

An LTL Proof System for Runtime Verification

Clare Cini¹ and Adrian Francalanza¹

Computer Science, ICT, University of Malta
{clare.cini.08,adrian.francalanza}@um.edu.mt

Abstract. We propose a local proof system for LTL formalising deductions within the constraints of Runtime Verification (RV), and show how such a system can be used as a basis for the construction of online runtime monitors. Novel soundness and completeness results are proven for this system. We also prove decidability and incrementality properties for a monitoring algorithm constructed from it. Finally, we relate its expressivity to existing symbolic analysis techniques used in RV.

1 Introduction

Runtime verification (RV) is a lightweight verification technique that checks whether the current execution of a system under scrutiny satisfies or violates a given correctness property. It has its origins in model checking, as a more scalable (yet still formal) approach to program verification where state explosion problems (which are part and parcel of model checking) are mitigated [LS09]. Linear Temporal Logic, (LTL) [Pnu77] is prevalently used for formal expositions of RV [Gei01,SRA04,BLS07,BLS10,BLS11,BF12], because it has a pleasingly straightforward definition over strings, denoting execution traces.

Proof systems [Bus98,TS00] embody mechanical syntactic deductions similar to those made by monitors in RV. We propose a proof system for LTL attuned to the constraints of an RV setting, and show how it can be used as a basis for monitor construction. Although deductive systems for LTL exist, e.g., [MP91,KI11,BL08] they are geared towards reasoning about the full statespace of a system. By contrast, our proof system is *local* [SW91] focussing on checking whether a specific point lies within a property set, instead of interpreting a formula *wrt.* a set of points; this mirrors closely the runtime analysis in RV.

RV settings pose further constraints on our symbolic analysis. In online settings, deductions are often performed on the *partial* traces generated thus far, while the program is still executing. This has two important consequences: (a) *conclusive* deductions must be *consistent with any extension* leading to a complete trace (b) in order to keep RV overheads low, *inconclusive* deductions must be reusable, and contribute to deductions of subsequent extensions *i.e.*, the analysis must be *incremental*. In addition, monitors for partial traces typically reason about trace *satisfactions*, but also trace *violations* [BLS11,BLS10] so as to determine good/bad prefixes [KYV01]. Accordingly, proof system deductions should reason directly about both satisfactions and violations. Moreover, timely detections often require *synchronous* monitor instrumentation where monitored

system execute in lock-step with the respective monitor, producing a trace event and waiting for the monitor to terminate its (incremental) analysis before executing further. Thus, in order for such an instrumentation to be safe, it is important to ensure that incremental deductions are *decidable*.

We formally compare our system with other LTL symbolic techniques used in RV. In particular we consider Geilen’s work [Gei01], based on informative prefixes [KYV01], as well as Sen *et al.*’s work [SRA04] which is based on derivatives [HR01]. Apart from enabling a better understanding of each approach, facilitating comparisons between seemingly different formalisations, this study enables cross fertilisation of techniques from one formalisation to the other.

The paper is structured as follows. After introducing the logic, §2, we present our proof system in §3. §4 presents the associated monitor algorithm. §5 details formal comparisons with other symbolic analyses and §6 concludes.

2 The Logic: An LTL Primer

Syntax. Fig. 1 defines the *core* syntax of LTL as used in [BLS11,EFH⁺03], parameterised by a set of predicates $p \in \text{PRED}$. It consists of two base cases, *i.e.*, the true formula, **tt**, and a predicate formula, p , standard negation and conjunction constructors, $\neg\psi$ and $\psi_1 \wedge \psi_2$, and the characteristic *next* and *until* formulas, $X\psi$ and $\psi_1 \text{ U } \psi_2$ *resp.* Other studies of LTL (*e.g.*, [Gei01,BL08]) prefer to work with formulas in *negation normal form* (nnf), where negations are pushed to the leaves of a formula. To accommodate this, we also consider an extended LTL syntax in Fig. 1, that also includes base formulas for falsity, **ff**, and constructors such as disjunctions, $\varphi_1 \vee \varphi_2$, and release formulas, $\varphi_1 \text{ R } \varphi_2$. Our extended syntax also employs an extended predicate notation that includes *co-predicates*, *i.e.*, for any predicate¹ $p = S \subseteq \Sigma$, its co-predicate, denoted as \bar{p} , represents its *dual* and is defined as $\Sigma \setminus S$. This allows us to eliminate negations from normalised formulas; because of this we sometimes refer to an nnf formula as *negation-free*. Fig. 1 also defines a translation function, $\langle - \rangle :: \text{LTL} \rightarrow \text{ELTL}$ from formulas of the core LTL to a negation-free formula in the extended syntax.

Model. The logic semantics is also given in Fig. 1. It assumes an alphabet, Σ (with element variables σ), over which predicates are defined, $p :: \Sigma \rightarrow \text{BOOL}$. As in other RV studies [Gei01,SRA04,BLS11], the logic is defined over *infinite* strings, $s \in \Sigma^\omega$; *finite* strings over the same alphabet are denoted by the variable $t \in \Sigma^*$. A string with element σ at its head is denoted as σs (*resp.* σt). For indexes $i, j \in \text{NAT}$, s_i denotes the i^{th} *element* in the string (starting from index 0) and $[s]^i$ denotes the *suffix* of s starting at index i ; note that for any s , $[s]^0 = s$. Infinite strings with a regular (finite) pattern t are sometimes denoted as t^* , whereas the shorthand $t \dots$ represents infinite strings with a (finite) prefix t .

Semantics. The denotational semantic function $\llbracket - \rrbracket :: \text{ELTL} \rightarrow \mathcal{P}(\Sigma^\omega)$ is defined by induction over the structure of LTL formulas; in Fig. 1 we define the

¹ Predicates are sometimes denoted as sets over Σ .

Core LTL Syntax

$$\psi \in \text{LTL} ::= \text{tt} \mid p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \mathbf{X}\psi \mid \psi_1 \mathbf{U} \psi_2$$

Extended LTL Syntax

$$\begin{aligned} \varphi \in \text{ELTL} ::= & \text{tt} \mid p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathbf{U} \varphi_2 \mid \mathbf{X}\varphi \mid \neg\varphi \\ & \mid \text{ff} \mid \bar{p} \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \mathbf{R} \varphi_2 \end{aligned}$$

Formula Translation (Normalisation)

$$\begin{aligned} \langle \text{tt} \rangle &\stackrel{\text{def}}{=} \text{tt} & \langle \neg\text{tt} \rangle &\stackrel{\text{def}}{=} \text{ff} & \langle p \rangle &\stackrel{\text{def}}{=} p & \langle \neg p \rangle &\stackrel{\text{def}}{=} \bar{p} \\ \langle \mathbf{X}\psi \rangle &\stackrel{\text{def}}{=} \mathbf{X}\langle \psi \rangle & \langle \neg\mathbf{X}\psi \rangle &\stackrel{\text{def}}{=} \mathbf{X}\langle \neg\psi \rangle & \langle \neg\neg\psi \rangle &\stackrel{\text{def}}{=} \langle \psi \rangle \\ \langle \psi_1 \wedge \psi_2 \rangle &\stackrel{\text{def}}{=} \langle \psi_1 \rangle \wedge \langle \psi_2 \rangle & \langle \neg(\psi_1 \wedge \psi_2) \rangle &\stackrel{\text{def}}{=} \langle \neg\psi_1 \rangle \vee \langle \neg\psi_2 \rangle \\ \langle \psi_1 \mathbf{U} \psi_2 \rangle &\stackrel{\text{def}}{=} \langle \psi_1 \rangle \mathbf{U} \langle \psi_2 \rangle & \langle \neg(\psi_1 \mathbf{U} \psi_2) \rangle &\stackrel{\text{def}}{=} \langle \neg\psi_1 \rangle \mathbf{R} \langle \neg\psi_2 \rangle \end{aligned}$$

