



Specification-Based Monitoring of Cyber-Physical Systems: A Survey on Theory, Tools and Applications

Ezio Bartocci¹(✉), Jyotirmoy Deshmukh², Alexandre Donzé³,
Georgios Fainekos⁴, Oded Maler⁵, Dejan Ničković⁶,
and Sriram Sankaranarayanan⁷

¹ Technische Universität Wien, Vienna, Austria
`ezio.bartocci@tuwien.ac.at`

² University of Southern California, Los Angeles, USA

³ University of California at Berkeley, Berkeley, USA

⁴ Arizona State University, Tempe, USA

⁵ VERIMAG, CNRS and University of Grenoble-Alpes (UGA),
Saint Martin d'Hères, France

⁶ AIT Austrian Institute of Technology GmbH, Vienna, Austria

⁷ University of Colorado, Boulder, USA

Abstract. The term Cyber-Physical Systems (CPS) typically refers to engineered, physical and biological systems monitored and/or controlled by an embedded computational core. The behaviour of a CPS over time is generally characterised by the evolution of physical quantities, and discrete software and hardware states. In general, these can be mathematically modelled by the evolution of continuous state variables for the physical components interleaved with discrete events. Despite large effort and progress in the exhaustive verification of such hybrid systems, the complexity of CPS models limits formal verification of safety of their behaviour only to small instances. An alternative approach, closer to the practice of simulation and testing, is to monitor and to predict CPS behaviours at simulation-time or at runtime. In this chapter, we summarise the state-of-the-art techniques for qualitative and quantitative monitoring of CPS behaviours. We present an overview of some of the important applications and, finally, we describe the tools supporting CPS monitoring and compare their main features.

1 Introduction

Dynamic Behaviours and Their Evaluation. The world around us is in a constant flux with “things” changing dynamically. Planets move, temperatures rise and fall, rivers flow, rocks break down. In addition to these physical dynamics, a large part of the changing world is due to the activities of living systems and in particular humans, their social constructs and the artefacts they build. Houses are illuminated and air-conditioned, power is generated, distributed and consumed, cars drive on roads and highways, plants manufacture

materials and objects, commercial transactions are made and recorded in information systems. Airplanes fly, continuously changing location and velocity while their controllers deal, via sensors and actuators, with various state variables in the engine and wings. Conceptually those processes can be viewed as temporal behaviours, waveforms or signals or time series or sequences, where continuous and discrete variables change their values over time and various types of events occur along the time axis.

The systems that generate these behaviours are evaluated to some extent as good or bad, efficient or worthless, excellent or catastrophic. Such an evaluation can apply to the system in question as a whole, to some of its components or to a particular period of operation. We use *monitoring* to denote the act of observing and evaluating such temporal behaviours. Behaviours can be very long, spanning over a large stretches of time, densely populated with observations. They can also be very wide, recording many variables and event types. As such they carry too much information by themselves to be easily and directly evaluated. What should be distilled out of these behaviours should somehow be expressed and specified. The mathematical objects that do this job are functions that map complex and information-rich behaviours into low dimensional vectors of bits and/or numbers that indicate satisfaction of logical requirements and the values of various performance indices, see Fig. 1.

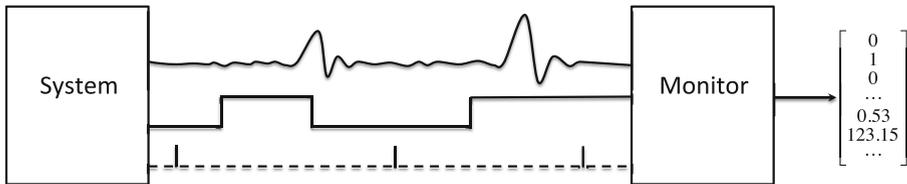


Fig. 1. Monitoring as reducing complex temporal behaviours into low-dimensional vectors of bits and numbers. The first of three behaviours is continuous and the two others are at the timed level of abstraction, state-based (signal) and event-based.

The evaluation can be based on a gross abstraction of the behaviour, for example the event of an airplane crash corresponds to a zero location on the z dimension and a large downward velocity at some point in time. Likewise the death of a patient can be specified by the stabilisation of his or her heart beat signal to a constant value. More often than not, those global catastrophic events may be related to (and preceded by) more detailed temporal behaviours that involve intermediate steps and variables, for example, some rise in the engine temperature which is not followed by certain actions such as turning on a cooling system. In less safety-critical contexts, systems are evaluated for performance, for example the time a client spends in a queue between requesting and being granted, or the energy consumption of a computer or a chemical plant along some segment of time.

