Trace Diagnostics using Temporal Implicants

Thomas Ferrère[1], Oded Maler[1], Dejan Ničković[2]

[1] Verimag, University of Grenoble / CNRS
[2] AIT Austrian Institute of Technology

Abstract. Runtime verification and model checking are two important methods for assessing correctness of systems. In both techniques, detecting an error is witnessed by an *execution* that violates the system *specification*. However, a faulty execution on its own may not provide sufficiently precise insight to the causes of the reported violation. Additional, often manual effort is required to properly diagnose the system. In this paper we present a method for analyzing such causes. The specifications we consider are expressed in LTL (Linear Temporal Logic) and MTL (Metric Temporal Logic), and the execution models are taken as *ultimately-periodic* words, and *finite variability* continuous signals respectively. The *diagnostics* problem is defined for the propositional case as the search for a small *implicant* of a formula which is satisfied by a given valuation, or equivalently a subset of that valuation sufficient to render the formula true. We propose a suitable notion of implicants in the temporal case, that are semantically based on signal subsets, and guarantee the existence of *prime* implicants for arbitrary temporal properties. An inductive procedure for finding temporal implicants is obtained by the introduction of selection functions that appear in a process equivalent to Skolemization in first order logic. Through the model restrictions we impose for LTL and MTL we are able to generate concise implicants of a property, describing a small fragment of the input signal that causes violation of a formula.

1 Introduction

Our work is concerned with the problem of temporal *monitoring*: given a *single* behavior w, either in discrete or dense time, and a temporal property φ check whether $w \models \varphi$. This problem is known as *runtime verification* in software and *assertion checking* in hardware. In addition to the yes/no answer, we would like to produce an informative *diagnostics*, a small fragment of the behavior which provides a sufficient condition for the violation of φ by w. This additional information helps localizing and explaining the causes of the fault. We solve the diagnostics problem for MTL [9], for which we assume that the input signal w has bounded variability. We further extend our results to LTL [15] under the assumption of an ultimately periodic input sequence. This makes our technique applicable to the analysis of counter-examples executions as produced by a *model checking* procedure.

Consider the temporal logic formula $\Box(p \rightarrow \Diamond_{[1,2]}\, q)$. It requires that for any instant in time where p holds, there exists another instant within 1 to 2 time units where q holds. The behavior depicted in Figure 1 violates this temporal property – the violation can be explained by the fact that p holds at time 1 and q does not hold throughout $[2, 3]$. Such a concise piece of information, compared to w which can be a very long signal can

increase confidence in monitoring and model-checking procedures, and promote their further acceptance in various application domains.

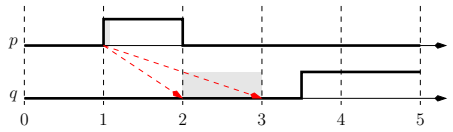

Fig. 1. A behavior that violates $\Box(p \rightarrow \Diamond_{[1,2]}\, q)$. Grey-shaded area gives one possible explanation.

Finding an explanatory sub-model in the propositional case, is strongly related to the concept of prime implicants of a formula. The problem that we pose in this paper, finding explanatory temporal sub-models, is novel and in order to solve it we had to overcome numerous issues that come from the infinitude of the temporal models. Our main result is an inductive explanation generation scheme for MTL which produces focused dense time sub-signals sufficient to explain violation. A crucial ingredient of the procedure is the elimination of disjunctive operations by the introduction of selection functions similar in spirit to Skolem functions used to eliminate existential quantification. Under a finite variability assumption we can show that explanations can be taken as finitely variable. For LTL, we show similarly that infinite ultimately-periodic sequences admit ultimately-periodic explanatory sub-models.

Related Work The problem of understanding a counter-example by finding the reason for the failure of a temporal logic formula in the trace itself was studied in [1]. This work differs from ours in several aspects. It adopts a different notion of failures based on Halpern and Pearl causality [6] and considers only LTL but not dense-time temporal logics. The explanations of ultimately periodic sequences are handled by unfolding the trace. Finally, the authors are interested in the detection of the first failure in a trace. In our work we provide more flexibility by means of selection functions, which allow to choose between several different failures. In [13], the authors propose a procedure that provides a minimal debugging window for traces that violate an MTL formula. The result can be seen as a coarse-grain diagnostic, providing a small segment of the input trace yet not discriminating signal segments that cause the violation.
There have been various studies on obtaining additional debugging information from counter-examples in LTL model checking. Tight automata [10] were introduced to find shortest finite counter-examples for safety properties, and extended in [17] to infinite words and full LTL. Comparing erroneous and correct traces with distance metrics in order to localize errors has been studied in the context of software checking in [5]. The problem of finding and repairing violations of LTL properties by sequential circuits was studied in [8], where a repair solution based on a game-oriented approach is proposed. A related problem is that of computing unsatisfiable cores for LTL, i.e. finding smaller unsatisfiable sub-formulas, as studied in [16, 14, 7]. At the syntactic level minimal unsatisfiable cores bear some similarity with prime implicants; they primarily address formal verification concerns.

2 Propositional Foundations

Consider the problem of explaining why a formula φ is *violated* by a given execution w of some system, seen as finding the part of the execution w that causes φ to be violated. Note that through negation this is equivalent to solving the dual problem of explaining why some formula is *satisfied*. We first introduce and study the problem in the simple setting of propositional logic.

2.1 Problem Statement

Let $\mathbb{P}$ be a *finite* set of propositional variables. A valuation w is taken to be a function $P \to \mathbb{B}$ with $P \subseteq \mathbb{P}$ its domain and $\mathbb{B} := \{0, 1\}$ the set of Boolean values. We define propositional formulas over $\mathbb{P}$ and the constant true the usual way. The set of models of a formula φ is noted $[\![\varphi]\!]$. For φ and ψ two formulas we write $\psi \Rightarrow \varphi$ when $[\![\psi]\!] \subseteq [\![\varphi]\!]$, and $\psi \Leftrightarrow \varphi$ when $[\![\psi]\!] = [\![\varphi]\!]$. Note that implication ($\Rightarrow$) induces a partial order over classes of equivalent ($\Leftrightarrow$) formulas.

Definition 1 (Terms, Implicants and Prime Implicants). *A* term γ *is defined as a conjunction of literals. If* $\gamma \Rightarrow \varphi$ *then the term* γ *is an* implicant *of formula* φ*. If moreover* γ *is maximal with respect to* $\Rightarrow$ *modulo equivalence we talk of* prime implicant*.*

We say that γ *explains* the satisfaction of φ by w, if γ is an implicant of φ and w is a model of γ. Note that the least general explanation of φ relative to w is a term representing the truth status of every variable in w. It is intuitively clear, however that we opt for explanations that are smaller and more general, omitting "don't care" variables. We aim at providing explanations that use small subsets of "do care" variables. The most general explanations are in particular the prime implicants of φ satisfied by w.

Problem (Diagnostics). *Given a valuation* w *and a formula* φ*, find a prime implicant* γ *of* φ *such that* $w \models \gamma$*.*

2.2 Syntactic and Semantic Formulations

Take φ a formula, w a model of φ and γ a solution to corresponding instance of the diagnostic problem. As $\gamma \Rightarrow \varphi$, there exists a proof of φ under hypothesis γ; a correct algorithm solving the diagnostics problem is implicitly constructing that proof. The more general the implicant is, the more complex the associated proof can be.