Semantics

$$\begin{aligned} \llbracket \text{tt} \rrbracket &\stackrel{\text{def}}{=} \Sigma^\omega & \llbracket \text{ff} \rrbracket &\stackrel{\text{def}}{=} \emptyset \\ \llbracket p \rrbracket &\stackrel{\text{def}}{=} \{s \mid p(s_0)\} & \llbracket \bar{p} \rrbracket &\stackrel{\text{def}}{=} \{s \mid \text{not } p(s_0)\} \\ \llbracket \varphi_1 \wedge \varphi_2 \rrbracket &\stackrel{\text{def}}{=} \llbracket \varphi_1 \rrbracket \cap \llbracket \varphi_2 \rrbracket & \llbracket \varphi_1 \vee \varphi_2 \rrbracket &\stackrel{\text{def}}{=} \llbracket \varphi_1 \rrbracket \cup \llbracket \varphi_2 \rrbracket \\ \llbracket \mathbf{X}\psi \rrbracket &\stackrel{\text{def}}{=} \{s \mid [s]^1 \in \llbracket \psi \rrbracket\} & \llbracket \neg\varphi \rrbracket &\stackrel{\text{def}}{=} (\Sigma^\omega) \setminus \llbracket \varphi \rrbracket \\ \llbracket \varphi_1 \mathbf{U} \varphi_2 \rrbracket &\stackrel{\text{def}}{=} \{s \mid \exists j \text{ such that } [s]^j \in \llbracket \varphi_2 \rrbracket \text{ and } (i < j \text{ implies } [s]^i \in \llbracket \varphi_1 \rrbracket)\} \\ \llbracket \varphi_1 \mathbf{R} \varphi_2 \rrbracket &\stackrel{\text{def}}{=} \{s \mid \forall j \text{ we have } ([s]^j \in \llbracket \varphi_2 \rrbracket) \text{ or } (\exists i < j \text{ such that } [s]^i \in \llbracket \varphi_1 \rrbracket)\} \end{aligned}$$

Fig. 1. Linear Temporal Logic Syntax and Semantics

semantics for the extended LTL syntax (of which the core syntax is a subset). Most cases are standard. For instance, $\llbracket \text{tt} \rrbracket$ (resp. $\llbracket \text{ff} \rrbracket$) returns the universal (resp. empty) set of strings, $\llbracket \neg\varphi \rrbracket$ returns the dual of $\llbracket \varphi \rrbracket$, whereas $\llbracket \varphi_1 \wedge \varphi_2 \rrbracket$ (resp. $\llbracket \varphi_1 \vee \varphi_2 \rrbracket$) denotes the intersection (resp. union) of the meaning of its subformulas, $\llbracket \varphi_1 \rrbracket$ and $\llbracket \varphi_2 \rrbracket$. The meaning of $\llbracket p \rrbracket$ (resp. $\llbracket \bar{p} \rrbracket$) contains all strings whose *first element* satisfies the predicate p (resp. \bar{p}). The temporal formulas are more involving. The denotation of $\llbracket \mathbf{X}\varphi \rrbracket$ contains all strings whose immediate suffix (i.e., at index 1) is included in $\llbracket \varphi \rrbracket$. Until formulas $\llbracket \varphi_1 \mathbf{U} \varphi_2 \rrbracket$ contain all strings that contain a suffix (at *some* index j) satisfying $\llbracket \varphi_2 \rrbracket$, and *all* the suffixes preceding j satisfy $\llbracket \varphi_1 \rrbracket$. Finally, release formulas, $\varphi_1 \mathbf{R} \varphi_2$ contain strings whose suffixes *always* satisfy φ_2 , as well as strings that contain a suffix satisfying both φ_1 and φ_2 and *all* the preceding suffixes satisfying φ_2 .

The denotational semantics allows us to observe the duality between the formulas tt , $\varphi_1 \wedge \varphi_2$ and $\psi_1 \mathbf{U} \psi_2$, and their counterparts ff , $\varphi_1 \vee \varphi_2$ and $\psi_1 \mathbf{R} \psi_2$. It also helps us understand the mechanics of the translation function, pushing negation to the leaves of a formula using negation propagation identities (e.g., DeMorgan's law), converting constructors to their dual constructor; at the leaves the function then performs direct translations from tt and p to ff and \bar{p}

resp. The semantics also allows us to prove Prop. 1, justifying the use of a corresponding negation-free formula instead of a core LTL formula, in order to reason exclusively in terms of positive interpretations. It states that (i) the translation function is *total* (otherwise the equality cannot be determined) but also that (ii) the translated formula *preserves the semantic meaning* of the original formula.

Proposition 1. *For any $\psi \in \text{LTL}$, $\llbracket \psi \rrbracket = \llbracket \langle \psi \rangle \rrbracket$*

Example 1. The behaviour of a traffic-light system may be described by observing its states consisting of green, g , orange, o , and red, r . Complete executions may thus be represented as traces (strings) over the alphabet $\Sigma = \{g, o, r\}$. Predicate definitions may be defined as sets over this alphabet Σ e.g., $\text{st} = \{o, r\}$ is true for *stopping* actions o and r ; singleton-set predicates are denoted by single-letter names e.g., $\text{g} = \{g\}$. We can specify the following properties:

- $(\neg r) \wedge Xr$ describes a trace where the system is not in a red state initially, but turns red at the next instant. e.g., $gr\dots$ or $or\dots$;
- $\text{g} \text{ U } \text{o}$ describes traces that eventually switch to the orange state from a green state e.g., $go\dots$ or $gggo\dots$;
- Gst , i.e., always st , which is shorthand for $\neg(\text{tt} \text{ U } \neg\text{st})$, describes traces that contain only stopping states, e.g., strings of the form $(or)^*$ or r^* .

To determine whether $r^* \in \llbracket \text{Gst} \rrbracket$, we can use $\llbracket \neg(\text{tt} \text{ U } \neg\text{st}) \rrbracket$ for which we would need to calculate $\llbracket \text{tt} \text{ U } \neg\text{st} \rrbracket$ and then take its dual. Alternatively, we can calculate the denotation of $\langle \neg(\text{tt} \text{ U } \neg\text{st}) \rangle$, which translates to ff R st , and check inclusion wrt. the translated negation-free formula, safe in the knowledge that $\llbracket \neg(\text{tt} \text{ U } \neg\text{st}) \rrbracket = \llbracket \text{ff R st} \rrbracket$. Using similar reasoning, to determine whether $r\dots \notin \llbracket (\neg r) \wedge Xr \rrbracket$, we can check whether $r\dots \in \llbracket r \vee Xr \rrbracket$ holds. ■

3 An Online Monitoring Proof System

Online² runtime verification of LTL properties consist in determining whether the current execution satisfies (or violates) a property from *the trace generated thus far*. We present a local proof system [SW91,BS92] that characterises such runtime analysis, and allows us to determine whether any (complete) trace ts with *finite* prefix t is included in (or excluded from) $\llbracket \varphi \rrbracket$. The proof system is defined as the least relation satisfying the rules in Fig. 2. These rules employ two, mutually dependent, judgements: the sequent $t \vdash^+ \varphi$ denotes a *satisfaction* judgement, whereas $t \vdash^- \psi$ denotes a *violation* judgement; note the polarity differentiating the two judgements, i.e., $+$ and $-$.