Monitoring Real Systems and Monitoring Simulated Models. Before going further, let us distinguish between two major contexts in which the monitoring of dynamic behaviours can take place (see a more elaborate discussion in [95]). The first is the monitoring of real systems during their execution via online measurements. Here the role of monitoring is to alert in *real time* in order to trigger corrective actions, either by a human operator or by a supervisory layer of control. A primitive form of this type of monitoring exists in many domains: indicators on the control panel of a car, airplane or electronic device, monitors for physiological conditions of patients in a hospital and SCADA (Supervisory Control and Data Acquisition) systems for controlling complex large-scale systems such as airports, railways or industrial plants. In fact, any information system can be viewed as performing some kind of a monitoring activity.

The other context is during model-based system design and development where all or some of the system components do not exist yet in flesh and blood and their models, as well as the model of the environment they are supposed to interact with, exist as virtual objects of mathematical and computational nature. The design process of such systems is typically accompanied by an extensive simulation and verification campaign where the response of the system to numerous scenarios is simulated and evaluated. Most of the work described in this chapter originates from the design-time monitoring context, where simulation traces constitute the input of the monitoring process. Many techniques and considerations are shared, nevertheless, with the monitoring of real systems.

The activity of simulating a system and checking its behaviour is part of the verification and validation process whose goal is to ensure, as much as possible, that the system behaves as expected and to avoid unpleasant surprises after its deployment. In some restricted contexts of simple programs or digital circuits, this process can be made exhaustive and “formal” in the sense that all possible classes of scenarios are covered. When dealing with cyber-physical systems, whose existence and interaction scope are not confined to the world inside a computer for which practically exact models exist, complete formal verification is impossible, if not meaningless. In this domain, simulation-based lightweight verification is the common practice, accompanied by the hope of providing a good finite coverage of the infinite space of behaviours.

Rigorous Specification Formalisms. Part of the runtime verification movement is coming from formal verification circles, attempting to export to the simulation-based verification domain another ingredient of formal verification, namely, the rigorous specification of the system requirements. In the context of discrete systems, software or digital hardware, formalisms such as *temporal logic* or *regular expressions* are commonly used. They can specify in a declarative manner which system behaviours, that is, sequences of states and events, conform with the intention of the designer in terms of system functionality, and which of these behaviours do not. Such specifications can be effectively translated into monitoring programs that observe behaviours and check whether the requirements are satisfied. As such they can replace or complement tedious manual inspection of simulation traces or ad hoc programming of property testers.

Let us give some intuitive illustrations of the nature of these formalisms. Linear-time temporal logic (LTL) provides a compact language for speaking of sequences and the relations between their values at different points in time. The semantics of an LTL formula φ is time dependent with $(w, t) \models \varphi$ indicating that formula φ holds for sequence w at position t . The simplest formulas are state formulas which are satisfied at t according to the value of the sequence at t . That is, writing p for the truth value of a logical variable, or $x > 0$ for a numerical variable, is interpreted at each t as $p[t]$ and $x[t] > 0$, respectively.

More complex formulas are built using Boolean and temporal operators. The latter are divided into two types, future and past operators. The satisfaction of a future operator at position t depends on the values of the sequence at some or all the positions from t onward, that is, the suffix of w from t to $|w|$. For example $\Box p$ (*always p*) is true at any t such that p holds at every $t' \geq t$. The analogous past formula $\Box p$ (*historically p*) holds at t if p holds at any position $t' \leq t$, in other words, along the prefix of w from 0 to t . The satisfaction of a future formula by the whole sequence w is defined as its satisfaction at position 0 while that of a past formula, by its satisfaction at $|w|$. Past formulas have some advantages such as causality, while future LTL is more commonly used and is considered by some to be more intuitive.

The formula $\Box p$ quantifies universally over all time instances. The dual formula $\Diamond p$ (*eventually p*) quantifies existentially. It holds at t if p holds at *some* t' in the future. The weakness of such a property from a practical standpoint is that there is no bound on the distance between t and t' , a fact that may upset some impatient clients waiting for a response during their lifetime. We should note that in verification, formulas are often interpreted over *infinite* sequences generated by automata, while in monitoring we deal with finite sequences and if p never becomes true until the end of w , formula $\Diamond p$ is falsified.

A more quantitative alternative to \Diamond can be expressed in discrete time using the *next* operator. Formula $\bigcirc p$ (*next p*) holds at t if p holds at $t + 1$. Thus, the requirement that each p is followed by q within 2 to 3 time steps is captured by the formula

$$\Box (p \rightarrow (\bigcirc (\bigcirc q)) \vee (\bigcirc (\bigcirc (\bigcirc q)))).$$

This formulation may become cumbersome for large delay constants and by extending the syntax we can write this formula as

$$\Box (p \rightarrow \Diamond_{[2,3]} q)$$

with $\Diamond_{[a,b]} p$ being satisfied at t if p is satisfied at some $t' \in [t + a, t + b]$. In discrete time, this can be viewed as a syntactic sugar, but in dense time where *next* is anyway meaningless, this construct allows events to occur anywhere in an interval, not necessarily at sampling points or clock ticks.