Example 1. Take formula $\varphi := (p \wedge q) \vee (p \wedge \neg q)$ and valuation $w := \{p \mapsto 1, q \mapsto 0, r \mapsto 0\}$. The formulas $\alpha := p$ and $\beta := p \wedge \neg q$ are both implicants of φ, and satisfied by w with α a prime implicant of φ. In sequent calculus the proof $\beta \vdash \varphi$ is direct through a right disjunction rule, while the proof $\alpha \vdash \varphi$ requires the application of several rules, and uses non-intuitionistic reasoning.

We now sketch the semantic counter-parts of implicants, beginning with a refinement relation $\sqsubseteq$ between valuations.

Definition 2. *For two valuations* $u : P \to \mathbb{B}$ *and* $v : Q \to \mathbb{B}$ *we have* $u \sqsubseteq v$ *if and only if* $P \subseteq Q$ *and* $u(p) = v(p)$ *for all* $p \in P$*.*

The space of valuations is a semi-lattice with respect to $\sqsubseteq$ with meet operation $\sqcap$ and least element $\bot$. Let u and v be some valuations with domain P and Q respectively. The valuation $u \sqcap v$ has domain $\{p \in P \cap Q \ : \ u(p) = v(p)\}$ and value $u \sqcap v\,(p) = u(p)$ where defined. The least element $\bot$ is the nowhere-defined valuation with domain $\emptyset$. One can think of a valuation v over $P \subseteq \mathbb{P}$ as a compact representation for all valuations w over $\mathbb{P}$ such that $v \sqsubseteq w$. A valuation v corresponds to a term γ_v (the conjunction of literals true according to v) and reciprocally any satisfiable term γ corresponds to a valuation v_γ that assigns a value to variables according to the literals in γ.

Definition 3 (sub-model). *A valuation v is a sub-model of φ if for all valuations w over $\mathbb{P}$ such that $v \sqsubseteq w$ we have $w \models \varphi$; if moreover v is minimal with respect to $\sqsubseteq$ we talk of minimal sub-model.*

A valuation v is a (minimal) sub-model of φ if and only if γ_v is a (prime) implicant of φ. Hence the diagnostics of φ with respect to w can equivalently be seen as finding the minimal sub-model of φ contained in w.

2.3 Practical Solution

Note that the *minimal* diagnostics problem is at least as hard as a satisfiability query, since tautologies can be recognized by their unique prime implicant, the empty term true. However when only considering implicants that are satisfied by a given model w, knowing the truth value of each sub-formula of φ on w allows to construct sub-models $v \sqsubseteq w$ in a simple, top-down fashion. For every formula we take for implicant a combination of implicants for its sub-formulas that are satisfied by w, or violated by w when in the context of a negation. Accordingly we define an operator E (and its dual F) that for a given formula φ returns an implicant of φ (respectively of $\neg\varphi$) which under suitable assumptions is satisfied by w. The explanation of φ is then defined as

$$\mathrm{Exp}(\varphi) = \begin{cases} E(\varphi) & \text{if } w \models \varphi \\ F(\varphi) & \text{otherwise} \end{cases}$$

with

$$\begin{aligned} E(p) &= p & F(p) &= \neg p \\ E(\neg\varphi) &= F(\varphi) & F(\neg\varphi) &= E(\varphi) \\ E(\varphi_1 \vee \varphi_2) &= E(\xi(\varphi_1 \vee \varphi_2)) & F(\varphi_1 \vee \varphi_2) &= F(\varphi_1) \wedge F(\varphi_2) \end{aligned}$$

where ξ is a *selection function* satisfying $\xi(\varphi_1 \vee \varphi_2) \in \{\varphi_1, \varphi_2\}$. When for any formula $\varphi_1 \vee \varphi_2$ such that $w \models \varphi_1 \vee \varphi_2$ it holds $w \models \xi(\varphi_1 \vee \varphi_2)$, we say that ξ is *correct* with respect to w. We can take for example

$$\xi : \varphi_1 \vee \varphi_2 \mapsto \begin{cases} \varphi_1 & \text{if } w \models \varphi_1 \\ \varphi_2 & \text{otherwise} \end{cases}$$

This gives priority to the left disjunct. Under the assumption that ξ is correct with respect to w, the formula $\mathrm{Exp}(\varphi)$ is a solution to the diagnostics problem associated to φ and w. In the case of Example 1, applying the procedure on φ and w yields the explanation β.

3 Temporal Issues

We introduce the temporal logics LTL [15] and MTL [9] in a unified framework. Temporal formulas will be given by the grammar

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi\,\mathcal{U}_I\varphi$$

where I is a real interval with integer endpoints. Other temporal connectives are introduced through the abbreviations $\varphi_1\,\mathcal{U}\varphi_2 := \varphi_1\,\mathcal{U}_{(0,+\infty)}\varphi_2$ for *strict until*, $\bigcirc\,\varphi := \mathrm{false}\,\mathcal{U}\,\varphi$ for *next*, $\varphi_1\,\tilde{\mathcal{U}}\,\varphi_2 := \varphi_2 \vee (\varphi_1 \wedge \varphi_1\,\mathcal{U}\varphi_2)$ for *non-strict until*, $\Diamond\varphi := \mathrm{true}\,\tilde{\mathcal{U}}\,\varphi$ for *eventually*, $\Box\varphi := \neg(\Diamond\neg\varphi)$ for *always* and $\Diamond_I\varphi := \mathrm{true}\,\mathcal{U}_I\varphi$, $\Box_I\varphi := \neg\Diamond_I\neg\varphi$ for their timed versions. LTL formulas are then constructed using temporal connectives $\bigcirc$ and $\tilde{\mathcal{U}}$, while MTL formulas are constructed using temporal connectives $\Diamond_I$ and $\mathcal{U}$.

A temporal behavior is defined as a function $\mathbb{T} \times P \to \mathbb{B}$, for $P \subseteq \mathbb{P}$ a set of propositions and $\mathbb{T}$ a linearly-ordered time domain. Given a behavior w we note $w[t]$ its value at time $t \in \mathbb{T}$ taken to be a vector of Boolean values, and w_p the behavior $\mathbb{T} \to \mathbb{B}$ that is the projection of w on the component $p \in \mathbb{P}$. The models for both logics are defined over infinite time domains, $\mathbb{N}$ and $[0, d)$ respectively. In the following, we use the term *signal* to refer both to discrete continuous time behaviors.

We denote by $I \oplus J = \{t + t' \mid t \in I \text{ and } t'' \in J\}$ and $I \ominus J = \{t - t' \mid t \in I \text{ and } t'' \in J\}$ the Minkowski sum and difference of two intervals I and J, that we may simply note $t \oplus J$ and $t \ominus J$ when I is the punctual interval $[t, t]$. The semantics of a temporal logic formula φ with respect to a signal $w : \mathbb{T} \to \mathbb{B}^{\mathbb{P}}$ and time $t \in \mathbb{T}$ are given as follows:

$$\begin{array}{lcl}
(w,t) \models p & \leftrightarrow & w_p[t] = 1 \\
(w,t) \models \neg\varphi & \leftrightarrow & (w,t) \not\models \varphi \\
(w,t) \models \varphi_1 \vee \varphi_2 & \leftrightarrow & (w,t) \models \varphi_1 \text{ or } (w,t) \models \varphi_2 \\
(w,t) \models \varphi_1\,\mathcal{U}_I\varphi_2 & \leftrightarrow & \exists t' \in (t \oplus I) \cap \mathbb{T},\ (w,t') \models \varphi_2 \text{ and} \\
 & & \forall t'' \in (t,t') \cap \mathbb{T},\ (w,t'') \models \varphi_1
\end{array}$$

We say that w is a model of φ and write $w \models \varphi$ when $(w, 0) \models \varphi$. A signal can be "projected" for any formula φ to its *satisfaction signal* $w_\varphi : \mathbb{T} \to \mathbb{B}$ such that $w_\varphi[t] = 1$ if and only if $(w, t) \models \varphi$. We extend the notion of satisfaction signal w_φ to sets of formulas Ψ by letting $w_\Psi : \mathbb{T} \times \Psi \to \mathbb{B}$ be a multi-dimensional signal featuring the corresponding $|\Psi|$ satisfaction signals w_ψ for $\psi \in \Psi$. The satisfaction signals of φ and of all its sub-formulas ψ are given as the result of applying a monitoring procedure such as the one from [11] to w and φ.

