Monitoring Usage-control Policies in Distributed Systems

David Basin, Matúš Harvan, Felix Klaedtke, and Eugen Zălinescu
ETH Zurich, Computer Science Department, Switzerland

***Abstract*—We have previously presented a monitoring algorithm for compliance checking of policies formalized in an expressive metric first-order temporal logic. We explain here the steps required to go from the original algorithm to a working infrastructure capable of monitoring an existing distributed application producing millions of log entries per day. The main challenge is to correctly and efficiently monitor the trace interleavings obtained by totally ordering actions that happen at the same time. We provide solutions based on formula transformations and monitoring representative traces. We also report, for the first time, on statistics on the performance of our monitor on real-world data, providing evidence of its suitability for nontrivial applications.**

I. Introduction

Determining whether the usage of sensitive data complies with regulations and policies is a growing concern for companies, administrations, and end users alike. In the context of IT systems, this question amounts to whether one can implement processes that monitor other processes. In previous work [1], [2], we have demonstrated that metric first-order temporal logic (MFOTL) is a good candidate for monitoring data usage to determine policy compliance. In particular, the metric temporal operators allow one to formalize both qualitative and quantitative temporal relationships between actions and, as the logic is first-order, we can also formulate dependencies between the finite but unbounded number of agents and data elements in IT systems. We have given a monitoring algorithm for MFOTL [1] and many usage-control policies can be naturally formulated in the fragment that the monitor handles efficiently [2].

In this paper, we extend our previous work by deploying and evaluating our monitoring approach in a real-world concurrent and distributed setting. This is in contrast to our previous analysis [2], which we carried out in a non-distributed setting where we used log files filled with synthetically generated actions. In the following, we describe our monitoring setup and the challenges we faced. We begin with an abstract description of the systems that we handle.

System Model. The types of entities in our systems are *data*, *(data) stores*, *agents*, and *actions*. Data is stored in distributed data stores such as databases and repositories and created, read, modified, propagated, combined, and deleted by actions initiated by agents. Agents are either humans or applications, including database triggers.

Agents always access the data directly from a store and never indirectly from another agent. Whenever an agent wants to use some data, it accesses the appropriate store, uses the data, and discards it afterwards. For subsequent usage, it must access the store again. Before discarding the data, the agent may write it, possibly after processing it in some way, into the same or a different store. In this way, data can propagate between stores. A consequence of this restriction on the interaction between system entities is that the use of data is always observable at the data stores.

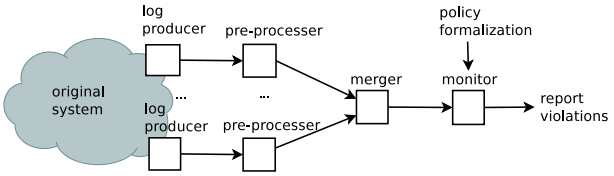

Figure 1. System Extension

Systems are governed by *(usage-control) policies*, which state requirements on the usage of the data. For example, only agents with particular credentials may modify data, or data must be deleted after two years from a given store. Agents may or may not comply with policies.

Logging and Monitoring. Given a system that is an instance of the above system model, we must extend it to support logging and monitoring. To determine whether a policy is violated we usually need to relate actions that are carried out in different parts of the system. Moreover, the ordering of actions and the time elapsed between them is important. To relate actions and the times when they happen, we log them locally, annotating each action with a timestamp, and merge these logs after some pre-processing. We then monitor this merged stream of logged actions. These system extensions are depicted in Figure 1.

Challenges and Contributions. Individual logs are totally ordered and timestamped using local clocks. However, even assuming clock synchronization [3], we have only a partial order on system actions [4] as multiple actions with the same timestamp may occur in different logs. Our main theoretical challenge is to monitor such a partially ordered set of actions, which is, in general, an intractable problem. In Section III, we identify a subclass of formulas that describe properties that are insensitive to the ordering of actions labeled by the same timestamp and for which it suffices to monitor a particular merging of the logs, namely, the merging that assumes that actions with equal timestamps happen simultaneously. Furthermore, in case

the given formula is outside this class we provide means to meaningfully monitor this merge by approximating the described property.

A practical challenge is to deploy adequate logging mechanisms. The mechanisms should be complete in that they log all occurrences of policy-relevant system actions. They should also be accurate in that if an action is logged then it has happened in the system and the corresponding log entry accurately describes the action, e.g, it describes the involved data and the associated timestamp is the actual time when the action happened. Incomplete or inaccurate logging may lead to false positives and false negatives when monitoring the system.

In Section IV, we explain how we handle these practical challenges in our case study. Where possible, we use existing logging mechanisms and extract policy-relevant information from the produced log entries. For system components where no logging was available, we either added logging directly to the components or we extended the components with proxy mechanisms that logged actions. However, proxies have limitations: agents do not necessarily access a store over a proxy and proxies see requested actions but not necessarily all the effects on the involved data. In our case, the interactions could be accurately observed but not for all agents, which led to accurate but incomplete logs.

Summarizing, we see our contributions as follows. We provide solutions for efficiently monitoring partially ordered logs, which is a central problem in monitoring real-time concurrent distributed systems. Moreover, we evaluate the performance of our monitoring approach and demonstrate its effectiveness on a real-world application.

Organization. The remainder of this paper is structured as follows. In Section II, we give background on MFOTL and our monitor. In Section III, we show how we handle the interleavings of multiple streams of logged actions from different log producers. In Section IV, we report on our case study. In Section V, we discuss related work and in Section VI, we draw conclusions. The Appendices A–D contain additional proof details. Additional details on the case study are given in Appendix E.

II. Preliminaries

We briefly review metric first-order temporal logic (MFOTL) and describe how we use it to monitor systems.

Syntax and Semantics. Let $\mathbb{I}$ be the set of nonempty intervals over $\mathbb{N}$. We will write an interval $I \in \mathbb{I}$ as $[b, b') := \{a \in \mathbb{N} \mid b \le a < b'\}$, where $b \in \mathbb{N}$, $b' \in \mathbb{N} \cup \{\infty\}$, and $b < b'$. A *signature* S is a tuple (C, R, ι), where C is a finite set of constant symbols, R is a finite set of predicates disjoint from C, and the function $\iota : R \to \mathbb{N}$ associates each predicate $r \in R$ with an arity $\iota(r) \in \mathbb{N}$. In the following, let $S = (C, R, \iota)$ be a signature and V a countably infinite set of variables, assuming $V \cap (C \cup R) = \emptyset$.

$$
\begin{array}{ll}
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models t \approx t' & \text{iff } v(t) = v(t') \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models t < t' & \text{iff } v(t) < v(t') \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models r(t_1, \ldots, t_{\iota(r)}) & \text{iff } (v(t_1), \ldots, v(t_{\iota(r)})) \in r^{\mathcal{D}_i} \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\neg \phi) & \text{iff } (\bar{\mathcal{D}}, \bar{\tau}, v, i) \not\models \phi \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\phi \vee \psi) & \text{iff } (\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \phi \text{ or } (\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \psi \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\exists x.\, \phi) & \text{iff } (\bar{\mathcal{D}}, \bar{\tau}, v[x/d], i) \models \phi, \text{for some } d \in |\bar{\mathcal{D}}| \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\bullet_I \, \phi) & \text{iff } i > 0,\ \tau_i - \tau_{i-1} \in I, \text{ and } (\bar{\mathcal{D}}, \bar{\tau}, v, i-1) \models \phi \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\bigcirc_I \, \phi) & \text{iff } \tau_{i+1} - \tau_i \in I \text{ and } (\bar{\mathcal{D}}, \bar{\tau}, v, i+1) \models \phi \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\phi \, \mathsf{S}_I \, \psi) & \text{iff for some } j \le i,\ \tau_i - \tau_j \in I, (\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \psi, \\
 & \quad \text{and } (\bar{\mathcal{D}}, \bar{\tau}, v, k) \models \phi, \text{for all } k \in [j+1, i+1) \\
(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models (\phi \, \mathsf{U}_I \, \psi) & \text{iff for some } j \ge i,\ \tau_j - \tau_i \in I,\ (\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \psi, \\
 & \quad \text{and } (\bar{\mathcal{D}}, \bar{\tau}, v, k) \models \phi, \text{ for all } k \in [i, j)
\end{array}
$$

Figure 2. Semantics of MFOTL

Formulas over the signature S are given by the grammar

$$\phi ::= t_1 \approx t_2 \mid t_1 < t_2 \mid r(t_1, \ldots, t_{\iota(r)}) \mid (\neg \phi) \mid (\phi \vee \phi) \mid (\exists x.\, \phi) \mid (\bullet_I \, \phi) \mid (\bigcirc_I \, \phi) \mid (\phi \, \mathsf{S}_I \, \phi) \mid (\phi \, \mathsf{U}_I \, \phi),$$

where $t_1, t_2, \ldots$ range over the elements in $V \cup C$, and r, x, and I range over the elements in R, V, and $\mathbb{I}$, respectively.

To define MFOTL's semantics, we need the following notions. A *structure* $\mathcal{D}$ over S consists of a domain $|\mathcal{D}| \neq \emptyset$ and interpretations $c^{\mathcal{D}} \in |\mathcal{D}|$ and $r^{\mathcal{D}} \subseteq |\mathcal{D}|^{\iota(r)}$, for each $c \in C$ and $r \in R$. A *temporal structure* over S is a pair $(\bar{\mathcal{D}}, \bar{\tau})$, where $\bar{\mathcal{D}} = (\mathcal{D}_0, \mathcal{D}_1, \ldots)$ is a sequence of structures over S and $\bar{\tau} = (\tau_0, \tau_1, \ldots)$ is a sequence of natural numbers (i.e., timestamps), where:

(1) The sequence $\bar{\tau}$ is monotonically increasing (i.e., $\tau_i \le \tau_{i+1}$, for all $i \ge 0$) and makes progress (i.e., for every $i \ge 0$, there is some $j > i$ such that $\tau_j > \tau_i$).
(2) $\bar{\mathcal{D}}$ has constant domains, i.e., $|\mathcal{D}_i| = |\mathcal{D}_{i+1}|$, for all $i \ge 0$. We denote the domain by $|\bar{\mathcal{D}}|$ and require that $|\bar{\mathcal{D}}|$ is strict linearly ordered by a relation $<$.
(3) Each constant symbol $c \in C$ has a rigid interpretation, i.e., $c^{\mathcal{D}_i} = c^{\mathcal{D}_{i+1}}$, for all $i \ge 0$. We denote c's interpretation by $c^{\bar{\mathcal{D}}}$.

A *valuation* is a mapping $v : V \to |\bar{\mathcal{D}}|$. We abuse notation by applying a valuation v also to constant symbols $c \in C$, with $v(c) = c^{\bar{\mathcal{D}}}$. For a valuation v, a variable x, and $d \in |\bar{\mathcal{D}}|$, $v[x/d]$ is the valuation mapping x to d and not altering the other variables' valuation.

The semantics of MFOTL, $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \phi$, is given in Figure 2, where $(\bar{\mathcal{D}}, \bar{\tau})$ is a temporal structure over the signature S, with $\bar{\mathcal{D}} = (\mathcal{D}_0, \mathcal{D}_1, \ldots)$, $\bar{\tau} = (\tau_0, \tau_1, \ldots)$, v a valuation, $i \in \mathbb{N}$, and ϕ a formula over S. Note that the temporal operators are labeled with intervals I and a formula of the form $(\bullet_I \, \phi)$, $(\bigcirc_I \, \phi)$, $(\phi \mathsf{S}_I \psi)$, or $(\phi \mathsf{U}_I \psi)$ is only satisfied in $(\bar{\mathcal{D}}, \bar{\tau})$ at the time point i, if it is satisfied within the bounds given by the interval I of the respective temporal operator, which are relative to the current timestamp τ_i.

Terminology and Notation. We use standard syntactic sugar such as $\blacksquare_I \, \phi := \neg(\mathit{true} \mathsf{S}_I \neg \phi)$ and $\square_I \, \phi := \neg(\mathit{true} \mathsf{U}_I \neg \phi)$, where $\mathit{true} := \exists x.\, x \approx x$. We also use non-metric operators like $\square \, \phi := \square_{[0,\infty)} \, \phi$. We omit parentheses where possible,

e.g., unary operators (temporal and Boolean) bind stronger than binary ones. A formula ϕ is *bounded* if the interval I of every temporal operator U_I occurring in ϕ is finite. We use standard terminology like *atomic formula* and *subformula*.

System Monitoring. We illustrate our use of MFOTL and our monitoring algorithm [1] for compliance checking by the simple policy stating that reports must be approved within at most 10 time units before they are published:

$$\Box \forall x.\, publish(x) \rightarrow \blacklozenge_{[0,11)}\, approve(x)\,.$$

We assume that the actions for publishing and approving reports are logged in relations. Specifically, for each time point $i \in \mathbb{N}$, we have the unary relations $PUBLISH_i$ and $APPROVE_i$ such that (1) $x \in PUBLISH_i$ iff report f is published at time point i and (2) $x \in APPROVE_i$ iff report x is approved at time point i. Observe that there can be multiple approvals at the same time point for different reports. Furthermore, every time point i has a timestamp $\tau_i \in \mathbb{N}$.

The corresponding temporal structure $(\bar{\mathcal{D}}, \bar{\tau})$ with $\bar{\mathcal{D}} = (\mathcal{D}_0, \mathcal{D}_1, \dots)$ and $\bar{\tau} = (\tau_0, \tau_1, \dots)$ of a sequence of logged publishing and approval actions is as follows. The only relational symbols in $\bar{\mathcal{D}}$'s signature are *publish* and *approve*, both of arity 1. The domain of $\bar{\mathcal{D}}$ consists of all reports. The ith structure in $\bar{\mathcal{D}}$ is timestamped with τ_i and contains the relations $PUBLISH_i$ and $APPROVE_i$.