Sequential composition is realised in future LTL using the *until* operator. The formula $p \mathcal{U} q$ (*p until q*) is satisfied at t if q occurs at some later point in time while p holds continuously until then. Using this operator, for which

◇ and □ are degenerate cases, one can require that some process should not start as long as another process has not terminated. The semantics of *until* is defined below.¹

$$(w, t) \models \varphi_1 \mathcal{U} \varphi_2 \quad \text{iff} \quad \exists t' \geq t ((w, t') \models \varphi_2 \wedge \forall t'' \in [t, t'] (w, t'') \models \varphi_1) \quad (1)$$

The past counter-part of *until* is the *since* operator with $q \mathcal{S} p$ (q since p) meaning that p occurred in the past and q has been holding continuously since then. The semantics of *since* is given below.

$$(w, t) \models \varphi_2 \mathcal{S} \varphi_1 \quad \text{iff} \quad \exists t' \leq t ((w, t') \models \varphi_1 \wedge \forall t'' \in [t', t] (w, t'') \models \varphi_2) \quad (2)$$

It is interesting to compare these operators with the *concatenation* operation used in regular expressions.

Regular expressions constitute a fundamental and popular formalism in computer science, conceived initially to express the dynamic behaviour of neural networks, and later applied to lexical and grammatical analysis. Traditionally, such expressions are defined over a monolithic alphabet of symbols but in order to present them in the same style as LTL, we will use product alphabets such as $\{0, 1\}^n$, defined and accessed via variables. Thus an expression p in our approach would be interpreted in the traditional approach as the set of all Boolean vectors in a global alphabet in which the entry corresponding to p is 1.

In discrete time, p is satisfied by any sequence of length one in which p holds. Sequential composition is realised by the concatenation operation where $\varphi_1 \cdot \varphi_2$ is satisfied by any sequence w that admits a factorisation $w = w_1 \cdot w_2$ such that w_1 satisfies φ_1 and w_2 satisfies φ_2 . This is best illustrated by defining the semantics using the satisfaction relation $(w, t, t') \models \varphi$ which holds whenever the subsequence of w starting at t and ending in t' satisfies expression φ :

$$(w, t, t') \models \varphi_1 \cdot \varphi_2 \quad \text{iff} \quad \exists t'' \in [t, t'] (w, t, t'') \models \varphi_1 \wedge (w, t'', t') \models \varphi_2 \quad (3)$$

The Kleene star allows to repeat concatenation for an indefinite but finite number times, with φ^* being satisfied by any sequence that admits a finite factorisation in which all factors satisfy φ . As an example, expression $(\neg p)^* \cdot q \cdot p$ specifies sequences in which a finite (possibly empty) time segment where p does not hold is followed by the occurrence of q followed by p .

Note that unlike LTL, regular expressions are more symmetric with respect to the arrow of time, as can be seen by the difference between their respective semantics definitions. The definitions of \models in (1) and (2) go recursively from t to the future or the past, respectively. When they come up from the recursion they do it in the opposite direction: for future LTL, satisfaction is computed backwards and that of past LTL is computed forward. For concatenation, in contrast, the semantics of \models in (3) is defined by a double recursion which takes the whole sequence and splits it into two parts which are the arguments for the two recursive calls. The semantics is collected from both ends while coming up from the recursion.

¹ Variants of *until* may differ on whether φ_2 is required to occur or whether φ_1 can cease to hold at the moment φ_2 starts or only after that.

Going Cyber-Physical. The exportation of these formalisms and their monitoring algorithms to the cyber-physical world has to cope with the hybrid nature of such systems. The dynamics of digital systems is captured by discrete event systems such as automata, generating discrete sequences of logical states and events. Physical systems are modelled using formalisms such as differential equations, producing behaviours viewed as continuous signals and trajectories. Specification formalisms and monitoring algorithms should then be extended so as to express and check temporal properties of such behaviours. This topic is the focus of the present chapter, centred around *Signal Temporal Logic* (STL), first presented in [96], along with a monitoring algorithm, further elaborated in the thesis [107] and explained from first principles in [98].

STL is a straightforward extension of (propositional) LTL along two orthogonal dimensions, namely, moving from discrete to dense time and using predicates on numerical values in addition to basic (atomic) propositions. The first feature is present in real-time variants of temporal logic such as MTL/MITL while the second has been explored in various first-order extensions of discrete-time LTL. We believe that some of the popularity of STL comes from the smooth and simple integration of these two features. This popularity, as attested by numerous publications that apply it to application domains ranging from analog circuits, via robotics, control systems and engineering education, down to biomedical and biochemical domains, justifies the role STL plays in this chapter, although in principle other variants of logic could do the job as well.