Fig. 2 includes three satisfaction axioms (PTRU, PPRD and PCOP) and three violation axioms (NFLS, NPRD and NCOP); the system is parametric wrt. the pre-computation of predicates and co-predicates, p and \bar{p} . The conjunction and disjunction rules, PAND, POR1 and POR2 (resp. NAND1, NAND2 and NOR) *decompose* the composite formula of the judgement for their premises. The negation

² By contrast, *offline* monitoring typically works on *complete* execution traces. [RH05]

Satisfaction Rules

$$\begin{array}{l}
\text{PTRU} \frac{}{t \vdash^+ \text{tt}} \quad \text{PPRD} \frac{p(\sigma)}{\sigma t \vdash^+ p} \quad \text{PCOP} \frac{\bar{p}(\sigma)}{\sigma t \vdash^+ \bar{p}} \quad \text{PNEG} \frac{t \vdash^- \varphi}{t \vdash^+ \neg \varphi} \\
\text{PAND} \frac{t \vdash^+ \varphi_1 \quad t \vdash^+ \varphi_2}{t \vdash^+ \varphi_1 \wedge \varphi_2} \quad \text{PNXT} \frac{t \vdash^+ \varphi}{\sigma t \vdash^+ X\varphi} \\
\text{POR1} \frac{t \vdash^+ \varphi_1}{t \vdash^+ \varphi_1 \vee \varphi_2} \quad \text{POR2} \frac{t \vdash^+ \varphi_2}{t \vdash^+ \varphi_1 \vee \varphi_2} \\
\text{PUNT1} \frac{t \vdash^+ \varphi_2}{t \vdash^+ \varphi_1 \text{ U } \varphi_2} \quad \text{PUNT2} \frac{\sigma t \vdash^+ \varphi_1 \quad t \vdash^+ \varphi_1 \text{ U } \varphi_2}{\sigma t \vdash^+ \varphi_1 \text{ U } \varphi_2} \\
\text{PREL1} \frac{t \vdash^+ \varphi_1 \quad t \vdash^+ \varphi_2}{t \vdash^+ \varphi_1 \text{ R } \varphi_2} \quad \text{PREL2} \frac{\sigma t \vdash^+ \varphi_2 \quad t \vdash^+ \varphi_1 \text{ R } \varphi_2}{\sigma t \vdash^+ \varphi_1 \text{ R } \varphi_2}
\end{array}$$

Violation Rules

$$\begin{array}{l}
\text{NFLS} \frac{}{t \vdash^- \text{ff}} \quad \text{NPRD} \frac{\bar{p}(\sigma)}{\sigma t \vdash^- p} \quad \text{NCOP} \frac{p(\sigma)}{\sigma t \vdash^- \bar{p}} \quad \text{NNEG} \frac{t \vdash^+ \varphi}{t \vdash^- \neg \varphi} \\
\text{NOR} \frac{t \vdash^- \varphi_1 \quad t \vdash^- \varphi_2}{t \vdash^- \varphi_1 \vee \varphi_2} \quad \text{NNXT} \frac{t \vdash^- \psi}{\sigma t \vdash^- X\psi} \\
\text{NAND1} \frac{t \vdash^- \varphi_1}{t \vdash^- \varphi_1 \wedge \varphi_2} \quad \text{NAND2} \frac{t \vdash^- \varphi_2}{t \vdash^- \varphi_1 \wedge \varphi_2} \\
\text{NUNT1} \frac{t \vdash^- \varphi_1 \quad t \vdash^- \varphi_2}{t \vdash^- \varphi_1 \text{ U } \varphi_2} \quad \text{NUNT2} \frac{\sigma t \vdash^- \varphi_2 \quad t \vdash^- \varphi_1 \text{ U } \varphi_2}{\sigma t \vdash^- \varphi_1 \text{ U } \varphi_2} \\
\text{NREL1} \frac{t \vdash^- \varphi_2}{t \vdash^- \varphi_1 \text{ R } \varphi_2} \quad \text{NREL2} \frac{\sigma t \vdash^- \varphi_1 \quad t \vdash^- \varphi_1 \text{ R } \varphi_2}{\sigma t \vdash^- \varphi_1 \text{ R } \varphi_2}
\end{array}$$

Fig. 2. Satisfaction and Violation Proof Rules

rules PNEG and NNEG also decompose the formula, but *switch the modality* of the sequents for their premises, *transitioning* from one judgement form to the other. Specifically, in the case of PNEG, the satisfaction sequent $t \vdash^+ \neg \varphi$ is defined in terms of the *violation* sequent $t \vdash^- \varphi$ (and dually for NNEG).

The rules for the temporal formulas may decompose judgement formulas, e.g., PUNT1, PREL1, NUNT1, NREL1, but may also analyse suffixes of the trace *in incremental fashion*. For instance, in order to prove $\sigma t \vdash^+ X\varphi$, rule PNXT requires the satisfaction judgement to hold for the *immediate* suffix t and the subformula φ , i.e., $t \vdash^+ \varphi$. Similarly, to prove the satisfaction sequent $\sigma t \vdash^+ \varphi_1 \text{ U } \varphi_2$, rule PUNT2 requires a satisfaction proof of the current trace σt and the subformula φ_1 , as well as a satisfaction proof of the immediate suffix t wrt. $\varphi_1 \text{ U } \varphi_2$. Since this suffix premise is wrt. to the same composite formula $\varphi_1 \text{ U } \varphi_2$, it may well be the case that PUNT2 is applied again for suffix t . In fact, satisfaction proofs for until formulas are characterised by a series of PUNT2 applications, followed by an application of rule PUNT1 (the satisfaction proofs for $\varphi_1 \text{ R } \varphi_2$ and violation

proofs for $\varphi_1 \text{ U } \varphi_2$ and $\varphi_1 \text{ R } \varphi_2$ follow an analogous structure). This incremental analysis structure mirrors that of RV algorithms for LTL [Gei01,SRA04,BLS11] and contrasts with the descriptive nature of the *resp.* semantic definition for $\varphi_1 \text{ U } \varphi_2$ (Fig. 1) (which merely stipulates the *existence* of some index j at which point φ_2 holds without stating *how* to find this index).

We note the inherent symmetry between the satisfaction and violation rules, internalising the negation-propagation mechanism of the normalisation function $\langle - \rangle$ of §2 through rules PNEG and NNEG. For instance, there are no satisfaction (*resp.* violation) proof rules for the formula **ff** (*resp.* **tt**). The *resp.* predicate axioms for satisfactions and violations are dual to one another, as are the rules for conjunctions and disjunctions. More precisely, following $\langle \neg(\psi_1 \wedge \psi_2) \rangle \stackrel{\text{def}}{=} \langle \neg\psi_1 \rangle \vee \langle \neg\psi_2 \rangle$ from Fig. 1, the violation rules for conjunctions (NAND1 and NAND2) have the same structure as the satisfaction rules for the *resp.* disjunctions (POR1 and POR2). The symmetric structure carries over to the temporal proof rules as well, *e.g.*, violation rules NUNT1 and NUNT2 have an analogous structure to that of rules PREL1 and PREL2.

