ,  $\varphi \otimes \psi \in \mathcal{L}_3^c$ ,  $\varphi \text{S}_I \psi \in \mathcal{L}_3^c$ , and  $\varphi \text{U}_I \psi \in \mathcal{L}_3^c$ .

Note that  $\mathcal{L}_3^c$  allows universal quantification and, by using  $\neg$ , also existential quantification of free variables, and both quantifiers can be nested freely. But if an  $\mathcal{L}_3^c$  formula contains a sub-formula with no quantifiers and two or more predicates, they must have the same free variables. As all of  $\mathcal{L}_3$ 's connectives are retained and their application is not restricted,  $\mathcal{L}_3^c$  can still express all monotonic finitary operators. However, they cannot be used as liberally as in  $\mathcal{L}_3$ .

The first and second policy examples in Section 3 fall within  $\mathcal{L}_3^c$ . However, due to the free-variable restriction, the following formula is not in  $\mathcal{L}_3^c$ :

$$\square \forall s. \forall r. \forall m. \text{send}(s, r, m) \rightarrow \diamond_I \text{authorize}(m).$$

It says that all messages  $m$ , sent by  $s$  to  $r$  must be subsequently authorized. This is a typical compliance policy from the HIPAA Privacy Rule [1]. By pushing the quantification of  $s$  and  $r$  inside the antecedent, we obtain a formula in  $\mathcal{L}_3^c$ :

$$\square \forall m. (\exists s. \exists r. \text{send}(s, r, m)) \rightarrow \diamond_I \text{authorize}(m).$$

One can check that evaluating  $\forall x. \varphi \rightarrow \psi$  and  $(\exists x. \varphi) \rightarrow \psi$ , as well as  $\exists x. \varphi \wedge \psi$  and  $(\exists x. \varphi) \wedge \psi$ , where  $x \notin \text{fv}(\psi)$ , over an arbitrary logging knowledge base and an arbitrary time point yields the same truth value.

It is not always possible to rewrite a formula such that the result falls into  $\mathcal{L}_3^c$ . Recall the third example (the separation-of-duty requirement) from Section 3. Clearly, it does not fall within  $\mathcal{L}_3^c$ . However, if there are finitely many datasets, we can partially ground the formula, obtaining a family of formulas  $\varphi_{d,d'}$ , where  $d$  and  $d'$  range over the datasets. Each is in  $\mathcal{L}_3^c$  after similar rewriting as above:

$$\varphi_{d,d'} := \square (\exists s. (\exists o. \text{access}(s, o, d)) \wedge \exists o'. \blacklozenge \text{access}(s, o', d')) \rightarrow \text{C}_{[0,2]} \neg \text{conflict}(d, d').$$

Syntactic rewriting and partial grounding cannot always be applied. Still,  $\mathcal{L}_3^c$  is an expressive fragment that captures a wide-range of compliance policies.

Finally, we state our result on the algorithm’s completeness on  $\mathcal{L}_3^c$  formulas. To do so, we define the stronger invariant  $Inv_c(\varphi, \tau, E)$  which, in addition to  $Inv(\varphi, \tau, E)$ , requires that there are  $v', v'' \in \mathbf{3}$  with  $v' \neq v''$  such that  $S_{v'} = \text{Fin } V'$  and  $S_{v''} = \text{Fin } V''$  for some sets  $V', V''$ , where  $E = (S_t, S_f, S_\perp)$ .

**Theorem 11.** *Let  $\bar{\mathcal{D}}$  be a logging knowledge base,  $\varphi$  a formula in  $\mathcal{L}_3^c$ , and  $\tau \in \mathbb{N}$  a time point. The procedure  $\text{eval}(\varphi, \Gamma_\tau, \tau)$  returns a value  $E$  that satisfies the invariant  $Inv_c(\varphi, \tau, E)$ , whenever  $\text{init}(\varphi)$ ,  $\text{eval}(\varphi, \Gamma_0, 0)$ ,  $\dots$ ,  $\text{eval}(\varphi, \Gamma_{\tau-1}, \tau-1)$  were called previously in this order, where  $\Gamma_{\tau'} = \text{triples}(\mathcal{D}_{\tau'})$ , for  $\tau' \leq \tau$ .*

