

Explainable Online Monitoring of Metric First-Order Temporal Logic

Leonardo Lima¹, Jonathan Julián Huerta y Munive¹, and Dmitriy Traytel¹

Department of Computer Science, University of Copenhagen, Copenhagen, Denmark

Abstract. Metric first-order temporal logic (MFOTL) is an expressive formalism for specifying temporal and data-dependent constraints on streams of time-stamped, data-carrying events. It serves as the specification language of several runtime monitors. These monitors input an MFOTL formula and an event stream prefix and output satisfying assignments to the formula’s free variables. For complex formulas, it may be unclear why a certain assignment is output. We propose an approach that accompanies assignments with detailed explanations, in the form of proof trees. We develop a new monitor that outputs such explanations. Our tool incorporates a formally verified checker that certifies the explanations and a visualization that allows users to interactively explore and understand the outputs.

1 Introduction

Runtime monitoring is concerned with the analysis of events produced by a system during its execution. An online monitor searches for given complex patterns in event streams, processing the stream incrementally, i.e., one event at a time. If it finds a pattern match, the monitor outputs a verdict to its user. The nature of a verdict depends on both the monitor and its pattern specification language. For propositional specification languages, such as metric temporal logic (MTL) [6, 21], typical verdicts are streams of Booleans [8, 28, 31], where each Boolean signifies the presence or the absence of a pattern match, i.e., the satisfaction or violation of the MTL formula at every position in the input stream.

Users might find Boolean outputs difficult to interpret, especially when complex patterns like nesting temporal operators are involved. In particular, Boolean verdicts give no insight into how monitors produce them—we have to trust their correctness. Even when assuming infallible monitors, verdict justifications can help us to ensure that we expressed correctly our intentions in the specification and, e.g., that it is not vacuously true [23].

Lima et al. [25] propose the use of richer verdicts in an MTL monitor. Specifically, they use proof trees in a dedicated proof system resembling MTL’s semantics to explain why a formula is satisfied or violated. They develop the EXPLANATOR2 monitor, which outputs a stream of size-minimal proof trees, and design an interactive graphical user interface for exploring and understanding these informative verdicts. In addition, they formally verify, in the Isabelle/HOL proof assistant, a proof tree checker certifying that their proof system rules were correctly applied. Thus proof tree verdicts serve a two-fold purpose: as machine-checkable certificates and human-readable explanations.

In this work, we significantly widen the scope of the “proof tree verdicts” approach. We provide certifiable and explainable monitoring verdicts for metric first-order temporal logic (MFOTL) [14] with bounded future operators and without equality between variables. MFOTL extends MTL with data parameters and first-order quantification and is an

expressive formalism with many practical applications [7, 10–13]. We extend Lima et al.’s MTL proof system to MFOTL with the expected rules for quantifiers (Section 2): e.g., the universally quantified formula $\forall x. \alpha$ is satisfied if α with x replaced by d is satisfied *for all* domain values d . The key challenge here is that the domain is typically infinite, which results in the above proof rule for \forall to be infinitely branching. This is problematic because it is unclear how to validate a correct application of the \forall rule in a proof tree checker.

A crucial observation is that without equality between variables, proof trees cannot distinguish values outside of the *active domain*, i.e., the finite set of data values from the monitored event stream prefix and from the formula’s constants. Thus, the active domain’s size plus one bounds the number of choices for d requiring *different* proof trees, and we can reuse them—with the extra “plus one” representing values outside the active domain. Thus, to represent the \forall rule it suffices to store a finite partition of the domain and one subproof for each part. We obtain finite proof objects, develop a checker for them, and formally verify the checker’s correctness in Isabelle/HOL (Section 3).

The proof system explains how to deal with closed MFOTL formulas. A Boolean verdict for a formula with free variables only makes sense relative to a variable assignment. Hence, traditional MFOTL monitors compute sets of satisfying variable assignments [15, 30] instead of Boolean verdicts. In our setting, an explanation for a formula with free variables must provide a proof tree for any variable assignment (satisfying or violating). For infinite domains, there are infinitely many assignments, but the same idea that worked for quantifiers comes to our rescue: it suffices to consider a finite partition of the domain for each variable. Inspired by binary decision diagrams (BDDs) [16], we organize the partitions for different variables hierarchically in *partitioned decision trees (PDTs)*. PDTs are trees where each leaf stores a generic data item and each node (representing a variable) branches on a finite partition of the domain (Section 4). The partitions may change from one node to the other. PDTs can be compacted (or reduced in BDD terminology).

We thus have arrived at our notion of *explainable verdicts* for MFOTL formulas: PDTs whose leaves are proof objects. We extend our verified checker from proof objects to such verdicts and Lima et al.’s algorithm for MTL [25] to MFOTL (Section 5). Our algorithm extension is modular in the sense that it merely adds a layer of PDTs, but keeps Lima et al.’s algorithms for temporal operators unchanged. We implement the extended algorithm in a new monitor and also extend Lima et al.’s interactive visualization of proof objects. We demonstrate the effectiveness of our new tool on MFOTL policies from the literature (Section 6). In summary, we make the following contributions:

- We develop a proof system for MFOTL satisfaction and violation at a time-point for a given event stream and verify its soundness and completeness in Isabelle/HOL.
- We finitely represent our proof system’s proof trees and formally verify a checker for them. The key idea is that finite partitions of infinite domains are sufficient.
- We design partitioned decision trees (PDTs) to represent functions from variable assignments to generic data items in a way that enables sharing and compression.
- We develop an algorithm computing explanations: PDTs with proof objects as leaves. We implement the algorithm in a new monitor, along with an interactive visualization of explanations and integrated with the verified proof tree checker for certification.

Our tool, called WHYMON, is publicly available [2].

$v, i \models tt$	$v, i \models \exists x. \alpha$ iff $v[x \mapsto d], i \models \alpha$ for some $d \in \mathbb{D}$
$v, i \not\models ff$	$v, i \models \forall x. \alpha$ iff $v[x \mapsto d], i \models \alpha$ for all $d \in \mathbb{D}$
$v, i \models p(\bar{i})$ iff $p(\llbracket \bar{i} \rrbracket_v) \in \Gamma_i$	$v, i \models \bullet_I \alpha$ iff $i > 0$, $\tau_i - \tau_{i-1} \in I$, and $v, i-1 \models \alpha$
$v, i \models x \approx c$ iff $v(x) = c$	$v, i \models \circ_I \alpha$ iff $\tau_{i+1} - \tau_i \in I$ and $v, i+1 \models \alpha$
$v, i \models \neg \alpha$ iff $v, i \not\models \alpha$	$v, i \models \blacklozenge_I \alpha$ iff $v, j \models \beta$ for some $j \leq i$ with $\tau_i - \tau_j \in I$
$v, i \models \alpha \wedge \beta$ iff $v, i \models \alpha$ and $v, i \models \beta$	$v, i \models \blacklozenge_I \alpha$ iff $v, j \models \beta$ for some $j \geq i$ with $\tau_j - \tau_i \in I$
$v, i \models \alpha \vee \beta$ iff $v, i \models \alpha$ or $v, i \models \beta$	$v, i \models \blacksquare_I \alpha$ iff $v, j \models \beta$ for all $j \leq i$ with $\tau_i - \tau_j \in I$
$v, i \models \alpha \rightarrow \beta$ iff $v, i \not\models \alpha$ or $v, i \models \beta$	$v, i \models \square_I \alpha$ iff $v, j \models \beta$ for all $j \geq i$ with $\tau_j - \tau_i \in I$
$v, i \models \alpha \mathcal{S}_I \beta$ iff $v, j \models \beta$ for some $j \leq i$ with $\tau_i - \tau_j \in I$ and $v, k \models \alpha$ for all $j < k \leq i$	
$v, i \models \alpha \mathcal{U}_I \beta$ iff $v, j \models \beta$ for some $j \geq i$ with $\tau_j - \tau_i \in I$ and $v, k \models \alpha$ for all $i \leq k < j$	

Fig. 1: Semantics of MFOTL for a fixed stream $\sigma = \langle \tau_i, \Gamma_i \rangle_{i \in \mathbb{N}}$.

Further Related Work Lima et al.’s work [25], which we extend, is based on the work by Basin et al. [9] that employed proof trees as explanations in the context of understanding counterexamples of LTL model checkers. We refer to these works for a discussion of related proof systems for propositional temporal logics and regular expressions.

In the first-order monitoring setting, we are on unexplored territory with verdicts that go beyond satisfying assignments. Nonetheless our work incorporates ideas from existing first-order monitors. Most closely related is Havelund et al.’s DEJAVU monitor [18], which uses BDDs to represent sets of satisfying assignments. Our work generalizes BDDs to branching over partitions of the domain and storing generic data (e.g., proof objects) instead of Booleans in the leaves. In addition, the DEJAVU authors make use of the fact that without equality between variables the formula’s satisfaction cannot be influenced by different values outside the active domain. We generalize this observation so that not only the satisfaction but rather entire proof trees can be reused when exchanging values outside the active domain. Finally, DEJAVU only supports past temporal operators and closed formulas, whereas our algorithm supports both past and future operators and free variables.

Havelund et al.’s key observation fails for equalities between variables. For example, the formula $x \approx y \rightarrow p(x, y)$ is satisfied for any pair of distinct values $c \neq d$ outside of the predicate p ’s interpretation, but it is violated if we pick the same value c for both x and y . A classic result by Ailamazyan et al. [5, 19] shows that for the relational calculus (MFOTL without temporal operators) it suffices to distinguish a finite number of equivalence classes of values outside of the active domain. While it is conceivable that this result generalizes to MFOTL with equality, we leave this generalization as future work.

The MFOTL monitor MonPoly [14, 15] and its formally verified counterpart VeriMon [30] output streams of satisfying assignments for formulas in the so-called monitorable fragment. The fragment ensures that all subformulas always evaluate to finite sets of satisfying assignments. Our monitor does not suffer from this limitation; even more it returns all satisfying and violating assignments (labeled and explained as such).

Outside of first-order monitoring, our visualization takes some inspiration from the stream runtime verification tool TeSSLa [24], which can provide output for all intermediate streams. Similarly, we provide output for all subformulas, but our proof trees allow us to focus on the relevant dependencies between a formula and its subformulas.

Metric first-order temporal logic (MFOTL) We recall MFOTL’s syntax and semantics. We fix an infinite domain \mathbb{D} (e.g., containing integers and strings). Terms $t \in \mathbb{T}$ are either variables $x, y, z \in \mathbb{V}$ or constants $c, d \in \mathbb{D}$. Overlines indicate lists (finite sequences), e.g.,

if t is a term, then \bar{t} is a list of terms. The grammar below specifies MFOTL's syntax, where $p \in \mathbb{E}$ is a predicate name (e.g., a string) and $I \in \mathbb{I} \subseteq 2^{\mathbb{N}}$ is a non-empty interval.

$$\alpha ::= tt \mid \text{ff} \mid p(\bar{t}) \mid x \approx c \mid \neg \alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \rightarrow \alpha \mid \exists x. \alpha \mid \forall x. \alpha \mid \\ \bullet_I \alpha \mid \circ_I \alpha \mid \blacklozenge_I \alpha \mid \blacklozenge_I \alpha \mid \blacksquare_I \alpha \mid \square_I \alpha \mid \alpha \mathcal{S}_I \alpha \mid \alpha \mathcal{U}_I \alpha$$