3.1 Syntactic Rewritings

The fragment of temporal logic based on operators $\neg, \vee, \Diamond_I$ and $\mathcal{U}$ as introduced, has the same expressiveness as the fragment $\neg, \vee$ and $\mathcal{U}_I$, often taken as primitive MTL operators. This equivalence is based on the observation that the timed until operator

admits a decomposition into a timing part, and a sequential part [4]. For instance, we have $\varphi \mathcal{U}_{(a,b)} \psi \Leftrightarrow \Box_{(0,a]} \varphi \wedge \Box_{[a,a]} (\varphi \mathcal{U} \psi) \wedge \Diamond_{(a,b)} \psi$.

For the purpose of handling the negation of an until formula we introduce its dual operation *release*, with non-strict and strict versions as follows.

$$\varphi \tilde{\mathcal{R}} \psi := \psi \tilde{\mathcal{U}} (\psi \wedge \varphi) \ \vee \ \Box \psi$$
$$\varphi \mathcal{R} \psi := \varphi \mathcal{U} \, \mathrm{true} \ \vee \ \psi \mathcal{U} (\psi \wedge \varphi) \ \vee \ \psi \mathcal{U} (\psi \wedge \varphi \mathcal{U} \, \mathrm{true}) \ \vee \ \Box_{(0,\infty)} \psi$$

In the release property, the right-argument may possibly never occur. In discrete time $\neg(\varphi \tilde{\mathcal{U}} \, \psi) \Leftrightarrow \neg \varphi \tilde{\mathcal{R}} \, \neg \psi$, while in continuous time $\neg(\varphi \mathcal{U} \psi) \Leftrightarrow \neg \varphi \mathcal{R} \, \neg \psi$. We explain the MTL negation of an until as follows: $\varphi \mathcal{U} \psi$ does not hold if φ is immediately false, or if φ becomes false before (or immediately when) ψ becomes true, or if ψ never holds in the future.

3.2 Semantic Restrictions

We now introduce some definitions allowing us to place restrictions on the kind of signals we consider. Given a some constant in $\mathbb{T}$ we note $w^{a\cdot\cdot}$ the shifted sequence such that $w^{a\cdot\cdot}[t] = w[t + a]$ for all $t \in \mathbb{T}$. For a and b constants in $\mathbb{T}$ we say that some sequence w is *ultimately periodic* with *period* a and *prefix* b if $w^{a+b\cdot\cdot} = w^{b\cdot\cdot}$ holds. Some real interval I is said to be *uniform* with respect to signal w when $w[t] = w[t']$ for all t and t' in I; if moreover I is maximal with respect to $\subseteq$ we talk of a *maximally uniform* interval. The *variability* of a signal is taken to be the largest number of its maximally uniform segments in any unit length interval.

In what follows on one hand we assume that all continuous signals have finite variability, that is *MTL finite variability semantics*. On the other hand we consider arbitrary discrete signals, that is *LTL unrestricted semantics*. However we will always assume that the input signal to the LTL diagnostic is ultimately-periodic.

3.3 Sub-Models of a Formula

We define, similarly to the propositional case, *sub-signals* with domain $T \subseteq \mathbb{T} \times \mathbb{P}$ and a partial order relation $\sqsubseteq$ over them. Sub-signals u and v with respective domains R and S verify $u \sqsubseteq v$ if and only if $R \subseteq S$ and $u_p[t] = v_p[t]$ for all $(t, p) \in R$. Given formula φ and sub-signal v, we say that v is a *sub-model* of φ if for all signals $w \sqsupseteq v$ it holds $w \models \varphi$.

To ensure finite representation we introduce corresponding semantic restrictions (ultimate periodicity, finite variability) on sub-signals. The notions of uniform segment and shifting operation extend to sub-signals in a natural way. A finite variability sub-signal has a domain $\bigcup_{p \in P} T_p \times \{p\}$, where $T_p \subseteq \mathbb{T}$ is the domain of p, such that the number of segments in the intersection of each T_p with a unit interval admits a maximum. An ultimately-periodic signal v with period a and prefix b has a domain T such that $(t + b, p) \in T$ if and only if $(t + b + a, p) \in T$.

In the discrete case, a relative ultimate-periodicity hypothesis does not guarantee the existence of a minimal sub-model.

Example 2 (LTL). The formula $\varphi := \Diamond \Box\, p$ has no minimal ultimately-periodic sub-model over the discrete time domain $\mathbb{N}$. Consider the monotone sequence (v_i) of ultimately-periodic sub-models of φ with period 1 and prefix i, and domain $[i, \infty) \times \{p\}$. The sub-signal $\bigsqcap_{i \in \mathbb{N}} v_i = \bot$ is not a sub-model of φ.

We thus fix the period a and prefix b as given by the input signal, and restrict our analysis to sub-models with corresponding ultimate periodicity. For representation convenience we define the domain $\mathbb{T}_{a,b} = \{0, 1, \ldots, b-1, b^\infty, (b+1)^\infty, \ldots, (a+b-1)^\infty\}$ featuring *recurrent time* symbols t^∞ for $t = b, b+1, \ldots, a+b-1$. For an arbitrary signal w and element $t^\infty \in \mathbb{T}_{a,b}$, we have that $w[t^\infty] = 1$ iff $w[t + a\,n] = 1$ for all $n \in \mathbb{N}$. Any ultimately-periodic signal over $\mathbb{T} = \mathbb{N}$ with corresponding period and prefix may be seen without loss of generality as a signal over $\mathbb{T}_{a,b}$.

In the continuous case, uniformly bounding the variability does not even guarantee the existence of a minimal sub-model.

Example 3 (MTL). The formula $\varphi = p\,\mathcal{U}\,\mathrm{true}$ has no minimal sub-model over the dense time domain $[0, 1)$. Consider the monotone sequence (v_i) of sub-models of φ with variability 1, and domain $(0, \frac{1}{i+1}] \times \{p\}$. The sub-signal $\bigsqcap_{i \in \mathbb{N}} v_i = \bot$ is not a sub-model of φ.

To overcome limit problems we extend the temporal domain $\mathbb{T} = [0, d)$ to non-standard reals taken in $\mathbb{T}^+ = \{t^+ \,:\, t \in \mathbb{T}\}$ and $\mathbb{T}^- = \{t^- \,:\, t \in (0, d]\}$, with $\mathbb{T}^* = \mathbb{T} \cup \mathbb{T}^+ \cup \mathbb{T}^-$. We note $w[t^+]$ the right limit of some signal w at time t and $w[t^-]$ its left limit. Any finite variability signal over $\mathbb{T} = [0, d)$ may be extended to a signal over $\mathbb{T}^*$.