## 6 Related Work

The only work we are aware of that addresses the problem of compliance checking with incomplete knowledge is Garg et al. [21]. Their policy language is a restricted first-order logic. It has a more liberal usage of free variables compared to  $\mathcal{L}_3^c$ , but it does not consider  $\perp$  at the object-level and cannot express the  $\otimes$  operator. They adopt a weaker logging assumption, whereby a finite or an infinite number of event occurrences can be unknown. However, their compliance algorithm is not suitable for on-line monitoring and, more importantly, it is incomplete, even with our logging assumption. Recall our first policy example in Section 3. If the web-server’s logger crashes and there are no denials, their algorithm does not report that there are no violations. Instead, it wrongly reports that there may be potential violations, where in fact there are none. Similarly, it may also fail to report violations. For example, given a specification of the form

$$\square \forall \bar{x}. c(\bar{x}) \rightarrow \exists \bar{y}. c'(\bar{x}, \bar{y}) \wedge \forall \bar{z}. \varphi(\bar{x}, \bar{y}, \bar{z}),$$

then all  $\bar{x}$  that violate the policy by making  $c$  true and  $\varphi$  false, but for which all  $c'$  events are missing, are not reported. This is because their algorithm evaluates formulas in a top-down fashion: it first finds all  $\bar{x}$  that satisfy  $c$ , then it partially grounds<sup>2</sup> the consequent, then it finds all  $\bar{y}$  that satisfy  $c'$ , and then partially grounds  $\varphi$ , and so forth. However, if there are no partial groundings, the algorithm stops further evaluations. In contrast, since our algorithm works in a bottom-up fashion, it does not have this problem.

The problem of incompleteness and disagreements is also present in other fields, and some approaches there are also based on many-valued logics. Some access-control policy languages [15, 18] use multiple truth values to represent different access-control decisions. These languages are propositional and do not support temporal reasoning. Several model-checking approaches [13, 14, 16] also consider a many-valued truth space. However, their many-valued semantics do not guarantee policy-compliance monotonicity. Furthermore, their specification languages only have the classical Boolean and temporal connectives.

<sup>2</sup> Their logging assumption and language restrictions guarantee that there are always only finitely many satisfying ground instances.

Bauer et al. [8] extend the classical LTL semantics by also assigning non-Boolean truth values to finite and complete prefixes of infinite traces. Their semantics differentiate whether all or some extensions of a finite trace satisfy a property. However, the Boolean and temporal operators are not extended over the additional truth values. Furthermore, they do not consider the ordering  $\leq_k$  of the truth values in knowledge.

Another approach to dealing with incompleteness is to make quantitative statements, e.g., how certain it is whether a property is violated. Stoller et al. [26] present such an approach for monitoring traces with gaps. Their solution first assigns probabilities to whether events happened during gaps, and then computes the overall probability that a temporal property is violated. This solution is orthogonal to ours. It requires a reliable training set to derive appropriate probability assignments for different event occurrences.

## 7 Conclusions

In complex IT systems, logging failures happen and knowledge about the occurrence of system actions is incomplete when monitoring the system. Furthermore, system components can disagree on whether actions took place. Approaches for checking system compliance based on the classical Boolean setting are insufficient since they may incorrectly report policy violations. A three-valued truth space allows us to correctly distinguish between violations and potential violations. The solution presented in this paper carefully adopts a three-value truth space so that policy evaluations are correct regardless of how knowledge gaps are resolved. The presented monitoring algorithm shows that policy violations and potential violations can be soundly and completely determined.

As future work we will investigate how to efficiently resolve potential violations as prior knowledge gaps are incrementally resolved. We also plan case studies to evaluate our monitoring algorithm in real-world settings. Finally, we would like to explore different truth spaces to distinguish between different kinds of knowledge gaps and disagreements.

**Acknowledgments.** We thank Germano Caronni and Matúš Harvan for fruitful discussions on this topic.

## References

1. The Health Insurance Portability and Accountability Act of 1996 (HIPAA), 1996. Public Law 104-191.
2. Gramm-Leach-Bliley Act of 1999 (GLBA), 1999. Public Law 106-102.
3. S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.
4. H. Barringer, A. Groce, K. Havelund, and M. Smith. Formal analysis of log files. *J. Aero. Comput. Inform. Comm.*, 7:365–390, 2010.
5. D. Basin, M. Harvan, F. Klaedtke, and E. Zălinescu. Monitoring usage-control policies in distributed systems. In *TIME'11*, pages 88–95, 2011.