Besides the first-order logic operators, the syntax includes the past \bullet (previous), \blacklozenge (once), \blacksquare (historically), \mathcal{S} (since) and future \circ (next), \blacklozenge (eventually), \square (always), \mathcal{U} (until) temporal operators. We use \wedge for universal and \vee for existential quantification at the metalanguage level to avoid confusion with MFOTL formulas. We also use common interval notation $[a, b) = \{n \mid a \leq n < b\}$ or $[a, c] = \{n \mid a \leq n \leq c\}$, for $a, c \in \mathbb{N}$ and $b \in \mathbb{N} \cup \{\infty\}$, and omit intervals when $I = [0, \infty) = \mathbb{N}$. Whenever we write $[a, c]$, we exclusively denote the range $[a, \dots, c]$ (rather than the two element sequence $[a, c]$). Furthermore, we assume that the future operators (\blacklozenge , \square , and \mathcal{U}) intervals are finite (also called bounded). We write $a + I$ for $\{a + x \mid x \in I\}$ and $a \mathcal{R} I$ for $\bigwedge x \in I. a \mathcal{R} x$ (where $\mathcal{R} \in \{<, \leq, >, \geq\}$). We interpret formulas over streams σ : infinite sequences of time-stamped sets of events $\sigma = \langle \tau_i, \Gamma_i \rangle_{i \in \mathbb{N}}$. We call the indices $i \in \mathbb{N}$ time-points, so that Γ_i is the set of events and $\tau_i \in \mathbb{N}$ is the time-stamp at time-point i . The time-stamps τ_i must be monotone ($\bigwedge i. j. i \leq j \longrightarrow \tau_i \leq \tau_j$) and eventually increasing ($\bigwedge \tau. \bigvee i. \tau_i > \tau$). Each event has the form $p(d_1, \dots, d_n)$ where p is the event name and $d_i \in \mathbb{D}$. Given a total assignment v mapping variables to values in \mathbb{D} , we define $\llbracket x \rrbracket_v = v(x)$ and $\llbracket c \rrbracket_v = c$. The notation $\llbracket \bar{t} \rrbracket_v = \bar{c}$ lifts this operation to lists of terms. We define the satisfaction relation $v, i \models_{\sigma} \alpha$ in the usual way (Figure 1). Finally, the *earliest time-point* $\text{ETP}_{\sigma}(\tau)$ of $\tau \in \mathbb{N}$ on σ is the smallest time-point i such that $\tau_i \geq \tau$. Analogously, the *latest time-point* $\text{LTP}_{\sigma}(\tau)$ of $\tau \geq \tau_0$ on σ is the largest i such that $\tau_i \leq \tau$. We omit the stream σ (e.g., \models , $\text{ETP}(\tau)$ and $\text{LTP}(\tau)$) if it is clear from the context.

2 Proof System

We introduce a local proof system for MFOTL (Figure 2). "Local" means here that the proof system does not talk about satisfiability in general, but rather about the formula's satisfaction or violation for a fixed stream, assignment, and time-point.

Our proof system consists of two mutually dependent judgments, \vdash_{σ}^{+} and \vdash_{σ}^{-} (again σ is omitted when clear), that characterize a formula's satisfaction $v, i \vdash_{\sigma}^{+} \alpha$ and violation $v, i \vdash_{\sigma}^{-} \alpha$ relations for assignment v , stream σ , and time-point i . The rules of our proof system closely follow the MFOTL semantics (Figure 1) and extend the proof system used by Lima et al. [25] with assignments (that are mostly passed around without modification) and the rules for quantifiers (which modify the assignments). The rules for atomic predicates and Boolean constants and operators are self-explanatory: e.g., predicates are satisfied if a matching event is present in the trace; a conjunction is satisfied if both conjuncts are satisfied; a conjunction is violated if either of the conjuncts is violated.

The rule \exists^{+} states that for v to satisfy $\exists x. \alpha$ at i , it suffices to provide a domain value d such that the updated assignment $v[x \mapsto d]$ setting x to d satisfies α at i . Conversely, \exists^{-} asserts that the violation of $\exists x. \alpha$ under v at i requires showing that all domain values make $v[x \mapsto d]$ violate α at i . Since the universal quantifier is dual to the existential one, the rules \forall^{-} and \forall^{+} exchange the relations \vdash_{σ}^{+} and \vdash_{σ}^{-} compared to \exists^{+} and \exists^{-} .

$$\begin{array}{c}
 \frac{}{v, i \vdash^+ tt} tt^+ \quad \frac{}{v, i \vdash^- ff} ff^- \quad \frac{p(\llbracket \bar{t} \rrbracket_v) \in \Gamma_i}{v, i \vdash^+ p(\bar{t})} p^+ \quad \frac{p(\llbracket \bar{t} \rrbracket_v) \notin \Gamma_i}{v, i \vdash^- p(\bar{t})} p^- \quad \frac{v, i \vdash^+ \alpha \quad v, i \vdash^- \beta}{v, i \vdash^- \alpha \rightarrow \beta} \rightarrow^- \\
 \frac{v, i \vdash^- \alpha}{v, i \vdash^+ \neg \alpha} \neg^+ \quad \frac{v, i \vdash^- \alpha}{v, i \vdash^- \alpha \wedge \beta} \wedge_L^- \quad \frac{v, i \vdash^- \beta}{v, i \vdash^- \alpha \wedge \beta} \wedge_R^- \quad \frac{v, i \vdash^+ \alpha \quad v, i \vdash^+ \beta}{v, i \vdash^+ \alpha \wedge \beta} \wedge^+ \quad \frac{v, i \vdash^- \alpha}{v, i \vdash^+ \alpha \rightarrow \beta} \rightarrow_L^+ \\
 \frac{v, i \vdash^+ \alpha}{v, i \vdash^- \neg \alpha} \neg^- \quad \frac{v, i \vdash^+ \alpha}{v, i \vdash^+ \alpha \vee \beta} \vee_L^+ \quad \frac{v, i \vdash^+ \beta}{v, i \vdash^+ \alpha \vee \beta} \vee_R^+ \quad \frac{v, i \vdash^- \alpha \quad v, i \vdash^- \beta}{v, i \vdash^- \alpha \vee \beta} \vee^- \quad \frac{v, i \vdash^+ \beta}{v, i \vdash^+ \alpha \rightarrow \beta} \rightarrow_R^+ \\
 \frac{v[x \mapsto d], i \vdash^+ \alpha}{v, i \vdash^+ \exists x. \alpha} \exists^+ \quad \frac{\wedge d. v[x \mapsto d], i \vdash^+ \alpha}{v, i \vdash^+ \forall x. \alpha} \forall^+ \quad \frac{\wedge d. v[x \mapsto d], i \vdash^- \alpha}{v, i \vdash^- \exists x. \alpha} \exists^- \quad \frac{v[x \mapsto d], i \vdash^- \alpha}{v, i \vdash^- \forall x. \alpha} \forall^- \\
 \frac{}{v, 0 \vdash^- \bullet_I \alpha} \bullet_0^- \quad \frac{i > 0 \quad \tau_i - \tau_{i-1} < I}{v, i \vdash^- \bullet_I \alpha} \bullet_{<I}^- \quad \frac{i > 0 \quad \tau_i - \tau_{i-1} > I}{v, i \vdash^- \bullet_I \alpha} \bullet_{>I}^- \quad \frac{i > 0 \quad v, i-1 \vdash^- \alpha}{v, i \vdash^- \bullet_I \alpha} \bullet^- \\
 \frac{\tau_{i+1} - \tau_i \in I \quad v, i+1 \vdash^+ \alpha}{v, i \vdash^+ \circ_I \alpha} \circ^+ \quad \frac{\tau_{i+1} - \tau_i < I}{v, i \vdash^- \circ_I \alpha} \circ_{<I}^- \quad \frac{\tau_{i+1} - \tau_i > I}{v, i \vdash^- \circ_I \alpha} \circ_{>I}^- \quad \frac{v, i+1 \vdash^- \alpha}{v, i \vdash^- \circ_I \alpha} \circ^- \\
 \frac{j \leq i \quad \tau_i - \tau_j \in I \quad v, j \vdash^+ \alpha}{v, i \vdash^+ \blacklozenge_I \alpha} \blacklozenge^+ \quad \frac{\tau_i < \tau_0 + I}{v, i \vdash^- \blacklozenge_I \alpha} \blacklozenge_{<I}^- \quad \frac{j \geq i \quad \tau_j - \tau_i \in I \quad v, j \vdash^+ \alpha}{v, i \vdash^+ \blacklozenge_I \alpha} \blacklozenge^+ \\
 \frac{\tau_i \geq \tau_0 + I \quad \wedge j \in [E_i^p(I), L_i^p(I)]. v, j \vdash^- \alpha}{v, i \vdash^- \blacklozenge_I \alpha} \blacklozenge^- \quad \frac{\wedge j \in [E_i^f(I), L_i^f(I)]. v, j \vdash^- \beta}{v, i \vdash^- \blacklozenge_I \alpha} \blacklozenge^- \\
 \frac{j \leq i \quad \tau_i - \tau_j \in I \quad v, j \vdash^- \alpha}{v, i \vdash^- \blacksquare_I \alpha} \blacksquare^- \quad \frac{\tau_i < \tau_0 + I}{v, i \vdash^+ \blacksquare_I \alpha} \blacksquare_{<I}^+ \quad \frac{j \geq i \quad \tau_j - \tau_i \in I \quad v, j \vdash^- \alpha}{v, i \vdash^+ \square_I \alpha} \square^- \\
 \frac{\tau_i \geq \tau_0 + I \quad \wedge j \in [E_i^p(I), L_i^p(I)]. v, j \vdash^+ \alpha}{v, i \vdash^+ \blacksquare_I \alpha} \blacksquare^+ \quad \frac{\wedge j \in [E_i^f(I), L_i^f(I)]. v, j \vdash^+ \beta}{v, i \vdash^+ \square_I \alpha} \square^+ \\
 \frac{j \leq i \quad \tau_i - \tau_j \in I \quad v, j \vdash^+ \beta \quad \wedge k \in (j, i]. v, k \vdash^+ \alpha}{v, i \vdash^+ \alpha \mathcal{S}_I \beta} \mathcal{S}^+ \quad \frac{i > 0 \quad \tau_i - \tau_{i-1} \in I \quad v, i-1 \vdash^+ \alpha}{v, i \vdash^+ \bullet_I \alpha} \bullet^+ \\
 \frac{i \leq j \quad \tau_j - \tau_i \in I \quad v, j \vdash^+ \beta \quad \forall k \in [i, j]. v, k \vdash^+ \alpha}{v, i \vdash^+ \alpha \mathcal{U}_I \beta} \mathcal{U}^+ \quad \frac{\tau_i < \tau_0 + I}{v, i \vdash^- \alpha \mathcal{S}_I \beta} \mathcal{S}_{<I}^- \quad \frac{v(x) = c}{v, i \vdash^+ x \approx c} \approx^+ \\
 \frac{\tau_i \geq \tau_0 + I \quad \wedge k \in [E_i^p(I), L_i^p(I)]. v, k \vdash^- \beta}{v, i \vdash^- \alpha \mathcal{S}_I \beta} \mathcal{S}_{\infty}^- \quad \frac{\wedge k \in [E_i^f(I), L_i^f(I)]. v, k \vdash^- \beta}{v, i \vdash^- \alpha \mathcal{U}_I \beta} \mathcal{U}_{\infty}^- \\
 \frac{E_i^p(I) \leq j \quad j \leq i \quad \tau_i \geq \tau_0 + I \quad v, j \vdash^- \alpha \quad \wedge k \in [j, L_i^p(I)]. v, k \vdash^- \beta}{v, i \vdash^- \alpha \mathcal{S}_I \beta} \mathcal{S}^- \\
 \frac{i \leq j \quad j < L_i^f(I) \quad v, j \vdash^- \alpha \quad \forall k \in [E_i^f(I), j]. v, k \vdash^- \beta}{v, i \vdash^- \alpha \mathcal{U}_I \beta} \mathcal{U}^- \quad \frac{v(x) \neq c}{v, i \vdash^- x \approx c} \approx^-
 \end{array}$$

 Fig. 2: Local proof system for MFOTL on a fixed stream $\sigma = \langle \tau_i, \Gamma_i \rangle_{i \in \mathbb{N}}$.

