Combining the Temporal and Epistemic Dimensions for MTL Monitoring

Eugene Asarin[1], Oded Maler[2], Dejan Nickovic[3], and Dogan Ulus[2]

[1] IRIF, Paris Diderot University
[2] Verimag, CNRS / University of Grenoble-Alpes, France
[3] Austrian Institute of Technology

Abstract. We define a new notion of satisfaction of a temporal logic formula φ by a behavior w. This notion, denoted by $(w, t, t') \models \varphi$, is characterized by two time parameters: the position t from which satisfaction is considered, and the end of the (finite) behavior t' which indicates how much do we *know* about the behavior. We define this notion in dense time where φ is a formula in the future fragment of metric temporal logic (MTL) and w is a Boolean signal of bounded variability. We show that the set of all pairs (t, t') such that $(w, t, t') \models \varphi$ can be expressed as a finite union of two-dimensional zones and give give an effective procedure to compute it.

1 Introduction and Motivation

Within the traditional use of temporal logic (TL) in verification, formulas are interpreted over non-terminating[1] behaviors, viewed mathematically as ω-words. These are sequences which are infinite in one dimension with a time domain order-isomorphic to $\mathbb{N}$ (to $\mathbb{R}_+$ or $\mathbb{Q}_+$, if we consider dense time). In this setting, the availability of a generative model of the system dynamics is assumed in a form of a transition system (automaton) where all those behaviors are represented by infinite runs that go through cycles. Likewise, the TL specification can also be translated to an ω-automaton and the verification problem reduces to a test of inclusion between two ω-regular languages [31]. This problem can be solved by reasoning about cycles in finite-state automata.

Historical Remark: This was not always the point of view in the early works of logicians on tense logic, before the importation of TL to verification by Pnueli [24, 25]. Kamp [13] who added the *until* and *since* operators to the original tense logic of Prior [26], and showed expressive equivalence to the first-order theory of sequences, considered arbitrary time structures satisfying order axioms, that could be infinite in both directions. Regular languages over bi-infinite words, indexed by $\mathbb{Z}$ rather than $\mathbb{N}$, were considered by Nivat and Perrin in [23]. The current ω-view has been nailed down in the *anchored* interpretation of Manna and Pnueli [21] which associated an initial state with every computation and, moreover, considered satisfaction from this initial state as

[1] In the context of reactive systems, finite behaviors are sometimes even considered anomalous, representing deadlocks.

having a special status compared to satisfaction from an arbitrary point in time. Readers interested in more historical and technical details are advised to consult [30] and the references therein. ⌟

There are several contemporary motivations to consider finite, time bounded behaviors as the semantic model for TL. In many (if not most) real-life situations, especially in the hybrid cyber-physical world, exhaustive verification is impossible and one resorts to simulation-based (runtime, dynamic, lightweight) verification, where behaviors are generated individually. Each of these behaviors is checked for property satisfaction, or using a language-theoretic terminology, the inclusion test of model checking is replaced by numerous membership tests. We use the term *monitoring* for this activity. An important advantage of monitoring is that it can be applied to systems models not admitting a clean description (programs, simulators, black boxes) and hence not amenable to formal reasoning. For a behavior to be observed and checked by a mortal agent (or analyzed by a terminating program), it should be finite and the semantics of the specifications should be adapted to yield answers based on such finite behaviors.

This problem had to be (and has been) addressed by anyone developing such monitoring tools [1]. One way to tackle this issue is to provide finitary interpretation of TL. The truncated semantics for future TL is rigorously studied in [9], where weak, strong and neutral interpretation of the temporal specifications are proposed. This work is further developed in [8], providing the topological characterization of the weakness and strength of temporal formulas. In [6], the authors study LTL interpreted over finite behaviors and show the limited expressiveness of the logic in the finitary setting. They propose linear dynamic logic over finite traces (LDL_f) that significantly increases the expressiveness of the logic without additional computational cost. Another way to address the interpretation of TL over finite behaviors is to employ a 3-valued semantics ranging over $\{0, 1, \bot\}$ where $\bot$ is viewed as *unknown* to model the fact that the finite behavior does not contain sufficient information needed to determine the satisfaction or a violation of a temporal formula at a given instant in time [3, 27, 2]. We finally mention [20], which discusses various options of handling temporal logic over finite behaviors.

Another motivation comes from the application of specification formalisms outside the traditional design-time verification framework. After all, monitoring can be applied to data measured from real physical systems, not only to models [16]. In monitoring real systems during their execution we would like to detect some alarming patterns of behavior as they occur, so as to do something about them. In not so safety-critical situations, we would like to analyze a given behavior and distinguish, say, periods in which some bounded-response property has been satisfied from periods when it was not. All these application domains call for an approach where finite segments of a behavior, not necessarily starting at time zero, and certainly not ending at the "end", are the major objects of study. In this setting, monitoring has often to be performed *online*, with the values of the monitored behavior being disclosed progressively as time goes by.

The major contribution of this work is in defining a two-dimensional notion of satisfiability, denoted by $(w, t, t') \models \varphi$, where t indicates, as usual, the position from which satisfaction is considered, and t' indicates the endpoint of the signal, the limit of our current knowledge about it. This work is partially inspired by [28] where the relation $(w, t, t') \models \varphi$ means that the segment $w[t, t')$ of a Boolean signal w matches a timed

regular expression. In that paper, the match set $\mathcal{M}(\varphi, w) = \{(t, t') : (w, t, t') \models \varphi\}$ was shown to be computable and to consist of a finite union of zones. While we borrow some of the two-dimensional techniques from [28], it turns out that for TL, which is less symmetric than regular expressions with respect to the direction of time, these notions are trickier and require a distinction between the temporal and epistemic components. Note that unlike other works that combine knowledge and time [12, 29, 10], where knowledge is relative to different agents in a distributed system who may observe different variables and events at different times, our notion is centralized and is focused on the knowledge associated with the unfolding of time.

The rest of the paper is organized as follows. In Section 2 we present the algorithm for MTL monitoring of Boolean signals as developed in [17, 22]. It is based on two major operation, interval back-shifting to treat the *timed eventually* operator and another operation on intervals to handle the *untimed until*. In Section 3 we define the two-dimensional satisfaction relation for future MTL and its associated match-set computation problem. We show that the latter can be solved by extending the above two interval-based operations to deal with zones. In Section 4 we illustrate how our implementation of the algorithm works on MTL formulas of a practical interest. Section 5 is devoted to conclusions and suggestions for future work. Needless to say, the results and insights obtained for dense time and MTL hold, as a degenerate case, for the discrete time setting of LTL and sequences.