3.4 Temporal Implicants

We now introduce sentences based on (possibly infinite) conjunctions of unary predicates $p[t]$ and their negation $\neg p[t]$ for t in some domain $\mathbb{D}$, that we will take to be $\mathbb{T}_{a,b}$ or $\mathbb{T}^*$.

Definition 4 (Terms, Implicants and Prime Implicants). *Temporal terms are are defined using the grammar*

$$\gamma := p[t] \mid \neg p[t] \mid \gamma \wedge \gamma \mid \bigwedge_{t \in D} \theta[t]$$

where $p \in \mathbb{P}$ is a propositional variable, t is a time in $\mathbb{D}$, D is a subset of $\mathbb{D}$, and θ a function from $\mathbb{D}$ to terms. The semantics $\models$ of temporal terms relative to a signal w are as expected for literals and binary conjunctions, and for the case of general conjunctions are given by

$$w \models \bigwedge_{t \in D} \theta[t] \;\leftrightarrow\; \forall t \in D,\, w \models \theta[t]$$

An implicant *of some temporal formula φ is a temporal term γ such that $\gamma \Rightarrow \varphi$. We talk of* prime implicant *when γ is maximal with respect to $\Rightarrow$ modulo equivalence.*

The above definition of temporal terms is very general, allowing arbitrary functions θ under infinite conjunctions. However temporal terms can always be written in a simpler normal form as follows, which is straightforward to prove by structural induction.

Proposition 1 (Normal Form). *For every temporal term γ there exists a temporal term of the form $\bigwedge_{\ell\in L}\bigwedge_{t\in T_\ell}\ell[t]$ equivalent to γ, with L the set of propositional literals over $\mathbb{P}$. Assuming L is ordered this normal form is unique.*

For any term γ we will write $\bigwedge_{\ell\in L}\bigwedge_{t\in V_\ell^\gamma}\ell[t]$ its normal form. It is clear that normal form temporal terms are analogous to sub-signals over temporal domains $\mathbb{T}_{a,b}$ and $\mathbb{T}^*$. Notably $\Rightarrow$ defines a partial order over normal form terms; given arbitrary terms α and β, it holds $\alpha\Rightarrow\beta$ if and only if $V_\ell^\beta\subseteq V_\ell^\alpha$ for all $\ell\in L$.

By considering such terms we obtain the existence of at least one prime implicant for every satisfiable LTL and MTL formula. Recall that we consider full discrete semantics over $\mathbb{T}=\mathbb{N}$, and finite variability continuous semantics over $\mathbb{T}=[0,d)$.

Proposition 2 (Existence of Prime Implicants). *For any LTL formula φ and sequence w with period a and prefix b such that $w\models\varphi$ there exists a prime implicant γ of φ over $\mathbb{T}_{a,b}$ such that $w\models\gamma$. For any MTL formula φ and signal w such that $w\models\varphi$ there exists a prime implicant γ of φ over $\mathbb{T}^*$ such that $w\models\gamma$.*

Proof. Let us note Γ the set of implicants γ of φ such that $w\models\gamma$, that we may assume in normal form. Note that in both discrete and continuous cases, w seen as a temporal term is itself an implicant of φ. This gives us $\Gamma\neq\emptyset$. In the discrete case Γ is finite, so that there exists of a maximal element of Γ relative to $\Rightarrow$, which proves our proposition. In the continuous case, we demonstrate the existence of a maximal element of Γ relative to $\Rightarrow$ by direct application of Zorn's Lemma. Consider Δ an arbitrary totally ordered subset of Γ. We show that Δ has a maximum α in Γ, which we will identify as $\alpha=\bigwedge_{\ell\in L}\bigwedge_{t\in U_\ell}\ell[t]$, where each $U_\ell=\bigcap_{\gamma\in\Delta}\overline{V_\ell^\gamma}$ is the intersection over all $\gamma\in\Delta$ of the closure of V_ℓ^γ in $\mathbb{T}^*$. For this we need to establish (1) $\gamma\Rightarrow\alpha$ for all $\gamma\in\Delta$; (2) $\alpha\Rightarrow\beta$ for any upper bound β of Δ; and (3) $\alpha\in\Gamma$.
The facts (1) and (2) do not pose any difficulty. For (3) to hold, we need to show $w\models\alpha$ which is trivial, and to show $\alpha\Rightarrow\varphi$. To demonstrate that latter fact we take w' an arbitrary model of α and show that there exists some $\gamma\in\Delta$ such that $w'\models\gamma$. As $\Delta\subseteq\Gamma$ we will then have $w'\models\varphi$ by definition of Γ.
Assume, in search of a contradiction that $w'\not\models\gamma$ for all $\gamma\in\Delta$. For each $\gamma\in\Delta$ there exists $\ell\in L$ and $t\in V_\ell^\gamma$ such that $w'_\ell[t]=0$. We may construct a sequence (γ_i,ℓ_i,t_i) of $\Delta\times L\times\mathbb{T}^*$ such that $t_i\in V_{\ell_i}^\gamma$ and $w'_{\ell_i}[t_i]=0$ for all $i\in\mathbb{N}$, and such that (γ_i) is monotone and diverging, that is $\gamma_i\Rightarrow\gamma_j$ if $i\leq j$, and for all $\gamma\in\Delta$ there exists $i\in\mathbb{N}$ such that $\gamma\Rightarrow\gamma_i$. We take $s_i\in\mathbb{T}$ the standard part of t_i, that is $t_i\in\{s_i^+,s_i^-,s_i\}$. As L is finite, we can safely assume that the sequence (ℓ_i) is constant. As $\mathbb{T}$ is bounded, by Bolzano-Weierstrass Theorem we may in turn assume that the sequence (s_i) is monotone and convergent, an assumption that we extend to (t_i). Let us note ℓ the value of (ℓ_i), and t the limit of (t_i). As (γ_i) is monotone, the subsequence of times $(t_j)_{j\geq i}$ has all its values in $V_\ell^{\gamma_i}$, so that $t\in\overline{V_\ell^{\gamma_i}}$. In particular $t\in\bigcap_{i\in\mathbb{N}}\overline{V_\ell^{\gamma_i}}=\bigcap_{\gamma\in\Delta}\overline{V_\ell^\gamma}=U_\ell$ as (γ_i) is diverging. Then $t\in U_\ell$ gives us $w'_\ell[t]=1$. If t is a standard real, there exists i such that $t_i=t$, and in particular $w'_\ell[t_i]=1$. Else t is non-standard, and as (t_i) converges to t, by finite variability of w' there also exists i such that $w'[t_i]=1$. Yet $w'[t_i]=0$ by hypothesis. Contradiction! Therefore there exists $\gamma\in\Delta$ such that $w'\models\gamma$.