To detect policy violations, our monitor [1] iteratively processes the temporal structure $(\bar{\mathcal{D}}, \bar{\tau})$ representing the stream of logged actions. This can be done offline or online. At each time point i, it outputs the valuations satisfying the negation of the formula $publish(x) \rightarrow \blacklozenge_{[0,11)}\, approve(x)$. Note that we drop the outermost quantifier since we are not only interested in whether the policy is violated but also which data is responsible for the reported violations.

In general, we assume that policies formalized in MFOTL are of the form $\Box\psi$, where ψ is bounded. Since ψ is bounded, the monitor need only take into account a finite prefix of $(\bar{\mathcal{D}}, \bar{\tau})$ when determining the satisfying valuations of $\neg\psi$ at a time point i. To effectively determine all these valuations, we also assume here that predicates have finite interpretations in $(\bar{\mathcal{D}}, \bar{\tau})$, i.e., the relation $r^{\mathcal{D}_j}$ is finite, for every predicate r and every $j \in \mathbb{N}$. Furthermore, we require that $\neg\psi$ can be rewritten to a temporal-subformula-domain-independent formula, a generalization of the standard notion of domain-independent database queries [5].

III. Monitoring Concurrently Logged Actions

In this section, we first prove the intractability of monitoring where logs are produced in a concurrent setting. We then show how to partially overcome this obstacle by monitoring a single log where all actions with equal timestamps are assumed to have happened at the same point in time. Proof details are given in the Appendices A–D.

Log Interleavings. Intuitively, an interleaving of logs preserves the ordering of the logged actions with respect to their timestamps, but allows for all possible orderings of actions with equal timestamps that are recorded by different log producers. To define this, let $\mathrm{img}(f)$ denote the set $\{y \in Y \mid f(x) = y, \text{ for some } x \in X\}$, for a function $f : X \rightarrow Y$. Furthermore, we assume in this section that all temporal structures have the same signature (C, R, ι), equal domains, and that constant symbols are equally interpreted. Note that any two temporal structures in which the common constant symbols are equally interpreted can easily be extended so that their extensions fulfill this requirement.

Definition 1. *Let* $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$, $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$, *and* $(\bar{\mathcal{D}}, \bar{\tau})$ *be temporal structures.* $(\bar{\mathcal{D}}, \bar{\tau})$ *is an* interleaving *of* $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ *and* $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *if there are strictly monotonic functions* $f_1, f_2 : \mathbb{N} \rightarrow \mathbb{N}$ *with*
(1) $\mathrm{img}(f_1) \cup \mathrm{img}(f_2) = \mathbb{N}$,
(2) $\mathrm{img}(f_1) \cap \mathrm{img}(f_2) = \emptyset$, *and*
(3) $\tau_i^k = \tau_{f_k(i)}$ *and* $r^{\mathcal{D}_i^k} = r^{\mathcal{D}_{f_k(i)}}$, *for all* $k \in \{1, 2\}$, $i \in \mathbb{N}$, $r \in R$.
We denote by $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \,⧓\, (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *the set of all interleavings of the temporal structures* $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ *and* $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$.

Since there are usually multiple interleavings of two temporal structures, we formulate policy violations in terms of a set of temporal structures.

Definition 2. *Let* $\mathbf{T}$ *be a set of temporal structures.*
(1) $\mathbf{T}$ weakly violates *the formula* ϕ *at time point* $i \in \mathbb{N}$ *if for some* $(\bar{\mathcal{D}}, \bar{\tau}) \in \mathbf{T}$ *and some valuation* v, *it holds that* $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \not\models \phi$.
(2) $\mathbf{T}$ strongly violates *the formula* ϕ *at time point* $i \in \mathbb{N}$ *if for all* $(\bar{\mathcal{D}}, \bar{\tau}) \in \mathbf{T}$, *there is some valuation* v *such that* $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \not\models \phi$.

Unfortunately, even in a propositional setting, determining whether the set of interleavings weakly or strongly violates a formula is intractable.

Theorem 3. *Let* $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ *and* $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *be temporal structures,* $i \in \mathbb{N}$, *and* ϕ *a quantifier-free sentence with only Boolean and non-metric past operators that neither contains the equality symbol* $\approx$ *nor the ordering symbol* $\prec$.
1. *Determining whether the set of interleavings* $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \,⧓\, (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *weakly violates* ϕ *at* i *is NP-complete.*
2. *Determining whether the set of interleavings* $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \,⧓\, (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *strongly violates* ϕ *at* i *is coNP-complete.*

Note that both decision problems are well defined as ϕ does not contain future operators. We therefore only need to examine the finite prefixes with length $i + 1$ of the interleavings to determine whether ϕ is weakly or strongly violated at the given time point i.

Collapsing Interleaved Logs. We first give conditions with respect to an arbitrary set of temporal structures for when it suffices to monitor a single temporal structure. We then identify a natural temporal structure for the set of interleavings of two temporal structures, which we use for

monitoring.

Definition 4. *The temporal structure* $(\bar{\mathcal{C}}, \bar{\kappa})$ *is* sufficient *for the formula* ϕ *on the set* $\mathbf{T}$ *of temporal structures if for all valuations* v*, the following conditions are fulfilled:*
(C1) If $(\bar{\mathcal{C}}, \bar{\kappa}, v, 0) \models \phi$ *then* $(\bar{\mathcal{D}}, \bar{\tau}, v, 0) \models \phi$*, for all* $(\bar{\mathcal{D}}, \bar{\tau}) \in \mathbf{T}$.
(C2) If $(\bar{\mathcal{C}}, \bar{\kappa}, v, 0) \not\models \phi$ *then* $(\bar{\mathcal{D}}, \bar{\tau}, v, 0) \not\models \phi$*, for all* $(\bar{\mathcal{D}}, \bar{\tau}) \in \mathbf{T}$.

In the following, the set $\mathbf{T}$ in the above definition will be the set of interleavings of two temporal structures. For the temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$, we will use the so-called collapse:

Definition 5. *Let* $(\bar{\mathcal{D}}, \bar{\tau})$ *and* $(\bar{\mathcal{C}}, \bar{\kappa})$ *be temporal structures.* $(\bar{\mathcal{C}}, \bar{\kappa})$ *is a* collapse *of* $(\bar{\mathcal{D}}, \bar{\tau})$ *if there is a monotonic surjective function* $f : \mathbb{N} \to \mathbb{N}$ *such that*
(1) if $\tau_i = \tau_j$ *then* $f(i) = f(j)$*, for all* $i, j \in \mathbb{N}$*,*
(2) $\kappa_{f(i)} = \tau_i$*, for all* $i \in \mathbb{N}$*, and*
(3) $r^{\mathcal{C}_j} = \bigcup_{i \in f^{-1}(j)} r^{\mathcal{D}_i}$*, for all* $j \in \mathbb{N}$ *and* $r \in R$.

Intuitively, the structures of the temporal structure $(\bar{\mathcal{D}}, \bar{\tau})$ with equal timestamps are collapsed into a single structure. The collapse is uniquely defined and we denote it by $col(\bar{\mathcal{D}}, \bar{\tau})$. Furthermore, the collapses of temporal structures in the set of interleavings of two given temporal structures are all isomorphic.

Before we identify formulas for which the collapse of an interleaving of given temporal structures can be correctly used for monitoring, we give practical reasons that justify its use for monitoring. First, observe that the collapse can be incrementally obtained from an arbitrary interleaving of two given temporal structures. Hence, monitoring the collapse can be done efficiently. Second, note that the actual ordering of actions logged with equal timestamps in a concurrent system cannot be known. Hence, it does not make sense to consider just one arbitrary interleaving. Assuming that equally timestamped actions have happened at the same point in time naturally "hides" the differences between interleavings. Moreover, reasonable policies for a concurrent system should not care about the ordering of equally timestamped actions in case of accurate and precise clocks. In other words, if the collapsed temporal structure is not sufficient for the policy on the set of interleavings, then the policy might not be the intended one for the system. Finally, monitoring the collapsed temporal structure is practically more efficient than monitoring an interleaving. This is because the monitor is invoked less often since time points with equal timestamps are merged to a single one. Hence, the monitor processes the logged actions with equal timestamp in a single invocation.

Monitoring the Collapse. Intuitively, collapse-sufficient formulas are formulas that do not yield false positives and false negatives when monitoring the collapse of an interleaving:

Definition 6. *Let* ϕ *be a formula. For* $k \in \{1, 2\}$*, we say that* ϕ *has the property (Ck) if* $(\bar{\mathcal{C}}, \bar{\kappa})$ *fulfills the condition (Ck) in Definition 4 with respect to* ϕ *and* $(\bar{\mathcal{D}}, \bar{\tau}) \between (\bar{\mathcal{D}}', \bar{\tau}')$*, for every* $(\bar{\mathcal{D}}, \bar{\tau})$*,* $(\bar{\mathcal{D}}', \bar{\tau}')$*, and* $(\bar{\mathcal{C}}, \bar{\kappa})$*, where* $(\bar{\mathcal{C}}, \bar{\kappa})$ *is the collapse of an interleaving of* $(\bar{\mathcal{D}}, \bar{\tau})$ *and* $(\bar{\mathcal{D}}', \bar{\tau}')$*. Moreover,* ϕ *is* collapse-sufficient *if it has the properties (C1) and (C2).*

Monitoring the collapse of a collapse-sufficient formula is correct with respect to strong violations. Since the formula has property (C2), violations found in $(\bar{\mathcal{C}}, \bar{\kappa})$ imply that the set of interleavings strongly violates the formula. The converse is ensured by the property (C1): if no violation is found in $(\bar{\mathcal{C}}, \bar{\kappa})$ then all interleavings are policy compliant. Furthermore, by monitoring $(\bar{\mathcal{C}}, \bar{\kappa})$ we also detect when the set of interleavings weakly violates the given formula. The reason is that if a formula is strongly violated by the set of interleavings then it also weakly violated, since the set of interleavings is always nonempty.

Example 7. *The formula* $\Box \forall x.publish(x) \rightarrow \blacklozenge_{[0,11)} approve(x)$ *is not collapse-sufficient. Suppose that a report* x *is published in* $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ *at time point* i*, i.e.,* $x \in publish^{\mathcal{D}^1_i}$ *and only approved in* $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *at the equally timestamped time point* j*, i.e.,* $x \in approve^{\mathcal{D}^2_j}$ *with* $\tau^2_j = \tau^1_i$*. Then there is an interleaving* $(\bar{\mathcal{D}}, \bar{\tau}) \in (\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ *where the approval action comes (pointwise) strictly after the publish action. As a result, we cannot handle this formula correctly by monitoring the collapsed temporal structure* $(\bar{\mathcal{C}}, \bar{\kappa})$ *of an interleaving of the given temporal structures* $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ *and* $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$.

A slightly stronger policy can be efficiently monitored. Namely, the policy that requires that an approval action must happen timewise strictly before the publish action, i.e., $\Box \forall x.\, publish(x) \rightarrow \blacklozenge_{[1,11)} approve(x)$*. This formula is collapse-sufficient. Similarly,* $\Box \forall x.\, publish(x) \rightarrow \lozenge_{[0,1)} \blacklozenge_{[0,11)} approve(x)$ *is also collapse-sufficient. It formalizes the slightly weaker policy where publish actions must be timewise but not pointwise previously approved.*

Note that stutter-invariance [6] is a necessary condition for collapse-sufficiency. However, it is not a sufficient condition. For example, the formula $\Box \forall x.\, p(x) \wedge q(x)$ is stuttering-invariant but not collapse-sufficient.

A Collapse-sufficient Fragment. In the following, we present a fragment of collapse-sufficient formulas. Our fragment is defined in terms of an algorithm that identifies formulas that have property (C1) or property (C2).

The algorithm labels the atomic subformulas of the given formula and propagates these labels bottom-up to the formula's root using a fixed set of inference rules. The labels represent invariants, which capture the relation between violations found in a collapsed temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$ at some time point and violations found in its pre-images $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ at a time point with an equal timestamp, where $col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ denotes the set of temporal structures $(\bar{\mathcal{D}}', \bar{\tau}')$ with $col(\bar{\mathcal{D}}', \bar{\tau}') = (\bar{\mathcal{C}}, \bar{\kappa})$. Note that $(\bar{\mathcal{D}}, \bar{\tau}) \between (\bar{\mathcal{D}}', \bar{\tau}') \subsetneq col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$, where $(\bar{\mathcal{C}}, \bar{\kappa})$ is the collapse of an interleaving of

$$\frac{}{t \approx t' : (\models \forall)} \qquad \frac{}{t \approx t' : (\not\models \forall)}$$

$$\frac{}{r(t_1, \ldots, t_{\iota(r)}) : (\models \exists)} \qquad \frac{}{r(t_1, \ldots, t_{\iota(r)}) : (\not\models \forall)}$$

$$\frac{\psi : (\models \forall)}{\psi : (\models \exists)} \qquad \frac{\psi : (\not\models \forall)}{\psi : (\not\models \exists)}$$

$$\frac{\psi : (\models \forall)}{\blacklozenge_I \psi : (\models \forall)} \qquad \frac{\psi : (\models \exists)}{\blacklozenge_I \psi : (\models \exists)} \qquad \frac{\psi : (\models \exists)}{\blacklozenge_I \psi : (\models \forall)}\ 0 \notin I \qquad \frac{\psi : (\not\models \forall)}{\blacklozenge_I \psi : (\not\models \forall)}$$

$$\frac{\psi : (\models \exists)}{\blacklozenge_I \Diamond_J \psi : (\models \forall)}\ 0 \in I \cap J$$

Figure 3. Selection of Inference Rules

the temporal structures $(\bar{\mathcal{D}}, \bar{\tau})$ and $(\bar{\mathcal{D}}', \bar{\tau}')$.