2 Preliminaries

A Boolean signal w is a function from an interval $dom(w) = [0, \ell)$ to $\mathbb{B}^n$. The signal is infinite when $\ell = \infty$, and finite, otherwise. We restrict ourselves to signals that satisfy the sanity condition of *bounded variability*, which for finite signals means that $dom(w)$ can be partitioned into finitely many intervals, and w is constant in each interval. Such an interval is said to be *maximal* if it is not strictly contained in another interval where the signal is constant.

The syntax of the future fragment of metric temporal logic (MTL) as defined in [15] is given by

$$\varphi := p \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \mathcal{U}_I \varphi_2$$

where $p \in \{p_1, \ldots, p_n\}$ is a propositional variable and I is any non-empty interval of the form $[a, b]$, $[a, b)$, $(a, b]$ or (a, b) with a and b being integers. To avoid tedious case analysis and focus on the new features introduced by the two-dimensional notion of satisfaction, we treat only the case $I = [a, b]$. It has been shown in [19] that with this restriction, if w decomposes into unions of maximal intervals which are left-closed right-open, all the other signals generated during the monitoring procedure admit such a decomposition without singular points, a fact that will simplify the presentation. For the same reason, we do not explore all variants of the timed *until* operator and focus on the non-strict version, whose semantics is given using the standard satisfaction relation $(w, t) \models \varphi$ indicating that w satisfies φ *from* position t:

$$(w, t) \models \varphi_1 \mathcal{U}_{[a,b]} \varphi_2 \text{ iff } \exists r \in [t + a, t + b](w, r) \models \varphi_2 \wedge \forall r' \in [t, r]\ (w, r') \models \varphi_1$$

The timed *eventually* operator $F_{[a,b]}$ is a degenerate case where φ_1 is replaced by *true*, $F_{[a,b]}\varphi = \top \mathcal{U}_{[a,b]}\varphi$. It only requires that t will occur sometime in $[t+a, t+b]$. Its dual, the timed *always*, which require φ to hold throughout the interval, is defined as $G_{[a,b]}\varphi = \neg(F_{[a,b]}\neg\varphi)$.

It has been shown in [7, 22] that the timed until operator can be rewritten as a combination of $F_{[a,b]}$ and the *untimed* until $\mathcal{U}$ which does not put any restriction on the future time point r:

$$\varphi_1 \mathcal{U}_{[a,b]} \varphi_2 = G_{[0,a]}\varphi_1 \mathcal{U} \varphi_2 \wedge F_{[a,b]}\varphi_2$$

Hence from now on we consider the syntax

$$\varphi := p \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid F_{[a,b]} \mid \varphi_1 \mathcal{U} \varphi_2$$

The monitoring procedure of [17, 22] is based on reformulating the time-dependent satisfaction relation in terms of *satisfaction signals*. A satisfaction signal for a formula φ relative to signal w is a one-dimensional Boolean signal[2] $\varphi(.)$ such that

$$\varphi(t) = \begin{cases} 1 \text{ if } (w,t) \models \varphi \\ 0 \text{ if } (w,t) \not\models \varphi \end{cases}$$

The standard semantics of MTL can be reformulated in terms of such signals. We use w_p to denote the projection of w on variable p.

Definition 1 (MTL Semantics with Satisfaction Signals). *The semantics of MTL formulas with respect to a Boolean signal w is defined inductively:*

$$\begin{aligned}
&p(t) = w_p(t) \\
&(\neg\varphi)(t) = \neg(\varphi(t)) \\
&(\varphi \vee \psi)(t) = \varphi(t) \vee \psi(t) \\
&(F_{[a,b]}\varphi)(t) = \bigvee_{r \in [t+a, t+b]} \varphi(r) \\
&(\varphi_1 \mathcal{U} \varphi_2)(t) = \bigvee_{r \geq t} (\varphi_2(r) \wedge \bigwedge_{r' \in [t,r]} \varphi_1(r'))
\end{aligned}$$

We say that w satisfies φ from t if $\varphi(t) = 1$. We use $M(\varphi, w)$, or simply $M(\varphi)$ when w is clear from the context, to denote the time points from which φ is satisfied. For every φ, $M(\varphi)$ admits a *canonical representation* as a minimal set $\mathcal{I}$ of maximal intervals. The above semantics can be viewed as specifying recursive calls that descend the parse tree of φ down to the atomic propositions whose satisfaction signals are just the appropriate projections of w. Then, while climbing up, it combines the lower-level satisfaction signals until it gets to the top formula. The crucial procedures are those that compute the satisfaction signals of $F_{[a,b]}\varphi$ and $\varphi_1 \mathcal{U} \varphi_2$ from those of their sub-formulas.

Let $\varphi' = F_{[a,b]}\varphi$. The back-shifting method of [17, 22] for computing φ' from φ is based on the following simple concepts that generalize time shifts to the non-deterministic setting (more on these operations and their relation to determinism can be found in [18]).

[2] By a slight abuse of notation we use the same symbol for a formula and its satisfaction signal.

Definition 2 (Forward and Backward Cones, Back-Shift). *Let t be a time point and let $I = [c, d]$ be an interval.*

1. *The $[a, b]$-forward cone of t is the interval $[t + a, t + b]$;*
2. *The $[a, b]$-back cone of t is the interval $[t - b, t - a]$;*
3. *The $[a, b]$-back shift of interval I is $I' = \sigma_{[a,b]}(I) = [c - b, d - a]$.*

The forward cone consists of all time points r such that $\varphi(r)$ may influence $\varphi'(t)$. The back cone specifies the points r such that $\varphi'(r)$ can be influenced by $\varphi(t)$, those that t is in their forward cone. The back-shift of I is the union of the back cones of its elements, the set of all time points t such that $\varphi'(t)$ is influenced by $\varphi(r)$, $r \in I$. The following observation underlies the monitoring procedure of [17, 22].