Example 2. Recall property $\mathbf{g} \text{ U } \mathbf{o}$ from Ex. 1. We can construct the satisfaction proof for trace go and the violation proof for trace gr below:

$$\begin{array}{c} \text{PPRD} \frac{\mathbf{g}(g)}{go \vdash^+ \mathbf{g}} \quad \text{PUNT1} \frac{\text{PPRD} \frac{\mathbf{o}(o)}{o \vdash^+ \mathbf{o}}}{o \vdash^+ \mathbf{g} \text{ U } \mathbf{o}} \quad \text{NPRD} \frac{\bar{\mathbf{o}}(g)}{gr \vdash^- \mathbf{o}} \quad \text{NUNT1} \frac{\text{NPRD} \frac{\bar{\mathbf{g}}(r)}{r \vdash^- \mathbf{g}} \quad \text{NPRD} \frac{\bar{\mathbf{o}}(r)}{r \vdash^- \mathbf{o}}}{r \vdash^- \mathbf{g} \text{ U } \mathbf{o}} \\ \text{PUNT2} \frac{\quad}{go \vdash^+ \mathbf{g} \text{ U } \mathbf{o}} \quad \text{NUNT2} \frac{\quad}{gr \vdash^- \mathbf{g} \text{ U } \mathbf{o}} \end{array}$$

Crucially, however, we are *unable* to construct *any* proof for the trace gg . For instance, attempting to construct a satisfaction proof fails because we hit the end-of-trace, ϵ , before completing the proof tree. Intuitively, we do not have enough information from the trace generated thus far to conclude that the complete trace satisfies the property. For instance, the next state may be o , in which case we can infer satisfaction for ggo , or it can be r , in which case we infer a violation for ggr ; if it is g , we postpone any conclusive judgement once again.

$$\begin{array}{c} \text{PPRD} \frac{\mathbf{g}(g)}{gg \vdash^+ \mathbf{g}} \quad \text{PUNT2} \frac{\text{PPRD} \frac{\mathbf{g}(g)}{g \vdash^+ \mathbf{g}} \quad ?? \frac{\quad}{\epsilon \vdash^+ \mathbf{g} \text{ U } \mathbf{o}}}{g \vdash^+ \mathbf{g} \text{ U } \mathbf{o}} \\ \text{PUNT2} \frac{\quad}{gg \vdash^+ \mathbf{g} \text{ U } \mathbf{o}} \end{array} \quad \blacksquare$$

Properties. Our proof system is sound, in the following sense.

Theorem 1 (Soundness). *For arbitrary t, φ :*

$$(t \vdash^+ \varphi \text{ implies } \forall s. ts \in \llbracket \varphi \rrbracket) \quad \text{and} \quad (t \vdash^- \varphi \text{ implies } \forall s. ts \notin \llbracket \varphi \rrbracket)$$

Proof. By rule induction on $t \vdash^+ \varphi$ and $t \vdash^- \varphi$. □

Example 3. The satisfaction and violation proofs of Ex. 2 suffice to prove $gos \in \llbracket \mathbf{g} \text{ U } \mathbf{o} \rrbracket$ and $grs \notin \llbracket \mathbf{g} \text{ U } \mathbf{o} \rrbracket$ for any (infinite) suffix s . Moreover, to determine whether $r \dots \notin \llbracket (\neg r) \wedge Xr \rrbracket$ from Ex. 1, it suffices to consider the prefix r and

either construct a violation proof directly, or else normalise the negation of the formula, $\langle \neg((\neg r) \wedge Xr) \rangle = r \vee X\bar{r}$ and construct a satisfaction proof:

$$\begin{array}{c} \text{PPRE} \frac{r(r)}{r \vdash^+ r} \\ \text{PNEG} \frac{r \vdash^+ r}{r \vdash^- \neg r} \\ \text{NAND1} \frac{r \vdash^- \neg r}{r \vdash^- (\neg r) \wedge Xr} \end{array} \qquad \begin{array}{c} \text{PPRE} \frac{r(r)}{r \vdash^+ r} \\ \text{POR1} \frac{r \vdash^+ r}{r \vdash^+ r \vee X\bar{r}} \end{array} \quad \blacksquare$$

Remark 1. The apparent redundancy (Ex. 3) allows us to use the proof system as a unifying framework that embeds other approaches (cf. §5), which may handle negation directly [SRA04], or work exclusively with formulas in nnf [Gei01].

Our proof system handles empty strings ϵ , as these arise naturally from the incremental analysis of finite traces discussed above.

Example 4. We can prove $oo\dots \in \llbracket X \text{tt} \rrbracket$ from the prefix oo , by constructing the proof tree below; the leaf node relies on being able to deduce $\epsilon \vdash^+ \text{tt}$:

$$\begin{array}{c} \text{PTRU} \frac{}{\epsilon \vdash^+ \text{tt}} \\ \text{PNXT} \frac{\epsilon \vdash^+ \text{tt}}{o \vdash^+ X \text{tt}} \\ \text{PNXT} \frac{o \vdash^+ X \text{tt}}{oo \vdash^+ XX \text{tt}} \end{array} \quad \blacksquare$$

Theorem 2 (Incompleteness). *For arbitrary t, φ :*

$(\forall s. ts \in \llbracket \varphi \rrbracket \text{ does not imply } t \vdash^+ \varphi) \quad \text{and} \quad (\forall s. ts \notin \llbracket \varphi \rrbracket \text{ does not imply } t \vdash^- \varphi)$

Proof. By counter example. For the positive case, consider $t = \epsilon$. We have $\forall s. ts \in \llbracket X \text{tt} \rrbracket$ but $t \not\vdash^+ X \text{tt}$, $\forall s. ts \in \llbracket p \vee \bar{p} \rrbracket$ but $t \not\vdash^+ p \vee \bar{p}$, and $\forall s. ts \in \llbracket \text{ff R tt} \rrbracket$ but $t \not\vdash^+ \text{ff R tt}$. Curiously, whenever $p(\sigma)$ holds for *all* $\sigma \in \Sigma$, we also have $\forall s. ts \in \llbracket p \rrbracket$ but $t \not\vdash^+ p$. Analogous examples can be drawn up for the negative case. \square

We can however prove completeness for a syntactic subset of the logic, limiting ourselves to *discriminating* predicates³, i.e., predicates p where $\exists \sigma_1, \sigma_2 \in \Sigma$ such that $\sigma_1 \neq \sigma_2, p(\sigma_1)$ and $\neg p(\sigma_2)$. We define the following syntactic subset:

$$\begin{aligned} \phi \in \text{PLTL} &::= \text{tt} \mid \text{ff} \mid p \mid \bar{p} \mid \phi_1 \wedge \phi_2 \mid \phi_1 \text{U} \phi_2 \mid \neg \gamma \\ \gamma \in \text{NLTL} &::= \text{tt} \mid \text{ff} \mid p \mid \bar{p} \mid \gamma_1 \vee \gamma_2 \mid \gamma_1 \text{R} \gamma_2 \mid \neg \phi \end{aligned}$$

Theorem 3 (Partial Completeness). *For arbitrary t, ϕ, γ :*

$(\forall s. ts \in \llbracket \phi \rrbracket \text{ implies } t \vdash^+ \phi) \quad \text{and} \quad (\forall s. ts \notin \llbracket \gamma \rrbracket \text{ implies } t \vdash^- \gamma)$

Proof. By induction on the structure of ϕ and γ . \square

4 An Automation

An automated proof search using the rules in Fig. 2 can be *syntax directed* by the formula (and the polarity) since, for most formulas, there is only *one* applicable rule. Moreover, the exception cases have *at most* two applicable rules.