In order to be relevant to real applications, we should keep in mind that continuous dynamical systems are the object of study of various branches of mathematics and engineering, in particular, control and signal processing. These domains have developed over the years a variety of ways to measure and evaluate such systems and their behaviours, which are appropriate to their physical and mathematical nature. There is a variety of mathematical norms that reduce such behaviours into single numbers. There are transformations like Fourier's that extract the spectral properties of signals for the purpose of classification or noise removal. There are many statistical ways to assess signals and time series and detect occurring patterns. The challenge in monitoring cyber-physical systems is to integrate these traditional performance measures with those provided by the newly developed verification-inspired formalisms which are more suitable for capturing sequential aspects of behaviours.

2 Specification Languages

In this section, we present Signal Temporal Logic (STL) [96] as the specification language that we use in this document for expressing properties of CPS. We introduce the syntax of the formalism, together with its qualitative and quantitative semantics.

2.1 Signal Temporal Logic

STL [96] extends the continuous-time Metric Temporal Logic (MTL) [87] with numerical predicates over real-valued variables. In particular, STL enables reasoning about real-time properties at the interface between components that exhibit both discrete and continuous dynamics.

We denote by X and P finite sets of *real* and *propositional* variables. We let $w : \mathbb{T} \rightarrow \mathbb{R}^m \times \mathbb{B}^n$ be a multi-dimensional *signal*, where $\mathbb{T} = [0, d] \subseteq \mathbb{R}$, $m = |X|$ and $n = |P|$. Given a variable $v \in X \cup P$ we denote by $\pi_v(w)$ the *projection* of w on its component v .

We now define the variant of STL that contains both *past* and *future* temporal operators. The syntax of an STL formula φ over $X \cup P$ is defined by the grammar

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

where $p \in P$, $x \in X$, $\sim \in \{<, \leq\}$, $c \in \mathbb{Q}$, and $I \subseteq \mathbb{R}^+$ is an interval. We define the semantics of STL as the *satisfiability relation* $(w, t) \models \varphi$, indicating that the signal w satisfies φ at the time point t , according to the following definition. Given that we interpret the logic only over the finite traces, we let the satisfaction relation to be defined only for $t \in \mathbb{T}$.

$$\begin{aligned} (w, t) \models p &\leftrightarrow \pi_p(w)[t] = \text{true} \\ (w, t) \models x \sim c &\leftrightarrow \pi_x(w)[t] \sim c \\ (w, t) \models \neg\varphi &\leftrightarrow (w, t) \not\models \varphi \\ (w, t) \models \varphi_1 \vee \varphi_2 &\leftrightarrow (w, t) \models \varphi_1 \text{ or } (w, t) \models \varphi_2 \\ (w, t) \models \varphi_1 \mathcal{U}_I \varphi_2 &\leftrightarrow \exists t' \in (t + I) \cap \mathbb{T} : (w, t') \models \varphi_2 \text{ and } \forall t'' \in (t, t') (w, t'') \models \varphi_1 \\ (w, t) \models \varphi_1 \mathcal{S}_I \varphi_2 &\leftrightarrow \exists t' \in (t - I) \cap \mathbb{T} : (w, t') \models \varphi_2 \text{ and } \forall t'' \in (t', t) (w, t'') \models \varphi_1 \end{aligned}$$

We say that a signal w *satisfies* a STL formula φ , denoted by $w \models \varphi$, iff $(w, 0) \models \varphi$. In the remainder of this section, we discuss several specific aspects of the STL syntax and semantics.

Finitary interpretation. Specification formalisms with future temporal operators are typically defined over *infinite* behaviours. In particular, we can have a specification that is satisfied at time t iff a future obligation is fulfilled at some future time instant $t' > t$. Consequently, observing a finite prefix of a behaviour may not be sufficient to determine the satisfaction or the violation of a temporal specification according to its standard semantics. The finitary interpretation of future temporal logics is a well-studied problem in the monitoring research field. In order to tackle this issue, we adapt the semantics of \mathcal{U} and restrict the existential quantification of time to the (possibly bounded) signal domain. This altered semantics provides a natural interpretation of STL over finite signals. In particular, the *eventually* operator has a so-called *strong* interpretation – $\diamond \varphi$ is satisfied iff φ holds at any time before the signals ends. Similarly, the *always* operator has a *weak* interpretation under this semantics – $\square \varphi$ is satisfied iff φ is not violated during the signal duration.