Observation 1 (Back Shifting) *The back-shift of interval I consists of all points whose forward cone intersects I: $\sigma_{[a,b]}(I) = \{t : [t + a, t + b] \cap I \neq \emptyset\}$.*

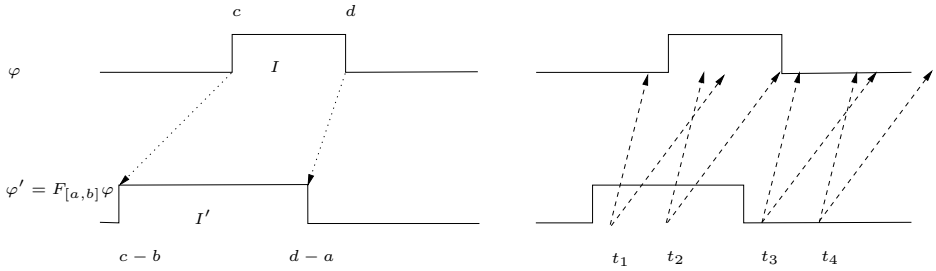

Fig. 1. A finite satisfaction signal φ which is true at interval $I = [c, d]$. (left): Back-shifting $I' = \sigma_{[a,b]}(I)$; (right): the forward cones of points $t_1, t_2 \in I'$ do indeed intersect I hence $\varphi' = F_{[a,b]}\varphi$ holds there. On the other hand, $\varphi = 0$ throughout all the forward cone of t_3 and hence $\varphi'(t_3) = 0$. The forward cone of t_4 does not intersect I either but part of it goes outside $dom(\varphi)$. We use in this paper a semantics where $\varphi'(t_4) = 0$.

The satisfaction signal of $\varphi' = F_{[a,b]}\varphi$ is thus computed by back-shifting all maximal intervals in $M(\varphi)$ and their union characterizes φ', see Figure 1. This procedure is obviously correct for points t such that $[t + a, t + b] \subseteq dom(\varphi)$. For other points like t_4 in the figure, the question is how to evaluate the disjunction (existential quantification) over the values of φ in that cone. A common approach, the one used implicitly in [22], is to consider $\varphi'(t) = 0$ if $\varphi(r) = 0$ for all $r \in [t + a, t + b] \cap dom(\varphi)$. We will use this semantics but our results can be easily adapted to an alternative 3-valued semantics where $\varphi(t) = \perp$ (unknown) if some possible completion of the signal lead to satisfaction and some others, to violation.

To illustrate the computation of $\varphi = \varphi_1 \mathcal{U} \varphi_2$ observe first that $(\varphi_1 \mathcal{U} \varphi_2)$ holds at t when φ_1 holds continuously between t and some future point r where φ_2 holds. This motivates the following operation between intervals $I_1 = [c, d)$ and $I_2 = [c', d')$:

$$\Omega(I_1, I_2) = \begin{cases} \emptyset & \text{if } d \leq c' \\ [c, d) & \text{if } c' < d \wedge d \leq d' \\ [c, d') & \text{if } c' < d \wedge d' < d \end{cases}$$

The three cases are illustrated Figure 2. The following observation justifies the computation of the set $M(\varphi)$ of positive intervals in the satisfaction signal of φ, by applying this operation to all pairs of maximal intervals in $M(\varphi_1)$ and $M(\varphi_2)$.

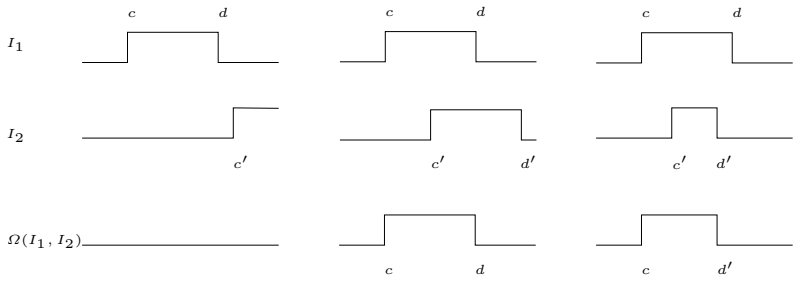

Fig. 2. Computing $\varphi_1 \mathcal{U} \varphi_2$ by computing $\Omega(I_1, I_2)$ for two maximal intervals. (a) φ_1 does not hold until φ_2; (b) it does but stops holding before φ_2 stops; (c) it does but φ_2 stops holding before and hence some parts of I_1 have no future where φ_2 holds.

Observation 2 *Let $M(\varphi_1)$ and $M(\varphi_2)$ be represented, respectively, by sets $\mathcal{I}_1$ and $\mathcal{I}_2$ of maximal intervals. Then*

$$M(\varphi) = \bigcup_{I_1 \in \mathcal{I}_1} \bigcup_{I_2 \in \mathcal{I}_2} \Omega(I_1, I_2).$$

The fact that $\mathcal{I}_1$ consists of maximal intervals is crucial here. If an interval $[c, d)$ satisfying $d' < d$ is split into non-maximal intervals $[c, e)$ and $[e, d)$ with $e < c'$, the points in $[c, e)$ will be wrongly considered as not satisfying $\varphi_1 \mathcal{U} \varphi_2$.

3 Satisfaction in Two Dimensions

The essence of our definition is to consider the end of the signal as an additional parameter t'. We would like to know what can be said about satisfaction at t after observing a prefix $w[0, t')$. Although this characterization of the pair (t, t') is different here from the matching property used for regular expressions in [28], we will use a similar terminology, partly because we have not yet found a simple name for this relation.

Let w be a Boolean signal defined over a bounded time domain $dom(w) = [0, \ell)$. Any sub-interval $[t, t')$ of $dom(w)$ defines a sub-segment of w that we denote by $w' = w[t, t')$. The set of all non-empty sub-segments of w can be represented by the triangle $T_w = \{(t, t') : 0 \leq t < t' < \ell\}$ (Figure 3-(a)).

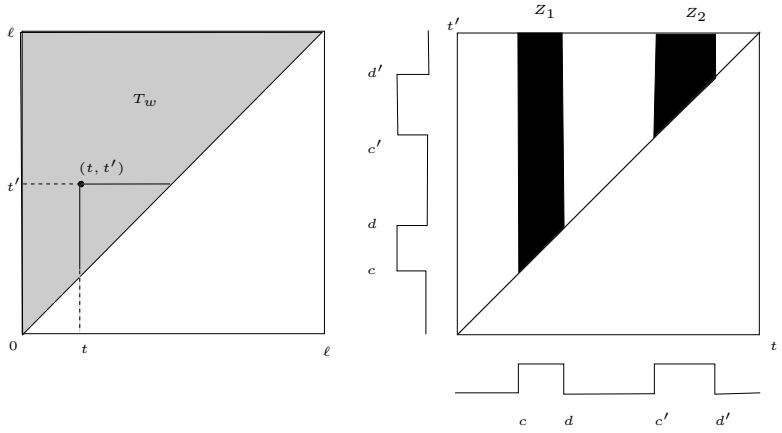

Fig. 3. (a) The triangle T_w associated with a signal defined over $[0, \ell)$. The length of the horizontal or vertical line from a point (t, t') to the diagonal is $t' - t$, the length of the segment $[t, t')$ that it represents; (b) A proposition which is true at the intervals $[c, d]$ and $[c', d']$. Its match set is $Z_1 \cup Z_2$.