4 MTL Diagnostics

In this section, we propose an effective procedure to compute implicants of an MTL formula φ relative to a multi-dimensional signal $w : [0, d) \times \mathbb{P} \to \mathbb{B}$ of length d. First, note that the satisfaction signal w_φ of a given formula φ relative to a finite variability signals w has itself finite variability, the variability of satisfaction signals growing at most quadratically with the size of the formula. Like satisfaction, an explanation for a temporal formula is *time dependent* and should be a function from the time domain to formulas that explain satisfaction or violation at some time t. Analogously to the notion of satisfaction signal $w_\varphi : \mathbb{T} \to \mathbb{B}$ we define the notion of *explanation signal* noted $E(\varphi)$ such that $E(\varphi)[t]$ explains the satisfaction of φ by w at time $t \in \mathbb{T}$. We then construct explanations through definitions of $E(\varphi)[t]$ and its dual $F(\varphi)[t]$, which are inductive on the structure of formula φ, and on the times t at which explanations of its sub-formulas are required. We are able to guarantee finite representation by producing finitely variable explanation signals. We use selection functions ξ_φ to relate the truth of some formula φ at time t with the truth of its sub-formulas at some time $\xi_\varphi[t]$. Arbitrary selection functions may yet lead to explanations which are almost as large as the signal itself, yet we can find selection functions that allow best "explanation sharing". For instance given a non-singular interval I the same t' may belong to $t + I$ for every t in some interval T. Hence a selection function satisfying $\xi_\varphi[t] = t'$ for every $t \in T$ will use only one point to witness the satisfaction of $\varphi = \Diamond_I \psi$ throughout T.

4.1 Non-Standard MTL Semantics

The relation $<$ over the reals of $\mathbb{T}$ naturally extends to $\mathbb{T}^*$ with $t^- < t < t^+$. We then define over $\mathbb{T}^*$ the relation $\ll$ with $t \ll t'$ if and only if $t < t'$ or $t = t' \notin \mathbb{R}$. Interval notations using angled parentheses are introduced with $\langle t, t' \rangle := \{t'' \,:\, t \ll t'' \ll t'\}$. The sum of symbolic limit t^+, respectively t^-, and a real number a is taken as $(t + a)^+$, respectively $(t + a)^-$. The sum $t \boxplus I$ of some $t \in \mathbb{T}^*$ and a real interval I is then defined as the closure in $\mathbb{T}^*$ of $t \oplus I$, and similarly for the difference $t \boxminus I$.
We extend the satisfaction relation for temporal formulas to *non-standard* reals by writing $(w, t) \models \varphi$ if $w_\varphi[t] = 1$. Now conditions induced by timed eventually, and until operators can be expressed in terms of the *closed* intervals[1] of $\mathbb{T}^*$ we introduced.

Lemma 1. *For any formula φ, ψ, (finite variability) signal w, and time or symbolic limit $t \in \mathbb{T}^*$ we have*

- $(w, t) \models \Diamond_I \varphi$ *if and only if there exists* t' *in* $t \boxplus I$ *such that* $(w, t') \models \varphi$*;*
- $(w, t) \models \varphi \,\mathcal{U}\, \psi$ *if and only if there exists* $t' \gg t$ *such that* $(w, t') \models \psi$ *and for all* $t'' \in \langle t, t' \rangle$ *it holds* $(w, t'') \models \varphi$*.*

4.2 Explanation Operators

We may now formally define operators $E(\varphi)[t]$ and $F(\varphi)[t]$ providing explanations of φ, or $\neg\varphi$ relative to signal w at time $t \in \mathbb{T}^*$ in the form of temporal terms. The explanation

[1] A closed interval of $\mathbb{T}^*$ has the form $[t, t']$, where $t, t' \in \mathbb{T}^*$.

of a formula φ relative to a signal w is then given by the application of E or F at time 0. We let

$$\mathrm{Exp}(\varphi) = \begin{cases} E(\varphi)[0] & \text{if } (w,0) \models \varphi \\ F(\varphi)[0] & \text{otherwise} \end{cases}$$

with

$$\begin{aligned}
E(p)[t] &= p[t] & F(p)[t] &= \neg p[t] \\
E(\neg\varphi)[t] &= F(\varphi)[t] & F(\neg\varphi)[t] &= E(\varphi)[t] \\
E(\varphi \vee \psi)[t] &= E(\xi_{\varphi\vee\psi}[t])[t] & F(\varphi \vee \psi)[t] &= F(\varphi)[t] \wedge F(\psi)[t] \\
E(\Diamond_I \varphi)[t] &= \begin{cases} E(\varphi)[t+a] & \text{if } I = [a,a] \\ E(\varphi)[\xi_{\Diamond_I\varphi}[t]] & \text{otherwise} \end{cases} & F(\Diamond_I \varphi)[t] &= \bigwedge_{t' \in t \boxplus I} F(\varphi)[t'] \\
E(\varphi \,\mathcal{U}\, \psi)[t] &= E(\psi)[\xi_{\varphi\,\mathcal{U}\,\psi}[t]] \wedge \bigwedge_{\substack{t' \in \\ \langle t, \xi_{\varphi\,\mathcal{U}\,\psi}[t]\rangle}} E(\varphi)[t'] & F(\varphi \,\mathcal{U}\, \psi)[t] &= E(\neg\varphi \mathcal{R} \neg\psi)[t]
\end{aligned}$$

where $\xi_{\varphi\vee\psi}$ from $\mathbb{T}^*$ to formulas, $\xi_{\Diamond_I\varphi}$ and $\xi_{\varphi\,\mathcal{U}\,\psi}$ from $\mathbb{T}^*$ to $\mathbb{T}^*$ are *selection functions* such that for all $t \in \mathbb{T}^*$, it holds $\xi_{\varphi\vee\psi}[t] \in \{\varphi, \psi\}$, $\xi_{\Diamond_I\varphi}[t] \in t \boxplus I$ and $\xi_{\varphi\,\mathcal{U}\,\psi}[t] \gg t$.

We say that a selection function ξ_φ is correct with respect to w if for all $t \in \mathbb{T}^*$ such that $(w,t) \models \varphi$ we have

- $(w,t) \models \xi_\varphi[t]$ when φ is of the form $\varphi_1 \vee \varphi_2$,
- $(w, \xi_\varphi[t]) \models \varphi_1$ when φ is of the form $\Diamond_I \varphi_1$,
- $(w, \xi_\varphi[t]) \models \varphi_2$ and $\forall t' \in \langle t, \xi[t]\rangle,\ (w,t') \models \varphi_1$ when φ is of the form $\varphi_1 \,\mathcal{U}\, \varphi_2$.

The following result is straightforward to prove from Lemma 1.

Theorem 1 (Soundness). *A term* $\mathrm{Exp}(\varphi)$*, correct with respect to signal* w *is a solution to the diagnostics problem of* φ *with respect to* w*.*

Moreover, given finite variability selection functions our explanations can be effectively represented.

Proposition 3 (Finite Representation). *Assuming all selection functions are finitely variable, the term* $\mathrm{Exp}(\varphi)$ *has a normal form* $\bigwedge_{\ell\in L} \bigwedge_{t'\in T_\ell} \ell[t']$ *such that each* T_ℓ *a* finite union of intervals *of* $\mathbb{T}^*$.

4.3 Computation of Selection Functions

We describe procedures that define explicit instances of selection functions, and satisfying the correctness and finite variability criteria. The explanation operators can be made constructive when the normalization of the terms they produce is interleaved with the instantiation of selection functions over intervals appearing in the normalization process. It is indeed sufficient to define selection functions piecewise on closed intervals T of $\mathbb{T}^*$. Furthermore we can assume that for such intervals T, formula φ holds for all $t \in T$ as the correctness assumption is void outside such intervals, with the finite variability assumption then trivial to match.