The rules \blacklozenge^+ and \blacklozenge^- are mere restatements of the MFOTL semantics. Since the operators \blacksquare_I and \square_I are respectively dual to \blacklozenge_I and \blacklozenge_I , their violation rules \blacksquare^- and \square^- once again exchange \vdash_σ^+ and \vdash_σ^- compared to \blacklozenge^+ and \blacklozenge^- . The rule $\blacksquare_{<I}^+$ accounts for the vacuous truth of the operator \blacksquare_I near the start of the stream (when no time-points fall within the interval I). Dually, the rule $\blacklozenge_{<I}^-$ asserts the violation of \blacklozenge_I near the start of the stream. The remaining rules \blacklozenge^- , \blacklozenge^+ , \blacksquare^+ , and \square^+ use notation $E_i^p(I)$, $L_i^p(I)$, $E_i^f(I)$, and $L_i^f(I)$ to refer to time-points of particular interest relative to the current time-point i . Specifically,

for a future formula $\varphi = \mathcal{F}_I \alpha$ with $\mathcal{F} \in \{\circ, \diamond, \square\}$ and interval $I = [a, b]$ or $I = [a, b)$ such that $b \neq \infty$, the formula's semantics at time-point i may need to refer to any time-point with time-stamp in $[\tau_i + a, \dots, \tau_i + b]$. The latest such time-point is $L_i^f(I) = \text{LTP}(\tau_i + b)$ while the earliest one is $E_i^f(I) = \max(i, \text{ETP}(\tau_i + a))$. For past operators $\mathcal{P} \in \{\bullet, \blacklozenge, \blacksquare\}$, the relevant time-stamp interval is $[\tau_i - b, \dots, \tau_i - a]$ and the interval's earliest time-point is $E_i^p(I) = \text{ETP}(\tau_i - b)$ and its latest time-point is $L_i^p(I) = \min(i, \text{LTP}(\tau_i - a))$.

Proof trees emerging from repeated application of the rules in our proof system contain all the necessary information to explain why a formula is satisfied or violated. In other words, our proof system is sound and complete, i.e., the following result holds.

Theorem 1. *Let α be a formula, v a variable assignment, $i \in \mathbb{N}$ a time-point, and $\sigma = \langle \tau_i, \Gamma_i \rangle_{i \in \mathbb{N}}$ a trace. Then $v, i \vdash_{\sigma}^+ \alpha \iff v, i \models_{\sigma} \alpha$ and $v, i \vdash_{\sigma}^- \alpha \iff v, i \not\models_{\sigma} \alpha$.*

We have formalized and verified this result in Isabelle/HOL.

Example 1. Consider the standard publish–approve example [14] requiring that any file f published by an author a , must first be approved by a manager m of a within the previous seven days. The formalization of this policy as a closed MFOTL formula is:

$$\varphi = \forall a. \forall f. \text{publish}(a, f) \rightarrow (\blacklozenge_{[0,7]} \exists m. (\neg \text{mgr}_F(m, a) \mathcal{S} \text{mgr}_S(m, a)) \wedge \text{approve}(m, f)).$$

Here, the events $\text{mgr}_S(m, a)$ and $\text{mgr}_F(m, a)$ mark m starting and finishing being a 's manager. Formally, m is currently a manager of a if m started being a 's manager in the past and has not finished being a 's manager since. Thus, the manager relation changes over time. Consider the stream $\langle \tau_i, \Gamma_i \rangle_{i \in \mathbb{N}}$, where $\tau_0 = \tau_1 = 0$, $\tau_2 = 4$, $\tau_3 = 10$, and

$$\begin{aligned} \Gamma_0 &= \{\text{mgr}_S(\text{Mallory}, \text{Alice}), \text{mgr}_S(\text{Merlin}, \text{Bob}), \text{mgr}_S(\text{Merlin}, \text{Charlie})\}, \text{ and} \\ \Gamma_1 &= \{\text{approve}(\text{Mallory}, 152)\}, \text{ and} \\ \Gamma_2 &= \{\text{approve}(\text{Merlin}, 163), \text{publish}(\text{Alice}, 160), \text{mgr}_F(\text{Merlin}, \text{Charlie})\}, \text{ and} \\ \Gamma_3 &= \{\text{approve}(\text{Merlin}, 187), \text{publish}(\text{Bob}, 163), \text{publish}(\text{Alice}, 163), \\ &\quad \text{publish}(\text{Charlie}, 163), \text{publish}(\text{Charlie}, 152)\}. \end{aligned}$$

In the following we abbreviate the subformulas of φ as follows: $\varphi_L = \text{publish}(a, f)$, $\varphi_1 = \neg \text{mgr}_F(m, a) \mathcal{S} \text{mgr}_S(m, a)$, $\varphi_2 = \text{approve}(m, f)$, $\varphi_3 = \exists m. \varphi_1 \wedge \varphi_2$, $\varphi_R = \blacklozenge_{[0,7]} \varphi_3$, and $\varphi' = \varphi_L \rightarrow \varphi_R$. The following proof tree shows that φ is violated at time-point 3 for any v :

$$\frac{\frac{\frac{\text{approve}(d, 152) \notin \Gamma_i}{v[a \mapsto \text{Charlie}, f \mapsto 152, m \mapsto d], i \vdash^- \varphi_2} P^-}{v[a \mapsto \text{Charlie}, f \mapsto 152, m \mapsto d], i \vdash^- \varphi_1 \wedge \varphi_2} \wedge_R^-}{\frac{\text{publish}(\text{Charlie}, 152) \in \Gamma_3}{v[a \mapsto \text{Charlie}, f \mapsto 152], 3 \vdash^+ \varphi_L} P^+ \quad \frac{v[a \mapsto \text{Charlie}, f \mapsto 152], i \vdash^- \varphi_3}{v[a \mapsto \text{Charlie}, f \mapsto 152], 3 \vdash^- \varphi_R} \blacklozenge^-}{\frac{v[a \mapsto \text{Charlie}, f \mapsto 152], 3 \vdash^- \varphi_L \rightarrow \varphi_R}{v[a \mapsto \text{Charlie}], 3 \vdash^- \forall f. \varphi'} \rightarrow^-} \forall^-} \forall^-$$

Given φ_R 's temporal constraint, we note that $\tau_3 \geq 0$ and need to check $v, i \vdash^- \varphi_3$ for the time-points $i \in \{2, 3\}$ (as $[E_3^p([0, 7]), L_3^p([0, 7])] = \{2, 3\}$). Both subproofs are identical, so we parameterize them over i . In addition, the \exists^- subproofs are valid for an arbitrary manager $d \in \mathbb{D}$ (abbreviating infinite branching over all possible domain values). \square

$$\begin{aligned}
 \mathfrak{sp} = & \mathit{tt}^+(\mathbb{N}) \mid p^+(\mathbb{N}, \mathbb{E}, \bar{t}) \mid \neg^+(\mathfrak{vp}) \mid \wedge^+(\mathfrak{sp}, \mathfrak{sp}) \mid \vee_L^+(\mathfrak{sp}) \mid \vee_R^+(\mathfrak{sp}) \mid \rightarrow_L^+(\mathfrak{vp}) \mid \rightarrow_R^+(\mathfrak{sp}) \\
 & \mid \forall^+(\mathbb{V}, \mathfrak{U}_{\mathbb{D}}(\mathfrak{sp})) \mid \exists^+(\mathbb{V}, \mathbb{D}, \mathfrak{sp}) \mid \bullet^+(\mathfrak{sp}) \mid \circ^+(\mathfrak{sp}) \mid \blacklozenge^+(\mathbb{N}, \mathfrak{sp}) \mid \blacklozenge^+(\mathbb{N}, \mathfrak{sp}) \\
 & \mid \blacksquare_{<I}^+(\mathbb{N}) \mid \blacksquare^+(\mathbb{N}, \overline{\mathfrak{sp}}) \mid \square^+(\mathbb{N}, \overline{\mathfrak{sp}}) \mid \mathcal{S}^+(\mathfrak{sp}, \overline{\mathfrak{sp}}) \mid \mathcal{U}^+(\mathfrak{sp}, \overline{\mathfrak{sp}}), \\
 \mathfrak{vp} = & \mathit{ff}^-(\mathbb{N}) \mid p^-(\mathbb{N}, \mathbb{E}, \bar{t}) \mid \neg^-(\mathfrak{sp}) \mid \wedge_R^-(\mathfrak{vp}) \mid \wedge_L^-(\mathfrak{vp}) \mid \vee^-(\mathfrak{vp}, \mathfrak{vp}) \mid \rightarrow^-(\mathfrak{sp}, \mathfrak{vp}) \\
 & \mid \forall^-(\mathbb{V}, \mathbb{D}, \mathfrak{vp}) \mid \exists^-(\mathbb{V}, \mathfrak{U}_{\mathbb{D}}(\mathfrak{vp})) \mid \bullet^-(\mathfrak{vp}) \mid \bullet_{<I}^-(\mathbb{N}) \mid \bullet_{>I}^-(\mathbb{N}) \mid \bullet_{\emptyset}^- \mid \circ^-(\mathfrak{vp}) \mid \circ_{<I}^-(\mathbb{N}) \\
 & \mid \circ_{>I}^-(\mathbb{N}) \mid \blacklozenge^-(\mathbb{N}, \overline{\mathfrak{vp}}) \mid \blacklozenge_{<I}^-(\mathbb{N}) \mid \blacklozenge^-(\mathbb{N}, \overline{\mathfrak{vp}}) \mid \blacksquare^-(\mathbb{N}, \mathfrak{vp}) \mid \square^-(\mathbb{N}, \mathfrak{vp}) \\
 & \mid \mathcal{S}_{<I}^-(\mathbb{N}) \mid \mathcal{S}^-(\mathbb{N}, \mathfrak{vp}, \overline{\mathfrak{vp}}) \mid \mathcal{S}_{\infty}^-(\mathbb{N}, \overline{\mathfrak{vp}}) \mid \mathcal{U}^-(\mathbb{N}, \mathfrak{vp}, \overline{\mathfrak{vp}}) \mid \mathcal{U}_{\infty}^-(\mathbb{N}, \overline{\mathfrak{vp}})
 \end{aligned}$$

Fig. 3: Grammar for our proof objects.

3 Proof Object Checker

This section introduces our proof objects and their checker: finite data-representations of our proof system’s trees, and an algorithm that certifies if a given proof object faithfully proves the satisfaction or violation of a formula under a given assignment and stream. We discuss the soundness, completeness, and executability of these constructions.

To algorithmically manipulate proof trees, we define an explicit representation of *satisfactions* \mathfrak{sp} and *violations* \mathfrak{vp} via the grammar in Figure 3, where each constructor corresponds to a proof rule of our proof system (Figure 2), and its arguments represent subproofs and parameters that are part of a rule. The disjoint union $\mathfrak{p} = \mathfrak{sp} \uplus \mathfrak{vp}$ is our type of *proof objects*. The proof object \forall^+ requires information about satisfactions for all domain elements $d \in \mathbb{D}$ which we finitely represent with our *valued partitions* $P \in \mathfrak{U}_{\mathbb{D}}(\mathfrak{sp})$. Recall that a partition P of a set A is a collection of non-empty, pair-wise disjoint subsets of A that cover A . That is, $D_i \cap D_j = \emptyset$ for $D_i, D_j \in P$ with $D_i \neq D_j$ and $\bigcup P = A$. Partitions enable us to finitely represent all elements of the domain using finitely many finite sets and the co-finite complement of their union. In valued partitions $P \in \mathfrak{U}_{\mathbb{D}}(\mathfrak{sp})$, each set in the partition is tagged with a satisfaction explaining why its elements satisfy the argument of a universally quantified formula. Formally, our valued partitions $P \in \mathfrak{U}_{\mathbb{D}}(Z)$ are lists of pairs of a set D_i and a value $z \in Z$ from a given set Z such that the sets D_i form a partition of \mathbb{D} . Similarly, \exists^- stores a valued partition $P \in \mathfrak{U}_{\mathbb{D}}(\mathfrak{vp})$ of violations.

Our proof objects $p \in \mathfrak{p}$ represent satisfactions or violations at a certain time-point. We define a function $\text{tp}(p)$ (omitted) to compute this time-point. Either this information can be obtained recursively (e.g., $\text{tp}(\circ^+(p)) = \text{tp}(p) - 1$) or, in cases where it cannot, it is stored directly in the proof objects (e.g., $\text{tp}(\mathit{tt}^+(i)) = i$). We lift tp to sequences (yielding sequences of time-points) and valued partitions as $\text{tp}(P) = \text{tp}(p_1)$, where (D_1, p_1) is the partition P ’s first entry. To characterize *valid* proof objects, we define the relation \vdash_{σ} (Figure 4) that checks that proof objects constitute correct applications of our proof system’s rules. Here, \vdash is not an executable algorithm yet since the proof objects \forall^+ and \exists^- require a recursive call for each element of each set in the partition, and at least one of such sets is infinite for infinite domains. We will improve on this aspect after an example.