Definition 3 (Matching and Match Sets). *A segment (t, t') of signal w matches an MTL formula φ, denoted as $(w, t, t') \models \varphi$, if $(w[0, t'), t) \models \varphi$. The match-set of φ in w is the set of all matching segments:*

$$\mathcal{M}(\varphi, w) = \{(t, t') : (w, t, t') \models \varphi\}.$$

The relation between matching in one and two dimensions can be expressed as

$$\mathcal{M}(\varphi, w) = \bigcup_{t' \in [0,\ell)} M(\varphi, w[0, t')) \times \{t'\} \tag{1}$$

We will use notation $\mathcal{M}(\varphi)$ when w is clear from the context. To compute $\mathcal{M}(\varphi)$ we first define the two-dimensional analog of satisfaction signal, the satisfaction map, where $\varphi(t, t')$ indicates the satisfaction status of φ by $w[t, t')$.

Definition 4 (MTL Matching Semantics with Satisfaction Maps). *The matching semantics of MTL formulas with respect to a Boolean signal w is defined inductively as follows:*

$$\begin{aligned}
&p(t, t') = w_p(t) \wedge t < t' < \ell \\
&(\neg\varphi)(t, t') = \neg(\varphi(t, t')) \\
&(\varphi \vee \psi)(t, t') = \varphi(t, t') \vee \psi(t, t') \\
&(F_{[a,b]}\varphi)(t, t') = \bigvee_{r \in [t+a, t+b]} \varphi(r, t') \\
&(\varphi_1 \mathcal{U} \varphi_2)(t, t') = \bigvee_{r \geq t} (\varphi_2(r, t') \wedge \bigwedge_{r' \in [t, r]} \varphi_1(r', t'))
\end{aligned}$$

In the sequel we will show that for a bounded-variability signal w, $\mathcal{M}(\varphi, w)$ can be expressed as a finite union of zones.

Definition 5 (Two-dimensional Zones). *A two-dimensional zone Z is a subset of $\mathbb{R}^2_+$ which is defined via a conjunction of orthogonal and difference inequalities of the following form*

$$\begin{array}{c} \underline{\alpha} \prec t \prec \overline{\alpha} \\ \underline{\beta} \prec t' \prec \overline{\beta} \\ \underline{\gamma} \prec t' - t \prec \overline{\gamma} \end{array} \tag{2}$$

where $\prec$ is either $<$ or $\leq$. The representation of a zone by the intervals $[\underline{\alpha}, \overline{\alpha}]$, $[\underline{\beta}, \overline{\beta}]$ and $[\underline{\gamma}, \overline{\gamma}]$ can be tightened and brought into a normal form where no inequality is implied by any combination of the others, except possibly in a marginal way.[3] We assume that we always work with such normalized zones where the constants satisfy the following constraints.

$$\begin{array}{c} \underline{\beta} - \overline{\gamma} \leq \underline{\alpha} \leq \overline{\alpha} \leq \overline{\beta} - \underline{\gamma} \\ \underline{\alpha} + \underline{\gamma} \leq \underline{\beta} \leq \overline{\beta} \leq \overline{\alpha} + \overline{\gamma} \\ \underline{\beta} - \overline{\alpha} \leq \underline{\gamma} \leq \overline{\gamma} \leq \overline{\beta} - \underline{\alpha} \end{array}$$

As convex sets, zones are not closed under union and complementation and these operations yield the class of sets that we will call timed polyhedra. Non-convex timed polyhedra can be expressed as a finite union of zones but this representation is not unique, and there is no canonical minimal representation as in the case of intervals. Moreover, the choice of the zones in the representation may affect the correctness of the procedure we propose in the sequel for the until operator. For this reason we define explicitly the notion of a representation of a timed polyhedron.

Definition 6 (Timed Polyhedra, Representation). *A timed polyhedron* $\mathbf{Z}$ *is a set expressible as a Boolean combination of orthogonal and difference constraints as in (2). A set of zones $\mathcal{Z} = \{Z_1, \ldots, Z_k\}$ is a representation of* $\mathbf{Z}$ *if*

$$\mathbf{Z} = \bigcup_i Z_i$$

To characterize match sets we will first show that those of propositions are timed polyhedra and that the latter are closed under the operations in Definition 4. This is trivial for disjunction and negation and we focus on $F_{[a,b]}$ and $\mathcal{U}$ for which we provide constructive proofs.

When a proposition p holds throughout an interval $I = [c, d]$, $\mathcal{M}(p)$ contains segments of w whose starting point t is in I. The role of t' is just to ensure, in addition, that $[t, t')$ is a well-defined segment, a sub-interval of $dom(w)$. This can be written explicitly as $(c \leq t \leq d) \wedge (t < t' < \ell)$, or using a zone-like notation, $(c \leq t \leq d) \wedge (0 < t' - t) \wedge (0 \leq t' < \ell)$. This will hold for any p-interval and consequently $\mathcal{M}(p)$ is a timed polyhedron (see Figure 3-(b)). The concepts of cones and back-shifting can be adapted to points and zones in $\mathbb{R}^2_+$. Recall that the satisfaction of $F_{[a,b]}\varphi$ at t is a function of the satisfaction of φ throughout $[t + a, t + b]$. The role of t' is to determine whether parts of the forward cone go outside $dom(\varphi)$ and should not be considered. Figure 4 illustrates the effect of t and t' on satisfaction.

[3] It means that if a constraint $f(t, t') \leq c$ is implied by other constraints, the constraint $f(t, t') \leq c - \varepsilon$ is not implied by them for any $\varepsilon > 0$.