Disjunction Consider the formula $\varphi \vee \psi$ and the signal w. A finitely variable selection function $\xi_{\varphi\vee\psi}$ and correct with respect to w can be constructed as follows. By finite variability hypothesis on w, the satisfaction signal $w_{\varphi,\psi} : \mathbb{T} \to \mathbb{B}$ has finite variability over any interval T where $\varphi \vee \psi$ holds. We partition T in k maximally uniform intervals T_i, and take $\xi_{\varphi\vee\psi}[t] = \varphi$ over intervals T_i where $(w, t) \models \varphi$, and $\xi_{\varphi\vee\psi}[t] = \psi$ over other intervals. The function $\xi_{\varphi\vee\psi}$ is uniform over all T_i, so has finite variability.

Timed Eventually Now consider the formula $\Diamond_I \varphi$ for I non-singular and the signal w, and assume that the formula is satisfied over some interval T. we want to build a procedure that generates a small set of witnesses of φ that account for the satisfaction of $\Diamond_I \varphi$ by w over T. The satisfaction of $\Diamond_I \varphi$ over T can be explained by the satisfaction of φ at some time points within the interval $T \boxplus I = \bigcup_{t\in T} t \boxplus I$, and in particular the satisfaction of φ at some $s \in T + I$ provides a sufficient explanation for the satisfaction of $\Diamond_I \varphi$ for all $t \in (s \boxminus I) \cap T$. We use these two observations to generate a piecewise constant selection function $\xi_{\Diamond_I\varphi}$ defined over T and correct relative to a signal w.

1: $\xi_{\Diamond_I\varphi} \leftarrow \emptyset$
2: **while** $T \neq \emptyset$ **do**
3: $\quad t \leftarrow \min(T)$
4: $\quad S \leftarrow (t \boxplus I) \cap \{t' \ : \ (w, t') \models \varphi\}$
5: $\quad s \leftarrow \max(S)$
6: $\quad R \leftarrow (s \boxminus I) \cap T$
7: $\quad \xi_{\Diamond_I\varphi} \leftarrow \xi_{\Diamond_I\varphi} \cup (R \times \{s\})$
8: $\quad T \leftarrow T \setminus R$
9: **return** $\xi_{\Diamond_I\varphi}$

(a)

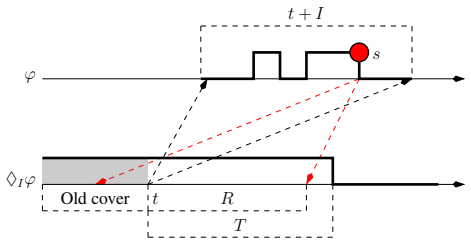

(b)

Fig. 2. (a) Algorithm to find a correct instance of $\xi_{\Diamond_I\varphi}$ over T relative to signal w; (b) Example of R and s computation for $\Diamond_I \varphi$.

We present the procedure in Figure 2-(a); it works as follows. The selection function is initialized (line 1) as nowhere defined. In every iteration of the main while loop (line 2), we find a time domain $S = (t \boxplus I) \cap \{t' \ : \ (w, t') \models \varphi\}$ such that φ is satisfied inside S and any point in S provides a sufficient (Lemma 1) explanation for the satisfaction of $\Diamond_I \varphi$ at t taken as the earliest time of T. Such set S is obtained directly from the satisfaction signal w_φ, that we suppose already computed by the monitoring procedure. We then take s the latest time of S, which constitutes a minimal subset of S sufficient to explain the satisfaction of $\Diamond_I \varphi$ throughout the domain $s \boxminus I$; when intersected with T it gives R, a prefix of T. At the end of the iteration, the definition of $\xi_{\Diamond_I\varphi}$ over the interval R is taken as s, which we may write $R \times \{s\}$ identifying selection functions with subsets of $\mathbb{T}^* \times \mathbb{T}^*$ (line 7). The covered prefix R can be removed from T (line 8). The procedure terminates when T the domain remaining to cover becomes empty.

Untimed Until Consider the formula $\varphi \mathcal{U} \psi$ and the signal w and assume that the formula is satisfied over T, taken without loss of generality to be a closed interval of $\mathbb{T}^*$. For $t \in \mathbb{T}$, similarly to the case of timed eventually a single witness $t' > t$ of ψ along with a uniform interval (t, t') where φ holds is sufficient to explain the satisfaction of $\varphi \mathcal{U} \psi$ over the whole interval $[t, t')$. With such observations we generate a piecewise constant selection function $\xi_{\varphi \mathcal{U} \psi}$ correct with respect to some signal w and defined over T. We make use of a subroutine $W(\varphi, \psi, t)$ that returns the set of witnesses of ψ in signal w that are sufficient to explain $\varphi \mathcal{U} \psi$ at time $t \in \mathbb{T}^*$ where $\varphi \mathcal{U} \psi$ holds. Thanks to Lemma 1 we have $W(\varphi, \psi, t) = \{t' \gg t \ : \ (w, t') \models \psi \ \text{ and } \forall t'' \in \langle t, t' \rangle, \ (w, t'') \models \varphi\}$. Assuming the satisfaction signals w_φ and w_ψ given by the monitoring algorithm, the procedure $W(\varphi, \psi, t)$ can be realized as follows. First decompose the domain $\{t' \in \mathbb{T}^* \ : \ t' \gg t\}$ as a finite partition into uniform intervals of $\mathbb{T}^*$ with respect to $w_{\varphi,\psi}$ that we can assume of the form $[t_i, t_i]$ and (t_i, t_{i+1}) with t_i an ordered sequence of times. Start from the interval containing $t_0 = t$ and iterate through intervals, remembering the latest interval where ψ holds, until φ stops holding at $[t_i, t_i]$ or (t_i, t_{i+1}). The latest interval witness of ψ is then either (t_{i-1}, t_i) or $[t_i, t_i]$, and we take $W(\varphi, \psi, t) = [t, t_i)$.

1: $\xi_{\varphi \mathcal{U} \psi} \leftarrow \emptyset$
2: **while** $T \neq \emptyset$ **do**
3: $\quad t \leftarrow \min(T)$
4: $\quad S \leftarrow W(\varphi, \psi, t)$
5: $\quad$ **if** $S \cap T = \emptyset$ **then** $s \leftarrow \min(S)$
6: $\quad$ **else** $\quad s \leftarrow \max(S \cap T)$
7: $\quad R \leftarrow [t, s) \cap T$
8: $\quad \xi_{\varphi \mathcal{U} \psi} \leftarrow \xi_{\varphi \mathcal{U} \psi} \cup (R \times \{s\})$
9: $\quad T \leftarrow T \setminus R$
10: **return** $\xi_{\varphi \mathcal{U} \psi}$

(a)

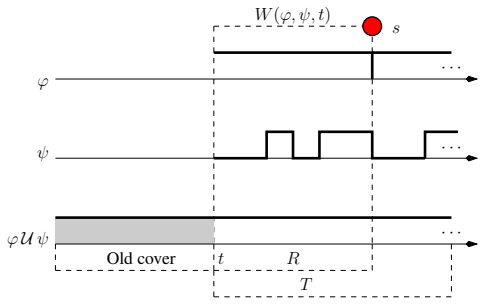

(b)

Fig. 3. (a) Algorithm to find a correct instance of $\xi_{\varphi \mathcal{U} \psi}$ over T relative to signal w; (b) Example of R and s computation for $\varphi \mathcal{U} \psi$.