The labels and their corresponding invariants are as follows for a formula ϕ:

($\models \forall$): For all valuations v and all $i \in \mathbb{N}$, if $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$ then for every $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ and every $j \in \mathbb{N}$ with $\kappa_i = \tau_j$, it holds that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$.

($\models \exists$): For all valuations v and all $i \in \mathbb{N}$, if $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$ then for every $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$, there is some $j \in \mathbb{N}$ with $\kappa_i = \tau_j$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$.

($\not\models \forall$): For all valuations v and all $i \in \mathbb{N}$, if $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \phi$ then for every $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ and every $j \in \mathbb{N}$ with $\kappa_i = \tau_j$, it holds that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \not\models \phi$.

($\not\models \exists$): For all valuations v and all $i \in \mathbb{N}$, if $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \phi$ then for every $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$, there is some $j \in \mathbb{N}$ with $\kappa_i = \tau_j$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \not\models \phi$.

The first symbol ($\models$ or $\not\models$) in a label states whether the formula is satisfied in the collapsed temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$. The second symbol ($\forall$ or $\exists$) states whether the formula is satisfied at some equally timestamped time point or at all equally timestamped time points in all temporal structures $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$.

Due to space limitations, Figure 3 shows only some of our inference rules. All rules can be found in Appendix C, where we also prove their soundness.

First, consider the rules in Figure 3 for atomic formulas. An atomic formula $t \approx t'$ depends only on the valuation and therefore can be labeled ($\models \forall$) and ($\not\models \forall$). An atomic formula of the form $r(t_1, \ldots, t_{\iota(r)})$ can be labeled ($\models \exists$) and ($\not\models \forall$). We only explain the labeling ($\models \exists$). The explanation for the label ($\not\models \forall$) is analogous. The interpretation of a predicate in a collapsed temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$ at a time point i is the union of the predicate's interpretations at all time points j in a temporal structure $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ for which τ_j equals κ_i. Therefore, if $\bar{a} \in r^{\mathcal{C}_i}$ then $\bar{a} \in r^{\mathcal{D}_j}$, for some $j \in \mathbb{N}$ with $\tau_j = \kappa_i$. Note that $\bar{a} \in r^{\mathcal{D}_j}$ does not necessarily hold for all these js; hence, we cannot label $r(t_1, \ldots, t_{\iota(r)})$ with ($\not\models \forall$).

The next two rules in Figure 3 express that the invariants corresponding to the labels ($\models \forall$) and ($\not\models \forall$) imply the invariants corresponding to ($\models \exists$) and ($\not\models \exists$), respectively.

Next, we consider the inference rules for the temporal operator $\blacklozenge_I$. We first justify the inference rule that allows us to propagate the label ($\models \forall$) from ψ to $\blacklozenge_I \psi$. If $\blacklozenge_I \psi$ is satisfied in the collapsed temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$ at time point i then ψ is satisfied at some previous time point $j \leq i$ in $(\bar{\mathcal{C}}, \bar{\kappa})$ with $\kappa_i - \kappa_j \in I$. Because ψ is labeled with ($\models \forall$), all time points with timestamp κ_j in the temporal structure $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ also satisfy ψ, and hence, all time points with timestamp κ_i satisfy $\blacklozenge_I \psi$ in $(\bar{\mathcal{D}}, \bar{\tau})$. When ψ is labeled with ($\models \exists$), possibly only a single time point k in $(\bar{\mathcal{D}}, \bar{\tau})$ with $\tau_k = \kappa_j$ satisfies ψ. If $0 \in I$ then $\blacklozenge_I \psi$ might not be satisfied at time points before k, even if these time points have the timestamp κ_i. So, we can label $\blacklozenge_I$ with ($\models \exists$) but not with ($\models \forall$). However, if $0 \notin I$ then ψ is satisfied in $(\bar{\mathcal{C}}, \bar{\kappa})$ at a time point j with the timestamp $\kappa_j < \kappa_i$. Hence $\blacklozenge_I \psi$ is satisfied in $(\bar{\mathcal{D}}, \bar{\tau})$ at all time points with the timestamp κ_i. This allows us to label $\blacklozenge_I \psi$ with ($\models \forall$). Finally, consider the rule where ψ is labeled ($\not\models \forall$). If $\blacklozenge_I \psi$ is violated in the collapsed temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$ at timestamp κ_i then ψ is violated at all previous points in the temporal structure $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ that satisfy the metric constraints given by I. But then $\blacklozenge_I \psi$ is also violated in $(\bar{\mathcal{D}}, \bar{\tau})$ at all time points with the timestamp κ_i. Hence we can label $\blacklozenge_I \psi$ with ($\not\models \forall$).

We can try to label a formula solely based on inference rules that involve only a single Boolean or temporal operator. However, with more specialized inference rules like the one for $\blacklozenge_I \Diamond_J \psi$ given in Figure 3, we are more likely to succeed in propagating labels to the root of the formula. Intuitively, with the nesting of the operators $\blacklozenge_I$ and $\Diamond_J$, and when $0 \in I \cap J$, the ordering of equally timestamped time points becomes irrelevant since from a given time point, we can freely choose any of these time points that satisfy the metric constraints given by the intervals I and J. Hence, a labeling ($\models \exists$) for ψ allows us to label $\blacklozenge_I \Diamond_J \psi$ with ($\models \forall$).

Finally, we remark that there are no inference rules for the temporal operators $\bullet_I$ and $\circ_I$ because these operators inherently rely on the relative ordering of the structures in a temporal structure.

Based on the labels at the root of the formula, we can determine if the formula has the property (C1) or the property (C2). The conclusions we can draw are stated in the following lemma, which follows from the soundness of the inference rules.

Lemma 8. *Let ϕ be a formula.*
1. *If ϕ can be labeled by ($\models \forall$) then ϕ has property (C1).*
2. *If ϕ can be labeled by ($\not\models \forall$) then ϕ has property (C2).*
3. *If ϕ can be labeled by ($\models \exists$) then $\Diamond \phi$ has property (C1).*
4. *If ϕ can be labeled by ($\not\models \exists$) then $\Box \phi$ has property (C2).*

Based on this lemma, we obtain the following theorem.

Theorem 9. *If the formula ϕ can be labeled by ($\models \forall$) and ($\not\models \forall$) then it is collapse-sufficient. Moreover, we can determine in linear time in the formula's length whether ϕ can be labeled by ($\models \forall$), ($\models \exists$), ($\not\models \forall$), and ($\not\models \exists$).*

Note that formulas of the form $\Box \psi$ are already collapse-

sufficient if ψ can be labeled by $(\not\models \exists)$ and $\Box\psi$ can be labeled by $(\models \forall)$. Even if only one of these labellings can be derived, monitoring $\Box\psi$ on the collapsed temporal structure of an interleaving is still useful. For example, if ψ is labeled by $(\not\models \exists)$ then violations that are found on the collapsed temporal structure relate to strong violations on the set of interleavings. However, we might miss some violations.

Example 10. *We illustrate our algorithm and its inference rules by applying it to the formula* $\Box\forall x.\, publish(x) \rightarrow \blacklozenge_{[0,11)}\, approve(x)$. *We first remove some syntactic sugar and obtain the formula* $\Box\forall x.\, \neg publish(x) \vee \blacklozenge_{[0,11)}\, approve(x)$. *We start by labeling the atomic subformulas. Both* $publish(x)$ *and* $approve(x)$ *are labeled with* $(\models \exists)$ *and* $(\not\models \forall)$. *According to the inference rules for the temporal operator* $\blacklozenge_I$ *we label* $\blacklozenge_{[0,11)}\, approve(x)$ *with* $(\models \exists)$ *and* $(\not\models \forall)$. *We cannot label it with* $(\models \forall)$ *since the interval contains* 0. *Moreover, the subformula* $\neg publish(x)$ *is labeled with* $(\not\models \exists)$ *and* $(\models \forall)$. *The subformulas* $\neg publish(x) \vee \blacklozenge_{[0,11)}\, approve(x)$ *and* $\forall x.\, \neg publish(x) \vee \blacklozenge_{[0,11)}\, approve(x)$ *are labeled* $(\models \exists)$ *and* $(\not\models \exists)$. *We conclude that the formula* $\Box\forall x.\, \neg publish(x) \vee \blacklozenge_{[0,11)}\, approve(x)$ *has the property (C2). It does not have the property (C1), as shown in Example 7.*

The formula $\Box\forall x.\, publish(x) \rightarrow \blacklozenge_{[1,11)}\, approve(x)$ *has both properties (C1) and (C2). The labeling starts similarly but* $\blacklozenge_{[1,11)}\, approve(x)$ *is additionally labeled with* $(\models \forall)$ *since the interval of the temporal operator does not contain* 0. *This label propagates to the root of the formula. We conclude that* $\Box\forall x.\, \neg publish(x) \vee \blacklozenge_{[1,11)}\, approve(x)$ *also has property (C1).*

Policy Approximation. In Example 7, we have seen that we can obtain collapse-sufficient policies by strengthening or weakening the original policy. In the following, we present a systematic approach along these lines by over-approximating and under-approximating policies.

Let ϕ be a formula in positive normal form. We obtain the *weakened* formula ϕ^w by replacing each atomic subformula $r(t_1, \ldots, t_{\iota(r)})$ that occurs positively in ϕ by $\blacklozenge_I \lozenge_{I'}\, r(t_1, \ldots, t_{\iota(r)})$, for some intervals I and I' with $0 \in I \cap I'$. Analogously, in the *strengthened* formula ϕ^s, we replace each negative occurrence of an atomic subformula $r(t_1, \ldots, t_{\iota(r)})$ by $\blacklozenge_I \lozenge_{I'}\, r(t_1, \ldots, t_{\iota(r)})$.

Theorem 11. *Let* ϕ^w *and* ϕ^s *be weakened and strengthened formulas of the formula* ϕ *in positive normal form. The formulas* $\phi \rightarrow \phi^w$ *and* $\phi^s \rightarrow \phi$ *are valid. Moreover,*
1. *if* ϕ^s *is collapse-sufficient then* ϕ *has property (C1), and*
2. *if* ϕ^w *is collapse-sufficient then* ϕ *has property (C2).*

Weakened and strengthened formulas are more likely to be collapse-sufficient, since their subformulas of the form $\blacklozenge_I \lozenge_{I'}\, r(t_1, \ldots, t_{\iota(r)})$ can be labeled with $(\models \forall)$, while $r(t_1, \ldots, t_{\iota(r)})$ can only be labeled with the weaker label $(\models \exists)$. Simultaneously weakening and strengthening always results in a collapse-sufficient formula. However, the resulting formula does not necessarily relate to the original formula.

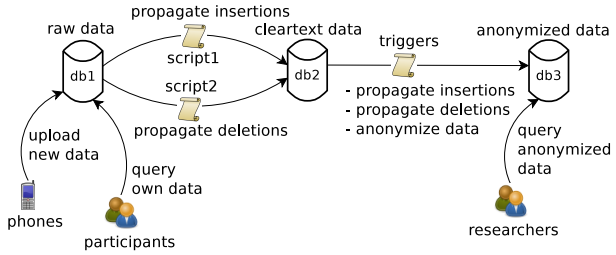

Figure 4. Nokia's Data-collection Campaign

Finally, note that by inserting the temporal operators $\blacklozenge_{[0,1)}$ and $\lozenge_{[0,1)}$ around positively occurring atomic subformulas, the ordering of equally timestamped actions becomes irrelevant. This is desirable in systems where the clocks used to timestamp the actions are synchronized but too coarse-grained. Taking this idea further, by putting temporal operators $\blacklozenge_{[0,b)}$ and $\lozenge_{[0,b)}$ around these subformulas with $b \geq 1$, we take into account that the timestamps in a temporal structure are inaccurate and might differ from their actual value by the threshold b—a situation that occurs in practice.

IV. Practical Experience

In this section, we describe the implementation of our monitoring approach within Nokia's Data-collection Campaign [7], which is a real-world application with realistic usage-control policies. Furthermore, we report on the monitor's performance and our findings.

Scenario. The campaign,[1] which was launched in 2009, collects contextual information from cell phones of about 180 participants. This sensitive data includes phone locations, call and SMS information, and the like. The data collected by a participant's phone is propagated into the databases db1, db2, and db3. The phones use WLAN to periodically upload their data to database db1. Every night, the synchronization script script1 copies the data from db1 to db2. Furthermore, triggers running on db2 anonymize and copy the data to db3, where researchers can access and analyze the anonymized data. The participants can access and delete their own data using a web interface to db1. Deletions are propagated to all databases: from db1 to db2 by the synchronization script script2, which also runs every night, and from db2 to db3 by database triggers. Figure 4 summarizes the various usages of data in the campaign.

Within the campaign, data is organized by records and can easily be identified. When uploading data from a phone into db1, a unique identifier is generated for each record. This identifier together with an identifier of the participant who contributed the data is attached to the record.

Policies. The collected data is subject to various policies in order to protect the participants' privacy. For example, there are access control rules and policies governing

[1] See http://research.nokia.com/page/11367 for details.