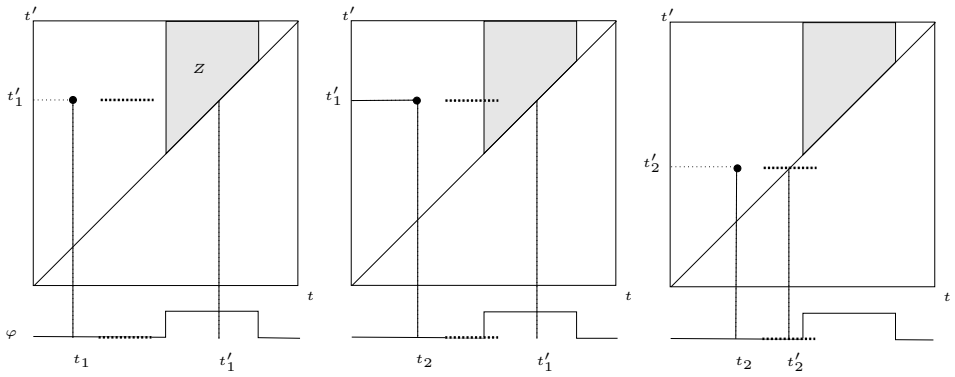

Fig. 4. The effect of t and t' on satisfaction of $\varphi' = F_{[a,b]}\varphi$ with respect to a given satisfaction signal φ whose match-set is the zone Z. The thick dashed lines indicate the forward cones of the respective values of t. (a) segment (t_1, t'_1) does not satisfy φ' because the forward cone of t_1 does not intersect Z; (b) segment (t_2, t'_1) which starts later does satisfy φ' because the forward cone of (t_2, t'_1) intersects Z. In other words $[t_2+a, t_2+b]$ intersects φ^1 before t'_1; (c) segment (t_2, t'_2) which ends earlier than (t_2, t'_1) does not satisfy φ' because it ends before φ becomes true.

Definition 7 (Cones and Back-Shifting in the Plane). *Let (t, t') be a point in $\mathbb{R}^2_+$ and let Z be a zone represented by the following (normalized) inequalities:*

$$\begin{array}{c} \underline{\alpha} \le t \le \overline{\alpha} \\ \underline{\beta} \le t' \le \overline{\beta} \\ \underline{\gamma} \le t' - t \le \overline{\gamma} \end{array}$$

Then

1. *The $[a, b]$-forward cone of (t, t') is $[t + a, t + b] \times \{t'\}$;*
2. *The $[a, b]$-back cone of (t, t') is $[t - b, t - a] \times \{t'\}$;*
3. *The $[a, b]$-back shift of zone Z is $Z' = \sigma_{[a,b]}(I)$, a zone defined by*

$$\begin{array}{c} \underline{\alpha} - b \le t \le \overline{\alpha} - a \\ \underline{\beta} \le t' \le \overline{\beta} \\ \underline{\gamma} + a \le t' - t \le \overline{\gamma} + b \end{array} \tag{3}$$

Claim (Zone Back Shifting). The back-shift of a zone consists of all points whose forward cone intersects Z:

$$\sigma_{[a,b]}(Z) = \{(t, t') : [t + a, t + b] \times \{t'\} \cap Z \neq \emptyset\}$$

Proof. Given that Z is normalized, (3) is what you get by applying quantifier elimination to the following formula

$$\exists r \begin{pmatrix} a \le r \le b \\ \underline{\alpha} \le t + r \le \overline{\alpha} \\ \underline{\beta} \le t' \le \overline{\beta} \\ \underline{\gamma} \le t' - (t + r) \le \overline{\gamma} \end{pmatrix}$$

∎

Figure 5 illustrates zone back-shifting. Perhaps the simplest way to view it is to back-shift the vertices of Z along the horizontal t dimension. The left vertices are shifted by b and the right ones by a.

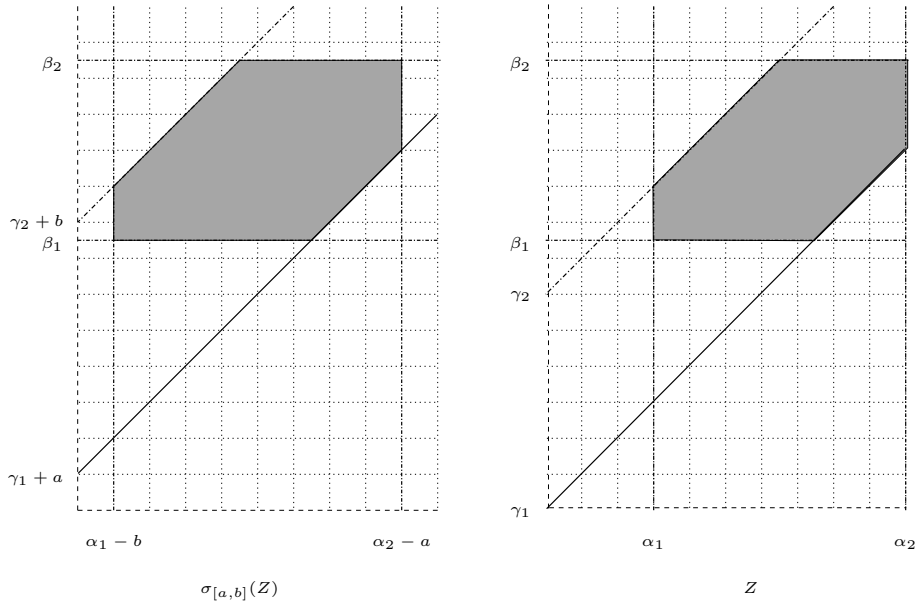

Fig. 5. An illustration of zone back-shifting.

What remains to be shown is that if both $\mathcal{M}(\varphi_1)$ and $\mathcal{M}(\varphi_2)$ are timed polyhedra, so is $\mathcal{M}(\varphi_1 \mathcal{U} \varphi_2)$. Recalling the relation between one-dimensional matching by intervals and two-dimensional matching as expressed in (1), we will associate with every $\mathbf{Z} \subseteq T_w$ and every t', a one-dimensional object, the t'-slice (projection) of $\mathbf{Z}$, defined as $I_{\mathbf{Z},t'} = \{t : (t,t') \in \mathbf{Z}\}$. For a convex zone Z, $I_{Z,t'}$ is a single interval and for this reason we first prove our result for the case where both match-sets are single zones.

Claim. Let $\mathcal{M}(\varphi_1) = Z_1$ and $\mathcal{M}(\varphi_2) = Z_2$ be zones, then $\mathcal{M}(\varphi_1 \mathcal{U} \varphi_2)$ is also a zone.