Example 2. The following violation proof object p at time-point 3 (i.e., $\text{tp}(p) = 3$) is valid for formula φ on stream σ from Example 1 (i.e., $v, p \vdash_{\sigma} \varphi$ for any assignment v):

$$\begin{aligned}
 p &= \forall^-(a, \text{Charlie}, \forall^-(f, 152, p_{\rightarrow}^-)), \text{ where} \\
 p_{\rightarrow}^- &= \rightarrow^-(p_L^+, p_{\blacklozenge}^-), p_L^+ = p^+(3, \text{publish}, [a, f]), \\
 p_{\blacklozenge}^- &= \blacklozenge^-(3, [\exists^-(x, [(\mathbb{D}, p_2^-)]), \exists^-(x, [(\mathbb{D}, p_3^-)])]), \text{ and} \\
 p_i^- &= \wedge_R^-(p^-(i, \text{approve}, [m, f])) \text{ for } i \in \{2, 3\}.
 \end{aligned}$$

$v, \pi^+(i) \vdash \pi$		$v, \text{ff}^-(i) \vdash \text{ff}$	
$v, p^+(i, \bar{p}, \bar{i}) \vdash p(\bar{i})$	iff $p(\llbracket \bar{i} \rrbracket_v) \in \Gamma_i$	$v, p^-(i, \bar{p}, \bar{i}) \vdash p(\bar{i})$	iff $p(\llbracket \bar{i} \rrbracket_v) \notin \Gamma_i$
$v, \approx^+(i, x, c) \vdash x \approx c$	iff $v(x) = c$	$v, \approx^-(i, x, c) \vdash x \approx c$	iff $v(x) \neq c$
$v, \neg^+(vp) \vdash \neg \alpha$	iff $v, vp \vdash \alpha$	$v, \neg^-(sp) \vdash \neg \alpha$	iff $v, sp \vdash \alpha$
$v, \rightarrow_L^+(vp) \vdash \alpha \rightarrow \beta$	iff $v, vp \vdash \alpha$	$v, \wedge_L^-(vp) \vdash \alpha \wedge \beta$	iff $v, vp \vdash \alpha$
$v, \rightarrow_R^+(sp) \vdash \alpha \rightarrow \beta$	iff $v, sp \vdash \beta$	$v, \wedge_R^-(vp) \vdash \alpha \wedge \beta$	iff $v, vp \vdash \beta$
$v, \exists^+(x, d, sp) \vdash \exists x. \alpha$	iff $v[x \mapsto d], sp \vdash \alpha$	$v, \vee_L^+(sp) \vdash \alpha \vee \beta$	iff $v, sp \vdash \alpha$
$v, \forall^-(x, d, vp) \vdash \forall x. \alpha$	iff $v[x \mapsto d], vp \vdash \alpha$	$v, \vee_R^+(sp) \vdash \alpha \vee \beta$	iff $v, sp \vdash \beta$
$v, \wedge^+(sp_1, sp_2) \vdash \alpha \wedge \beta$	iff $v, sp_1 \vdash \alpha$ and $v, sp_2 \vdash \beta$ and $\text{tp}(sp_1) = \text{tp}(sp_2)$		
$v, \vee^-(vp_1, vp_2) \vdash \alpha \vee \beta$	iff $v, vp_1 \vdash \alpha$ and $v, vp_2 \vdash \beta$ and $\text{tp}(vp_1) = \text{tp}(vp_2)$		
$v, \rightarrow^-(sp_1, vp_2) \vdash \alpha \rightarrow \beta$	iff $v, sp_1 \vdash \alpha$ and $v, vp_2 \vdash \beta$ and $\text{tp}(sp_1) = \text{tp}(vp_2)$		
$v, \forall^+(x, P) \vdash \forall x. \alpha$	iff $\wedge (D_k, sp_k) \in P. \text{tp}(sp_k) = \text{tp}(P)$ and $\wedge d \in D_k. v[x \mapsto d], sp_k \vdash \alpha$		
$v, \exists^-(x, P) \vdash \exists x. \alpha$	iff $\wedge (D_k, vp_k) \in P. \text{tp}(vp_k) = \text{tp}(P)$ and $\wedge d \in D_k. v[x \mapsto d], vp_k \vdash \alpha$		
$v, \bullet^+(sp) \vdash \bullet_I \alpha$	iff $v, sp \vdash \alpha$ and $\text{tp}(\bullet^+(sp)) = \text{tp}(sp) + 1$ and $\tau_{\text{tp}(\bullet^+(sp))} - \tau_{\text{tp}(sp)} \in I$		
$v, \circ^+(sp) \vdash \circ_I \alpha$	iff $v, sp \vdash \alpha$ and $\text{tp}(\circ^+(sp)) + 1 = \text{tp}(sp)$ and $\tau_{\text{tp}(sp)} - \tau_{\text{tp}(\circ^+(sp))} \in I$		
$v, \blacklozenge^+(i, sp) \vdash \blacklozenge_I \alpha$	iff $v, sp \vdash \alpha$ and $i \geq \text{tp}(sp)$ and $\tau_i - \tau_{\text{tp}(sp)} \in I$		
$v, \blacktriangleright^+(i, sp) \vdash \blacktriangleright_I \alpha$	iff $v, sp \vdash \alpha$ and $i \leq \text{tp}(sp)$ and $\tau_{\text{tp}(sp)} - \tau_i \in I$		
$v, \blacksquare^+(i, \bar{sp}) \vdash \blacksquare_I \alpha$	iff $(\wedge sp \in \bar{sp}. v, sp \vdash \alpha)$ and $\text{tp}(\bar{sp}) = [\mathbf{E}_i^p(I), \mathbf{L}_i^p(I)]$ and $\tau_i \geq \tau_0 + I$		
$v, \square^+(i, \bar{sp}) \vdash \square_I \alpha$	iff $(\wedge sp \in \bar{sp}. v, sp \vdash \alpha)$ and $\text{tp}(\bar{sp}) = [\mathbf{E}_i^f(I), \mathbf{L}_i^f(I)]$		
$v, \mathcal{S}^+(sp, \bar{sp}) \vdash \alpha \mathcal{S}_I \beta$	iff $(\wedge sp' \in \bar{sp}. v, sp' \vdash \alpha)$ and $v, sp \vdash \beta$ and $\text{tp}(\mathcal{S}^+(sp, \bar{sp})) \geq \text{tp}(sp)$ and $\text{tp}(\bar{sp}) = [\text{tp}(sp) + 1, \text{tp}(\mathcal{S}^+(sp, \bar{sp}))]$ and $\tau_{\text{tp}(\mathcal{S}^+(sp, \bar{sp}))} - \tau_{\text{tp}(sp)} \in I$		
$v, \mathcal{U}^+(sp, \bar{sp}) \vdash \alpha \mathcal{U}_I \beta$	iff $(\wedge sp' \in \bar{sp}. v, sp' \vdash \alpha)$ and $v, sp \vdash \beta$ and $\text{tp}(\mathcal{U}^+(sp, \bar{sp})) \leq \text{tp}(sp)$ and $\text{tp}(\bar{sp}) = [\text{tp}(\mathcal{U}^+(sp, \bar{sp})), \text{tp}(sp)]$ and $\tau_{\text{tp}(sp)} - \tau_{\text{tp}(\mathcal{U}^+(sp, \bar{sp}))} \in I$		
$v, \bullet_0^- \vdash \bullet_I \alpha$	iff $\text{tp}(\bullet_0^-) = 0$	$v, \bullet_{<I}^-(i) \vdash \bullet_I \alpha$	iff $i > 0$ and $\tau_i - \tau_{i-1} < I$
$v, \circ_{<I}^-(i) \vdash \circ_I \alpha$	iff $\tau_{i+1} - \tau_i < I$	$v, \bullet_{>I}^-(i) \vdash \bullet_I \alpha$	iff $i > 0$ and $\tau_i - \tau_{i-1} > I$
$v, \circ_{>I}^-(i) \vdash \circ_I \alpha$	iff $\tau_{i+1} - \tau_i > I$	$v, \blacksquare_{<I}^-(i) \vdash \blacksquare_I \alpha$	iff $\tau_i < \tau_0 + I$
$v, \blacklozenge_{<I}^-(i) \vdash \blacklozenge_I \alpha$	iff $\tau_i < \tau_0 + I$	$v, \mathcal{S}_{<I}^-(i) \vdash \alpha \mathcal{S}_I \beta$	iff $\tau_i < \tau_0 + I$
$v, \bullet^-(vp) \vdash \bullet_I \alpha$	iff $v, vp \vdash \alpha$ and $\text{tp}(\bullet^-(vp)) = \text{tp}(vp) + 1$		
$v, \circ^-(vp) \vdash \circ_I \alpha$	iff $v, vp \vdash \alpha$ and $\text{tp}(\circ^-(vp)) + 1 = \text{tp}(vp)$		
$v, \blacklozenge^-(i, \bar{vp}) \vdash \blacklozenge_I \alpha$	iff $(\wedge vp \in \bar{vp}. v, vp \vdash \alpha)$ and $\text{tp}(\bar{vp}) = [\mathbf{E}_i^p(I), \mathbf{L}_i^p(I)]$ and $\tau_i \geq \tau_0 + I$		
$v, \blacktriangleright^-(i, \bar{vp}) \vdash \blacktriangleright_I \alpha$	iff $(\wedge vp \in \bar{vp}. v, vp \vdash \alpha)$ and $\text{tp}(\bar{vp}) = [\mathbf{E}_i^f(I), \mathbf{L}_i^f(I)]$		
$v, \blacksquare^-(i, vp) \vdash \blacksquare_I \alpha$	iff $v, vp \vdash \alpha$ and $i \geq \text{tp}(vp)$ and $\tau_i - \tau_{\text{tp}(vp)} \in I$		
$v, \square^-(i, vp) \vdash \square_I \alpha$	iff $v, vp \vdash \alpha$ and $i \leq \text{tp}(vp)$ and $\tau_{\text{tp}(vp)} - \tau_i \in I$		
$v, \mathcal{S}_\infty^-(i, \bar{vp}) \vdash \alpha \mathcal{S}_I \beta$	iff $(\wedge vp \in \bar{vp}. v, vp \vdash \beta)$ and $\text{tp}(\bar{vp}) = [\mathbf{E}_i^p(I), \mathbf{L}_i^p(I)]$ and $\tau_i \geq \tau_0 + I$		
$v, \mathcal{S}^-(i, vp, \bar{vp}) \vdash \alpha \mathcal{S}_I \beta$	iff $(\wedge vp \in \bar{vp}. v, vp \vdash \beta)$ and $v, vp \vdash \alpha$ and $\mathbf{E}_i^p(I) \leq \text{tp}(vp) \leq i$ and $\text{tp}(\bar{vp}) = [\text{tp}(vp), \mathbf{L}_i^p(I)]$ and $\tau_i \geq \tau_0 + I$		
$v, \mathcal{U}_\infty^-(i, \bar{vp}) \vdash \alpha \mathcal{U}_I \beta$	iff $(\wedge vp \in \bar{vp}. v, vp \vdash \beta)$ and $\text{tp}(\bar{vp}) = [\mathbf{E}_i^f(I), \mathbf{L}_i^f(I)]$		
$v, \mathcal{U}^-(i, vp, \bar{vp}) \vdash \alpha \mathcal{U}_I \beta$	iff $(\wedge vp \in \bar{vp}. v, vp \vdash \beta)$ and $v, vp \vdash \alpha$ and $i \leq \text{tp}(vp) < \mathbf{L}_i^f(I)$ and $\text{tp}(\bar{vp}) = [\mathbf{E}_i^f(I), \text{tp}(vp)]$		

Fig. 4: Proof checker for a fixed stream $\sigma = \langle \tau_i, \Gamma_i \rangle_{i \in \mathbb{N}}$.