³ This does not decrease expressivity, since **tt** and **ff** can be used for the other cases. Note also that the co-predicate of a discrimination predicate is also discriminating.

$$\begin{aligned}
\mathbf{exp}(d) &\stackrel{\text{def}}{=} \begin{cases} \{\langle \rangle\} & \text{if } \langle \rangle \in d \\ \{\} & \text{if } d = \{\} \\ d & \text{if } c \in d \text{ implies } \mathbf{sat}(c) \\ \mathbf{exp}(\bigcup_{c \in d} \mathbf{expC}(c)) & \text{otherwise} \end{cases} \\
\mathbf{expC}(c) &\stackrel{\text{def}}{=} \bigoplus_{o \in c} \mathbf{expO}(o) \\
\mathbf{expO}(o) &\stackrel{\text{def}}{=} \begin{cases} \{c \mid r \in \mathbf{rls}(\varphi, q), c = \mathbf{prm}(r, t, \varphi)\} & \text{if } o = (t, \varphi)^q \\ \{\langle \lceil \epsilon, \varphi \rceil^q \rangle\} & \text{if } o = \lceil \epsilon, \varphi \rceil^q \end{cases}
\end{aligned}$$

Fig. 3. A breadth-first incremental search algorithm

Notation. In what follows, $(t, \varphi)^+$ and resp. $(t, \varphi)^-$ denote the resp. outstanding proof obligations $t \vdash^+ \varphi$ and $t \vdash^- \varphi$. Since our algorithm works on partial traces, $\lceil \epsilon, \varphi \rceil^+$ and $\lceil \epsilon, \varphi \rceil^-$ are used to denote *saturated* proof obligations, where the string ϵ does not yield enough information to complete the proof search (e.g., $\epsilon \vdash^+ \mathbf{g} \cup \mathbf{o}$ in Ex. 2). A *conjunction set* $\langle o_1, \dots, o_n \rangle$ denotes a conjunction of proof obligations; metavariables o_i range over obligations of the form $(t, \varphi)^q$ or $\lceil t, \varphi \rceil^q$ for $q \in \{+, -\}$. A *disjunction set* $\{c_1, \dots, c_n\}$, where c_i range over conjunction sets, denotes a disjunction of conjunction sets.⁴ We employ a merge operation over disjunction sets, \oplus , defined below:

$$d \oplus d' \stackrel{\text{def}}{=} \{c \cup c' \mid c \in d, c' \in d'\}$$

The disjunction set $\{\langle \rangle\}$ acts as the *identity*, i.e., $\{\langle \rangle\} \oplus d = d \oplus \{\langle \rangle\} = d$, whereas the disjunction set $\{\}$ *annihilates* such sets, i.e., $\{\} \oplus d = d \oplus \{\} = \{\}$.

Algorithm. A breadth-first proof search algorithm is described in Fig. 3. Disjunction sets encode the alternative proof derivations that may lead to a completed proof-tree (resulting from multiple proof rules that can be applied at certain stages of the search), and conjunction sets represent the outstanding obligations within each potential derivation. Thus, a disjunction set with an element $\langle \rangle$, denotes a *successful* search, whereas an empty disjunction set $\{\}$ represents a *failed* search. Another terminating condition for the search algorithm of Fig. 3 is when a disjunction set contains only *saturated* conjunction sets: these containing *only* saturated obligations of the form $\lceil \epsilon, \varphi \rceil^q$ (the predicate $\mathbf{sat}(c)$ denotes this).

To verify whether the judgement $t \vdash^q \varphi$ holds, we initiate the function $\mathbf{exp}(-)$ with the disjunction set $\{\langle (t, \varphi)^q \rangle\}$. If none of the terminating conditions in Fig. 3 are met, $\mathbf{exp}(-)$ expands each conjunction set using $\mathbf{expC}(-)$, and recurses. Conjunction set expansion consists in expanding and merging *every* proof obligation using $\mathbf{expO}(-)$ and \oplus . Obligation expansion returns a disjunction set, where each conjunction set denotes the proof obligations resulting from the premises of the rules applied. It uses two auxilliary functions:

⁴ For clarity, conjunction set notation, $\langle - \rangle$, differs from that of disjunction sets, $\{-\}$.

- $\mathbf{rls}(\varphi, q)$ returns a set of rule names r from Fig. 2 that can be applied to obligations with the formula φ and polarity qualifier q (e.g., $\mathbf{rls}(\varphi_1 \mathbf{U} \varphi_2, +) = \{\mathbf{PUNT1}, \mathbf{PUNT2}\}$ and $\mathbf{rls}(\mathbf{X}\varphi, -) = \{\mathbf{NNXT}\}$).
- $\mathbf{prm}(r, t, \varphi)$ returns a conjunction set with the premises of rule r instantiated to the conclusion with string t and formula φ (e.g., $\mathbf{prm}(\mathbf{PUNT2}, go, \mathbf{g} \mathbf{U} \mathbf{o}) = \langle (go, \mathbf{g})^+, (o, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle$ and $\mathbf{prm}(\mathbf{PTRU}, go, \mathbf{tt}) = \langle \rangle$). Importantly:
 - (i) For cases such as $\mathbf{prm}(\mathbf{PUNT2}, \epsilon, \mathbf{g} \mathbf{U} \mathbf{o})$ the function returns $\langle [\epsilon, \mathbf{g} \mathbf{U} \mathbf{o}]^+ \rangle$ since the string ϵ prohibits the function from generating *all* the premises for the rule (one premise requires the string to be of length ≥ 1).
 - (ii) The function is *undefined* when rule conditions are not satisfied (e.g., $\mathbf{prm}(\mathbf{PPRD}, g, \mathbf{o})$ is undefined since $\mathbf{o}(g)$ does not hold).

Example 5. Recall the inconclusive judgement $gg \vdash^+ \mathbf{g} \mathbf{U} \mathbf{o}$ from Ex. 2.

$$\begin{aligned}
 \mathbf{exp}(\{\langle (gg, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\}) &= \mathbf{exp}(\{\langle (gg, \mathbf{o})^+ \rangle, \langle (gg, \mathbf{g})^+ \rangle, \langle (g, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\}) \\
 &= \mathbf{exp}(\{\} \cup (\{\langle \rangle\} \oplus \{\langle (g, \mathbf{o})^+ \rangle, \langle (g, \mathbf{g})^+ \rangle, \langle (\epsilon, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\})) \\
 &= \mathbf{exp}(\{\langle (g, \mathbf{o})^+ \rangle, \langle (g, \mathbf{g})^+ \rangle, \langle (\epsilon, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\}) = \{\langle [\epsilon, \mathbf{g} \mathbf{U} \mathbf{o}]^+ \rangle\} \quad \blacksquare
 \end{aligned}$$

Properties. An execution of $\mathbf{exp}(\{\langle (t, \varphi)^q \rangle\})$ may yield either of *three* verdicts. Apart from success, $\{\langle \rangle\}$, meaning that a full proof tree was derived, the algorithm partitions negative results as either a definite fail, $\{\}$, or an *inconclusive* verdict, consisting of a saturate disjunction set d (where $c \in d$ implies $\mathbf{sat}(c)$).

Saturated disjunction sets make the algorithm *incremental*, in the following sense. When a further suffix t' is learnt to a judgement $t \vdash^q \varphi$ with an inconclusive verdict, we can *reuse* the saturated disjunction set returned for $t \vdash^q \varphi$, instead of processing $tt' \vdash^q \varphi$ from scratch. This is done by converting each obligation of the form $[\epsilon, \varphi]^q$ in each saturated conjunction set to the *active* obligation $(t', \varphi)^q$ using an auxilliary “append” function $\mathbf{app}(-)$.

Example 6. To determine whether $ggo \vdash^+ \mathbf{g} \mathbf{U} \mathbf{o}$ holds, we can take the inconclusive outcome of $\mathbf{exp}(\{\langle (gg, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\})$ from Ex. 5, convert the saturated obligations using suffix o , $\mathbf{app}(\{\langle [\epsilon, \mathbf{g} \mathbf{U} \mathbf{o}]^+ \rangle\}, o) = \{\langle (o, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\}$, and calculate from that point onwards, $\mathbf{exp}(\{\langle (o, \mathbf{g} \mathbf{U} \mathbf{o})^+ \rangle\}) = \{\langle \rangle\}$. \blacksquare

Theorem 4 (Incrementality). $\mathbf{sat}(\mathbf{exp}(\{\langle (t_1, \varphi)^q \rangle\}))$ implies

$$\mathbf{exp}(\{\langle (t_1 t_2, \varphi)^q \rangle\}) = \mathbf{exp}(\mathbf{app}(\mathbf{exp}(\{\langle (t_1, \varphi)^q \rangle\}), t_2))$$