Proof. Following the semantics of until we have

$$\begin{array}{rll} (t,t') \in \mathcal{M}(\varphi_1 \mathcal{U} \varphi_2) & \text{iff} & \exists r \in [t,t'] \;\; (r,t') \in Z_2 \\ & \text{and} & \forall r' \in [t,r] \;\; (r',t') \in Z_1 \end{array}$$

which translates to

$$\exists r \in [t,t'] \left\{ \left\{ \begin{array}{ccc} \underline{\alpha}_2 \prec & r & \prec \overline{\alpha}_2 \\ \underline{\beta}_2 \prec & t' & \prec \overline{\beta}_2 \\ \underline{\gamma}_2 \prec & t'-r & \prec \overline{\gamma}_2 \end{array} \right\} \text{ and } \forall r' \in [t,r] \left\{ \begin{array}{ccc} \underline{\alpha}_1 \prec & r' & \prec \overline{\alpha}_1 \\ \underline{\beta}_1 \prec & t' & \prec \overline{\beta}_1 \\ \underline{\gamma}_1 \prec & t'-r' & \prec \overline{\gamma}_1 \end{array} \right\} \right\}$$

First we eliminate the universal quantification by taking the dual and applying the Fourier-Motzkin procedure and then eliminate the existential quantifier to finally obtain

$$\mathcal{M}(\varphi_1 \mathcal{U} \varphi_2) = \left\{ \begin{array}{ccccc} \underline{\alpha}_1 & \prec & t & \prec & \min\{\overline{\alpha}_1, \overline{\alpha}_2\} \\ \max \left\{ \begin{array}{l} \underline{\beta}_1, \underline{\beta}_2, \\ \underline{\alpha}_2 + \underline{\gamma}_1 \end{array} \right\} & \prec & t' & \prec & \min \left\{ \begin{array}{l} \overline{\beta}_1, \overline{\beta}_2, \\ \overline{\alpha}_1 + \overline{\gamma}_2 \end{array} \right\} \\ \max\{\underline{\gamma}_1, \underline{\gamma}_2\} & \prec & t' - t & \prec & \overline{\gamma}_1 \end{array} \right\} \tag{4}$$

Let us denote this operation on zones as $\mathbf{\Omega}(Z_1, Z_2)$. It can be viewed as performing in a symbolic manner an uncountable number of interval-based until computations:

$$\mathbf{\Omega}(Z_1, Z_2) = \bigcup_{t' \in [0,\ell)} \Omega(I_{Z_1,t'}, I_{Z_2,t'}) \times \{t'\}.$$

Consider now the more general case where $\mathcal{M}(\varphi_1)$ is a non-convex timed polyhedron $\mathbf{Z}$, represented as a set of zones $\mathcal{Z}$. Trying to apply the procedure to every pair of zones in the respective representations, we may face the following problem. There might be a maximal interval I in $I_{\mathbf{Z},t'}$ which is not fully included in a single zone in $\mathcal{Z}$ but is spread over two or more zones. This is illustrated in Figure 6-(a) with $\mathcal{Z} = \{Z_1, Z_2\}$, such that $I_{Z_1,t'} = I_1$ and $I_{Z_2,t'} = I_2$, both strict sub-intervals of I. On the other hand, adding to the representation the zone Z_3, shown in Figure 6-(b) will remedy this problem for I and for a bunch of other slices associate with a range of t' values. This motivates the following definition.

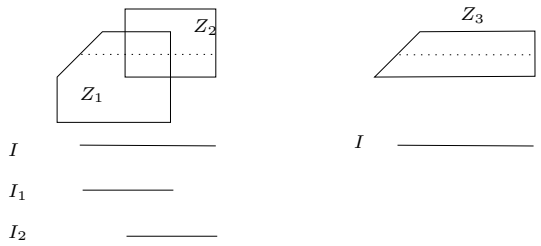

Fig. 6. (a) A representation $\mathcal{Z} = \{Z_1, Z_2\}$ of a timed polyhedron such that neither zone contains a maximal interval of a t'-slice. (b) Adding Z_3 to the representation fixes the problem.

Definition 8 (Maximal Zones, Maximal Normal Form). *Let* $\mathbf{Z}$ *be a timed polyhedron. A zone* $Z \subseteq \mathbf{Z}$ *is maximal in* $\mathbf{Z}$ *if there is no other zone* Z' *such that* $Z \subset Z' \subseteq \mathbf{Z}$. *A representation* $\mathcal{Z}$ *of* $\mathbf{Z}$ *is maximal if contains all maximal zones. A representation is reduced maximal if it consists of the set of all maximal zones.*

The notion of maximal representation is an adaptation of the concept of a *syllogistic form* of a Boolean function, a DNF representation that contains all maximal cubes. The reduced maximal representation corresponds to what is known as Blake normal form. Both were introduced in [4] and the reader can find more about it in [5]. In a maximal representation $\mathcal{Z}$ of $\mathbf{Z}$, every zone included in $\mathbf{Z}$ is included in some $Z \in \mathcal{Z}$, a fact which implies the following claim.

Claim (Pairwise Operation on Maximal Representation). Let $\mathcal{M}(\varphi_1) = \mathbf{Z}_1$ and $\mathcal{M}(\varphi_2) = \mathbf{Z}_2$ be timed polyhedra, represented by $\mathcal{Z}_1$ and $\mathcal{Z}_2$, respectively, with $\mathcal{Z}_1$ being maximal. Then $\mathcal{M}(\varphi_1 \mathcal{U} \varphi_2)$ is also a timed polyhedron computed as

$$\bigcup_{Z_1 \in \mathcal{Z}_1} \bigcup_{Z_2 \in \mathcal{Z}_2} \mathbf{\Omega}(Z_1, Z_2).$$

Proof. The inclusion

$$\mathcal{M}(\varphi_1 \mathcal{U} \varphi_2) \supseteq \bigcup_{Z_1, Z_2} \mathbf{\Omega}(Z_1, Z_2)$$

is trivial. Let us prove the opposite ($\subseteq$) inclusion. Consider any $(t, t') \in \mathcal{M}(\varphi_1 \mathcal{U} \varphi_2)$. By definition of the until satisfaction map, there exists $r \geq t$, such that

$$(r, t') \in \mathbf{Z}_2 \wedge \bigwedge_{r' \in [t,r]} ((r', t') \in \mathbf{Z}_1)$$

Applying the representation of $\mathbf{Z}_2$, we deduce from the first conjunct that $(r, t') \in Z_2$ for some zone $Z_2 \in \mathcal{Z}_2$. We rewrite the second conjunct as $I \subseteq \mathbf{Z}_1$, where I is the interval with extremities (t, t') and (r, t'). Using maximality of $\mathcal{Z}_1$, we deduce that $I \subseteq Z_1$ for some $Z_1 \in \mathcal{Z}_1$ or, in pointwise notation,

$$\bigwedge_{r' \in [t,r]} ((r', t') \in Z_1).$$

Gathering everything, we get that for some $r \geq t, Z_1 \in \mathcal{Z}_1, Z_2 \in \mathcal{Z}_2$, it holds that

$$(r, t') \in Z_2 \wedge \bigwedge_{r' \in [t,r]} ((r', t') \in Z_1),$$

in other words $(t, t') \in \mathbf{\Omega}(Z_1, Z_2)$. ⌟

This concludes the proof of our main result.