6. D. Basin, M. Harvan, F. Klaedtke, and E. Zălinescu. MONPOLY: Monitoring usage-control policies. In *RV'11*, volume 7186 of *LNCS*, pages 360–364, 2012.
7. D. Basin, F. Klaedtke, S. Müller, and B. Pfitzmann. Runtime monitoring of metric first-order temporal properties. In *FSTTCS'08*, volume 2 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49–60, 2008.
8. A. Bauer, M. Leucker, and C. Schallhart. Comparing LTL semantics for runtime verification. *J. Logic Comput.*, 20(3):651–674, 2010.
9. N. D. Belnap, Jr. A useful four-valued logic. In J. M. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*, volume 2 of *Episteme*, pages 7–37. D. Reidel Publishing Company, 1977.
10. S. Blamey. Partial logic. In D. M. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 5, pages 261–353. Kluwer Academic Publishers, 2002.
11. G. Boella, J. Broersen, and L. van der Torre. Reasoning about constitutive norms, counts-as conditionals, institutions, deadlines and violations. In *PRIMA'08*, volume 5357 of *LNCS*, pages 86–97, 2008.
12. J. Broersen. On the logic of 'being motivated to achieve  $\rho$ , before  $\delta'$ '. In *JELIA'04*, volume 3229 of *LNCS*, pages 334–346, 2004.
13. G. Bruns and P. Godefroid. Model checking partial state spaces with 3-valued temporal logics. In *CAV'99*, volume 1633 of *LNCS*, pages 274–287, 1999.
14. G. Bruns and P. Godefroid. Model checking with multi-valued logics. In *ICALP'04*, volume 3142 of *LNCS*, pages 281–293, 2004.
15. G. Bruns and M. Huth. Access control via Belnap logic: Intuitive, expressive, and analyzable policy composition. *ACM Trans. Inform. Syst. Secur.*, 14(1), 2011.
16. M. Chechik, B. Devereux, S. Easterbrook, and A. Gurfinkel. Multi-valued symbolic model-checking. *ACM Trans. Softw. Eng. Meth.*, 12(4):371–408, 2003.
17. J. Chomicki. Efficient checking of temporal integrity constraints using bounded history encoding. *ACM Trans. Database Syst.*, 20(2):149–186, 1995.
18. J. Crampton and C. Morisset. PTaCL: A language for attribute-based access control in open systems. In *POST'12*, volume 7215 of *LNCS*, pages 390–409, 2012.
19. N. Dinesh, A. K. Joshi, I. Lee, and O. Sokolsky. Checking traces for regulatory conformance. In *RV'08*, volume 5289 of *LNCS*, pages 86–103, 2008.
20. M. Fitting. Kleene's logic, generalized. *J. Log. Comput.*, 1(6):797–810, 1991.
21. D. Garg, L. Jia, and A. Datta. Policy auditing over incomplete logs: Theory, implementation and applications. In *CCS'11*, pages 151–162, 2011.
22. A. Groce, K. Havelund, and M. Smith. From scripts to specification: The evaluation of a flight testing effort. In *ICSE'10*, volume 2, pages 129–138, 2010.
23. S. Hallé and R. Villemaire. Runtime enforcement of web service message contracts with data. *IEEE Trans. Serv. Comput.*, 5(2):192–206, 2012.
24. T. Hvitved, F. Klaedtke, and E. Zălinescu. A trace-based model for multiparty contracts. *J. Log. Algebr. Program.*, 81(2):72–98, 2012.
25. S. C. Kleene. *Introduction to Metamathematics*. D. Van Nostrand, Princeton, 1950.
26. S. D. Stoller, E. Bartocci, J. Seyster, R. Grosu, K. Havelund, S. A. Smolka, and E. Zadok. Runtime verification with state estimation. In *RV'11*, volume 7186 of *LNCS*, pages 193–207, 2012.