Proof. By induction on t_2 . □

The algorithm of Fig. 3 is also *decidable* for the proof rules of Fig. 2. Intuitively, the main reason for this is because the proof system is *cut-free*, where rule premises are either defined in terms of string suffixes or subformula. Formally, we define a rank function $|-|$ mapping proof obligations to pairs of naturals, for which we assume a lexicographical ordering $(n_1, m_1) \geq (n_2, m_2) \stackrel{\text{def}}{=} n_1 \geq n_2 \vee (n_1 = n_2 \wedge m_1 \geq m_2)$ and the obvious function $\mathbf{max}(-)$ returning the

greatest element from a set of such pairs. Apart from $|t|$, we also assume $|\varphi|$ returning the *maximal depth* of the formula (e.g., $|p \cup \neg(\bar{p} \vee p)| = 3$ and $|\bar{p}| = 0$).

$$\begin{aligned} |(t, \varphi)^q| &\stackrel{\text{def}}{=} (|t|, |\varphi|) & |[t, \varphi]^q| &\stackrel{\text{def}}{=} (0, 0) & |c| &\stackrel{\text{def}}{=} \mathbf{max}(\{|o| \mid o \in c\} \cup \{(0, 0)\}) \\ |d| &\stackrel{\text{def}}{=} \text{if } \langle \rangle \in d \text{ then } (0, 0) \text{ else } \mathbf{max}(\{|c| \mid c \in d\} \cup \{(0, 0)\}) \end{aligned}$$

Above, the rank function maps saturated obligations to the bottom element $(0, 0)$. We overload the function to conjunction sets, where we add $(0, 0)$ to the $\mathbf{max}(-)$ calculation to cater for the case where c is empty. Following a similar pattern, we also extend the rank function to disjunction sets, but equate all sets with an empty conjunction set to the bottom element $(0, 0)$; this mirrors the termination condition of the algorithm in Fig. 3 which terminates the search as soon as the shortest proof tree is detected.

Theorem 5 (Decidability). *$\mathbf{exp}(\{\langle (t, \varphi)^q \rangle\})$ always terminates.*

Proof. Follows from the fact that when $|d| = (0, 0)$, $\mathbf{exp}(d)$ terminates immediately, and when $|d| \neq (0, 0)$, we have $|d| > |\bigcup_{c \in d} \mathbf{expC}(c)|$. \square

Runtime Monitoring. We can obtain a setup akin to the three-valued monitors of [BLS11] with outcome **Y**, denoting satisfaction, outcome **N**, denoting violation, and outcome **?**, denoting an inconclusive outcome. Following [BLS11], given a finite trace t and a property φ , we attempt to construct a deduction for *both* the satisfaction, $t \vdash^+ \varphi$ and violation, $t \vdash^- \varphi$, by *concurrently* running $\mathbf{exp}(\{\langle (t, \varphi)^+ \rangle\})$ and $\mathbf{exp}(\{\langle (t, \varphi)^- \rangle\})$ with the following possible outcomes:

1. We are able to construct a proof for $t \vdash^+ \psi$, corresponding to **Y**.
2. We are able to construct a proof for $t \vdash^- \psi$, corresponding to **N**.
3. We are unable to construct proofs for either case, corresponding to **?**.

Remark 2. The fourth possible outcome, *i.e.*, constructing a proof for *both* $t \vdash_3^+ \psi$ and $t \vdash_3^- \psi$, is ruled out by soundness (Thm. 1), which implicitly guarantees that our analysis is consistent (since the semantics is defined in terms of sets).

Like in most online RV setups, Thm. 4 allows for *incremental* monitor, as soon as individual trace elements are received. Moreover, Thm. 5 allows for a *safe synchronous* instrumentation, where the monitor and system execute in lock-step (*i.e.*, the system is paused after producing each monitored event so as to allow the monitor to carry out its analysis and perform timely detections). Since the monitoring analysis always terminates, the monitored system is guaranteed to be able to progress normally under a synchronous instrumentation.

5 Alternative RV Symbolic Techniques for LTL

We relate our deductive system to two prominent, but substantially distinct, symbolic techniques for LTL in the context of RV, namely [Gei01] and [SRA04].

5.1 Informative Prefixes

Intuitively, an *informative prefix* for a formula *explains* why a trace satisfies that formula [KYV01]. In [Gei01], trace satisfactions are monitored wrt. LTL formulas in *nnf*, by checking whether a trace contains an informative prefix.

Example 7. Recall $\mathbf{g} \mathbf{U} \mathbf{o}$ (Ex. 1). Prefix go is informative because (i) although the head, g , does not satisfy $\mathbf{g} \mathbf{U} \mathbf{o}$ in a definite manner, it allows the possibility of its suffix to satisfy the formula conclusively ($\mathbf{g}(g)$ holds); (ii) the immediate suffix, o , satisfies $\mathbf{g} \mathbf{U} \mathbf{o}$ conclusively ($\mathbf{o}(o)$ holds). In [Gei01], both go and o are deemed to be *locally-informative* wrt. $\mathbf{g} \mathbf{U} \mathbf{o}$ but go generates satisfaction obligations for the immediate suffix (*temporal informative successor*). ■

The algorithm in [Gei01] formalises the notion of locally informative by converting formulas to their *informative normal forms*. Moreover, temporal informative successors are formalised through the function $\text{next}(-)$, returning a set of formulas from a given formula and a trace element. For instance, in Ex. 7 $\text{next}(g, \mathbf{g} \mathbf{U} \mathbf{o}) = \{\mathbf{g} \mathbf{U} \mathbf{o}\}$ whereas $\text{next}(o, \mathbf{g} \mathbf{U} \mathbf{o}) = \{\}$. These functions are then used to construct automata that check for these properties over string prefixes.

$$\begin{array}{l}
\text{GTRU} \frac{}{\text{linf}(t, \mathbf{tt}, \emptyset)} \quad \text{GPRES1} \frac{p(\sigma)}{\text{linf}(\sigma t, p, \emptyset)} \quad \text{GPRES2} \frac{\bar{p}(\sigma)}{\text{linf}(\sigma t, \bar{p}, \emptyset)} \\
\text{GOR1} \frac{\text{linf}(t, \varphi_1, m)}{\text{linf}(t, \varphi_1 \vee \varphi_2, m)} \quad \text{GOR2} \frac{\text{linf}(t, \varphi_2, m)}{\text{linf}(t, \varphi_1 \vee \varphi_2, m)} \\
\text{GAND} \frac{\text{linf}(t, \varphi_1, m_1) \quad \text{linf}(t, \varphi_2, m_2)}{\text{linf}(t, \varphi_1 \wedge \varphi_2, m_1 \cup m_2)} \quad \text{GNXT} \frac{}{\text{linf}(t, \mathbf{X}\varphi, \{\varphi\})} \\
\text{GUNT1} \frac{\text{linf}(t, \varphi_2, m)}{\text{linf}(t, \varphi_1 \mathbf{U} \varphi_2, m)} \quad \text{GUNT2} \frac{\text{linf}(t, \varphi_1, m)}{\text{linf}(t, \varphi_1 \mathbf{U} \varphi_2, m \cup \{\varphi_1 \mathbf{U} \varphi_2\})} \\
\text{GREL1} \frac{\text{linf}(t, \varphi_1, m_1) \quad \text{linf}(t, \varphi_2, m_2)}{\text{linf}(t, \varphi_1 \mathbf{R} \varphi_2, m_1 \cup m_2)} \quad \text{GREL2} \frac{\text{linf}(t, \varphi_2, m)}{\text{linf}(t, \varphi_1 \mathbf{R} \varphi_2, m \cup \{\varphi_1 \mathbf{R} \varphi_2\})}
\end{array}$$