The problem of interpreting temporal logic over finite or truncated behaviours was extensively studied in [57], where *weak*, *strong* and *neutral* views

of the finitary semantics for LTL are proposed. In [56], the authors provide a topological characterisation of weak and strong temporal operators. An extensive discussion about different interpretations of temporal logics over finite traces is presented in [98]. Finally, we also mention the real-time monitoring framework from [29], where 3-valued $\{\text{true}, \text{false}, \text{inconclusive}\}$ semantics are used to provide a finitary interpretation of a real-time temporal logic.

Strict interpretation of temporal operators. We adopt the *strict* semantics of until and since as originally proposed in [8]. The strict interpretation of $\varphi\mathcal{U}\psi$ evaluated at time t requires that both ψ is satisfied at some t' *strictly greater* than t and that φ continuously holds in the interval (t, t') that *excludes* t . In other words, the satisfaction of $\varphi\mathcal{U}\psi$ at t depends only on the evaluation of φ and ψ at some times in the future of t . This is in contrast to the classical non-strict semantics of *until* and *since* in the discrete-time LTL [110], where $\varphi\mathcal{U}\psi$ is satisfied if ψ holds at t . Let us first denote by $\bar{\mathcal{U}}$ the non-strict until operator, and by \mathcal{U} its strict counterpart². We also recall that LTL contains the *next* operator \bigcirc in addition to $\bar{\mathcal{U}}$. We can show that in discrete time, a temporal logic with $\bar{\mathcal{U}}$ and \bigcirc such as LTL is equivalent to the logic that has \mathcal{U} only, by using the following rules.

$$\begin{aligned}\varphi_1\mathcal{U}\varphi_2 &\equiv \bigcirc(\varphi_1\bar{\mathcal{U}}\varphi_2) \\ \bigcirc\varphi &\equiv \text{false}\mathcal{U}\varphi \\ \varphi_1\bar{\mathcal{U}}\varphi_2 &\equiv \varphi_2 \vee (\varphi_1 \wedge (\varphi_1\mathcal{U}\varphi_2))\end{aligned}$$

In contrast to LTL, continuous-time temporal logics such as STL do not have the next operator. It turns out that the strict interpretation of the temporal operators strictly increases the expressiveness of the underlying logic in dense time, as it enables “forcing” the time to advance. The main practical consequence of strict interpretation of \mathcal{U} and \mathcal{S} is that it allows specification of instantaneous events in continuous time.

Derived operators. The syntactic definition of STL is minimal and includes only basic operators. We can derive other standard operators as follows:

$$\begin{aligned}\text{True constant: } \text{true} &\equiv p \vee \neg p \\ \text{False constant: } \text{false} &\equiv \neg \text{true} \\ \text{Conjunction: } \varphi_1 \wedge \varphi_2 &\equiv \neg(\neg\varphi_1 \vee \neg\varphi_2) \\ \text{Implication: } \varphi_1 \rightarrow \varphi_2 &\equiv \neg\varphi_1 \vee \varphi_2 \\ \text{Eventually: } \diamond_I\varphi &\equiv \text{true}\mathcal{U}_I\varphi \\ \text{Once: } \diamond_I\varphi &\equiv \text{true}\mathcal{S}_I\varphi \\ \text{Always: } \square_I\varphi &\equiv \neg\diamond_I\neg\varphi \\ \text{Historically: } \square_I\varphi &\equiv \neg\diamond_I\neg\varphi\end{aligned}$$

² We restrict our argument to the future operators for the sake of simplicity – the same reasoning can be applied to the past operators.

In addition to these derived operators, we can also define instantaneous *events* that have zero duration. Such events enable specification of *rising* and *falling edges* in boolean signals.

$$\begin{aligned} \text{Rising edge: } \uparrow \varphi &\equiv (\varphi \wedge (\neg\varphi \mathcal{S} \text{true})) \vee (\neg\varphi \wedge (\varphi \mathcal{U} \text{true})) \\ \text{Falling edge: } \downarrow \varphi &\equiv (\neg\varphi \wedge (\varphi \mathcal{S} \text{true})) \vee (\varphi \wedge (\neg\varphi \mathcal{U} \text{true})) \end{aligned}$$

2.2 Signal Temporal Logic with Quantitative Semantics

In Sect. 2.1, we introduced STL with *qualitative* semantics. This classical definition of STL enables to determine the *correctness* of a signal with respect to a specification. Specifically, it gives a binary pass/fail answer to the monitoring problem. When reasoning about hybrid systems that involve both discrete and continuous dynamics, the qualitative verdict may not be informative enough. After all, systems with continuous dynamics are usually expected to admit some degree of tolerance with respect to initial conditions, system parameters and environmental perturbations. Consequently, a quantitative degree of satisfaction/violation would be preferable to a simple yes/no output given by the qualitative interpretation of STL.