We present the main procedure to compute the selection function in Figure 3-(a). The procedure first assigns $\xi_{\varphi \mathcal{U} \psi}$ the empty function $\emptyset$ (line 1). In every iteration of the while loop, we compute an interval S whose elements s are witnesses of ψ providing sufficient explanation for the satisfaction of the $\varphi \mathcal{U} \psi$ throughout $[t, s)$. When S lies entirely outside T we take s to be the earliest suitable witness of ψ, so as not to impose a condition on φ beyond it (line 5). When S intersects with T on the contrary we look for the latest suitable witness of ψ in their intersection (line 6). As a direct corollary of Lemma 1 the interval $R = [t, s) \cap T$ is now accounted for, hence we define $\xi_{\varphi \mathcal{U} \psi}$ as taking the value s over interval R (line 8). Eventually R can be removed from T for the next iteration (line 9), and the procedure terminates when T becomes empty.

4.4 Discussion

Our procedure does not guarantee minimality, with for instance propositional tautologies not being recognized. However we obtain some form of temporal minimality through the proposed construction for selection functions. Intuitively each time a witness is required we select the furthest away, which maximizes the interval over which that witness is valid. Let φ be an *eventually* or *until* formula, w a trace such that φ holds for w on some domain T. We claim that selection functions ξ_φ constructed by our algorithms choose a set of witnesses $\xi_\varphi[T] = \{\xi_\varphi[t] \ : \ t \in T\}$ that is minimal relative to w.

The main advantage of our explanation principle is its hierarchical character: every sub-formula has its own explanation, which is then used in turn to account for the satisfaction or violation of its super-formulas. This makes the process of fault-finding transparent: if the fault lies in the specification then it can be localized syntactically, otherwise it lies in the system in which case the explanation of each sub-formula provides important insight on what went wrong. This also allows to solve the diagnostics problem efficiently. Under a uniform *bounded* variability assumption, computing $\mathrm{Exp}(\varphi, w)$ with our algorithms takes time quadratic in the size of the formula, and linear in the size of the input signal. The minimal diagnostics problem, that is to find a *prime* implicant of φ that is satisfied by w, is at least as hard as the satisfiability of MTL. For a bounded time domain the satisfiability of MTL is EXPSPACE-complete [3].

Let us now illustrate the overall process of deriving an explanation.

Example 4. We take φ the formula $\Diamond_{[1,2]} \Box(p \rightarrow \neg(q \,\mathcal{U}\, r))$, and $w : [0, 6) \rightarrow \mathbb{B}^3$ the right-continuous signal illustrated in Figure 4. The top-down computation of $\mathrm{Exp}(\varphi)$ is shown in terms of sub-signals, inductively extracted from satisfaction signals. We found

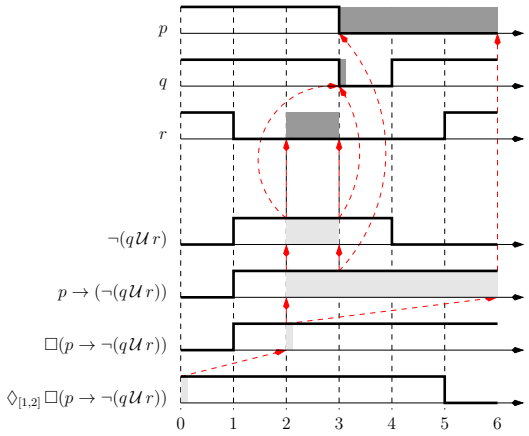

Fig. 4. Computing $\mathrm{Exp}(\varphi)$, for $\varphi = \Diamond_{[1,2]} \Box(p \rightarrow \neg(q \,\mathcal{U}\, r))$.

the diagnostics $\mathrm{Exp}(\varphi) \Leftrightarrow \bigwedge_{t \in (2,3]} \neg r[t] \wedge \bigwedge_{t \in [3,3]} \neg q[t] \wedge \bigwedge_{t \in [3,6)} \neg p[t]$. It reads as follows: r is false between 2 and 3, q is false at time 3 and p is false from 3 onwards.

5 LTL Ultimately-Periodic Diagnostics

Our explanation scheme for LTL formulas solves the diagnostics problem with respect to an ultimately-periodic sequence w with period a and prefix b. The first remark we make is that the satisfaction sequences w_φ of any LTL formula φ is then also ultimately-periodic, with same period and prefix. Let us say that a property φ is *future* if for any sequence w and time constant a it holds $(w, a) \models \varphi$ iff $w^{a\cdot\cdot} \models \varphi$. It is trivial to check that future properties have a satisfaction signal that preserves the prefix and period of its sub-formulas satisfaction signals. Following this remark we may assume that we dispose of the satisfaction sequences for each sub-formula; see [12] for the monitoring of ultimately-periodic sequences.

5.1 Recurrent LTL Semantics

We equip $\mathbb{T}_{a,b}$ with a pseudo-successor function $\widetilde{+}\,1$ that we define by $t \,\widetilde{+}\, 1 = t + 1$ for $t < b - 1$; $(b-1) \,\widetilde{+}\, 1 = b^\infty$; $t^\infty \,\widetilde{+}\, 1 = (t+1)^\infty$ for $b \leq t < a + b - 1$; and $(a + b - 1)^\infty \,\widetilde{+}\, 1 = b^\infty$. The n^{th} successor of a symbolic time $t \in \mathbb{T}_{a,b}$ is then given by $t \,\widetilde{+}\, 0 = t$ and $t \,\widetilde{+}\, n = (t \,\widetilde{+}\, 1) \,\widetilde{+}\, (n-1)$ for $n > 1$. We further define on $\mathbb{T}_{a,b}$ a preorder relation $\preceq$ such that $t \preceq t'$ if and only if there exists $n \geq 0$ such that $t' = t \,\widetilde{+}\, n$. The usual interval notations $[t, t')$ are extend to arbitrary $t, t' \in \mathbb{T}_{a,b}$ by letting $[t, t') = \{t'' \,:\, \exists n \geq 0,\, t'' = t \,\widetilde{+}\, n \text{ and } \forall k \leq n,\, t \,\widetilde{+}\, k \neq t'\}$. This coincides with the usual interval notation when t, t' are natural numbers, i.e. in the prefix. The semantics of LTL are extended to *recurrent* times by writing $(w, t^\infty) \models \varphi$ if $(w, t + a\, n) \models \varphi$ for all $n \in \mathbb{N}$. We then have the following equivalences.

Lemma 2. *For any formula φ, ψ, sequence w, and symbolic time $t \in \mathbb{T}_{a,b}$ we have*

- $(w, t) \models \bigcirc \varphi$ *if* $(w, t \,\widetilde{+}\, 1) \models \varphi$*;*
- $(w, t) \models \varphi \,\mathcal{U}\, \psi$ *if there exists* $t' \succeq t$ *such that* $(w, t') \models \psi$ *and* $(w, t'') \models \varphi$ *for all* $t'' \in [t, t')$.

Note that when w is ultimately-periodic with period a and prefix b, the preceding formulas are satisfied *if and only if* corresponding conditions hold.