In this section, we express the locally-informative predicate and the associated temporal informative successors as the single judgement $\text{linf}(t, \varphi, m)$, defined as the least relation satisfying the rules above. It states that t is locally informative for φ with obligations $m \in \mathcal{P}(\text{ELTL})$ for the succeeding suffix. For example, for a formula $\mathbf{X}\varphi$, *any* string t is locally informative, but requires the immediate suffix to satisfy φ (see GNXT); in the case of $\varphi_1 \mathbf{U} \varphi_2$, if t is locally informative for φ_1 with suffix obligations m , then t is also locally informative for $\varphi_1 \mathbf{U} \varphi_2$ with obligations $m \cup \{\varphi_1 \mathbf{U} \varphi_2\}$ (see GUNT2). Informative prefixes are formalised as the predicate $\text{inf}(t, \varphi)$ below. Note that recursion in $\text{inf}(t, \varphi)$ is employed on a substring of t and terminates when $m = \emptyset$.

$$\text{inf}(t, \varphi) \stackrel{\text{def}}{=} \exists m. \left(\text{linf}(t, \varphi, m) \quad \text{and} \quad (\varphi' \in m \text{ implies } (\text{inf}([t]^1, \varphi')) \right)$$

Example 8. We can deduce that $\text{inf}(go, g \cup o)$ because $\text{linf}(go, g \cup o, \{g \cup o\})$ and then that $\text{linf}(o, g \cup o, \emptyset)$. ■

We can formally show a correspondence between informative prefixes and our monitoring proof systems.

Theorem 6. *For all φ in nnf, $\text{inf}(t, \varphi)$ iff $t \vdash^+ \varphi$*

Proof. By structural induction on t , then by rule induction on $\text{linf}(t, \varphi, m)$ for the *only-if* case. By rule induction for the *if* case. □

In the *only-if* direction, Thm. 6 ensures that our system is as expressive as [Gei01]. In the *if* direction, Thm. 6 shows that every derivation in our proof system corresponds to an informative prefix as defined in [KYV01]⁵. This reinforces existing justifications as to why our system is unable to symbolically process certain prefix and formula pairs.

Example 9. The proof system of §3 is unable to deduce $\epsilon \vdash^+ \text{Xtt}$ (cf. proof for Thm. 2) even though, on a semantic level, this holds for any string continuation because ϵ is *not* an informative prefix of Xtt . ■

Thm. 6 has another important implication. One justification of informative prefixes is that any bad/good prefixes detected can be accompanied by an explanation [BLS11]. However, whereas in [Gei01] this explanation is given in terms of the algorithm implementation, delineating the proof system from its implementing monitor (as in our case) allows for better separation of concerns, by giving the explanation as a derivation⁶ in terms of the proof rules of Fig. 2.

5.2 Derivatives

In a derivatives approach [HR01,SRA04], LTL formulas are interpreted as *functions* that take a state *i.e.*, an element of the alphabet Σ , and return another LTL formula. The returned formula is then applied again to the next state in the trace, until either one of the *canonical* formulas tt or ff are reached; the trace analysis stops at canonical formulas, since tt (*resp.* ff) are idempotent, returning tt (*resp.* ff), irrespective of the state applied to. In [SRA04], coinductive deductive techniques are used on derivatives to establish LTL formula equivalences, which are then used to obtain optimal monitors for good/bad prefixes.

Example 10. Recall $\epsilon \not\vdash^+ \text{Xtt}$ from Ex. 9. In [SRA04], they establish that formulas tt and Xtt are equivalent *wrt.* good prefixes, $\text{tt} \equiv_G \text{Xtt}$, which allows them to reason symbolically about ϵ and Xtt in terms of ϵ and tt instead. ■

Formally, a derivative is a rewriting operator $-\{-\} :: \text{ALTL} \times \Sigma \longrightarrow \text{ALTL}$ (adapted from [SRA04]) defined on the structure of the formula as follows:

⁵ As a corollary, we also establish a correspondence between $t \vdash^- \varphi$ and $\text{inf}(t, \neg\varphi)$ as used in [Gei01] for bad prefixes.

⁶ The algorithm in §4 can be easily extended so as to record the rules used.

$$\begin{array}{ll}
\text{tt}\{\sigma\} \stackrel{\text{def}}{=} \text{tt} & \text{ff}\{\sigma\} \stackrel{\text{def}}{=} \text{ff} \\
\text{p}\{\sigma\} \stackrel{\text{def}}{=} \text{if } \text{p}(\sigma) \text{ then tt else ff} & \\
(\neg\psi)\{\sigma\} \stackrel{\text{def}}{=} (\text{tt} \oplus \psi)\{\sigma\} & \psi_1 \oplus \psi_2\{\sigma\} \stackrel{\text{def}}{=} \psi_1\{\sigma\} \oplus \psi_2\{\sigma\} \\
\psi_1 \wedge \psi_2\{\sigma\} \stackrel{\text{def}}{=} \psi_1\{\sigma\} \wedge \psi_2\{\sigma\} & \psi_1 \vee \psi_2\{\sigma\} \stackrel{\text{def}}{=} \psi_1\{\sigma\} \vee \psi_2\{\sigma\} \\
\text{X}\psi\{\sigma\} \stackrel{\text{def}}{=} \psi & \\
\psi_1 \text{U} \psi_2\{\sigma\} \stackrel{\text{def}}{=} \psi_2\{\sigma\} \vee (\psi_1\{\sigma\} \wedge \psi_1 \text{U} \psi_2) &
\end{array}$$

Above, we position rewriting definitions for core LTL formulas of Fig. 1 on the left; formula rewriting however also uses an *extended* set of formulas that include falsity, ff , disjunction, $\psi_1 \vee \psi_2$, and exclusive-or, $\psi_1 \oplus \psi_2$. The derivatives algorithm also works up to formula normalisations using the following equalities:

$$\begin{array}{llll}
\text{tt} \wedge \psi \equiv \psi & \text{ff} \wedge \psi \equiv \text{ff} & \text{ff} \vee \psi \equiv \psi & \text{tt} \vee \psi \equiv \text{tt} \\
\psi \wedge \psi \equiv \psi & \psi \vee \psi \equiv \psi & \text{ff} \oplus \psi \equiv \psi & (\psi_1 \wedge \psi_2) \oplus \psi_1 \oplus \psi_2 \equiv \psi_1 \vee \psi_2
\end{array}$$

Thus, for any finite trace t of the form $\sigma_1\sigma_2\dots\sigma_n$ we say:

- t is a good prefix for ψ iff $((\psi\{\sigma_1\})\{\sigma_2\})\{\dots\sigma_n\} \equiv \text{tt}$;
- t is a bad prefix for ψ iff $((\psi\{\sigma_1\})\{\sigma_2\})\{\dots\sigma_n\} \equiv \text{ff}$.

Example 11. The partial trace go is a *good prefix* for gUo (Ex. 1) because:

$$\begin{aligned}
(\text{gUo}\{g\})\{o\} &\stackrel{\text{def}}{=} \text{o}\{g\} \vee (\text{g}\{g\} \wedge \text{gUo})\{o\} \\
&\stackrel{\text{def}}{=} \text{ff} \vee (\text{g}\{g\} \wedge \text{gUo})\{o\} \equiv (\text{g}\{g\} \wedge \text{gUo})\{o\} \\
&\stackrel{\text{def}}{=} (\text{tt} \wedge \text{gUo})\{o\} \equiv \text{gUo}\{o\} \\
&\stackrel{\text{def}}{=} \text{o}\{o\} \vee (\text{g}\{o\} \wedge \text{gUo}) \stackrel{\text{def}}{=} \text{tt} \vee (\text{g}\{o\} \wedge \text{gUo}) \equiv \text{tt} \quad \blacksquare
\end{aligned}$$

Good (*resp.* bad) prefixes, as defined in[SRA04], correspond to finite traces with a satisfaction (*resp.* violation) proof in our system (§3), and vice-versa.