Fages and Rizk [113] and Fainekos and Pappas [61] proposed to tackle this issue by equipping the temporal logic with *quantitative* semantics. This extension replaces the binary satisfaction relation with the quantitative *robustness degree* function, while preserving the original syntax of the specification language. In essence, the robustness degree function gives a real value that indicates how far is a signal from satisfying or violating a specification. We illustrate the concept of the robustness degree function with a simple example on numerical predicates. Let $x < c$ be a numerical predicate. This predicate partitions the \mathbb{R} domain into the set of all real values that are strictly smaller than c and those that are greater or equal to c . Picking a concrete value for x , the robustness degree gives the relative position of x to c , instead of only indicating whether x is above or below the threshold. This idea is naturally extended to the logical and temporal operators that we now formalise.

Let φ be an STL formula, w a signal and t a time instant in \mathbb{T} . We then define the robustness degree function $\rho(\varphi, w, t)$ as follows.

$$\begin{aligned} \rho(p, w, t) &= \begin{cases} \infty & \text{if } \pi_p(w)[t] = \text{true} \\ -\infty & \text{otherwise} \end{cases} \\ \rho(x \sim c, w, t) &= c - \pi_x(w)[t] \\ \rho(\neg\varphi, w, t) &= -\rho(\varphi, w, t) \\ \rho(\varphi_1 \vee \varphi_2, w, t) &= \max\{\rho(\varphi_1, w, t), \rho(\varphi_2, w, t)\} \\ \rho(\varphi_1 \mathcal{U}_I \varphi_2, w, t) &= \sup_{t' \in (t+I) \cap \mathbb{T}} \min\{\rho(\varphi_2, w, t'), \inf_{t'' \in (t, t')} \rho(\varphi_1, w, t'')\} \\ \rho(\varphi_1 \mathcal{S}_I \varphi_2, w, t) &= \sup_{t' \in (t-I) \cap \mathbb{T}} \min\{\rho(\varphi_2, w, t'), \inf_{t'' \in (t', t)} \rho(\varphi_1, w, t'')\} \end{aligned}$$

We note that for a fixed formula φ and a given signal w , the quantity $\rho(\varphi, w, t)$ is a function of time, and can thus be treated as a signal. We refer to it as the robust satisfaction signal or robustness signal.

There is a couple of fundamental properties that relate the STL quantitative semantics to its qualitative counterpart. Consider an arbitrary STL formula φ , a signal w and time $t \in \mathbb{T}$. The first property says that for any $\rho(\varphi, w, t) \neq 0$, its sign determines whether $(w, t) \models \varphi$. The second property states that if $(w, t) \models \varphi$, then for any signal w' whose pointwise distance from w is smaller than $\rho(\varphi, w, t)$ we also have $(w', t) \models \varphi$.

We illustrate the difference between the qualitative and the quantitative semantics in Fig. 2. We can also observe the relation between the two semantics as stated in the previous paragraph.

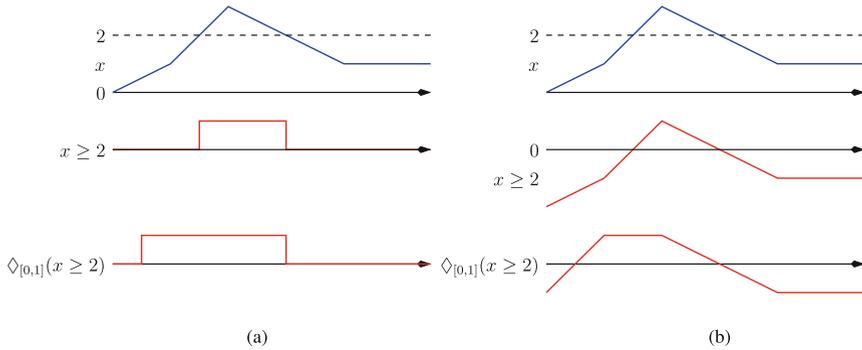


Fig. 2. Example for STL formula evaluation: (a) qualitative and (b) quantitative semantics.

Alternative quantitative semantics. In this section, we presented a quantitative semantics that allows measuring *spatial* robustness of STL specifications. This definition takes into account the spatial variations of signals when compared to STL specifications. Developing alternative notions of robustness degree for STL has been an active area of research in the recent years. In [51], the authors extend the quantitative semantics of STL by combining spatial with *time* robustness, thus also allowing to quantify temporal perturbations in signals. The idea of the combined space-time robustness for STL is further enhanced in [7] with *averaged* temporal operators. In [114], the authors identify that the bounded eventually $\diamond_{[a,b]}$ operator behaves like the *convolution* operator commonly used in filtering and digital signal processing. Following this surprising observation, one can develop various quantitative semantics for temporal logic, by defining the appropriate kernel window used for evaluating the formula. These additional operators enable reasoning not only about the worst-case but also the average-case behaviours. The Skorokhod metric provides an alternative way to measure mismatches between continuous signals in both space and time. An effective procedure for computing the Skorokhod distance between two behaviours is developed and presented in [42, 93]. This method is extended to estimate the Skorokhod distance between *reachpipes* in [94]. Nevertheless, there are no available methods yet to compute the Skorokhod distance between a signal and an STL formula.