Table I. Policy Formalizations in MFOTL

policy	MFOTL formalization
delete	$\Box\, \forall user.\, \forall data.\, delete(user, \mathsf{db2}, data) \rightarrow user \approx \mathsf{script2}$
ins-1-2	$\Box\, \forall user.\, \forall data.\, insert(user, \mathsf{db1}, data) \wedge data \not\approx \mathsf{unknown} \rightarrow \blacklozenge_{[0,1s)}\, \Diamond_{[0,30h]}\, \exists user'.\, insert(user', \mathsf{db2}, data) \vee delete(user', \mathsf{db1}, data)$
ins-2-3	$\Box\, \forall user.\, \forall data.\, insert(user, \mathsf{db2}, data) \wedge data \not\approx \mathsf{unknown} \rightarrow \blacklozenge_{[0,1s)}\, \Diamond_{[0,60s)}\, \exists user'.\, insert(user', \mathsf{db3}, data)$
del-1-2	$\Box\, \forall user.\, \forall data.\, delete(user, \mathsf{db1}, data) \wedge data \not\approx \mathsf{unknown} \rightarrow (\blacklozenge_{[0,1s)}\, \Diamond_{[0,30h)}\, \exists user'.\, delete(user', \mathsf{db2}, data)) \vee ((\Diamond_{[0,1s)}\, \blacklozenge_{[0,30h)}\, \exists user'.\, insert(user', \mathsf{db1}, data)) \wedge (\blacksquare_{[0,30h)}\, \Box_{[0,30h)}\, \neg \exists user'.\, insert(user', \mathsf{db2}, data)))$

Table II. Log Statistics

log	# time points	total # actions	# *insert* actions			# other actions
			db1	db2	db3	
1	29,672	1,462,700	82,486	678,840	678,840	22,534
2	10,870	969,520	23,828	472,369	472,369	954
3	6,601	1,019,428	33,229	492,411	492,411	1377
4	20,330	962,766	12,918	468,844	468,844	11,298
5	8,114	687,402	7,067	339,674	339,647	12,160
6	9,218	630,287	4,207	311,882	311,835	1,366
7	7,327	554,733	3,251	275,208	275,199	1,014
8	86,892	936,249	47,786	400,490	400,475	87,498
9	86,764	986,249	30,118	434,268	434,259	87,604

the process of propagating the data between databases. In particular, any insertion or deletion of data in db1 must be propagated to db2 within 30 hours, and from db2 to db3 within 1 minute. Furthermore, only the latest version of the synchronization scripts may be used and the scripts may not run longer than 6 hours. Finally, access to the databases is restricted to selected user accounts and the account script1 may be used only while the script script1 is running.

We present here just a few representative policies in Table I. Details about all the 14 policies are given in Appendix E. The predicates *insert* and *delete* correspond to the equally-named database commands. The arguments of these predicates are the agent that initiated the action, the name of the database where the action was carried out, and an identifier of the involved data.

Note that all policy formalizations in Table I are collapse-sufficient. However, some policies have slightly weaker or stronger variants that are not collapse-sufficient. For example, we obtained *ins-2-3* from the policy "all data inserted into db2 must also be inserted into db3 within 60 seconds" by weakening the formula $\Box\, \forall users.\, \forall data.\, insert(user, \mathsf{db2}, data) \wedge data \not\approx \mathsf{unknown} \rightarrow \Diamond_{[0,60s)}\, \exists user'.\, insert(user', \mathsf{db3}, data)$. Intuitively, *ins-2-3* is the policy formalization that we actually intended: we do not want to distinguish the relative ordering of the insertions into db2 and db3 when they are logged with the same timestamp. This is because the 1 second timestamp granularity that is used may not be fine-granular enough: the database triggers may be activated within milliseconds.

Logging Mechanisms. We extended the data-collection setup with mechanisms to log policy-relevant actions. We installed logging mechanisms for the three databases, the script script1, and the SVN repository, assuming synchronized clocks for timestamping. We now discuss details of these logging mechanisms.

As logs for the database db1 were not available, we implemented a proxy to inspect interactions of participants and phones with db1. The proxy logs what data is inserted and deleted. To observe the insertion of new data, we monitor the network traffic when the phone uploads data. For deletions, we use a custom front-end that logs the requests for deleting data. For practical reasons, we could deploy these mechanisms only for 2 out of the 180 participants. Hence, we have only partial logging for db1, which only affects 2 out of the 14 policies.

The databases db2 and db3 reside physically on a single PostgreSQL server, which logs the SQL queries. We extract relevant actions from these PostgreSQL logs. The main challenge is to determine what data is processed in a query since only the query itself is logged. Fortunately, most relevant queries are made by automated scripts or database triggers and contain enough information to determine what data is used. For example, an insert or delete query initiated by a synchronization script includes the identifier of the used data record. Hence, a simple syntactic analysis of these queries suffices to log the relevant actions in sufficient detail. When the analysis failed to extract the data, we identified the data with the constant unknown.

Evaluation. To evaluate the performance of our monitor on different data sets, we split the logs into smaller files, where each file corresponds to roughly 24 hours of log entries. Table II provides details about the collapsed temporal structures corresponding to these logs. Observe that the number of *insert* actions is significantly larger than the number of other actions. None of the log files used contains more than 100 *delete* actions. Table III shows the monitor's running times and memory usage for each policy and log file. For the experiments, we used a desktop computer with a 1150 MHz AMD Phenom 9600B Quad-Core CPU.

Monitoring invariants like the policy *delete* is fast: the monitor needed no more than 10 seconds for a 24-hours log file. More complex policies involving temporal operators with large time windows, take more time to monitor. For example, for the policy *ins-1-2*, the monitor took more than 4 hours in some cases. The policy *del-1-2* with an even larger time window, however, could be quickly monitored. The reason here is that the log files used contain only few *delete* actions. Although we monitored the logs offline, the running times indicate that an online monitoring approach is possible, since the running times are less than the time period covered by the logs. The memory requirements are also modest. For the policies *delete* and *ins-2-3*, the monitor does not require more than 10 MB of RAM. For *ins-1-2* and *del-1-2*, the monitor used under 200 MB of RAM, which is also acceptable due to the large time windows.

Findings. The monitor reported the following policy violations. First, some static access control policies like *delete* were violated. These violations were due to testing,

Table III. Monitor Performance — Running Times / Memory Usage

policy	log 1	log 2	log 3	log 4	log 5	log 6	log 7	log 8	log 9
delete	10 s / 4 MB	7 s / 4 MB	7 s / 4 MB	6 s / 4 MB	5 s / 4 MB	4 s / 4 MB	4 s / 4 MB	6 s / 4 MB	6 s / 4 MB
ins-1-2	231 m / 161 MB	44 m / 103 MB	67 m / 107 MB	24 m / 102 MB	9 m / 71 MB	5 m / 65 MB	3 m / 57 MB	73 m / 115 MB	48 m / 111 MB
ins-2-3	9 m / 8 MB	3 m / 7 MB	5 m / 8 MB	4 m / 8 MB	2 m / 8 MB	2 m / 7 MB	1 m / 7 MB	2 m / 8 MB	1 m / 6 MB
del-1-2	24 s / 176 MB	16 s / 139 MB	13 s / 87 MB	11 s / 79 MB	8 s / 58 MB	7 s / 53 MB	12 s / 111 MB	21 s / 184 MB	11 s / 102 MB

debugging, and other improvement activities going on while the system was running. Second, an earlier version of one of the synchronization scripts contained a bug, which was not detected in previous tests. Only a subset of the insertions were propagated between the databases. Third, while the campaign was running, the infrastructure was migrated to another server. After the migration, the deployment of the scripts was delayed, which caused policy violations.

Overall, the main reason for these violations is that we monitored an experimental system still under development. In this case study the monitor proved to be a powerful debugging tool. For commercial systems, it can detect policy violations thereby protecting the users' privacy and increasing users' trust in using the systems. Our findings also show that policy monitoring makes sense even in systems where users are honest and interested in honoring the policies.

V. Related Work

The usage-control architecture described by Pretschner et al. [8] and the UCON_{ABC} architecture of Park and Sandhu [9] both utilize monitoring techniques. However, the two architectures are only conceptual and have neither been deployed nor evaluated in a real-word setting.

Goodloe and Pike [10] recently surveyed the state of the art for monitoring distributed systems. We restrict ourselves here to the most related work. Bauer et al. [11] examine a setting where actions are totally ordered and system requirements are given in a propositional linear-time temporal logic. Both assumptions are too restrictive in our setting. However, their monitoring architecture additionally includes a component that analyzes the cause of a failure, which is fed back into the system. Genon et al. [12] present a monitoring algorithm for propositional LTL, where events are partially ordered. They use symbolic exploration methods to cope with the interleavings of events. It is unclear how their algorithm extends to a first-order setting. Moreover, in our approach, we consider formulas in a richer logic for which monitoring a single trace is sufficient. In contrast to these works and ours, Sen et al. [13] present a distributed monitoring approach, where multiple monitors are implemented locally and communicate with each other. These monitors are generated from a propositional past-time linear-time distributed temporal logic. A potential bottleneck is the monitors' communication overhead.

Finally, research on checking temporal integrity constrains [14], [15] of stored data and temporal triggers [16] in databases is related to our monitoring algorithm [1]. In fact, our monitoring algorithm extends Chomicki's monitor [14] by handling bounded future operators. These temporal operators are extremely useful for formalizing usage-control policies, which usually contain obligations. We are not aware of any implementation and experimental evaluation of Chomicki's monitoring algorithm.

VI. Conclusion

We theoretically and practically tackled the problem of monitoring the usage of data in concurrent distributed systems. We provided means to efficiently monitor concurrently generated logs. We also deployed and evaluated a monitoring architecture in a real-world application, Nokia's Data-collection Campaign. Our case study demonstrates the feasibility and benefits of monitoring the usage of sensitive data.

As future work we plan to develop monitoring techniques for more complex systems with more agents, actions, and databases. The challenges will be to handle less accurate and less complete logging, and to provide monitoring algorithms that scale up from millions to billions of log entries per day. Our future work also includes developing monitoring techniques that can also be used for policy enforcement, i.e., preventing policy violations.

Acknowledgments. This work was supported by the Nokia Research Center, Switzerland. The authors thank Imad Aad, Debmalya Biswas, Olivier Bornet, Olivier Dousse, Juha Laurila, and Valtteri Niemi for valuable input.

References

[1] D. Basin, F. Klaedtke, S. Müller, and B. Pfitzmann, "Runtime monitoring of metric first-order temporal properties," in *Proceedings of the 28th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, ser. Leibiz International Proceedings in Informatics (LIPIcs), vol. 2. Schloss Dagstuhl - Leibniz Center for Informatics, 2008, pp. 49–60.

[2] D. Basin, F. Klaedtke, and S. Müller, "Monitoring security policies with metric first-order temporal logic," in *Proceeding of the 15th ACM Symposium on Access Control Models and Technologies (SACMAT)*. ACM Press, 2010, pp. 23–34.

[3] A. S. Tanenbaum and M. van Steen, *Distributed Systems: Principles and Paradigms*. Prentice Hall, 2002.

[4] L. Lamport, "Time, clocks, and the ordering of events in a distributed system," *Commun. ACM*, vol. 21, no. 7, pp. 558–565, 1978.

[5] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases: The Logical Level*. Addison Wesley, 1994.

[6] L. Lamport, "What good is temporal logic?" in *Proceedings of the IFIP 9th World Computer Congress*, ser. Information Processing, vol. 83. North-Holland, 1983, pp. 657–668.

[7] I. Aad and V. Niemi, "NRC data collection campaign and the privacy by design principles," in *Proceedings of the International Workshop on Sensing for App Phones (PhoneSense)*, 2010.

[8] A. Pretschner, M. Hilty, and D. Basin, "Distributed usage control," *Commun. ACM*, vol. 49, no. 9, pp. 39–44, 2006.

[9] J. Park and R. Sandhu, "The UCON_{ABC} usage control model," *ACM Trans. Inform. Syst. Secur.*, vol. 7, no. 1, pp. 128–174, 2004.

[10] A. Goodloe and L. Pike, "Monitoring distributed real-time systems: A survey and future directions," NASA Langley Research Center, Tech. Rep. NASA/CR-2010-216724, July 2010.

[11] A. Bauer, M. Leucker, and C. Schallhart, "Model-based runtime analysis of distributed reactive systems," in *Proceedings of the 2006 Australian Software Engineering Conference (ASWEC)*. IEEE Computer Society, 2006.

[12] A. Genon, T. Massart, and C. Meuter, "Monitoring distributed controllers: When an efficient LTL algorithm on sequences is needed to model-check traces," in *Proceedings of the 14th International Symposium on Formal Methods (FM)*, ser. Lect. Notes Comput. Sci., vol. 4085. Springer, 2006, pp. 557–572.

[13] K. Sen, A. Vardhan, G. Agha, and G. Roşu, "Efficient decentralized monitoring of safety in distributed systems," in *Proceedings of the 26th International Conference on Software Engineering (ICSE)*. IEEE Computer Society, 2004, pp. 418–427.

[14] J. Chomicki, "Efficient checking of temporal integrity constraints using bounded history encoding," *ACM Trans. Database Syst.*, vol. 20, no. 2, pp. 149–186, 1995.

[15] U. W. Lipeck and G. Saake, "Monitoring dynamic integrity constraints based on temporal logic," *Inform. Syst.*, vol. 12, no. 3, pp. 255–269, 1987.

[16] A. P. Sistla and O. Wolfson, "Temporal triggers in active databases," *IEEE Trans. Knowl. Data Eng.*, vol. 7, no. 3, pp. 471–486, 1995.

[17] T. Massart, C. Meuter, and L. Van Begin, "On the complexity of partial order trace model checking," *Inform. Process. Lett.*, vol. 106, no. 3, pp. 120–126, 2008.

[18] N. Markey and P. Schnoebelen, "Model checking a path," in *Proceedings of the 14th International Conference on Concurrency Theory (CONCUR)*, ser. Lect. Notes Comput. Sci., vol. 2761. Springer, 2003, pp. 248–262.

Appendix

A. Additional Proof Details: Intractability Results

We remark that related intractability results for LTL on so-called partially ordered traces are given in [17]. However, the setting is different from ours. In particular, it is unclear how to describe the set of interleavings of two timestamped traces using partially ordered traces as defined in [17]. Moreover, we reduce SAT and TAUT, respectively, to the respective decision problem for proving its hardness. In [17], the global-predicate-detection decision problem is used.