5.2 Explanation Operators

The explanation scheme for LTL is derived from the propositional one by making operators E and F time dependent. Such operators take as input a formula φ and a symbolic time $t \in \mathbb{T}_{a,b}$, and return propositional terms over unary predicates $p[t']$ with $p \in \mathbb{P}$ and $t' \in \mathbb{T}_{a,b}$. The explanation Exp is then given by E or F at time 0. We let

$$\mathrm{Exp}(\varphi) = \begin{cases} E(\varphi)[0] & \text{if } w \models \varphi \\ F(\varphi)[0] & \text{otherwise} \end{cases}$$

with

$$
\begin{aligned}
E(p)[t] &= p[t] & F(p)[t] &= \neg p[t] \\
E(\neg\varphi)[t] &= F(\varphi)[t] & F(\neg\varphi)[t] &= E(\varphi)[t] \\
E(\bigcirc\varphi)[t] &= E(\varphi)[t \widetilde{+} 1] & F(\bigcirc\varphi)[t] &= F(\varphi)[t \widetilde{+} 1] \\
E(\varphi \vee \psi)[t] &= E(\xi_{\varphi\vee\psi}[t])[t] & F(\varphi \vee \psi)[t] &= F(\varphi)[t] \wedge F(\psi)[t] \\
E(\Box\varphi)[t] &= \bigwedge_{t' \succeq t} E(\varphi)[t'] & & \\
E(\varphi\,\tilde{\mathcal{U}}\,\psi)[t] &= E(\psi)[\xi_{\varphi\,\tilde{\mathcal{U}}\,\psi}[t]] \wedge \bigwedge_{t' \in [t, \xi_{\varphi\,\tilde{\mathcal{U}}\,\psi}[t])} E(\varphi)[t'] & F(\varphi\,\tilde{\mathcal{U}}\,\psi)[t] &= E(\neg\varphi\,\tilde{\mathcal{R}}\,\neg\psi)[t]
\end{aligned}
$$

where $\xi_{\varphi\vee\psi}$ from $\mathbb{T}_{a,b}$ to formulas, and $\xi_{\varphi\,\tilde{\mathcal{U}}\,\psi}$ from $\mathbb{T}_{a,b}$ to $\mathbb{T}_{a,b}$ are *selection functions* such that for all $t \in \mathbb{T}_{a,b}$, it holds $\xi_{\varphi\vee\psi}[t] \in \{\varphi, \psi\}$ and $\xi_{\varphi\,\tilde{\mathcal{U}}\,\psi}[t] \succeq t$.

We say that a selection function ξ_φ is correct with respect to w if for all $t \in \mathbb{T}_{a,b}$ such that $(w,t) \models \varphi$ we have

- $(w,t) \models \xi_\varphi[t]$ when φ is of the form $\varphi_1 \vee \varphi_2$,
- $(w, \xi_\varphi[t]) \models \varphi_2$ and $\forall t' \in [t, \xi[t]),\ (w,t') \models \varphi_1$ when φ is of the form $\varphi_1\,\tilde{\mathcal{U}}\,\varphi_2$.

Correct selection functions can easily be constructed, knowing that the domain $\mathbb{T}_{a,b}$ is finite. The following result is straightforward to prove from Lemma 2.

Theorem 2 (Soundness). *Under the assumption that selection functions are correct with respect to w, $\mathrm{Exp}(\varphi)$ is a solution to the diagnostic problem of φ and w.*

We now give an example of diagnostic produced by our explanation principle.

Example 5. Let φ be the formula $r \rightarrow \bigcirc\Box(p\,\tilde{\mathcal{U}}\,\neg q)$ and $w : \mathbb{N} \rightarrow \mathbb{B}^3$ be the sequence with period 4 and prefix 2 defined by the following ω-regular expression:

$$p\,q\,r \cdot p\,q\,r \cdot (\overline{p}\,\overline{q}\,\overline{r} \cdot p\,q\,\overline{r} \cdot \overline{p}\,\overline{q}\,r \cdot p\,q\,r)^\omega$$

We find $\mathrm{Exp}(\varphi) = p[1] \wedge \neg q[2^\infty] \wedge p[3^\infty] \wedge \neg q[4^\infty] \wedge p[5^\infty]$, which may be written as the ω-regular language $\mathrm{true} \cdot p \cdot (\overline{q} \cdot p \cdot \overline{q} \cdot p)^\omega$.

6 Conclusion and Perspectives

We have enriched MTL monitoring, and LTL model checking techniques with a focused analysis of the causes of satisfaction/violation of such a specification by a given temporal behavior. For monitoring applications we plan to develop an online version of our algorithms, which may then be integrated in the monitoring procedure to allow fault analysis on a simulation, or even a real execution without having to save the monitored signals. Looking more generally at temporal implicants it would be interesting to study alternative formulations, based on the desiderata of [2] that list atomicity of literals, closure under intersection, duality with implicates, etc. Our approach to diagnosis may then be transfered to other problems, such as fault localization where a system model is assumed to be available.

Acknowledgements This work was supported by the MISTRAL project A-1341-RT-GP coordinated by the European Defence Agency (EDA) and funded by 8 contributing Members (France, Germany, Italy, Poland, Austria, Sweden, Netherland and Luxenbourg) in the framework of the Joint Investment Programme on Second Innovative Concepts and Emerging Technologies (JIP-ICET 2) and the IKT der Zukunft of Austrian FFG project HARMONIA (nr. 845631).

References

1. Ilan Beer, Shoham Ben-David, Hana Chockler, Avigail Orni, and Richard J. Trefler. Explaining counterexamples using causality. In *Computer Aided Verification*, pages 94–108, 2009.
2. Meghyn Bienvenu. Prime implicates and prime implicants: From propositional to modal logic. *Journal of Artificial Intelligence Research*, 36(1):71–128, 2009.
3. Patricia Bouyer, Nicolas Markey, Joël Ouaknine, and James Worrell. On expressiveness and complexity in real-time model checking. In *Automata, Languages and Programming*, pages 124–135. Springer, 2008.
4. Deepak D'Souza and Nicolas Tabareau. On timed automata with input-determined guards. In *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, pages 68–83. Springer, 2004.
5. Alex Groce. Error explanation with distance metrics. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 108–122, 2004.
6. Joseph Y. Halpern and Judea Pearl. Causes and explanations: A structural-model approach: Part 1: Causes. In *Uncertainty in Artificial Intelligence*, pages 194–202, 2001.
7. François Hantry and Mohand-Said Hacid. Handling conflicts in depth-first search for LTL tableau to debug compliance based languages. In *Formal Languages and Analysis of Contract-Oriented Software*, pages 39–53, 2011.
8. Barbara Jobstmann, Stefan Staber, Andreas Griesmayer, and Roderick Bloem. Finding and fixing faults. *J. Comput. Syst. Sci.*, 78(2):441–460, 2012.
9. Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-time systems*, 2(4):255–299, 1990.
10. Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 19(3):291–314, 2001.
11. Oded Maler and Dejan Nickovic. Monitoring properties of analog and mixed-signal circuits. *STTT*, 15(3):247–268, 2013.
12. Nicolas Markey and Jean-François Raskin. Model checking restricted sets of timed paths. In *CONCUR*, volume 3170 of *Lecture Notes in Computer Science*, pages 432–447, 2004.
13. Subhankar Mukherjee and Pallab Dasgupta. Computing minimal debugging windows in failure traces of AMS assertions. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 31(11):1776–1781, 2012.
14. Ingo Pill, Simone Semprini, Roberto Cavada, Marco Roveri, Roderick Bloem, and Alessandro Cimatti. Formal analysis of hardware requirements. In *Design Automation Conference*, pages 821–826, 2006.
15. Amir Pnueli. The temporal logic of programs. In *Foundations of Computer Science*, pages 46–57. IEEE, 1977.
16. Viktor Schuppan. Towards a notion of unsatisfiable and unrealizable cores for LTL. *Sci. Comput. Program.*, 77(7-8):908–939, 2012.
17. Viktor Schuppan and Armin Biere. Shortest counterexamples for symbolic model checking of LTL with past. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 493–509, 2005.