3 Monitoring Algorithms

In this section, we present algorithms answering the following *monitoring* question: what is the qualitative and/or quantitative satisfaction of a formula φ by a signal w ? This problem is much easier than *model-checking*, i.e., proving that a system satisfies a formula, which is undecidable for STL even for simple classes of systems. Yet, it is desirable that efficient algorithms exist for monitoring, as this task can typically be repeated on large numbers of instances, or on signals of long durations. We consider two different settings: offline and online. In the offline setting, we assume that w is known before computing the satisfaction of φ . In the online setting, we assume only a partial knowledge of w , and compute successive estimates of φ satisfaction as new samples of w become available.

3.1 Offline Monitoring

For simplicity we restrict the presentation to the case of STL with future operators, and piecewise constant signals, i.e., we assume that w is completely defined by a sequence of time instants $t_0 < t_1 < \dots < t_i < \dots$ and values $w_0, w_1, \dots, w_i, \dots$ such that

$$\forall t \in [t_i, t_{i+1}), w[t] = w_i[t].$$

Here, we assume that the sequence of (time, value) pairs is finite, i.e., $i \leq N$. Moreover, we assume that the signal w holds its final value indefinitely, i.e., for all $t > t_N$, $w[t] = w[t_N]$. We present briefly an algorithm computing quantitative satisfaction of φ by w , adapted in a simpler form from [49]. For a purely Boolean monitoring algorithm, see [108]. Computing the quantitative satisfaction is not more complex than computing the Boolean, as both can be achieved in linear complexity in the size of signals.

The algorithm work by induction on the structure of the formula following the generic scheme presented in Algorithm 1.

Algorithm 1. Monitor(φ, w)

```

switch ( $\varphi$ )
  case  $p$ :
    return ComputeSatisfaction( $p, w$ )
  case  $x \sim c$ :
    return ComputeSatisfaction( $x \sim c, w$ )
  case  $* \varphi$ , where  $*$   $\in \{\neg, \diamond_I, \square_I\}$ :
     $w' := \text{Monitor}(\varphi, w)$ 
    return ComputeSatisfaction( $*$ ,  $w'$ )
  case  $\varphi * \psi$ , where  $*$   $\in \{\vee, \mathcal{U}_I\}$ :
     $w' := \text{Monitor}(\varphi, w)$ 
     $w'' := \text{Monitor}(\psi, w)$ 
    return ComputeSatisfaction( $*$ ,  $w', w''$ )
end switch

```

Table 1. ComputeSatisfaction for atomic predicate and Boolean operators. y is the signal returned, either Boolean or Real valued.

	Boolean	Quantitative
(p, w)	$y[t] = p$	if $p == \text{true}$, $y[t] = +\infty$, else $y[t] = -\infty$
$(x \sim c, w)$	$y[t] = (\pi_x(w)[t] \sim c)$	$y[t] = \pi_x(w)[t] - c$
(\neg, w')	$y[t] = \neg w[t]$	$y[t] = -w[t]$
(\vee, w', w'')	$y[t] = w'[t] \vee w''[t]$	$y[t] = \max(w'[t], w''[t])$

To implement Algorithm 1, we need to provide an implementation of the function ComputeSatisfaction for each instance of the switch statement. Instances which do not involve any temporal operator are straightforward and presented in Table 1. As can be expected, the only non-trivial case is with temporal operators.

In theory, we only need to handle the until operator (\mathcal{U}_I), as other temporal operators such as globally (\square_I) and eventually (\diamond_I) are derived from \mathcal{U}_I and Boolean operations. However, in practice, it is more efficient to deal with these operators separately. Also as we will see, \mathcal{U}_I is handled by combining specific algorithms for \square_I and \diamond_I and \mathcal{U}_I with $I = [0, \infty)$ (unbounded until). It turns out that for offline monitoring, unbounded operators are easier to handle than bounded time ones. Hence, we will present implementations for the unbounded cases before the bounded cases. In the following, we explain how to implement the ComputeSatisfaction function for unbounded eventually, timed eventually. The bounded globally case can be deduced from eventually using the equivalence $\square_I \varphi \Leftrightarrow \neg \diamond_I \neg \varphi$. Then we describe the algorithm for unbounded until and finally timed until operators.