The decision problem in Theorem 3(1) is in NP as a nondeterministic Turing machine can first guess the violating interleaving up to the given time point and then verify its guess in polynomial time [18]. Note that the Turing machine does not need to guess a valuation, as the input formula is a quantifier-free sentence and this contains no variables. Hardness is established by polynomially reducing SAT to the decision problem in Theorem 3(1) as shown below. Analogously, the coNP-hardness of the decision problem in Theorem 3(2) is shown by polynomially reducing TAUT to it, also explained below. This problem is in coNP since its complement is in NP.

Reduction from SAT. We show NP-hardness of the decision problem in Theorem 3(1) by reduction from SAT.

To fix notation, we recall that a propositional formula α over a set of atomic propositions P is satisfiable if there is an assignment θ of propositions to truth values $\bot$ (denoting false) and $\top$ (denoting true), i.e. $\theta : P \to \{\bot, \top\}$, such that $\theta(\alpha) = \top$, where θ is extended from atomic propositions to formulas as expected. The SAT problem asks whether a given propositional formula is satisfiable. SAT is NP-hard.

Suppose $P = \{p_0, \ldots, p_{n-1}\}$, with $n \geq 0$, is a set of atomic propositions. Let S be the signature (C, R, ι) with $C = \{c\}$, $R = \{q_0, r_0, \ldots, q_{n-1}, r_{n-1}\}$, and $\iota(q_i) = \iota(r_i) = 1$, for any $0 \leq i < n$. The two temporal structures $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ and $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$ over S are given by: $|\bar{\mathcal{D}}| = \{c\}$, $c^{\bar{\mathcal{D}}} = c$, $\tau_i^1 = \tau_i^2 = i$ for any $i \in \mathbb{N}$, and for any $k \in \{1, 2\}$ and $i, j \in \mathbb{N}$ with $0 \leq i < n$,

$$q_i^{\mathcal{D}_j^k} = \begin{cases} \{c\} & \text{if } k = 1 \text{ and } i = j, \\ \emptyset & \text{otherwise,} \end{cases}$$

$$r_i^{\mathcal{D}_j^k} = \begin{cases} \{c\} & \text{if } k = 2 \text{ and } i = j, \\ \emptyset & \text{otherwise.} \end{cases}$$

Given a propositional formula α over P, the MFOTL formula $\ulcorner\alpha\urcorner$ is obtained by replacing each occurrence of a proposition p_i in α with $\blacklozenge(r_i(c) \wedge \blacklozenge\, q_i(c))$. Thus, given a propositional formula α, the reduction constructs the two prefixes of length n of $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ and $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$ and the MFOTL formula $\ulcorner\alpha\urcorner$. This reduction is linear in the size of α. Its correctness is shown by Lemma 13. The following remarks and lemma will be needed.

Remark. For any interleaving $(\bar{\mathcal{D}}, \bar{\tau}) \in (\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$, the functions f_1 and f_2 in Definition 1 satisfy $f_k(i) \in \{2i, 2i + 1\}$ where $k \in \{1, 2\}$. Moreover, these functions are unique,

that is, if $g_1, g_2 : \mathbb{N} \to \mathbb{N}$ are strictly monotonic functions satisfying conditions (1)–(3) in Definition 1 then either $g_1 = f_1$ and $g_2 = f_2$, or $g_1 = f_2$ and $g_2 = f_1$. Furthermore, for any strictly monotonic functions f_1 and f_2 satisfying conditions (1) and (2) in Definition 1 and with $f_1(i), f_2(i) \in \{2i, 2i+1\}$ for $0 \leq i < n$, there is a unique temporal structure $(\bar{\mathcal{D}}, \bar{\tau})$ such that f_1 and f_2 also satisfy condition (3). In other words, the functions f_1, f_2 determine an interleaving of $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ and $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$

Lemma 12. *Let α be a propositional formula, θ a truth value assignment, v a valuation, and $(\bar{\mathcal{D}}, \bar{\tau})$ an interleaving of $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ given by the functions f_1 and f_2 such that $\theta(p_i) = \top$ iff $f_1(i) = 2i$, for any i with $0 \leq i < n$. It holds that $\theta(\alpha) = \top$ if and only if $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \models \ulcorner\alpha\urcorner$.*

Proof: We use structural induction on the form of α. The only interesting case is the base case, the other cases follow directly from the induction hypotheses. Thus let $\alpha = p_i \in P$.

Suppose that $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \models \blacklozenge(r_i(c) \wedge \blacklozenge\, q_i(c))$. That is, there is a time point $j \leq 2n$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models r_i(c)$ and such that there is a time point $j' \leq j$ for which $(\bar{\mathcal{D}}, \bar{\tau}, v, j') \models q_i(c)$. Then $c \in r_i^{\mathcal{D}_j}$ and $c \in q_i^{\mathcal{D}_{j'}}$. From the definition of an interleaving and the definitions of the interpretations of the predicates q_i and r_i, it follows that $j = f_2(i)$ and $j' = f_1(i)$. Then, as $f_1(i), f_2(i) \in \{2i, 2i+1\}$, $f_1(i) \neq f_2(i)$, and $j' \leq j$, we get that $f_1(i) = 2i$ and $f_2(i) = 2i + 1$. Thus $\theta(p_i) = \top$.

Suppose that $\theta(\alpha) = \top$. Then $f_1(i) = 2i$ and $f_2(i) = 2i + 1$. We have $(\bar{\mathcal{D}}, \bar{\tau}, v, 2i) \models q_i(c)$ and $(\bar{\mathcal{D}}, \bar{\tau}, v, 2i+1) \models r_i(c)$. Thus $(\bar{\mathcal{D}}, \bar{\tau}, v, 2i + 1) \models r_i(c) \wedge \blacklozenge\, q_i(c)$ and clearly $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \models \blacklozenge(r_i(c) \wedge \blacklozenge\, q_i(c))$. ∎

Lemma 13. *Let α be a propositional formula. It holds that α is satisfiable if and only if $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ weakly violates $\neg\ulcorner\alpha\urcorner$ at time point $2n$.*

Proof: Suppose first that α is satisfiable. Then there is a truth value assignment θ such that $\theta(\alpha) = \top$. Let $(\bar{\mathcal{D}}, \bar{\tau})$ be the interleaving determined by the functions f_1 and f_2 given by

$$f_1(i) = \begin{cases} 2i & \text{if } \theta(p_i) = \top, \\ 2i+1 & \text{otherwise,} \end{cases}$$

and

$$f_2(i) = \begin{cases} 2i & \text{if } \theta(p_i) = \bot, \\ 2i+1 & \text{otherwise.} \end{cases}$$

Let v be an arbitrary valuation. From Lemma 12, we obtain that $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \models \ulcorner\alpha\urcorner$, that is, $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \not\models \neg\ulcorner\alpha\urcorner$.

Suppose now that $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ weakly violates $\neg\ulcorner\alpha\urcorner$ at time point $2n$. Then there is an interleaving $(\bar{\mathcal{D}}, \bar{\tau})$ and a valuation v such $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \not\models \neg\ulcorner\alpha\urcorner$. Let f_1 and f_2 the be functions determined by $(\bar{\mathcal{D}}, \bar{\tau})$ as in Definition 1. Let θ be a truth value assignment such that $\theta(p_i) = \top$ if and only if $f_1(i) = 2i$. Using again Lemma 12, we get that θ is a satisfying assignment for α. ∎

Reduction from TAUT. We show coNP-hardness of the decision problem in Theorem 3(2) by reduction from TAUT.

We recall that a propositional formula α over a set of atomic propositions P is a tautology if $\theta(\alpha) = \top$ for any assignment θ of propositions to truth values. The TAUT problem asks whether a given propositional formula is a tautology. TAUT is coNP-hard.

We use the same reduction as for the decision problem in Theorem 3(1). The correctness of the reduction follows from the following lemma.

Lemma 14. *Let α be a propositional formula. It holds that α is a tautology if and only if $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ strongly violates $\neg\ulcorner\alpha\urcorner$ at time point $2n$.*

Proof: Suppose first that α is a tautology. Let $(\bar{\mathcal{D}}, \bar{\tau})$ be an arbitrary interleaving in $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ and f_1, f_2 be functions as in Definition 1. Let θ be a truth value assignment such that $\theta(p_i) = \top$ if and only if $f_1(i) = 2i$. Let v be an arbitrary valuation. Using Lemma 12, we obtain that $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \not\models \neg\ulcorner\alpha\urcorner$. Hence $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ strongly violates $\neg\ulcorner\alpha\urcorner$ at time point $2n$.

Suppose now that $(\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$ strongly violates $\neg\ulcorner\alpha\urcorner$ at time point $2n$. Let θ be an arbitrary truth value assignment. Let $(\bar{\mathcal{D}}, \bar{\tau})$ be the interleaving determined by the functions f_1 and f_2 given by

$$f_1(i) = \begin{cases} 2i & \text{if } \theta(p_i) = \top, \\ 2i+1 & \text{otherwise,} \end{cases}$$

and

$$f_2(i) = \begin{cases} 2i & \text{if } \theta(p_i) = \bot, \\ 2i+1 & \text{otherwise.} \end{cases}$$

There is a valuation v such $(\bar{\mathcal{D}}, \bar{\tau}, v, 2n) \not\models \neg\ulcorner\alpha\urcorner$. Using again Lemma 12, we get that θ is a satisfying assignment for α. Hence α is a tautology. ∎

B. Additional Proof Details: Derivation Rules

Figure 5 lists all the inference rules for label propagation. Lemma 15 (see below) shows the soundness of these rules.

When considering formulas in positive normal form, as required in Theorem 11, the Boolean operator $\vee$ and the temporal operators release R_I and trigger T_I are seen as primitives, instead of being defined as syntactic sugar. We recall that $\psi\ \mathsf{R}_I\ \chi$ abbreviates $\neg(\neg\psi\ \mathsf{S}_I\ \neg\chi)$ and $\psi\ \mathsf{T}_I\ \chi$ abbreviates $\neg(\neg\psi\, \mathsf{U}_I\, \neg\chi)$. Figure 6 lists propagation rules for formulas that use these operators. Their soundness follows from the soundness of rules in Figure 5 and the mentioned equivalences. For instance, the correctness of the rule

$$\frac{\psi : (\models \exists) \quad \chi : (\models \forall)}{(\psi\ \mathsf{R}_I\ \chi) \vee (\Diamond_J\, \psi) : (\models \forall)}\ 0 \notin I, 0 \in J$$

follows from unfolding the abbreviation $(\psi\ \mathsf{R}_I\ \chi) \vee (\Diamond_J\, \psi)$, which is $\neg((\neg\psi\ \mathsf{S}_I\ \neg\chi) \wedge (\Box_J\, \neg\psi))$, and the following deriva-

tion:

$$\frac{\dfrac{\psi : (\models \exists)}{\neg\psi : (\not\models \exists)} \quad \dfrac{\chi : (\models \forall)}{\neg\chi : (\not\models \forall)}}{\dfrac{(\neg\psi \,\mathsf{S}_I\, \neg\chi) \wedge (\Box_J\, \neg\psi) : (\not\models \forall)}{\neg((\neg\psi \,\mathsf{S}_I\, \neg\chi) \wedge (\Box_J\, \neg\psi)) : (\models \forall)}}\ 0 \notin I, 0 \in J$$

Finally, for convenience, Figure 7 lists some inference rules for formulas for which the main operator is one of the temporal operators $\blacklozenge_I$, $\Diamond_I$, $\blacksquare_I$, and $\Box_I$. These rules can be derived from the rules in Figure 5 by simply applying the definition of syntactic sugar. For instance, the rule

$$\frac{\psi : (\models \forall)}{\blacklozenge_I\, \psi : (\models \forall)}$$

can be derived from

$$\frac{\dfrac{\overline{x \approx x : (\models \forall)}}{\exists x.\, x \approx x : (\models \forall)} \quad \psi : (\models \forall)}{(\exists x.\, x \approx x)\, \mathsf{S}_I\, \psi : (\models \forall)}$$

Note that $\blacklozenge_I\, \psi$ is syntactic sugar for $(\exists x.\, x \approx x)\, \mathsf{S}_I\, \psi$.

We now show the soundness of the inference rules in Figure 5.

Lemma 15. *Let ϕ be a formula. If ϕ can be labeled with ℓ, then ϕ satisfies the invariant ℓ, where $\ell \in \{(\models \forall), (\not\models \forall), (\not\models \exists), (\models \exists)\}$.*

Proof: Let $(\bar{\mathcal{C}}, \bar{\kappa})$ be the collapse of an interleaving of two given temporal structures.

We proceed by induction on size of the derivation tree assigning label ℓ to ϕ. We make a case distinction based on the rule applied to label the formula, that is, the rule at the root of the tree. However, for clarity, we generally group cases by the formula's form.

For readability, and without loss of generality, we already fix an arbitrary valuation v, an arbitrary time point i, and an arbitrary temporal structure $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$.