Indeed, we use the definition in Figure 4 to certify that $v, p \vdash_\sigma \varphi$:

$$\begin{aligned}
v, p \vdash \varphi &\text{ iff } v[a \mapsto \text{Charlie}], \forall^-(f, 152, p_-^-) \vdash \forall f. \varphi_L \rightarrow \varphi_R \\
&\text{ iff } v[a \mapsto \text{Charlie}, f \mapsto 152], p_-^- \vdash \varphi_L \rightarrow \varphi_R \\
&\text{ iff } v[a \mapsto \text{Charlie}, f \mapsto 152], p_L^+ \vdash \varphi_L \text{ and } \text{tp}(p_L^+) = 3 = \text{tp}(p_\blacklozenge^+) \text{ and} \\
&\quad v[a \mapsto \text{Charlie}, f \mapsto 152], p_\blacklozenge^+ \vdash \varphi_R \\
&\text{ iff } v[a \mapsto \text{Charlie}, f \mapsto 152], \exists^-(x, [\llbracket \mathbb{D}, p_i^- \rrbracket]) \vdash \varphi_\exists \text{ for } i \in \{2, 3\} \\
&\text{ iff } v[a \mapsto \text{Charlie}, f \mapsto 152, x \mapsto d], p_i^- \vdash \varphi_1 \wedge \varphi_2 \text{ for all } d \in \mathbb{D}, i \in \{2, 3\} \\
&\text{ iff } \text{approve}(d, 152) \notin \Gamma_i \text{ for all } d \in \mathbb{D}, i \in \{2, 3\}, \text{ which is true.}
\end{aligned}$$

$$\begin{aligned}
 \text{LRTP } i \text{ tt} &= \text{LRTP } i \text{ ff} = \text{LRTP } i (\text{p}(\bar{t})) = \text{LRTP } i (x \approx c) = i, \\
 \text{LRTP } i (Qx. \alpha) &= \text{LRTP } i (\neg \alpha) = \text{LRTP } i \alpha \text{ for } Q \in \{\forall, \exists\}, \\
 \text{LRTP } i (\alpha \oplus \beta) &= \max(\text{LRTP } i \alpha) (\text{LRTP } i \beta) \text{ for } \oplus \in \{\vee, \wedge, \rightarrow\}, \\
 \text{LRTP } i (\alpha \oplus_I \alpha) &= \text{LRTP } (i-1) \alpha, \quad \text{LRTP } i (\alpha \circ_I \alpha) = \text{LRTP } (i+1) \alpha, \\
 \text{LRTP } i (\diamond_I \alpha) &= \text{LRTP } i (\square_I \alpha) = \text{LRTP } (L_i^f(I)) \alpha, \\
 \text{LRTP } i (\blacklozenge_I \alpha) &= \text{LRTP } i (\blacksquare_I \alpha) = \text{LRTP } (LTP_i^{\text{past}} I) \alpha, \\
 \text{LRTP } i (\alpha \mathcal{S}_I \beta) &= \max(\text{LRTP } i \alpha) (\text{LRTP } (LTP_i^{\text{past}} I) \beta), \\
 \text{LRTP } i (\alpha \mathcal{U}_I \beta) &= \max(\text{LRTP } (L_i^f(I) - 1) \alpha) (\text{LRTP } (L_i^f(I)) \beta), \text{ where} \\
 LTP_i^{\text{past}} I &= (\text{if } \tau_i \geq \tau_0 + I \text{ then } L_i^p(I) \text{ else } 0).
 \end{aligned}$$

Fig. 5: The formula’s latest relevant time-point at i for a fixed stream $\sigma = \langle \tau_i, F_i \rangle_{i \in \mathbb{N}}$.

We implicitly use in the above the true statements $\text{publish}(\text{Charlie}, 152) \in \Gamma_3$, $0 \leq \tau_3$, and $\text{tp}([\exists^-(x, [(\mathbb{D}, p_2^-)]), \exists^-(x, [(\mathbb{D}, p_3^-)])]) = [2, 3] = [E_i^p(I), L_i^p(I)]$. \square

Theorem 2. Fix a stream σ . The relation \vdash is sound and complete in the sense that $v, i \models \alpha$ iff there is a satisfaction sp such that $v, sp \vdash \alpha$ and $\text{tp}(sp) = i$. Similarly $v, i \not\models \alpha$ iff there is a violation vp such that $v, vp \vdash \alpha$ and $\text{tp}(vp) = i$.

We have established the above result in Isabelle. Below we sketch our overall approach and highlight the main challenge. We show both soundness and completeness by relating proof object validity (\vdash) to the proof system (\vdash^+ and \vdash^-), which we already know to be sound and complete, i.e., related to the semantics \models . Soundness is easy as the proof object directly provides the recipe for correctly applying the proof system rules. Formally, if $v, sp \vdash \alpha$ then $v, \text{tp}(sp) \vdash^+ \alpha$, and if $v, vp \vdash \alpha$, then $v, \text{tp}(vp) \vdash^- \alpha$. The proof follows immediately by mutual induction on the proof object structure.

Completeness of \vdash requires us to provide a valid proof object just from knowing $v, i \vdash^+ \alpha$ or $v, i \vdash^- \alpha$. We proceed by mutual induction on the derivations of \vdash^+ and \vdash^- . Only two of the quantifier cases are challenging. For the satisfaction of the universal quantifier (and similarly for the violation of \exists), we must construct a valued partition with finitely many subproofs. However, the induction hypothesis yields a separate proof object for every element of the domain \mathbb{D} , and all these proof objects may a priori be different. The crucial observation is that for all values that do not occur in the stream (or at least are not in reach of α with respect to a time-point i) we can reuse the same proof object. To formalize this observation, we first define a formula’s *active domain* at i , written $\text{AD}_i(\alpha)$, which formalizes the in “reach” intuition. To this end, we first define the *latest relevant time point* ($\text{LRTP } i \alpha$) of α at i (Figure 5). Intuitively, $\text{LRTP } i \alpha$ marks the largest time-point that may influence α ’s satisfiability at i . It exists, because we assume that future temporal operators have bounded intervals. Based on this, we define:

$$\text{AD}_i(\alpha) = \mathbb{D}(\alpha) \cup \bigcup_{k \leq \text{LRTP } i \alpha} \{d \mid d \text{ appears in some } \text{p}(d_1, \dots, d_n) \in \Gamma_k\}.$$

Here we write $\mathbb{D}(\alpha)$ for the set of constants $d \in \mathbb{D}$ occurring in subformulas of the form $x \approx d$ in α . (In contrast to constants occurring in atomic predicates, constants occurring in equalities may appear in α ’s satisfying assignments even if they are not part of the trace.) Note that $\text{AD}_i(\alpha)$ is finite. The active domain lets us formalize the key observation:

Lemma 1. Fix a stream σ , a formula α , a proof p , and two assignments v and v' . Let $i = \text{tp}(p)$, $\text{AD} = \text{AD}_i(\alpha)$, and V be the set of α ’s free variables. Assume that v and v'

may only disagree on V for values outside of the active domain at i , i.e.,

$$\forall x \in V. v(x) = v'(x) \vee (v(x) \notin \text{AD} \wedge v'(x) \notin \text{AD}).$$

Then, p 's validity status is the same for both assignments, i.e., $v, p \vdash \alpha$ iff $v', p \vdash \alpha$.

We now can finish the \forall^+ case of the completeness proof. By the induction hypothesis, there is a satisfaction $p(d) \in \mathfrak{sp}$ for each domain element $d \in \mathbb{D}$. Moreover, $\{\{d\} \mid d \in \text{AD}_i(\alpha)\} \cup \{\mathbb{D} \setminus \text{AD}_i(\alpha)\}$ is a finite partition of \mathbb{D} . Hence, the list of pairings $(\{d\}, p(d))$ for each $d \in \text{AD}_i(\alpha)$ and $(\mathbb{D} \setminus \text{AD}_i(\alpha), p(z))$ for some $z \in \mathbb{D} \setminus \text{AD}_i(\alpha)$ (which exists as \mathbb{D} is infinite) is a valued partition. Moreover, all subproofs are valid for all the values contained in the partition sets by combining the induction hypothesis with the above congruence Lemma 1 (for $p = p(z)$), and thus so is the overall \forall^+ proof object.

Lastly, we address the executability issue. The validity relation \vdash works with assignments v of values to variables. To avoid performing infinitely many recursive calls for the \forall^+ and \exists^- proof objects we now will work with *set assignments* V of sets of values to variables. We define a validity relation $V, p \vdash \alpha$ based on set assignments. The definition is the same as the one of $v, p \vdash \alpha$ except for the predicate and the quantifier cases:

$$\begin{aligned} V, p^+(i, \mathfrak{p}, \bar{i}) \vdash \mathfrak{p}(\bar{i}) & \quad \text{iff } \{\mathfrak{p}\} \times \llbracket \bar{i} \rrbracket_V \subseteq \Gamma_i \\ V, p^-(i, \mathfrak{p}, \bar{i}) \vdash \mathfrak{p}(\bar{i}) & \quad \text{iff } \{\mathfrak{p}\} \times \llbracket \bar{i} \rrbracket_V \subseteq \mathbb{D} \setminus \Gamma_i \\ V, \forall^+(x, P) \vdash \forall x. \alpha & \quad \text{iff } \bigwedge (D_k, sp_k) \in P. \text{tp}(sp_k) = \text{tp}(P) \text{ and } V[x \mapsto D_k], sp_k \vdash \alpha \\ V, \exists^+(x, d, sp) \vdash \exists x. \alpha & \quad \text{iff } V[x \mapsto \{d\}], sp \vdash \alpha \end{aligned}$$

and dually for \exists^- and \forall^- . Here, $\llbracket \bar{i} \rrbracket_V$ represents a transformation of the list of values $\llbracket \bar{i} \rrbracket$ to the set of all possible lists of values generated by V . Set assignments allow us to delay deciding values for quantifier subproofs to the predicate base case. Note that $\{\mathfrak{p}\} \times \llbracket \bar{i} \rrbracket_V \subseteq \Gamma_i$ and $\{\mathfrak{p}\} \times \llbracket \bar{i} \rrbracket_V \subseteq \mathbb{D} \setminus \Gamma_i$ are decidable because due to our partitions, we only encounter finite and co-finite sets. The set-assignment-based validity check is thus executable and thus provides the algorithm that we use as our formally verified proof object checker: $v, p \vdash \alpha = (\lambda x. \{v(x)\}), p \vdash \alpha$ (proved by induction on α using Lemma 1).

4 Partitioned Decision Trees

Our proof system is parameterized with an assignment, but in our monitoring approach we are interested in computing a proof object for every assignment. In this section, we introduce *partitioned decision trees* (PDTs), a specialized data structure for representing and efficiently manipulating variable assignments, inspired by the use of BDDs in runtime verification [17]. We want to represent functions of the form $f : \mathbb{D} \times \dots \times \mathbb{D} \rightarrow \mathfrak{p}$, i.e., mappings from tuples of domain elements to proof trees, where each tuple corresponds to a variable assignment to the formula's free variables. As argued in the previous section, we are only interested in such functions with a finite range. Thus, we organize the domain into a finite number of subsets $\mathbb{D} \times \dots \times \mathbb{D}$ such that each tuple element is partitioned separately (using valued partitions over the domain). As before, we work with finite and co-finite sets in the partition. PDTs $\mathbb{P}(A)$ are defined inductively as follows:

$$\mathbb{P}(A) = \text{Leaf } A \mid \text{Node } (\mathbb{V}, \bigoplus_{\mathbb{D}} (\mathbb{P}(A)))$$




Fig. 6: Resulting PDT for our running example at time-point 3.