Theorem 7. *For any finite trace $t = \sigma_1\dots\sigma_n$, and core LTL formula ψ :*

$$((\psi\{\sigma_1\})\{\dots\sigma_n\} \equiv \text{tt} \text{ iff } t \vdash^+ \psi) \text{ and } ((\psi\{\sigma_1\})\{\dots\sigma_n\} \equiv \text{ff} \text{ iff } t \vdash^- \psi)$$

Proof. For the *only-if* case both statements are proved simultaneously by numerical induction on n and then by structural induction on ψ . For the *if* case both statements are proved simultaneously by rule induction. \square

Apart from establishing a one-to-one correspondence between derivative prefixes and proof deductions for core LTL formulas in our system, Thm. 7 (together with Thm. 6) allows us to relate indirectly the informative prefixes of §5.1 to derivative prefixes. Moreover, Thm. 7 identifies from where the additional expressivity of the analysis in [SRA04] derives, namely through the deductive system for formula equivalence *wrt.* good/bad prefixed, $\vdash \psi_1 \equiv_G \psi_2$ and $\vdash \psi_1 \equiv_B \psi_2$. This opens up the possibility of merging the two approaches, perhaps by extending our deductive system with rules analogous to those show below:

$$\text{PEQ} \frac{t \vdash^+ \varphi_1 \quad \vdash \varphi_1 \equiv_G \varphi_2}{t \vdash^+ \varphi_2} \quad \text{NEQ} \frac{t \vdash^- \varphi_1 \quad \vdash \varphi_1 \equiv_B \varphi_2}{t \vdash^- \varphi_2}$$

6 Conclusion

We presented a proof system and the respective monitor generation for runtime-verifying LTL properties. One novel aspect of our approach is that we tease apart the specification of the symbolic analysis from its automation. This allows us to localise correctness results *e.g.*, soundness is determined for the proof system, independent of the subsequent automation. The higher level of abstraction used elucidates completeness studies, seldom tackled in other work in RV *e.g.*, syntactic class identified for Thm. 3 appears to be new. This separation of concerns also facilitates comparisons and cross-fertilisation with other symbolic techniques (§ 5) and leads to modular organisations that are easier to maintain *e.g.*, more efficient monitor automation may be considered without changing the proof rules. The concrete contributions are:

1. A *sound, local* LTL proof system inferring complete trace inclusion from finite prefixes, Thm. 1, together with completeness results, Thm. 2 and Thm. 3.
2. A mechanisation of proof derivations for this system that is formally incremental, Thm. 4, and decidable, Thm. 5.
3. An exposition of how the proof system can be used as a unifying framework where to relate different runtime monitoring formalisms, Thm. 7 and Thm. 6.

Related Work. Apart from the deductive system for LTL formula equivalence in [SRA04], there are other LTL proof systems [GPSS80,MP91,KI11,BL08]. Each differ substantially from ours. For instance, the model used in [GPSS80,MP91] is different from ours, *i.e.*, programs (*sets of traces*) instead of traces; the work in [GPSS80,KI11] is concerned with developing tableau methods for inferring the validity of a formula from a conjunction of formulas; [BL08] study cut-free sequent systems; importantly, none of these proof systems are local. In [MP91], they develop three *tailored* proof systems for separate classes of properties, namely safety, response and reactivity properties; crucially however, they do not consider aspects such as deductions from *partial* traces.

A substantial body of work studies alternative LTL semantics for partial traces [EFH⁺03,BLS07,BLS11]; consult [BLS10] for a comprehensive survey. Although complementary, the aim and methodology of this work is substantially different from ours. In particular, we keep the LTL semantics *constant*, and explore the soundness, completeness and expressivity aspects of our symbolic analysis *wrt.* to this fixed semantics.

Future Work. It would be fruitful to relate other LTL symbolic analyses to the ones discussed in §5. Our work may also be used as a point of departure for developing proof systems for other interpretations of LTL. For instance, a different LTL model to that of §2, consisting of both *finite* and infinite traces, alters the negation propagation identities used for the translation function (*e.g.*, $\neg X\psi \equiv X\neg\psi$ does not hold) which, amongst other things, would require tweaking to the proof rules. Similar issues arise in distributed LTL interpretations such as [BF12] where instead of having one execution trace, we have a *set of traces*

(one for each location). We also leave complexity analysis and the assessment of the runtime overheads introduced by our setup as future work.

References

- [BF12] Andreas Bauer and Ylis Falcone. Decentralised LTL Monitoring. In *FM*, volume 7436 of *LNCS*, pages 85–100. Springer, 2012.
- [BL08] Kai Brunner and Martin Lange. Cut-free sequent systems for temporal logic. *JLAP*, 76(2):216 – 225, 2008.
- [BLS07] Andreas Bauer, Martin Leucker, and Christian Schallhart. The good, the bad, and the ugly, but how ugly is ugly? In *RV*, volume 4839 of *LNCS*, pages 126–138, Berlin, Heidelberg, November 2007. Springer.
- [BLS10] Andreas Bauer, Martin Leucker, and Christian Schallhart. Comparing LTL semantics for runtime verification. *Logic and Comput.*, 20(3):651–674, 2010.
- [BLS11] Andreas Bauer, Martin Leucker, and Christian Schallhart. Runtime verification for LTL and TLTL. *TOSEM*, 20(4):14, 2011.
- [BS92] Julian Bradfield and Colin Stirling. Local model-checking for infinite state spaces. *TCS*, 96:157–174, 1992.
- [Bus98] Samuel R. Buss, editor. *Handbook of Proof Theory*. Elsevier, 1998.
- [EFH⁺03] Cindy Eisner, Dana Fisman, John Havlicek, Yoad Lustig, Anthony McIsaac, and David Van Campenhout. Reasoning with temporal logic on truncated paths. In *CAV*, volume 2725 of *LNCS*, pages 27–39. Springer, 2003.
- [Gei01] Marc Geilen. On the construction of monitors for temporal logic properties. *ENTCS*, 55(2):181–199, 2001.
- [GPSS80] Dov Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. On the temporal analysis of fairness. In *POPL*, pages 163–173, New York, NY, USA, 1980. ACM.
- [HR01] Klaus Havelund and Grigore Rosu. Monitoring programs using rewriting. In *ASE*, pages 135–143, Wash., DC, USA, 2001. IEEE.
- [KI11] Kensuke Kojima and Atsushi Igarashi. Constructive linear-time temporal logic: Proof systems and kripke semantics. *Inf. Comput.*, 209(12):1491–1503, 2011.
- [KYV01] Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. *Form. Methods Syst. Des.*, 19(3):291–314, October 2001.
- [LS09] Martin Leucker and Christian Schallhart. A brief account of Runtime Verification. *JLAP*, 78(5):293 – 303, 2009.
- [MP91] Zohar Manna and Amir Pnueli. Completing the Temporal Picture. *Theoretical Computer Science*, 83(1):97 – 130, 1991.
- [Pnu77] Amir Pnueli. The Temporal Logic of Programs. In *SFCS*, pages 46–57, Wash., DC, USA, 1977. IEEE.
- [RH05] Grigore Roşu and Klaus Havelund. Rewriting-based techniques for runtime verification. *Automated Software Engg.*, 12(2):151–197, April 2005.
- [SRA04] Koushik Sen, Grigore Rosu, and Gul Agha. Generating optimal linear temporal logic monitors by coinduction. In *ASIAN*, LNCS, pages 260–275. Springer, 2004.
- [SW91] Colin Stirling and David Walker. Local model-checking in the modal mu-calculus. *TCS*, 89:161–177, 1991.
- [TS00] A.S. Troelstra and H. Schwichtenberg. *Basic Proof Theory*. Cambridge University Press, 2000.