- We first consider the weakening rules.
  - ϕ is labeled with $(\models \forall)$ and $(\models \exists)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$. By the induction hypothesis, ϕ satisfies the invariant $(\models \forall)$, thus $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$ for any j with $\tau_j = \kappa_i$. By the definition of $(\bar{\mathcal{C}}, \bar{\kappa})$, there is at least one j with $\tau_j = \kappa_i$. Hence ϕ satisfies the invariant $(\models \exists)$.
  - ϕ is labeled with $(\not\models \forall)$ and with $(\not\models \exists)$. This case is analogous to the previous one.
- $\phi = t \approx t'$, where t and t' are variables or constants. In this case ϕ is labeled with $(\models \forall)$ and $(\not\models \forall)$.
  - ϕ is labeled with $(\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$. Then $v(t) = v(t')$. Clearly, $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$ for any time point j, as ϕ only depends on the valuation. The invariant $(\models \forall)$ is hence satisfied.
  - ϕ is labeled with $(\not\models \forall)$. This case is analogous to the previous one.
- $\phi = t \prec t'$, where t and t' are variables or constants. This case is analogous to the previous one.
- $\phi = r(t_1, \ldots, t_{\iota(r)})$, where $t_1, \ldots, t_{\iota(r)}$ are variables or constants. In this case ϕ is labeled with $(\models \exists)$ and $(\not\models \forall)$.
  - ϕ is labeled with $(\models \exists)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$. Then $(v(t_1), \ldots, v(t_{\iota(r)})) \in r^{\mathcal{C}_i}$. As $r^{\mathcal{C}_i} = \bigcup_{\{j \mid \tau_j = \kappa_i\}} r^{\mathcal{D}_j}$, there is a j with $\tau_j = \kappa_i$ such that $(v(t_1), \ldots, v(t_{\iota(r)})) \in r^{\mathcal{D}_j}$. Therefore $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$. Thus ϕ satisfies the invariant $(\models \exists)$.
  - ϕ is labeled with $(\not\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \phi$. Then for any j with $\tau_j = \kappa_i$ we have that $(v(t_1), \ldots, v(t_{\iota(r)})) \notin r^{\mathcal{D}_j}$, that is, $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$. Thus ϕ satisfies the invariant $(\not\models \forall)$.
- $\phi = \neg\psi$. If ψ is labeled with ℓ, then ϕ is labeled with $\neg\ell$, where $\neg\ell$ is $(\models \forall)$, $(\not\models \forall)$, $(\not\models \exists)$, or $(\models \exists)$ when ℓ is $(\not\models \forall)$, $(\models \forall)$, $(\models \exists)$, or $(\not\models \exists)$ respectively.
  - ϕ is labeled with $(\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \neg\psi$. By the induction hypothesis, ψ satisfies the invariant $(\not\models \forall)$. As $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \psi$, we have that $(\bar{\mathcal{D}}, \bar{\tau}, v, k) \not\models \psi$, that is, $(\bar{\mathcal{D}}, \bar{\tau}, v, k) \models \phi$, for all k with $\tau_k = \kappa_i$. Thus ϕ satisfies the invariant $(\models \forall)$.
  - The other cases are similar.
- $\phi = \psi \wedge \chi$. There are four rules to be analyzed.
  - ϕ, ψ, and χ are labeled with $(\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \psi \wedge \chi$. Then $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \psi$ and $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \chi$. By the induction hypothesis, ψ and χ satisfy the invariant $(\models \forall)$. Hence, for all j with $\tau_j = \kappa_i$, we have $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \psi$ and $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \chi$. Thus $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$ and $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \chi$ for all j with $\tau_j = \kappa_i$. Hence, ϕ satisfies the invariant $(\models \forall)$.
  - The other cases are similar.
- $\phi = \exists x.\psi$. There are four rules, one for each label: if ψ is labeled with ℓ, then ϕ is labeled with ℓ.
  - ℓ is $(\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \exists x.\psi$. Then there is a $d \in |\bar{\mathcal{D}}|$ such that $(\bar{\mathcal{C}}, \bar{\kappa}, v[x/d], i) \models \psi$. As ψ satisfies the invariant $(\models \forall)$, we have $(\bar{\mathcal{D}}, \bar{\tau}, v[x/d], j) \models \psi$ for all j with $\tau_j = \kappa_i$. That is, $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \exists x.\psi$ for all j with $\tau_j = \kappa_i$. Hence ϕ satisfies the invariant $(\models \forall)$.
  - The other cases are similar.
- $\phi = \psi\, \mathsf{S}_I\, \chi$. We have three rules to analyze.
  - ϕ, ψ, and χ are each labeled with $(\models \forall)$. By the induction hypothesis, ψ and χ satisfy the invariant $(\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$. Then, for some $j \leq i$ with $\kappa_i - \kappa_j \in I$, we have $(\bar{\mathcal{C}}, \bar{\kappa}, v, j) \models \chi$ and $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \models \psi$ for all $k \in [j+1, i+1)$. Let i' be an arbitrary time point such that $\tau_{i'} = \kappa_i$. As χ satisfies the invariant $(\models \forall)$, for the largest j' with $\tau_{j'} = \kappa_j$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, j') \models \chi$. Clearly, $\tau_{i'} - \tau_{j'} \in I$. From the definition of $(\bar{\mathcal{C}}, \bar{\kappa})$, for any $k' \in [j'+1, i'+1)$, there is a $k \in [j+1, i+1)$ such that $\tau_{k'} = \kappa_k$. Then, as ψ satisfies the invariant $(\models \forall)$, for any $k' \in [j'+1, i'+1)$, we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$. As ψ satisfies the invariant $(\models \forall)$, for all $k \in [j+1, i+1)$ and all k' with $\tau_{k'} = \kappa_k$,

$$\frac{\phi : (\models \forall)}{\phi : (\models \exists)} \qquad \frac{\phi : (\not\models \forall)}{\phi : (\not\models \exists)}$$

$$\frac{}{t \approx t' : (\models \forall)} \qquad \frac{}{t \approx t' : (\not\models \forall)} \qquad \frac{}{t \prec t' : (\models \forall)} \qquad \frac{}{t \prec t' : (\not\models \forall)}$$

$$\frac{}{r(t_1, \ldots, t_{\iota(r)}) : (\models \exists)} \qquad \frac{}{r(t_1, \ldots, t_{\iota(r)}) : (\not\models \forall)}$$

$$\frac{\psi : (\models \exists)}{\neg\psi : (\not\models \exists)} \qquad \frac{\psi : (\models \forall)}{\neg\psi : (\not\models \forall)} \qquad \frac{\psi : (\not\models \exists)}{\neg\psi : (\models \exists)} \qquad \frac{\psi : (\not\models \forall)}{\neg\psi : (\models \forall)}$$

$$\frac{\psi : (\models \forall) \quad \chi : (\models \forall)}{\psi \wedge \chi : (\models \forall)} \qquad \frac{\psi : (\models \forall) \quad \chi : (\models \exists)}{\psi \wedge \chi : (\models \exists)}$$

$$\frac{\psi : (\not\models \forall) \quad \chi : (\not\models \forall)}{\psi \wedge \chi : (\not\models \forall)} \qquad \frac{\psi : (\not\models \exists) \quad \chi : (\not\models \exists)}{\psi \wedge \chi : (\not\models \exists)}$$

$$\frac{\psi : (\models \forall)}{\exists x.\, \psi : (\models \forall)} \qquad \frac{\psi : (\models \exists)}{\exists x.\, \psi : (\models \exists)} \qquad \frac{\psi : (\not\models \forall)}{\exists x.\, \psi : (\not\models \forall)} \qquad \frac{\psi : (\not\models \exists)}{\exists x.\, \psi : (\not\models \exists)}$$

$$\frac{\psi : (\models \forall) \quad \chi : (\models \forall)}{\psi \,\mathsf{S}_I\, \chi : (\models \forall)} \qquad \frac{\psi : (\not\models \forall) \quad \chi : (\not\models \forall)}{\psi \,\mathsf{S}_I\, \chi : (\not\models \forall)} \qquad \frac{\psi : (\not\models \exists) \quad \chi : (\not\models \forall)}{\psi \,\mathsf{S}_I\, \chi : (\not\models \exists)} \qquad \frac{\psi : (\not\models \exists) \quad \chi : (\not\models \forall)}{(\psi \,\mathsf{S}_I\, \chi) \wedge (\Box_J\, \psi) : (\not\models \forall)}\ 0 \notin I, 0 \in J$$

$$\frac{\psi : (\models \forall) \quad \chi : (\models \forall)}{\psi \,\mathsf{U}_I\, \chi : (\models \forall)} \qquad \frac{\psi : (\not\models \forall) \quad \chi : (\not\models \forall)}{\psi \,\mathsf{U}_I\, \chi : (\not\models \forall)} \qquad \frac{\psi : (\not\models \exists) \quad \chi : (\not\models \forall)}{\psi \,\mathsf{U}_I\, \chi : (\not\models \exists)} \qquad \frac{\psi : (\not\models \exists) \quad \chi : (\not\models \forall)}{(\psi \,\mathsf{U}_I\, \chi) \wedge (\blacksquare_J\, \psi) : (\not\models \forall)}\ 0 \notin I, 0 \in J$$

$$\frac{\psi : (\models \exists)}{\blacklozenge_I\, \psi : (\models \exists)} \qquad \frac{\psi : (\models \exists)}{\blacklozenge_I\, \psi : (\models \forall)}\ 0 \notin I \qquad \frac{\psi : (\models \exists)}{\Diamond_I\, \psi : (\models \exists)} \qquad \frac{\psi : (\models \exists)}{\Diamond_I\, \psi : (\models \forall)}\ 0 \notin I$$

$$\frac{\psi : (\models \exists)}{\blacklozenge_I\, \Diamond_J\, \psi : (\models \forall)}\ 0 \in I \cap J$$

Figure 5. Inference Rules

$$\frac{\psi : (\not\models \forall) \quad \chi : (\not\models \forall)}{\psi \vee \chi : (\not\models \forall)} \qquad \frac{\psi : (\not\models \forall) \quad \chi : (\not\models \exists)}{\psi \vee \chi : (\not\models \exists)}$$

$$\frac{\psi : (\models \forall) \quad \chi : (\models \forall)}{\psi \vee \chi : (\models \forall)} \qquad \frac{\psi : (\models \exists) \quad \chi : (\models \exists)}{\psi \vee \chi : (\models \exists)}$$

$$\frac{\psi : (\not\models \forall) \quad \chi : (\not\models \forall)}{\psi \,\mathsf{R}_I\, \chi : (\not\models \forall)} \qquad \frac{\psi : (\models \forall) \quad \chi : (\models \forall)}{\psi \,\mathsf{R}_I\, \chi : (\models \forall)} \qquad \frac{\psi : (\models \exists) \quad \chi : (\models \forall)}{\psi \,\mathsf{R}_I\, \chi : (\models \exists)} \qquad \frac{\psi : (\models \exists) \quad \chi : (\models \forall)}{(\psi \,\mathsf{R}_I\, \chi) \vee (\Diamond_J\, \psi) : (\models \forall)}\ 0 \notin I, 0 \in J$$

$$\frac{\psi : (\not\models \forall) \quad \chi : (\not\models \forall)}{\psi \,\mathsf{T}_I\, \chi : (\not\models \forall)} \qquad \frac{\psi : (\models \forall) \quad \chi : (\models \forall)}{\psi \,\mathsf{T}_I\, \chi : (\models \forall)} \qquad \frac{\psi : (\models \exists) \quad \chi : (\models \forall)}{\psi \,\mathsf{T}_I\, \chi : (\models \exists)} \qquad \frac{\psi : (\models \exists) \quad \chi : (\models \forall)}{(\psi \,\mathsf{T}_I\, \chi) \vee (\blacklozenge_J\, \psi) : (\models \forall)}\ 0 \notin I, 0 \in J$$

Figure 6. Inference Rules for Formulas in Positive Normal Form

$$\frac{\psi : (\models \forall)}{\blacklozenge_I \psi : (\models \forall)} \qquad \frac{\psi : (\not\models \forall)}{\blacklozenge_I \psi : (\not\models \forall)}$$

$$\frac{\psi : (\models \forall)}{\lozenge_I \psi : (\models \forall)} \qquad \frac{\psi : (\not\models \forall)}{\lozenge_I \psi : (\not\models \forall)}$$

$$\frac{\psi : (\models \forall)}{\blacksquare_I \psi : (\models \forall)} \quad \frac{\psi : (\not\models \forall)}{\blacksquare_I \psi : (\not\models \forall)} \quad \frac{\psi : (\not\models \exists)}{\blacksquare_I \psi : (\not\models \exists)} \quad \frac{\psi : (\not\models \exists)}{\blacksquare_I \psi : (\not\models \forall)}\ 0 \notin I$$

$$\frac{\psi : (\models \forall)}{\square_I \psi : (\models \forall)} \quad \frac{\psi : (\not\models \forall)}{\square_I \psi : (\not\models \forall)} \quad \frac{\psi : (\not\models \exists)}{\square_I \psi : (\not\models \exists)} \quad \frac{\psi : (\not\models \exists)}{\square_I \psi : (\not\models \forall)}\ 0 \notin I$$

$$\frac{\psi : (\not\models \exists)}{\blacksquare_I \square_J \psi : (\not\models \forall)}\ 0 \in I \cap J$$

Figure 7. Derived Inference Rules

we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$. Hence $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \psi \,\mathsf{S}_I\, \chi$, and thus ϕ satisfies the invariant $(\models \forall)$.

– ϕ, ψ, and χ are each labeled with $(\not\models \forall)$. By the induction hypothesis, ψ and χ satisfy the invariant $(\not\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \phi$ and that, by absurdity, ϕ does not satisfy the invariant $(\not\models \forall)$. That is, there is an i' with $\tau_{i'} = \kappa_i$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \phi$. Then there is a $j' \leq i'$ with $\tau_{i'} - \tau_{j'} \in I$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j') \models \chi$ and for all $k' \in [j'+1, i'+1)$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k) \models \psi$. By the definition of $(\bar{\mathcal{C}}, \bar{\kappa})$, there is a j with $\kappa_j = \tau_{j'}$. As χ satisfies the invariant $(\not\models \forall)$, we have that $(\bar{\mathcal{C}}, \bar{\kappa}, v, j) \models \chi$. Similarly, we have that $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \models \psi$ for all $k \in [j+1, i+1)$. That is, $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$, which is a contradiction.