Theorem 1 (Match Sets for MTL). *For any MTL formula φ and a bounded variability Boolean signal w, $\mathcal{M}(\varphi, w)$ is a timed polyhedron represented as a finite union of zones.*

We sketch below how one can transform a representation of a timed polyhedron into a maximal one. We apply the multiplication (intersection) technique [4, 5] to complement a timed polyhedron $\mathbf{Z}$ represented by a set of zones. In essence it applies De Moragn rule to obtain a CNF representation of $\overline{\mathbf{Z}}$, and then opens the parentheses to

collect the terms (zones). The representation of $\overline{\mathbf{Z}}$ thus obtained is maximal by a direct extension of the results proved by Blake [5]. Applying this operation twice we obtain a maximal representation of $\mathbf{Z}$.

While a maximal representation is sufficient for proving the results, in our implementation we keep the representation reduced by making incremental inclusion tests and other optimization such as plane sweep techniques that may avoid intersections between zones that are far apart.

4 Case Study

In this section, we illustrate the computation of match sets on an example of a bounded recurrence property, taken from the catalog of commonly-used real-time properties [14]:

$$\varphi_1 \;:=\; (q \wedge \neg r \wedge Fr) \rightarrow (F_{[0,c]}(p \vee r)\, \mathcal{U}\, r)$$

Property φ_1 requires proposition p to hold at least every c time units between q and r. Such properties are commonly used to express periodic tasks to be performed between two events. Figure 7-(top left) depicts some input signals for propositions p, q, and r. The satisfaction maps for some sub-formulas are shown in Figure 7-(left) followed by the satisfaction map for the top-level formula φ_1 in Figure 7-(right). This figure illustrates the evolution of the formulas satisfaction with time and knowledge. Recall that a t'-section of the map gives us the satisfaction signal for the formula φ by $w[0, t')$. We depict t'-sections for time points t_1, t_2 and t_3 in Figure 7-(bottom-right). We can observe that φ_1 is satisfied at all times $t \in [0, t_1)$ based on the knowledge available at t_1. However, it turns out to be violated at some times $t \in [0, t_1)$ when additional knowledge about the input signals is provided at times t_2 and t_3.

Our techniques also open the way for using MTL for specifying local timed properties (patterns) that only hold at some segments of the signal. We illustrate this using three examples. First, consider a formula $\varphi_2 = q \wedge \varphi_1$ and its satisfaction map which filters away segments that satisfy φ_1 trivially due to $\neg q$. Second, we consider an MTL formula $\varphi_3 = GF_{[0,c]}(p \vee r)$ which describes time periods where p or r holds periodically at least every c time units. Third, in order to express a pattern describing time periods where p or r holds at least every c time units and starting with q, the formula φ_3 is intersected with q such that $\varphi_4 = q \wedge GF_{[0,c]}(p \vee r)$. Figure 8 depicts the satisfaction maps for φ_2, φ_3, and φ_4 using the same signals appearing in Figure 7.

5 Conclusions and Future Work

The major contribution of this work is in exporting and adapting the two-dimensional segment matching technology from timed regular expressions [28] to MTL. On the way to prove the main result, namely that the match sets for MTL are unions of zones, we had to cope with the alternating nature of the until operator, using the maximal representation for timed polyhedra. This concept, adapted from the syllogistic representation of Boolean functions, may have some other applications in the analysis of timed systems.

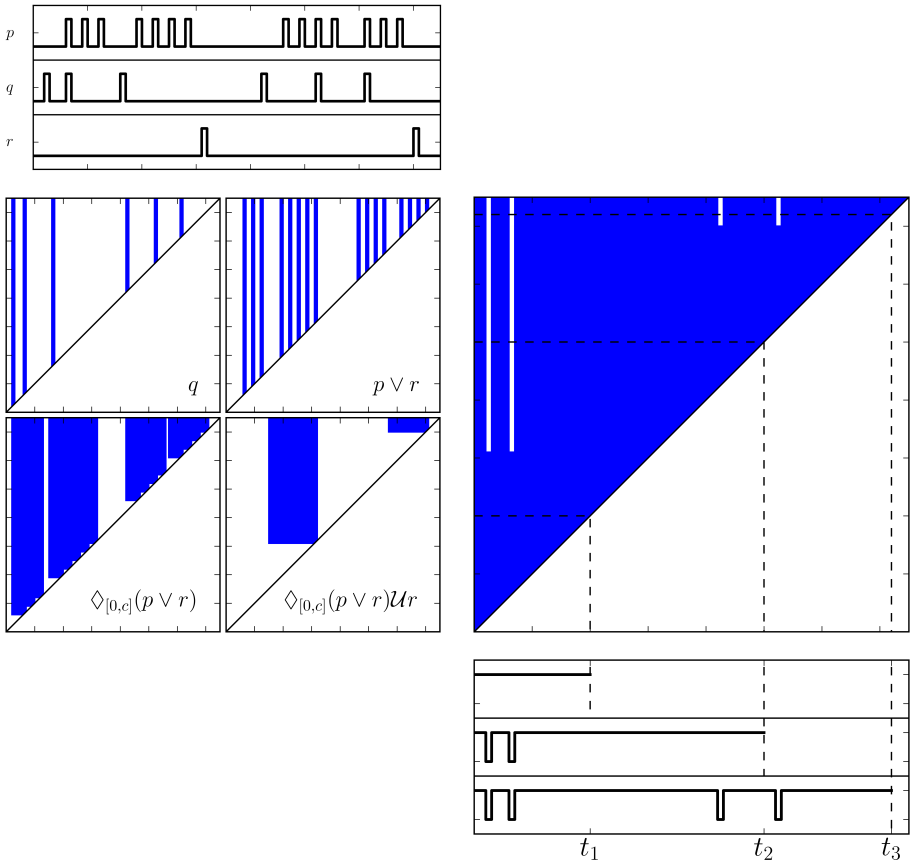

Fig. 7. Input signals p, q, and r, respectively (top left). Satisfaction maps for some subformulas (left). The satisfaction map for φ (right). Cross-sections of the satisfaction map for φ that corresponds to satisfaction signals at t_1, t_2, and t_3 (bottom right).

Our matching algorithm has been implemented and demonstrated on some non-trivial examples.

Regular expressions and temporal logic are inherently different due to various reasons, including the different nature of the major sequential operator (concatenation compared to *until*) and the positional and directed satisfaction relation of TL. Consequently, the MTL interpretation of the satisfaction map consists of separate positional and epistemological components. One way to go further in this direction is to consider a 3-dimensional satisfaction map defined on tuples (s, t, t') where $[t, t')$ stands for what is known about the signal and s is the position from which satisfaction is considered, not necessarily included in $[t, t')$. It looks a priori as if such an approach could handle full MTL with both future and past operators.