PDTs have leaves and nodes. Leaves store objects from the set A , while nodes store pairs of the form (x, P) , where x is a variable and P , a valued partition of the domain storing PDTs. PDTs generalize binary decision trees along two dimensions. First, the branching of their nodes is not binary but follows a given partition of the infinite domain \mathbb{D} . Second, their leaves do not store Boolean values. Instead, they store arbitrary objects, even though we will mostly use them with proof objects $A = \mathfrak{p}$. PDTs provide a way to organize the infinitely many possible variable assignments in a structured manner, storing only finitely many different proof objects. In monitoring, partitions will arise naturally, guided by the values occurring in the stream and assembled via operations that combine them.

Example 3. We continue the publish–approve example from Example 1. We consider the same stream but drop the top-level quantifiers from the formula φ : we only consider φ' with its free variables a and f . Figure 6 shows the PDT representing all assignments for φ' at time-point 3. The root node represents variable a , and the edges partition the values that a can take into the following domain subsets: $\{\text{Alice}\}$, $\{\text{Bob}\}$, $\{\text{Charlie}\}$, and $\mathbb{D} \setminus \{\text{Alice}, \text{Bob}, \text{Charlie}\}$. The second level is analogous for variable f . At every level of the PDT, the union of all choices cover the entire domain \mathbb{D} (by definition of partitions) and the partitions may differ at every node. The leaves of the PDT are different proof trees (formally, proof objects) which we represent by small black triangles. For example, \blacktriangle_3^- is the proof tree of φ' 's violation shown in Example 1. In contrast, \blacktriangle^+ (occurring in multiple leaves) is the proof tree shown in Figure 7 of φ' 's vacuous satisfaction: the left hand side of the implication ($\text{publish}(a, f)$) is violated for any assignment v updated by following the path from the PDT's root to the respective leaf (e.g., taking $a = \text{Alice}$ and $f = 42 \in \mathbb{D} \setminus \{163\}$). \square

$$\frac{\text{publish}(a, f) \notin \Gamma_3 \quad v, 3 \vdash^- \text{publish}(a, f) \quad P^-}{v, 3 \vdash^+ \varphi'} \rightarrow_L^+$$

Fig. 7: Proof tree \blacktriangle^+ .

Since PDTs are a generalization of BDDs, we use similar functions to manipulate them. We list the most important ones, for partitions and PDTs in Figure 8, but we only show and discuss the implementation of `apply2`, `merge2`, and `hide`. Most PDT-functions are parameterized by a variable list $vs :: \bar{\mathbb{V}}$ fixing the variable order. The functions `map_part` and `apply1` lift unary functions on objects to partitions and PDTs respectively.

The functions `merge2` and `apply2` do the same for binary functions; `apply2` generalizes the well-known `apply` function on BDDs [16]. On leaves, `apply2` maps f to the objects. When operating on a leaf and a node, `apply2` pushes f partially applied to the leaf to the node's leaves using `apply1`. Finally, on pairs of nodes, it proceeds recursively depending which of x , y , and z are equal. The most interesting case, $x = y = z$ occurs when both PDTs partition the domain values for z in different ways. Thus, we must

```

map_part :: (A ⇒ B) ⇒  $\mathfrak{P}_{\mathbb{D}}$ (A) ⇒  $\mathfrak{P}_{\mathbb{D}}$ (B)
merge2 :: (A ⇒ B ⇒ C) ⇒  $\mathfrak{P}_{\mathbb{D}}$ (A) ⇒  $\mathfrak{P}_{\mathbb{D}}$ (B) ⇒  $\mathfrak{P}_{\mathbb{D}}$ (C)
merge2 f [] P2 = []
merge2 f ((D1, v1) # P1) P2 =
  let P3 = map_filter (λ(D2, v2). if D1 ∩ D2 ≠ ∅ then Some (D1 ∩ D2, f v1 v2) else None) P2;
      P4 = map_filter (λ(D2, v2). if D2 \ D1 ≠ ∅ then Some (D2 \ D1, v2) else None) P2
  in P3 @ merge2 f P1 P4
pdt_of ::  $\bar{V}$  ⇒ A ⇒ A ⇒ 2(V→D) ⇒  $\mathbb{P}$ (A) split_prod ::  $\mathbb{P}(A \times B) \Rightarrow \mathbb{P}(A) \times \mathbb{P}(B)$ 
apply1 ::  $\bar{V} \Rightarrow (A \Rightarrow B) \Rightarrow \mathbb{P}(A) \Rightarrow \mathbb{P}(B)$  split_list ::  $\mathbb{P}(\bar{A}) \Rightarrow \overline{\mathbb{P}(A)}$ 
apply2 ::  $\bar{V} \Rightarrow (A \Rightarrow B \Rightarrow C) \Rightarrow \mathbb{P}(A) \Rightarrow \mathbb{P}(B) \Rightarrow \mathbb{P}(C)$ 
apply2 vs f (Leaf l1) (Leaf l2) = Leaf (f l1 l2)
apply2 vs f (Leaf l1) (Node (x, P2)) = Node (x, map_part (apply1 vs (λl2. f l1 l2)) P2)
apply2 vs f (Node (x, P1)) (Leaf l2) = Node (x, map_part (apply1 vs (λl1. f l1 l2)) P1)
apply2 (z # vs) f (Node (x, P1)) (Node (y, P2)) =
  if x = z and y = z then Node (z, merge2 (apply2 vs f) P1 P2)
  else if x = z then Node (x, map_part (λl. apply2 vs f l (Node (y, P2))) P1)
  else if y = z then Node (y, map_part (λr. apply2 vs f (Node (x, P1)) r) P2)
  else apply2 vs f (Node (x, P1)) (Node (y, P2))
apply3 ::  $\bar{V} \Rightarrow (A \Rightarrow B \Rightarrow C \Rightarrow D) \Rightarrow \mathbb{P}(A) \Rightarrow \mathbb{P}(B) \Rightarrow \mathbb{P}(C) \Rightarrow \mathbb{P}(D)$ 
hide ::  $\bar{V} \Rightarrow (A \Rightarrow A) \Rightarrow (\mathfrak{P}_{\mathbb{D}}(A) \Rightarrow A) \Rightarrow \mathbb{P}(A) \Rightarrow \mathbb{P}(A)$ 
hide vs leaf node (Leaf l) = Leaf (leaf l)
hide [z] leaf node (Node (x, P)) = Leaf (node (map_part unleaf P))
hide (z # vs) leaf node (Node (x, P)) =
  if x = z then Node (z, map_part (hide vs leaf node) P) else hide vs leaf node (Node (x, P))

```

Fig. 8: Selected functions on partitions and PDTs.

combine both partitions. For this, we use `merge2` that takes two valued partitions P_1 and P_2 , and iteratively “erodes” P_2 by intersecting its elements with the sets in P_1 while applying f . Since both P_1 and P_2 cover \mathbb{D} , the resulting set of intersections is a valued partition. The function `apply3` analogously combines three PDTs into one.

The function `hide` traverses the PDT similarly to `apply1`, while eliminating the last variable in the given variable list. It uses two higher-order arguments, in case the last layer is present (*node*) or absent (*leaf*). The function `pdt_of vs A B V` constructs a PDT from a finite set of partial assignments ($V :: 2^{(V \rightarrow \mathbb{D})}$) using A for leaves reached by paths from the set, and B for the other leaves. Finally, the `split_*` functions transpose a PDT storing pairs (lists of equal length) into a pair (list) of PDTs.

5 Monitoring Algorithm

We follow the typical online monitoring algorithm structure consisting of an initialization and a step (evaluation) function [25, 30]. The initializer `init` (omitted as standard) computes our monitor’s initial state $s \in \mathbb{S}$ from an MFOTL specification α . Figure 9 shows an excerpt of our monitor’s state, which recursively follows the formula structure and augments some operators with additional information, such as buffers storing verdicts from subformulas (\mathbb{B}_2 for \wedge and \mathbb{B}_3 for \mathcal{S}) or an operator-specific state (\mathbb{S}_{saux} for \mathcal{S}).

$$\begin{aligned}
 \mathbb{B}_2 &= \overline{\mathbb{P}(\mathbf{p})} \times \overline{\mathbb{P}(\mathbf{p})} & \mathbb{B}_3 &= \overline{\mathbb{P}(\mathbf{p})} \times \overline{\mathbb{P}(\mathbf{p})} \times \overline{\mathbb{N} \times \mathbb{N}} & \mathbb{S}_{saux} &= \dots \\
 \mathbb{S} &= \overline{\text{MPred } \mathbb{E} \overline{\mathbb{T}} \mid \text{MAnd } \mathbb{S} \mathbb{S} \mathbb{B}_2 \mid \text{MExists } \mathbb{E} \mathbb{S} \mid \text{MSince } \mathbb{I} \mathbb{S} \mathbb{S} \mathbb{B}_3 (\overline{\mathbb{P}(\mathbb{S}_{saux})}) \mid \dots \\
 \text{eval} &:: \overline{\mathbb{V}} \Rightarrow \mathbb{N} \Rightarrow \mathbb{N} \Rightarrow \mathbb{E} \times \overline{\mathbb{D}} \Rightarrow \mathbb{S} \Rightarrow \overline{\mathbb{P}(\mathbf{p})} \times \mathbb{S} \\
 \text{eval } vs \tau i \Gamma (\text{MPred } \mathbf{p} \ ts) &= \\
 &\text{let } e = \text{pdt_of} (\text{filter } (\lambda v. v \in \text{fv}(ts)) \ vs) \ (p^+(i, \mathbf{p}, ts)) \ (p^-(i, \mathbf{p}, ts)) \\
 &\quad \{\sigma \mid \exists ds. \mathbf{p}(ds) \in \Gamma \wedge \text{match } ts \ ds = \text{Some } \sigma\} \ \text{in } ([e], \text{MPred } \mathbf{p} \ ts) \\
 \text{eval } vs \tau i \Gamma (\text{MAnd } s_1 \ s_2 \ \text{buf}) &= \text{let } (es_1, s'_1) = \text{eval } vs \tau i \Gamma s_1; \quad (es_2, s'_2) = \text{eval } vs \tau i \Gamma s_2; \\
 &\quad (es, \text{buf}') = \text{buf2_take} (\text{apply2 } vs \ \text{do_and}) (\text{buf2_add } \text{buf} \ es_1 \ es_2) \ \text{in } (es, \text{MAnd } s'_1 \ s'_2 \ \text{buf}') \\
 \text{eval } vs \tau i \Gamma (\text{MExists } x \ s) &= \text{let } (es, s') = \text{eval } (vs@[x]) \ \tau i \Gamma s \\
 &\quad \text{in } (\text{map} (\text{hide } (vs@[x]) \ (\text{do_exists_leaf } x) \ (\text{do_exists_node } x)) \ es, \text{MExists } x \ s') \\
 \text{eval } vs \tau i \Gamma (\text{MSince } I \ s_1 \ s_2 \ \text{buf } \ \text{saux}) &= \\
 &\quad \text{let } (es_1, s'_1) = \text{eval } vs \tau i \Gamma s_1; \quad (es_2, s'_2) = \text{eval } vs \tau i \Gamma s_2; \\
 &\quad (es, \text{buf}', \text{saux}') = \text{buf2t_take} (\lambda e_1 \ e_2 \ (\tau, i) \ \text{saux}). \\
 &\quad \text{let } (\text{saux}', es') = \text{split_prod} (\text{apply3 } vs \ (\text{update_since } I \ \tau \ i) \ e_1 \ e_2 \ \text{saux}') \\
 &\quad \text{in } (\text{saux}', \text{split_list } es') (\text{buf2t_add } \text{buf} \ es_1 \ es_2 \ [(\tau, i)]) \ \text{saux}' \\
 &\quad \text{in } (es, \text{MSince } I \ s'_1 \ s'_2 \ \text{buf}' \ \text{saux}')
 \end{aligned}$$

Fig. 9: Involved types and selected cases of the monitor's eval function

$$\begin{aligned}
 \text{do_exists_leaf } x \ p &= \text{if } p \in \text{sp} \ \text{then } \exists^+(x, d \leftarrow \mathbb{D}, p) \ \text{else } \exists^-(x, [(\mathbb{D}, p)]) \\
 \text{do_exists_node } x \ P &= \text{if } \forall (D_i, p) \in P. p \in \text{sp} \\
 &\quad \text{then } \min (\text{map_filter } (\lambda (D_i, p). \text{if } p \in \text{sp} \ \text{then } \text{Some} (\exists^+(x, d \leftarrow D_i, p))) \ \text{else } \text{None}) \ P) \\
 &\quad \text{else } \exists^-(x, P)
 \end{aligned}$$

Fig. 10: Functions do_exists_leaf and do_exists_node.