– ϕ and ψ are labeled with $(\not\models \exists)$, and χ is labeled by $(\not\models \forall)$. By the induction hypothesis, ψ and χ satisfy the invariants $(\not\models \exists)$ and $(\not\models \forall)$ respectively. As before, suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \phi$ and that, by absurdity, ϕ does not satisfy the invariant $(\not\models \exists)$. That is, for all i' with $\tau_{i'} = \kappa_i$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \phi$. Consider the largest such i'. Then there is a $j' \leq i'$ with $\tau_{i'} - \tau_{j'} \in I$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j') \models \chi$ and for all $k' \in [j'+1, i'+1)$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$. By the definition of $(\bar{\mathcal{C}}, \bar{\kappa})$, there is a j with $\kappa_j = \tau_{j'}$. As χ satisfies the invariant $(\not\models \forall)$, we have that $(\bar{\mathcal{C}}, \bar{\kappa}, v, j) \models \chi$. Take $k \in [j+1, i+1)$ arbitrarily. If $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \not\models \psi$, as ψ satisfies the invariant $(\not\models \exists)$, then there is a k' with $\tau_{k'} = \kappa_k$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \not\models \psi$. This contradicts our assumption that $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \phi$, since such k' must be in the interval $[j'+1, i'+1)$. We thus have that $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \models \psi$ for all $k \in [j+1, i+1)$. Hence $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$, which is a contradiction.

- $\phi = \psi \,\mathsf{U}_I\, \chi$. This case is analogous to the previous one.
- $\phi = (\psi \,\mathsf{S}_I\, \chi) \wedge (\square_J \psi)$ with $0 \notin I$ and $0 \in J$. ϕ and χ are labeled with $(\not\models \forall)$, and ψ is labeled by $(\not\models \exists)$. By the induction hypothesis, ψ and χ satisfy the invariants $(\not\models \exists)$ and $(\not\models \forall)$ respectively. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \not\models \phi$ and that, by absurdity, ϕ does not satisfy the invariant $(\not\models \forall)$. That is, there is an i' with $\tau_{i'} = \kappa_i$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \phi$. Then there is a $j' \leq i'$ with $\tau_{i'} - \tau_{j'} \in I$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j') \models \chi$ and for all $k' \in [j'+1, i'+1)$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$; and for all $j'' \geq i'$ with $\tau_{j''} - \tau_{i'} \in J$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, j'') \models \psi$.
  By the definition of $(\bar{\mathcal{C}}, \bar{\kappa})$, there is a j with $\kappa_j = \tau_{j'}$. As χ satisfies the invariant $(\not\models \forall)$, we have that $(\bar{\mathcal{C}}, \bar{\kappa}, v, j) \models \chi$. Take $k \in [j+1, i)$ arbitrarily. If $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \not\models \psi$, as ψ satisfies the invariant $(\not\models \exists)$, then there is a k' with $\tau_{k'} = \kappa_k$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \not\models \psi$. This contradicts our assumption that $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \phi$. Indeed, such a k' must be in the interval $[j'+1, i''+1)$ where i'' is the largest such that $\tau_{i''} = \kappa_i$. If $k' \leq i'$ then $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \not\models \psi \,\mathsf{S}_I\, \chi$. If $k' > i'$ then $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \not\models \square_J \psi$, as $0 \in J$. We thus have that $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \models \psi$ for all $k \in [j+1, i+1)$. Hence $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \psi \,\mathsf{S}_I\, \chi$.
  As $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \square_J \psi$ and $0 \in J$, it follows that for all $k' \geq i'$ with $\tau_{k'} = \tau_{i'}$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$. We have seen that $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$ for all $k' \in [j'+1, i'+1)$. Because $\tau_{j'} < \tau_{i'}$ (as $0 \notin I$), it also follows that for all $k' \leq i'$ with $\tau_{k'} = \tau_{i'}$ we have $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$. Hence $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$ for all k' with $\tau_{k'} = \tau_{i'}$. As ψ satisfies the invariant $(\not\models \exists)$, we obtain that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \psi$. Similarly, we obtain that $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \models \psi$ for all $k > i$ such that $\kappa_k - \kappa_i \in J$. Hence $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \square_J \psi$.
  We showed that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$, which is a contradiction. Thus ϕ satisfies the invariant $(\not\models \forall)$.
- $\phi = (\psi \,\mathsf{U}_I\, \chi) \wedge (\blacksquare_J \psi)$ with $0 \notin I$ and $0 \in J$. This case is analogous to the previous one.
- $\phi = \blacklozenge_I \psi$. There are two rules to analyze. For both rules, ψ is labeled with $(\models \exists)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$. Then there is a $j \leq i$ with $\kappa_i - \kappa_j \in I$ such that $(\bar{\mathcal{C}}, \bar{\kappa}, v, j) \models \psi$. As, by the induction hypothesis, ψ satisfies the invariant $(\models \exists)$, there is a j' with $\tau_{j'} = \kappa_j$

such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j') \models \psi$.
  - ϕ is labeled with $(\models \exists)$. Take i' to be the largest k such that $\tau_k = \kappa_i$. Clearly, $\tau_{i'} - \tau_{j'} \in I$ and $j' \leq i'$. Hence $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \blacklozenge_I \psi$ and ϕ satisfies the invariant $(\models \exists)$.
  - $0 \notin I$ and ϕ is labeled with $(\models \forall)$. Take i' arbitrarily such that $\tau_{i'} = \kappa_i$. Clearly, $\tau_{i'} - \tau_{j'} \in I$ and, as $0 \notin I$, $\tau_{i'} - \tau_{j'} > 0$, thus $j' < i'$. Hence $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \blacklozenge_I \psi$. Thus ϕ satisfies the invariant $(\models \forall)$.
- $\phi = \Diamond_I \psi$. This case is analogous to the previous one.
- $\phi = \blacklozenge_I \Diamond_J \psi$ with $0 \in I \cap J$. There is only one rule to consider: ψ is labeled with $(\models \exists)$ and ϕ is labeled by $(\models \forall)$. Suppose that $(\bar{\mathcal{C}}, \bar{\kappa}, v, i) \models \phi$. Then there is a $j \leq i$ with $\kappa_i - \kappa_j \in I$ and there is a $k \geq j$ with $\kappa_k - \kappa_j \in I$ such that $(\bar{\mathcal{C}}, \bar{\kappa}, v, k) \models \psi$. As, by the induction hypothesis ψ satisfies the invariant $(\models \exists)$, there is a k' with $\tau_{k'} = \kappa_k$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, k') \models \psi$. Take i' arbitrarily such that $\tau_{i'} = \kappa_i$. If $k' \geq i'$ then $0 \leq \tau_{k'} - \tau_{i'} = \kappa_k - \kappa_i \leq \kappa_k - \kappa_j \in J$. As $0 \in J$, we have $\tau_{k'} - \tau_{i'} \in J$. Thus $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \Diamond_J \psi$ and, as $0 \in I$, $(\bar{\mathcal{D}}, \bar{\tau}, v, i') \models \blacklozenge_I \Diamond_J \psi$. The case when $k' < i'$ is similar. Hence ϕ satisfies the invariant $(\models \forall)$.

■

C. Additional Proof Details: Theorem 9

The implication in Theorem 9 follows directly from Lemma 8, which in turn follows the correctness of the derivation rules (Lemma 15) and from the following lemma.

Lemma 16. *Let ϕ be a formula.*

1. *If ϕ satisfies the invariant $(\models \forall)$, then ϕ has property (C1).*
2. *If ϕ satisfies the invariant $(\not\models \forall)$, then ϕ has property (C2).*
3. *If ϕ satisfies the invariant $(\models \exists)$, then $\Diamond \phi$ has property (C1).*
4. *If ϕ satisfies the invariant $(\not\models \exists)$, then $\Box \phi$ has property (C2).*

Proof: We fix a temporal structure $(\bar{\mathcal{C}}, \bar{\kappa})$.

1. Suppose ϕ satisfies the invariant $(\models \forall)$ and that $(\bar{\mathcal{C}}, \bar{\kappa}, v, 0) \models \phi$ for some valuation v. Then, for any $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$ and every $j \in \mathbb{N}$ with $\kappa_0 = \tau_j$, it holds that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$. By the definition of collapsed temporal structure, we have $\kappa_0 = \tau_0$. Hence ϕ satisfies (C1).
2. This case is analogous to the previous one.
3. Suppose ϕ satisfies the invariant $(\models \exists)$ and that $(\bar{\mathcal{C}}, \bar{\kappa}, v, 0) \models \Diamond \phi$ for some arbitrary valuation v. Then, for every $(\bar{\mathcal{D}}, \bar{\tau}) \in col^{-1}(\bar{\mathcal{C}}, \bar{\kappa})$, there is some $j \in \mathbb{N}$ with $\kappa_0 = \tau_j$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \phi$. It follows that $(\bar{\mathcal{D}}, \bar{\tau}, v, 0) \models \Diamond \phi$. Hence $\Diamond \phi$ satisfies (C1).
4. This case is analogous to the previous one.

■

Complexity of the Labeling Procedure. We now prove the other part of Theorem 9, which states that a formula ϕ can be labeled in time linear in its length, that is, in $O(|\phi|)$. We start with some definitions and then present a simple labeling algorithm and analyze its complexity.

For a formula ϕ, we define its immediate subformulas $isub(\phi)$ to be: (i) $\{\psi\}$ if $\phi = \neg\psi$, $\phi = \exists x.\, \psi$, $\phi = \bullet_I \psi$, or $\phi = \bigcirc_I \psi$; (ii) $\{\psi, \chi\}$ if $\phi = \psi \wedge \chi$, $\phi = \psi \,\mathsf{S}_I\, \chi$, or $\phi = \psi \,\mathsf{U}_I\, \chi$; and (iii) $\emptyset$ otherwise. For a rule r, we denote $\ell(r)$ the label of the conclusion of the rule.

We assume that the data structure used to represent formulas is a tree corresponding to the formula's syntax tree and that each node in the tree also stores 4 bits representing the 4 different labels. Initially these bits are set to 0, meaning that no label is associated with the corresponding subformula.

```
1  add_labels(φ)
2    foreach ψ ∈ isub(φ)
3      add_labels(ψ)
4    foreach rule r
5      if matches(φ, r) then
6        add_label(φ, ℓ(r))
```

The function $matches(\phi, r)$ checks if the formula ϕ pattern matches a rule r. The order of rules is arbitrary, with the exception that the weakening rules are checked last. So, for instance if ϕ received label $(\models \forall)$, then ϕ will match the appropriate weakening rule and it will also be labeled with $(\models \exists)$. As rules have constant size, and only at most the first two levels of the tree representing the formula ϕ need to be inspected, we conclude that the function executes in constant time.

The function $add\_label(\phi, \ell)$ simply adds the label ℓ to ϕ. Clearly, this operation can be performed in constant time.

Note that the execution of the lines 2 and 4–6 takes constant time: $|isub(\phi)| \leq 2$ for any ϕ, there is a fixed, constant number of rules, and the functions *matches* and *add_label* execute in constant time. Furthermore, the function *add_labels* is executed exactly $|\phi|$ times, once for each subformula of ϕ. Hence the whole labeling procedure of ϕ can be done in linear time in the size of ϕ.

D. Additional Proof Details: Theorem 11

We first show that ϕ^w is weaker than ϕ, or more precisely, that the formula $\phi \rightarrow \phi^w$ is valid. We proceed by structural induction on ϕ.

- $\phi = t \approx t'$, $\phi = t \prec t'$, $\phi = \neg(t \approx t')$, $\phi = \neg(t \prec t')$, or $\neg r(t_1, \ldots, t_{\iota(r)})$, where t, t', and t_i with $1 \leq i \leq \iota(r)$ are variables or constants. Then $\phi^w = \phi$, and the statement clearly holds.
- $\phi = r(t_1, \ldots, t_{\iota(r)})$. Then $\phi^w = \blacklozenge_J \Diamond_{J'} r(t_1, \ldots, t_{\iota(r)})$, for some intervals J and J' with $0 \in J \cap J'$. Let $(\bar{\mathcal{D}}, \bar{\tau})$ be a temporal structure, v a valuation, and i a time point. Suppose that $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \phi$. As $0 \in I \cap J$, we clearly have $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \blacklozenge_J \Diamond_{J'} \phi$, that is, $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \phi'$.
- $\phi = \psi \wedge \chi$, $\phi = \exists x.\, \psi$, $\phi = \bullet_I \psi$, $\phi = \bigcirc_I \psi$, $\phi = \psi \,\mathsf{S}_I\, \chi$, or $\phi = \psi \,\mathsf{U}_I\, \chi$. These cases follow directly from the induction hypotheses. We only present the case $\phi =$

$\psi \mathsf{S}_I \chi$. We have $\phi^w = \psi^w \mathsf{S}_I \chi^w$. Let $(\bar{\mathcal{D}}, \bar{\tau})$ be a temporal structure, v a valuation, and i a time point. Suppose that $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \phi$. Then there is a $j \leq i$ with $\tau_i - \tau_j \in I$ such that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \chi$ and $(\bar{\mathcal{D}}, \bar{\tau}, v, k) \models \psi$ for any $k \in [i+1, j+1)$. Using the induction hypotheses for ψ and χ, we obtain that $(\bar{\mathcal{D}}, \bar{\tau}, v, j) \models \chi^w$ and $(\bar{\mathcal{D}}, \bar{\tau}, v, k) \models \psi^w$ for any $k \in [i+1, j+1)$. Hence $(\bar{\mathcal{D}}, \bar{\tau}, v, i) \models \phi^w$.

The proof of the dual case, that is, that the formula $\phi^s \rightarrow \phi$ is valid, is similar. It is based on the remark that the formula $(\neg \blacklozenge_J \Diamond_{J'} r(t_1, \ldots, t_{\iota(r)})) \rightarrow \neg r(t_1, \ldots, t_{\iota(r)})$ is valid.