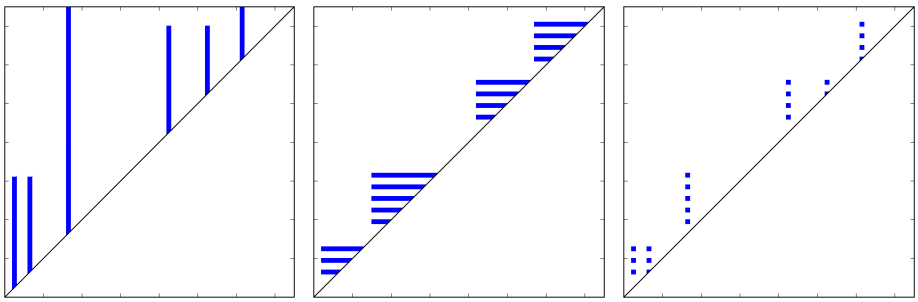

Fig. 8. Satisfaction maps for the formula $q \wedge \varphi_1$ (left), the formula $GF_{[0,c]}(p \vee r)$ (middle), and the formula $q \wedge GF_{[0,c]}(p \vee r)$ (right).

As mentioned, our technique can be adapted to a 3-values semantics with $\perp$ standing for unknown. To this end, the representation of the satisfaction map should be augmented with a second timed polyhedron $\mathcal{M}_\perp$, which should be shifted and manipulated in coordination with $\mathcal{M}$ and its complement.

Finally, the satisfaction of formulas in *interval temporal logics*, such as those studied in [11] and [32], is associated naturally with intervals. It might be the case that interval-based logics are more suited for defining patterns than point-based ones. We are currently working on the application of our techniques to handle metric extensions of such logics.

References

1. Yael Abarbanel, Ilan Beer, Leonid Gluhovsky, Sharon Keidar, and Yaron Wolfsthal. Focs–automatic generation of simulation checkers from formal specifications. In *CAV*, pages 538–542. Springer, 2000.
2. David A. Basin, Felix Klaedtke, and Eugen Zalinescu. Failure-aware runtime verification of distributed systems. In *FSTTCS*, pages 590–603, 2015.
3. Andreas Bauer, Martin Leucker, and Christian Schallhart. Monitoring of real-time properties. In *FSTTCS*, pages 260–272, 2006.
4. Archie Blake. *Canonical expressions in Boolean algebra*. PhD thesis, 1938.
5. Frank Markham Brown. *Boolean reasoning: the logic of Boolean equations*. Springer Science & Business Media, 2012.
6. Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *IJCAI*, pages 854–860, 2013.
7. Deepak D'Souza and Nicolas Tabareau. On timed automata with input-determined guards. In *FORMATS/FTRTFT*, pages 68–83, 2004.
8. Cindy Eisner, Dana Fisman, and John Havlicek. A topological characterization of weakness. In *PODC*, pages 1–8, 2005.
9. Cindy Eisner, Dana Fisman, John Havlicek, Yoad Lustig, Anthony McIsaac, and David Van Campenhout. Reasoning with temporal logic on truncated paths. In *CAV*, pages 27–39, 2003.

10. Dimitar P. Guelev, Catalin Dima, and Constantin Enea. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *Journal of Applied Non-Classical Logics*, 21(1):93–131, 2011.
11. Joseph Y Halpern and Yoav Shoham. A propositional modal logic of time intervals. *Journal of the ACM (JACM)*, 38(4):935–962, 1991.
12. Joseph Y Halpern and Moshe Y Vardi. The complexity of reasoning about knowledge and time. i. lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, 1989.
13. Hans Kamp. *Tense Logic and the Theory of Order*. PhD thesis, UCLA, 1968.
14. Sascha Konrad and Betty HC Cheng. Real-time specification patterns. In *ICSE*, pages 372–381, 2005.
15. Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.
16. Oded Maler. Some thoughts on runtime verification. In *International Conference on Runtime Verification*, pages 3–14, 2016.
17. Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In *FORMATS/FTRTFT*, pages 152–166, 2004.
18. Oded Maler, Dejan Nickovic, and Amir Pnueli. Real time temporal logic: Past, present, future. In *FORMATS*, pages 2–16, 2005.
19. Oded Maler, Dejan Nickovic, and Amir Pnueli. From MITL to timed automata. In *FORMATS*, pages 274–289, 2006.
20. Oded Maler, Dejan Nickovic, and Amir Pnueli. Checking temporal properties of discrete, timed and continuous behaviors. In *Pillars of Computer Science*, pages 475–505, 2008.
21. Zohar Manna and Amir Pnueli. The anchored version of the temporal framework. In *Workshop/School/Symposium of the REX Project*, pages 201–284. Springer, 1988.
22. Dejan Nickovic. *Checking timed and hybrid properties: Theory and applications*. PhD thesis, Université Joseph Fourier, Grenoble, France, 2008.
23. Maurice Nivat and Dominique Perrin. Ensembles reconnaissables de mots bi-infinis. In *STOC*, pages 47–59. ACM, 1982.
24. A. Pnueli. The temporal logic of programs. In *FOCS*, pages 46–57, 1977.
25. A. Pnueli. The Temporal Semantics of Concurrent Programs. *Theoretical Computer Science*, 13:45–60, 1981.
26. Arthur N Prior. *Past, present and future*, volume 154. Clarendon Press Oxford, 1967.
27. Thomas Reinbacher, Kristin Y. Rozier, and Johann Schumann. Temporal-logic based runtime observer pairs for system health management of real-time systems. In *TACAS*, pages 357–372, 2014.
28. Dogan Ulus, Thomas Ferrère, Eugene Asarin, and Oded Maler. Timed pattern matching. In *FORMATS*, pages 222–236, 2014.
29. Johan Van Benthem and Eric Pacuit. The tree of knowledge in action: Towards a common perspective. 2006.
30. Moshe Y. Vardi. From Church and Prior to PSL. In *25 Years of Model Checking*, volume 5000 of *Lecture Notes in Computer Science*, pages 150–171. Springer, 2008.
31. Moshe Y. Vardi and Pierre Wolper. An Automata-Theoretic Approach to Automatic Program Verification. In *LICS*, 1986.
32. Yde Venema. A modal logic for chopping intervals. *Journal of Logic and Computation*, 1(4):453–476, 1991.