Our function `eval`, partly shown in Figure 9, takes as inputs a new time-point i (along with its time-stamp τ and database Γ) and a monitor state s and outputs the next state s' and a list of PDTs of proof objects as verdicts. (In addition, `eval` keeps track of the variable ordering used in PDTs via the parameter vs .) Lists in the output are necessary because delays may occur for (bounded) future operators and a single time-point might trigger multiple outputs. Our algorithm extends Lima et al.'s algorithm [25] computing proof trees for MTL. We highlight our key additions to `eval` and the state Figure 9 in gray.

We focus on the predicate, conjunction, existential quantifier, and since cases. In the predicate case, we find all partial assignments σ mapping the predicate's variables to the values ds , so that $\mathbf{p}(ds) \in \Gamma$. We reuse VeriMon's match function [30] to compute such partial assignments. We convert this set of assignments to a PDT using `pdt_of`. In the resulting PDT, matching assignments lead to leaves using the satisfaction proof $p^+(i, \mathbf{p}, ts)$, whereas the others lead to the corresponding violation proof $p^-(i, \mathbf{p}, ts)$.

The conjunction case is taken almost without changes from Lima et al.'s [25] MTL algorithm. We reuse the buffering functions `buf2_add` and `buf2_take`. The first adds partial results to the buffer, while the second combines these results and dequeues them once both subformulas have produced results for a time-point. The only difference is that our buffers store PDTs of proof objects, whereas the MTL algorithm works with propositional proof objects. Accordingly, we reuse the Lima et al.'s function `do_and` $:: \mathbf{p} \Rightarrow \mathbf{p} \Rightarrow \mathbf{p}$ to combine two proof objects conjunctively, but lift it to PDTs using `apply2`.

The quantifier cases are a new addition of our work. As both cases proceed dually, we focus on $\exists x. \alpha$ formulas. Considering that α may have one more free variable than $\exists x. \alpha$, the recursive call appends x to the variable list ordering. The recursive call's output

is processed using our function `hide` to eliminate the quantified variable. The interesting cases occur near the leaves of α 's PDTs. If x is not present, `hide` will encounter a leaf, i.e., a proof object, and use the function `do_exists_leaf` (Figure 10) to perform a case distinction: satisfactions result in a satisfaction (`sp`) of $\exists x. \alpha$ with an arbitrary element d of the domain as the witness (we write $x \leftarrow X$ to denote an arbitrarily chosen element x of a non-empty set X); violations result in a violation (`vp`) with the trivial partition. If x is present as the last decision node, then `hide` will use `do_exists_node` (Figure 10) to construct the proof object for $\exists x. \alpha$. It performs a case distinction whether a satisfaction proof is contained in the partition of this last node. If it is, $\exists x. \alpha$ is satisfied and we compute the smallest (in proof size) such satisfaction proof, taking as our witness an arbitrary element of the respective partition set. Otherwise, all leaves are violations and we obtain a violation proof of $\exists x. \alpha$.

To reuse Lima et al.'s [25] temporal operator evaluation, our state stores a PDT whose leaves are the auxiliary state of these algorithms (instead of proof objects). This allows us to keep the complex auxiliary state and its update unchanged. For example, we use `apply3` to lift Lima et al.'s [25] `update_since` function to two PDTs storing proof objects for subformulas and a third one storing the auxiliary state. The resulting PDT has type $\mathbb{P}(\mathbb{S}_{\text{saux}} \times \bar{\mathbb{p}})$, which we transpose into the desired $\mathbb{P}(\mathbb{S}_{\text{saux}}) \times \mathbb{P}(\bar{\mathbb{p}})$ using `split_prod` and `split_list`.

6 Implementation and Case Study

We implement our algorithm in a new monitoring tool, called WHYMON [2]. Our implementation consists of 4 500 lines of OCaml code and incorporates an optimization of collapsing partition sets with the same stored values both in proof objects and in PDTs. Our formally verified checker contributes additional 1 700 lines of OCaml code generated from our Isabelle formalization, which itself comprises 6 400 lines of definitions and proofs. The checker's main function lifts the validity check of proof objects (\vdash) to PDTs, i.e., `check : trace \rightarrow formula \rightarrow pdt \rightarrow bool`, and is used to certify WHYMON's output. WHYMON includes a visualization [3] implemented in React [20] that consists of 2 400 lines of JavaScript and invokes a JavaScript version of our monitor, generated by `Js_of_ocaml` [32]. Here, we consider the data race policy [18] that captures possible concurrency issues in multithreaded programs on a stream prefix generated by Raszyk [27, Section 4.3]. Furthermore, we consider Nokia's Data-collection Campaign [4], which comes with a stream prefix of around 5 million time-points [1], for which we focus on the *del-2-3* policy [12] controlling data propagation between databases. We describe a violation for each scenario highlighting the advantages of our approach.

Example 4. We first return to Example 3 in our visualization tool, depicted in Figure 11. The table includes TP (time-points), TS (time-stamps), and Values columns. The following columns show the topmost operator of φ' 's subformulas or its predicate names (and their variables). In the Values column, for each of the already evaluated time-points, there is an associated button enclosing a \checkmark (for satisfactions), or a \times (for violations) or both. After clicking on this button, we are presented with a dropdown menu (as in Figure 12) that corresponds to a partition. The listed values are the (potentially multiple) variable assignments of the resulting PDT for that specific time-point. The formula φ' contains two free variables, a and f , and to single out a verdict we must select one value for

each. In particular, at time-point 3 we select $a = \text{Charlie}$ and $f = 152$. Note that in the visualization we focus on readability and omit set parentheses. Moreover, **Other** denotes the complement of the listed values. After choosing the assignments, a Boolean verdict appears in the next column matching the topmost operator of φ' , namely \rightarrow . Clicking on this Boolean verdict reveals and highlights the Boolean verdicts associated with its justification. The subformulas' columns of the current inspection are also highlighted. In this case, the implication is violated because the left side is satisfied, while the $\blacklozenge_{[0,7]}$ subformula is violated. We can explore this verdict further: the violation is justified by those of its subformula at time-points 2 and 3 (the time-points inside the interval are also highlighted). For each time-point, there is another dropdown menu where we can select an assignment for m . Here, the only listed value is **Any**, which corresponds to \mathbb{D} . Thus, the existential quantifier is violated because the subformula $\text{approve}(m, 152)$ is violated for all values that m can be assigned to (\mathbb{D}), and all justifications are identical.

Data Race Detection Multithreaded programs are pervasive and hard to debug. In particular, they are prone to data races, which occur when two threads access (read or write to) a shared address concurrently and at least one of these accesses is a write. Locking mechanisms that synchronize access to variables shared between threads are a plausible solution. We consider the following policy to detect data race potentials [18]:

$$\begin{aligned} \varphi_{dr} &= \text{data race}(t_1, t_2, x) \rightarrow \exists l. (\text{acq nrel}(t_1, x, l) \wedge \text{acq nrel}(t_2, x, l)), \text{ with} \\ \text{data race}(t_1, t_2, x) &= \blacklozenge (\text{read}(t_1, x) \vee \text{write}(t_1, x)) \wedge \blacklozenge \text{write}(t_2, x), \text{ and} \\ \text{acq nrel}(t, x, l) &= \blacksquare ((\text{read}(t, x) \vee \text{write}(t, x)) \rightarrow (\neg \text{rel}(t, l) \mathcal{S} \text{acq}(t, l))) \end{aligned}$$

where the predicates $\text{read}(t, x)$ and $\text{write}(t, x)$ specify read and write operations performed by thread t to shared address x , and $\text{acq}(t, l)$ and $\text{rel}(t, l)$ specify the acquisition and the release of lock l by thread t . Havelund et al. [18] consider a closed formula variant of this policy as their tool, DEJAVU, only supports closed formulas. In contrast, WHYMON supports open formulas. We consider the stream prefix:

$$\begin{aligned} & \langle (0, \{\text{acq}(9, 9)\}), (1, \{\text{read}(9, 3)\}), (2, \{\text{acq}(13, 19)\}), (3, \{\text{acq}(15, 3)\}), \\ & (4, \{\text{acq}(18, 15)\}), (5, \{\text{read}(13, 5)\}), (6, \{\text{write}(15, 4)\}), (7, \{\text{write}(15, 3)\}), \dots \rangle \end{aligned}$$

At time-point 7, WHYMON outputs a PDT with non-trivial assignments. We focus on the single violation in this PDT, which corresponds to the assignment $(\{9\}, \{15\}, \{3\})$ for (t_1, t_2, x) . This violation is shown in Figure 13. The topmost operator of φ_{dr} is an implication, and it is violated because the left side is satisfied (there was a data race), while the right side (the lock requirement) is violated. Specifically, the data race occurred because thread $t_1 = 9$ read address 3 at time-point 1, satisfying the $\blacklozenge_{[0,\infty]}$ subformula in the left conjunct, and thread $t_2 = 15$ wrote to address 3 at the current time-point 7, satisfying the $\blacklozenge_{[0,\infty]}$ subformula in the right conjunct. Moving to the right side of the implication, the violation of the existential indicates that its subformula is violated for every value of \mathbb{D} . In particular, the subsets of the domain $\{9\}$ and $\{9\}^c$ are each associated with a different violation. Here, we focus on the violation where $l = 9$. The subformula is a conjunction, and to be violated it suffices that one of the conjuncts is violated. This violation stems from the violation of the right conjunct $\blacksquare_{[0,\infty]}$ (note that $t_2 = 15$ is listed as the variable in the predicate columns). We omit the columns referring to the left conjunct, since all

TP	TS	Values	publish(a, f)	S10(a)	mgr_Firm(a)	mgr_Sim(a)	approve(m, f)
0	0	✓	✓	✓	✓	✓	✓
1	0	✓	✓	✓	✓	✓	✓
2	4	✓	✓	✓	✓	✓	✓
3	10	✗	✓	✓	✓	✓	✓

Fig. 11: Visualization of φ' 's violation at time-point 3 for ($\{\text{Charlie}\}, \{152\}$).

Fig. 12: Assignment selection for φ' at time-point 3.

TP	TS	Values	read(t1, x)	write(t2, x)	S10(a)	read(t2, x)	write(t2, x)	acquit(f)
0	0	✓	✓	✓	✓	✓	✓	✓
1	1	✓	✓	✓	✓	✓	✓	✓
2	2	✓	✓	✓	✓	✓	✓	✓
3	3	✓	✓	✓	✓	✓	✓	✓
4	4	✓	✓	✓	✓	✓	✓	✓
5	5	✓	✓	✓	✓	✓	✓	✓
6	6	✓	✓	✓	✓	✓	✓	✓
7	7	✗	✓	✓	✓	✓	✓	✓
8	8	✓	✓	✓	✓	✓	✓	✓
9	9	✓	✓	✓	✓	✓	✓	✓

Fig. 13: Visualization of φ_{del} 's violation at time-point 7 for ($\{9\}, \{15\}, \{3\}$).

TP	TS	Values	delete(x, dt2, y, data)	S10(a)	delete(u, dt3, v, data)	Formula
79	1304/2011, 13.57.02	✗	✓	✓	✓	✓
80	1304/2011, 13.57.03	✓	✓	✓	✓	✓
81	1304/2011, 13.57.05	✓	✓	✓	✓	✓
82	1304/2011, 13.57.13	✓	✓	✓	✓	✓
83	1304/2011, 13.57.30	✓	✓	✓	✓	✓
84	1304/2011, 13.57.32	✓	✓	✓	✓	✓
85	1304/2011, 13.59.32	✓	✓	✓	✓	✓
86	1304/2011, 14.08.01	✓	✓	✓	✓	✓
87	1304/2011, 14.08.34	✓	✓	✓	✓	✓

Fig. 14: Visualization of φ_{del} 's violation at time-point 79 for ($\{189810327\}, \{\text{user2}\}, \{\text{unknown}\}$).

entries are empty. Once again, the implication is violated because the left side is satisfied, i.e., thread $t_2 = 15$ wrote to address 3 at time-point 7, satisfying the disjunction, but $\mathcal{S}_{[0,\infty)}$ on the right side is violated, because thread $t_2 = 15$ never acquired the lock $l = 9$.