Finally, we prove statement (1). Statement (2) is similar. Let $(\bar{\mathcal{C}}, \bar{\kappa})$ be the collapse of two temporal structures $(\bar{\mathcal{D}}^1, \bar{\tau}^1)$ and $(\bar{\mathcal{D}}^2, \bar{\tau}^2)$. Suppose that ϕ^s is collapse-sufficient and that $(\bar{\mathcal{C}}, \bar{\kappa}, v, 0) \models \phi^s$, for some arbitrary valuation v. It follows that $(\bar{\mathcal{D}}, \bar{\tau}, v, 0) \models \phi^s$ for any $(\bar{\mathcal{D}}, \bar{\tau}) \in (\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$. As $\phi^s \rightarrow \phi$ is valid, we have that $(\bar{\mathcal{D}}, \bar{\tau}, v, 0) \models \phi$, for any $(\bar{\mathcal{D}}, \bar{\tau}) \in (\bar{\mathcal{D}}^1, \bar{\tau}^1) \between (\bar{\mathcal{D}}^2, \bar{\tau}^2)$.

E. Additional Details on Practical Experience

In this section, we describe in detail all our policies in Nokia's Data-collection Campaign, their MFOTL formalization, and the resources needed for monitoring.

Policies in Nokia's Data-collection Campaign. We first describe the domain and relations used for formalizing the policies. Then we describe the policies in natural language and give their formalization.

The domain, that is, the values that can occur as a parameter of a system actions are the databases db1, db2, db3, all database accounts, all data identifiers, the constant unknown, all possible names for the synchronization scripts, all possible subversion URLs, all possible subversion revision numbers, and the subversion status values latest, old, mod, and nosvn.

We represent actions in the system as elements in relations. We explain now the relations used. The elements of the relations for the predicates *select*, *insert*, *delete*, and *update* correspond to database operations with equally-named SQL commands. The parameters are the user executing the operation, the name of the database, and an identifier of the involved data. The elements in the relations for the predicates *start* and *stop* indicate the starting and finishing of a synchronization script and contain the name of the script as their only parameter. After the script script1 starts, it logs details about its SVN status in the relations for the predicate *svn*. The parameters are the name of the script, its SVN status determined by the command `svn status -u -v`, the SVN URL, and the SVN revision number. Possible values for SVN status are latest for the latest version, old for an older version, mod for a locally modified version, and nosvn if the script has not been checked out from the subversion repository. The relations for the predicate *commit* represent committing a new script version into the subversion repository. The parameters are the SVN URL and revision number.

Table IV. Policy Formalizations in MFOTL

policy	MFOTL formalization
delete	$\square\, \forall user.\, \forall data.\, delete(user, \mathsf{db2}, data) \rightarrow user \approx \mathsf{script2}$
insert	$\square\, \forall user.\, \forall data.\, insert(user, \mathsf{db2}, data) \rightarrow user \approx \mathsf{script1}$
select	$\square\, \forall user.\, \forall data.\, select(user, \mathsf{db2}, data) \rightarrow$ $user \approx \mathsf{script1} \vee user \approx \mathsf{script2} \vee user \approx \mathsf{triggers}$
update	$\square\, \forall user.\, \forall data.\, \neg update(user, \mathsf{db2}, data)$
script1	$\square\, \forall db.\, \forall data.\, select(\mathsf{script1}, db, data) \vee insert(\mathsf{script1}, db, data) \vee$ $delete(\mathsf{script1}, db, data) \vee update(\mathsf{script1}, db, data) \rightarrow$ $((\neg \blacklozenge_{[0,1s)} \Diamond_{[0,1s)} end(\mathsf{script1}))\ \mathsf{S}\ (\blacklozenge_{[0,1s)} \Diamond_{[0,1s)} start(\mathsf{script1}))) \vee$ $\blacklozenge_{[0,1s)} \Diamond_{[0,1s)} end(\mathsf{script1})$
runtime	$\square\, \forall script.\, start(script) \rightarrow (\neg \blacklozenge_{[0,1s)} \Diamond_{[0,1s)} end(script)) \wedge \Diamond_{[1s,6h)} end(script)$
svn	$\square\, \forall script.\, start(script) \rightarrow \blacklozenge_{[0,1s)} \Diamond_{[0,10s)} \exists url.\, \exists rev.\, svn(script, \mathsf{latest}, url, rev)$
svn2	$\square\, \forall script.\, \forall status.\, \forall url.\, \forall rev.\, svn(script, status, url, rev) \rightarrow$ $\blacksquare_{[1s,\infty)} (commit(url, rev') \rightarrow rev' \leq rev)$
ins-1-2	$\square\, \forall user.\, \forall data.\, insert(user, \mathsf{db1}, data) \wedge data \neq \mathsf{unknown} \rightarrow$ $\blacklozenge_{[0,1s)} \Diamond_{[0,30h]} \exists user'.\, insert(user', \mathsf{db2}, data) \vee delete(user', \mathsf{db1}, data)$
ins-2-3	$\square\, \forall user.\, \forall data.\, insert(user, \mathsf{db2}, data) \wedge data \neq \mathsf{unknown} \rightarrow$ $\blacklozenge_{[0,1s)} \Diamond_{[0,60s)} \exists user'.\, insert(user', \mathsf{db3}, data)$
ins-3-2	$\square\, \forall user.\, \forall data.\, insert(user, \mathsf{db3}, data) \wedge data \neq \mathsf{unknown} \rightarrow$ $\blacklozenge_{[0,60s)} \Diamond_{[0,1s)} \exists user'.\, insert(user', \mathsf{db2}, data)$
del-1-2	$\square\, \forall user.\, \forall data.\, delete(user, \mathsf{db1}, data) \wedge data \neq \mathsf{unknown} \rightarrow$ $(\blacklozenge_{[0,1s)} \Diamond_{[0,30h)} \exists user'.\, delete(user', \mathsf{db2}, data)) \vee$ $((\Diamond_{[0,1s)} \blacklozenge_{[0,30h)} \exists user'.\, insert(user', \mathsf{db1}, data)) \wedge$ $(\blacksquare_{[0,30h)} \square_{[0,30h)} \neg \exists user'.\, insert(user', \mathsf{db2}, data)))$
del-2-3	$\square\, \forall user.\, \forall data.\, delete(user, \mathsf{db2}, data) \wedge data \neq \mathsf{unknown} \rightarrow$ $\blacklozenge_{[0,1s)} \Diamond_{[0,60s)} \exists user'.\, delete(user', \mathsf{db3}, data)$
del-3-2	$\square\, \forall user.\, \forall data.\, delete(user, \mathsf{db3}, data) \wedge data \neq \mathsf{unknown} \rightarrow$ $\blacklozenge_{[0,60s)} \Diamond_{[0,1s)} \exists user'.\, delete(user', \mathsf{db2}, data)$

In the following, we informally state the policies in natural language and for the more involved policies, we provide additional explanations. The MFOTL formalization of the policies is shown in Table IV. The policies are:

- *delete:* Only user script2, representing the synchronization script script2, may delete data in db2 by executing the SQL delete command.
- *insert:* Only user script1, representing the synchronization script script1, may insert data in db2 by executing the SQL insert command.
- *select:* Only a limited set of users (script1, script2, triggers) may read data from db2 by executing the SQL select command.
- *update:* No SQL update commands are allowed in db2, only insertion and deletions.
- *script1:* Database operations may be executed under the user account script1 only while the script script1 is running. The motivation for this policy is that the account script1 should only be used by the script, so if the account is used while the script is not running, the account may have been compromised. The database operation can happen while the script is running, including the boundaries. That is, the time points when an operation happens and when the script starts or ends may have equal time stamps. The semantics of the S operator includes the script start, but excludes the script end. Therefore, the script end is allowed with the additional disjunct at the end of the formula.
- *runtime:* The synchronization scripts must run for at least 1 second and for no longer than 6 hours.
- *svn, svn2:* The synchronization scripts are maintained in an SVN repository. We require that when started, the synchronization scripts are the latest version available

Table V. Monitor Performance — Running Times / Memory Usage

policy	log 1	log 2	log 3	log 4	log 5	log 6	log 7	log 8	log 9
delete	10 s / 4 MB	7 s / 4 MB	7 s / 4 MB	6 s / 4 MB	5 s / 4 MB	4 s / 4 MB	4 s / 4 MB	6 s / 4 MB	6 s / 4 MB
insert	13 s / 4 MB	8 s / 4 MB	10 s / 4 MB	8 s / 4 MB	6 s / 4 MB	5 s / 4 MB	5 s / 4 MB	8 s / 4 MB	8 s / 4 MB
select	10 s / 4 MB	7 s / 4 MB	7 s / 4 MB	6 s / 4 MB	5 s / 4 MB	4 s / 4 MB	4 s / 4 MB	7 s / 4 MB	6 s / 4 MB
update	10 s / 4 MB	6 s / 4 MB	8 s / 4 MB	6 s / 4 MB	4 s / 4 MB	4 s / 4 MB	4 s / 4 MB	6 s / 4 MB	7 s / 4 MB
script1	14 s / 4 MB	9 s / 4 MB	10 s / 4 MB	9 s / 4 MB	6 s / 4 MB	6 s / 4 MB	5 s / 4 MB	9 s / 4 MB	8 s / 4 MB
runtime	12 s / 9 MB	8 s / 9 MB	8 s / 6 MB	8 s / 9 MB	5 s / 7 MB	5 s / 7 MB	4 s / 7 MB	7 s / 20 MB	7 s / 21 MB
svn	10 s / 4 MB	7 s / 4 MB	7 s / 4 MB	6 s / 4 MB	5 s / 4 MB	4 s / 4 MB	4 s / 4 MB	7 s / 4 MB	7 s / 4 MB
svn2	12 s / 16 MB	9 s / 16 MB	9 s / 16 MB	9 s / 16 MB	7 s / 16 MB	6 s / 16 MB	6 s / 16 MB	8 s / 16 MB	8 s / 16 MB
ins-1-2	231 m / 161 MB	44 m / 103 MB	67 m / 107 MB	24 m / 102 MB	9 m / 71 MB	5 m / 65 MB	3 m / 57 MB	73 m / 115 MB	48 m / 111 MB
ins-2-3	9 m / 8 MB	3 m / 7 MB	5 m / 8 MB	4 m / 8 MB	2 m / 8 MB	2 m / 7 MB	1 m / 7 MB	2 m / 8 MB	1 m / 6 MB
ins-3-2	7 m / 5 MB	3 m / 5 MB	5 m / 5 MB	4 m / 6 MB	2 m / 5 MB	2 m / 5 MB	1 m / 5 MB	2 m / 5 MB	1 m / 5 MB
del-1-2	24 s / 176 MB	16 s / 139 MB	13 s / 87 MB	11 s / 79 MB	8 s / 58 MB	7 s / 53 MB	12 s / 111 MB	21 s / 184 MB	11 s / 102 MB
del-2-3	10 s / 4 MB	6 s / 4 MB	7 s / 4 MB	6 s / 4 MB	5 s / 4 MB	4 s / 4 MB	4 s / 4 MB	6 s / 4 MB	6 s / 4 MB
del-3-2	10 s / 4 MB	6 s / 4 MB	7 s / 4 MB	6 s / 4 MB	4 s / 4 MB	4 s / 4 MB	4 s / 4 MB	6 s / 4 MB	6 s / 4 MB

in the repository (largest SVN revision number). We use two different formalizations, *svn* and *svn2*. The policy *svn* uses the status parameter of the relation *svn*. The policy *svn2* compares the revision number parameter of the relation *svn* with the committed revision numbers obtained from the subversion log via the *commit* relation. Computing the latest revision number is done by the logging mechanism for the policy *svn*, but by the monitor for the policy *svn2*. Monitoring both policies allows us to compare how efficiently the monitor copes with these different formalizations and to observe the impact of offloading the monitor by doing pre-computations in the logging mechanisms.

- *ins-*:* Data uploaded by the phone into db1 must be propagated to all databases. In particular, *ins-1-2* requires that data uploaded into db1 must be inserted into db2 within 30 hours after the upload, unless it has been deleted from db1 in between. Furthermore, *ins-2-3* and *ins-3-2* require that data may be inserted into db2 iff it is inserted into db3 within 1 minute. The time limit from db1 to db2 is 30 hours because the synchronization scripts run once every 24 hours and can run for up to 6 hours. The time limit from db2 to db3 is only 60 seconds as this synchronization is implemented by database triggers that start immediately upon a change in db2. Note that these policies require propagation of new data between db2 and db3 in both directions. However, between db1 and db2 only one direction is required. The reason is the incomplete logging for db1.
- *del-*:* Data deleted from db1 must be consistently deleted from all databases. The policies *del-2-3* and *del-3-2* are analogous to the policies *ins-2-3* and *ins-3-2*, respectively. The formalization of the policy *del-1-2* is more involved: If data is deleted from db1, then this data must also be deleted from db2 within 30 hours. However, if the data has just been uploaded to db1 and not yet propagated to db2, then it simply should not be propagated to db2 in the future either. Since the propagation would happen in at most 30 hours, we can simply consider the past and the future 30 hours to determine whether data has been and will be propagated to db2 or not.

Monitor Performance. Table V shows the monitor's running times and memory usage for all policies in Table IV and all log files in Table II.

Our reason for splitting the available stream of logged actions into smaller chunks (i.e., log files) is to evaluate our monitor on different data sets with different characteristics. Each of our chunks corresponds to a time span of approximately 24 hours. We point out that monitoring such chunks separately may reveal different violations than monitoring the whole stream of actions. This is because, policy conformance at a time point may depend on actions that have been logged in another (timewise subsequent or prior) chunk, as the time window of a temporal operator may overpass the time span of a chunk. Except for the policy *del-1-2*, all policy violations on the whole stream are also detected on a chunk. However, due to splitting, additional violations may be reported. We were not concerned about these issues, as our main focus was on evaluating the performance of the monitor. Moreover, we have manually checked that all violations reported in Section IV are indeed violations on the whole stream.