Data Propagation Nokia’s Data-collection Campaign [4] used three databases db_1 , db_2 and db_3 in the collection of sensitive information from mobile phones of participants. We focus on the policy φ_{del} [12], which controls the data propagation between databases db_2 and db_3 : if *data* is deleted from db_2 , then it must be deleted from db_3 within 1 minute.

$$\varphi_{del} = \text{delete}(x, db_2, y, data) \wedge data \not\approx [\text{unknown}] \rightarrow \diamond_{[0,60]} \exists u, v. \text{delete}(u, db_3, v, data)$$

where db_2 , db_3 , and $[\text{unknown}]$ are constants and $\text{delete}(db_{user}, db, p_{id}, data)$ specifies the deletion of *data* from participant p_{id} from database db using database user db_{user} . We used the REPLAYER tool [22] to convert the stream prefix to WHYMON’s format. We executed WHYMON’s command line interface with the entire prefix and found two violations. The following experiments were conducted on a computer with an Apple M1 Chip (8 cores) and 16GB of RAM. WHYMON took 17m51s to process the entire prefix. We also executed MONPOLY with a slightly modified yet equivalent policy (due to monitorability restrictions), and its running time amounted to 1m10s. MONPOLY outperforms WHYMON, but we must acknowledge the different outputs both monitors produce. MONPOLY only outputs variable assignments, whereas WHYMON outputs entire PDTs containing all assignments and a justification of the verdict in the form of a proof tree for each. We extract 100 time-points containing both violations and focus on the violation at time-point 79 for the assignment $(\{189810327\}, \{\text{user2}\}, \{[\text{unknown}]\})$ for $(data, x, y)$, depicted in Figure 14. Time-stamps are converted to actual dates (by enabling the option) and we omit time-points that do not contain relevant events for the violation. Let

$$\begin{aligned} \Gamma_{79} &= \{\text{delete}(\text{user2}, db_2, [\text{unknown}], 189810327), \\ \Gamma_{80} &= \{\text{delete}(\text{triggers}, db_3, [\text{unknown}], [\text{unknown}])\}, \\ \Gamma_{81} &= \{\text{delete}(\text{user2}, db_2, [\text{unknown}], 189810328)\}, \text{ and } \Gamma_{82} = \Gamma_{83} = \Gamma_{84} = \emptyset. \end{aligned}$$

The implication is violated because the left side is satisfied (there was a deletion at the current time-point 79), but $\diamond_{[0,59]}$ is violated. Note that $[0, 60)$ was replaced with the equivalent interval $[0, 59]$. For each time-point of $[E_{79}^f([0, 59]), L_{79}^f([0, 59])] = \{79, \dots, 84\}$, the subformula is violated. Regardless of the values we assign to u and v (all violations are identical), the subformula $\text{delete}(u, db_3, v, 189810327)$ is violated.

7 Conclusion

We describe an approach for MFOTL monitoring with verdicts in the form of proof objects for every free variable assignment. Such verdicts are useful for understandability and certification, which increases the monitor’s trustworthiness. We implement our approach in the tool WHYMON along with an interactive visualization for these verdicts, which we invite the reader to explore [3]. As future work, we plan to provide support for equality between variables and to improve our monitor’s performance by, e.g., stream slicing [29].

Data Availability Statement Our artifact [26] includes WHYMON’s source code at the artifact submission time together with instructions on how to set up WHYMON locally, extract our PDT checker, execute our examples, and replicate our case study.

Acknowledgements This research is supported by a Novo Nordisk Fonden start package grant (NNF20OC0063462). We thank David Basin, François Hublet, Srđan Krstić, Matthias Lott, Joshua Schneider for their suggestions on WHYMON’s and EXPLANATOR2’s user interfaces. We are also grateful to anonymous TACAS 2024 reviewers, who helped us improve the presentation of this paper with their valuable comments.

References

1. The Nokia case study log file (2014), <https://sourceforge.net/projects/monpoly/files/ldcc.tar/download>
2. WHYMON repository (2023), <https://github.com/runtime-monitoring/whymon>
3. WHYMON web interface (2023), <https://runtime-monitoring.github.io/whymon>
4. Aad, I., Niemi, V.: NRC data collection campaign and the privacy by design principles. In: Proceedings of the International Workshop on Sensing for App Phones (PhoneSense) (2010)
5. Ailamazyan, A.K., Gilula, M.M., Stolboushkin, A.P., Schwartz, G.F.: Reduction of a relational model with infinite domains to the case of finite domains. *Doklady Akademii Nauk SSSR* **286**(2), 308–311 (1986), <http://mi.mathnet.ru/dan47310>
6. Alur, R., Henzinger, T.A.: Real-time logics: Complexity and expressiveness. *Inf. Comput.* **104**(1), 35–77 (1993). <https://doi.org/10.1006/inco.1993.1025>
7. Arfelt, E., Basin, D.A., Debois, S.: Monitoring the GDPR. In: Sako, K., Schneider, S.A., Ryan, P.Y.A. (eds.) *ESORICS 2019*. LNCS, vol. 11735, pp. 681–699. Springer (2019). https://doi.org/10.1007/978-3-030-29959-0_33
8. Basin, D.A., Bhatt, B.N., Krstić, S., Traytel, D.: Almost event-rate independent monitoring. *Formal Methods Syst. Des.* **54**(3), 449–478 (2019). <https://doi.org/10.1007/s10703-018-00328-3>
9. Basin, D.A., Bhatt, B.N., Traytel, D.: Optimal proofs for linear temporal logic on lasso words. In: Lahiri, S.K., Wang, C. (eds.) *ATVA 2018*. LNCS, vol. 11138, pp. 37–55. Springer (2018). https://doi.org/10.1007/978-3-030-01090-4_3
10. Basin, D.A., Caronni, G., Ereth, S., Harvan, M., Klaedtke, F., Mantel, H.: Scalable offline monitoring of temporal specifications. *Formal Methods Syst. Des.* **49**(1-2), 75–108 (2016). <https://doi.org/10.1007/s10703-016-0242-y>
11. Basin, D.A., Dietiker, D.S., Krstić, S., Pignolet, Y., Raszyk, M., Schneider, J., Ter-Gabrielyan, A.: Monitoring the internet computer. In: Chechik, M., Katoen, J., Leucker, M. (eds.) *FM 2023*. LNCS, vol. 14000, pp. 383–402. Springer (2023). https://doi.org/10.1007/978-3-031-27481-7_22
12. Basin, D.A., Harvan, M., Klaedtke, F., Zalinescu, E.: Monitoring data usage in distributed systems. *IEEE Trans. Software Eng.* **39**(10), 1403–1426 (2013). <https://doi.org/10.1109/TSE.2013.18>
13. Basin, D.A., Klaedtke, F., Müller, S.: Monitoring security policies with metric first-order temporal logic. In: Joshi, J.B.D., Carminati, B. (eds.) *SACMAT 2010*. pp. 23–34. ACM (2010). <https://doi.org/10.1145/1809842.1809849>
14. Basin, D.A., Klaedtke, F., Müller, S., Zalinescu, E.: Monitoring metric first-order temporal properties. *J. ACM* **62**(2), 15:1–15:45 (2015). <https://doi.org/10.1145/2699444>
15. Basin, D.A., Klaedtke, F., Zalinescu, E.: The MonPoly monitoring tool. In: Reger, G., Havelund, K. (eds.) *RV-CuBES 2017*. Kalpa Publications in Computing, vol. 3, pp. 19–28. EasyChair (2017). <https://doi.org/10.29007/89hs>
16. Bryant, R.E.: Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers* **35**(8), 677–691 (1986). <https://doi.org/10.1109/TC.1986.1676819>

17. Havelund, K., Peled, D.: BDDs for representing data in runtime verification. In: Deshmukh, J., Nickovic, D. (eds.) RV 2020. LNCS, vol. 12399, pp. 107–128. Springer (2020). https://doi.org/10.1007/978-3-030-60508-7_6
18. Havelund, K., Peled, D., Ulus, D.: First-order temporal logic monitoring with BDDs. *Formal Methods Syst. Des.* **56**(1), 1–21 (2020). <https://doi.org/10.1007/s10703-018-00327-4>
19. Hull, R., Su, J.: Domain independence and the relational calculus. *Acta Informatica* **31**(6), 513–524 (1994). <https://doi.org/10.1007/BF01213204>
20. Hunt, P., O’Shannessy, P., Smith, D., Coatta, T.: React: Facebook’s functional turn on writing JavaScript. *ACM Queue* **14**(4), 40 (2016). <https://doi.org/10.1145/2984629.2994373>
21. Koymans, R.: Specifying real-time properties with metric temporal logic. *Real Time Syst.* **2**(4), 255–299 (1990). <https://doi.org/10.1007/BF01995674>
22. Krstić, S., Schneider, J.: A benchmark generator for online first-order monitoring. In: Deshmukh, J., Nickovic, D. (eds.) RV 2020. LNCS, vol. 12399, pp. 482–494. Springer (2020). https://doi.org/10.1007/978-3-030-60508-7_27
23. Kupferman, O., Vardi, M.Y.: Vacuity detection in temporal model checking. *Int. J. Softw. Tools Technol. Transf.* **4**(2), 224–233 (2003). <https://doi.org/10.1007/s100090100062>
24. Leucker, M., Sánchez, C., Scheffel, T., Schmitz, M., Schramm, A.: TeSSLa: runtime verification of non-synchronized real-time streams. In: Haddad, H.M., Wainwright, R.L., Chbeir, R. (eds.) SAC 2018. pp. 1925–1933. ACM (2018). <https://doi.org/10.1145/3167132.3167338>
25. Lima, L., Herasimau, A., Raszyk, M., Traytel, D., Yuan, S.: Explainable online monitoring of metric temporal logic. In: TACAS 2023. LNCS, vol. 13994, pp. 473–491. Springer (2023). https://doi.org/10.1007/978-3-031-30820-8_28
26. Lima, L., Huerta y Munive, J.J., Traytel, D.: Artifact for "Explainable online monitoring of metric first-order temporal logic" (2024). <https://doi.org/10.5281/zenodo.10439544>
27. Raszyk, M.: Efficient, Expressive, and Verified Temporal Query Evaluation. Ph.D. thesis, ETH Zürich (2022). <https://doi.org/10.3929/ethz-b-000553221>
28. Raszyk, M., Basin, D.A., Krstić, S., Traytel, D.: Multi-head monitoring of metric temporal logic. In: Chen, Y., Cheng, C., Esparza, J. (eds.) ATVA 2019. LNCS, vol. 11781, pp. 151–170. Springer (2019). https://doi.org/10.1007/978-3-030-31784-3_9
29. Schneider, J., Basin, D.A., Brix, F., Krstić, S., Traytel, D.: Scalable online first-order monitoring. *Int. J. Softw. Tools Technol. Transf.* **23**(2), 185–208 (2021). <https://doi.org/10.1007/s10009-021-00607-1>
30. Schneider, J., Basin, D.A., Krstić, S., Traytel, D.: A formally verified monitor for metric first-order temporal logic. In: Finkbeiner, B., Mariani, L. (eds.) RV 2019. LNCS, vol. 11757, pp. 310–328. Springer (2019). https://doi.org/10.1007/978-3-030-32079-9_18
31. Ulus, D.: Online monitoring of metric temporal logic using sequential networks. *CoRR abs/1901.00175* (2019). <https://doi.org/10.48550/arxiv.1901.00175>
32. Vouillon, J., Balat, V.: From bytecode to JavaScript: the Js_of_ocaml compiler. *Softw. Pract. Exp.* **44**(8), 951–972 (2014). <https://doi.org/10.1002/spe.2187>