ComputeSatisfaction($\diamond_{[0, \infty)}$, w) (*Unbounded Eventually*). From the quantitative semantics, recall that

$$\rho(\diamond_I(\varphi, w, t) = \sup_I \rho(\varphi, w, t).$$

If $I = [0, \infty)$, that means that

$$y[t] = \sup_{t' \geq t} \{w[t']\} = \max_{t_i \geq t} w[t_i]$$

It is easy to see that y is defined by a finite sequence satisfying the (backward) recurrence relation:

$$\begin{cases} y_N = w_N \\ y_k = \max(w_k, y_{k+1}) \end{cases}$$

Note that in the case of $I = [a, \infty)$, the result is obtained by shifting time by $-a$.

ComputeSatisfaction($\diamond_{[a, b)}$, w) (*Bounded Eventually*). From the quantitative semantics, we have that:

$$\rho(\diamond_{[a, b)}(\varphi, w, t) = \sup_{[a, b)} \rho(\varphi, w, t)$$

so that

$$y[t] = \max_{t_i \in [t+a, t+b)} \{w[t_i]\}.$$

In other words, the signal $t \rightarrow y[t]$ stores the sequence of maximums of signal w over a sliding finite window of size $b - a$. In [49], the authors observed that such a sliding window could be computed using an algorithm due to Daniel Lemire with linear complexity in the size of signal w [90]. We refer the reader to [49, 90] for details about this algorithm.

ComputeSatisfaction($\diamond_{[0, \infty)}$, w) (*Untimed Until*). To compute the satisfaction of untimed until, we are given two signals: w' and w'' and need to compute y . Since w' and w'' are of finite length N , y is also of finite length N . The computation goes backward starting from (y_N, t_N) . The time sequence t_0, t_1, \dots, t_{N-1} is obtained by merging and sorting the sequences t'_i and t''_i so that to simplify the notations, we assume that the sequences w'_i and w''_i are defined on the same time sequence, i.e., $w'_i = w'[t_i]$ and $w''_i = w''[t_i]$ for all i . Using standard min – max manipulations, we can show that the following recurrence is true:

$$\begin{cases} y_{N-1} = \min(w'_{N-1}, w''_{N-1}) \\ y_k = \max(\min(w'_k, w''_k), \min(w'_k, y_{k+1})), k \in \{0, \dots, N-2\} \end{cases}$$

Implementing this recurrence yields an algorithm with complexity in $O(2N)$.

ComputeSatisfaction($(\mathcal{U}_{[a, b)}, w', w'')$) (*Bounded Until*). To compute satisfaction signals for formulas involving timed operators, we make use of the following result:

Lemma 1. For two STL formulas φ, ψ ,

$$\varphi \mathcal{U}_{[a, b)} \psi \Leftrightarrow \diamond_{[a, b)} \psi \wedge \varphi \mathcal{U}_{[a, +\infty)} \psi \quad (4)$$

$$\varphi \mathcal{U}_{[a, +\infty)} \psi \Leftrightarrow \square_{[0, a)} (\varphi \mathcal{U} \psi) \quad (5)$$

Using this lemma, *ComputeSatisfaction*($\mathcal{U}_{[a, b)}, w', w''$) can be obtained by a sequence of intermediate computations using the algorithms above, as detailed in Algorithm 2. Note that all *ComputeSatisfaction* algorithm presented before run in linear time w.r.t. the number of samples in their input signals, so that for any STL formula, the computation of the robust satisfaction signal is also linear w.r.t. the size of the signals involved.

Algorithm 2. *ComputeSatisfaction*($\mathcal{U}_{[a, b)}, w', w''$) with $a < b < +\infty$

```

 $w_1 := \text{ComputeSatisfaction}(\mathcal{U}_{[0, +\infty)}, w', w'')$ 
 $w_2 := \text{ComputeSatisfaction}(\square_{[0, a)}, w_1)$  //  $w_2$  is the right hand side of (5)
 $w_3 := \text{ComputeSatisfaction}(\diamond_{[a, b)}, w'')$ 
return  $\text{ComputeSatisfaction}(\wedge, w_2, w_3)$ 

```

3.2 Online Monitoring

Offline algorithms assume that the entire trace is available to the monitoring procedure, and then run on the trace to produce either a Boolean satisfaction value or a quantitative (robust) satisfaction value. There are a number of situations where offline monitoring is unsuitable. Consider the case where the monitor is to be deployed in an actual system to detect erroneous behaviour. As embedded software is typically resource constrained, offline monitoring – which requires storing the entire observed trace – is impractical. Also, when a monitor is used in a simulation-based validation tool, a single simulation may run for several minutes or even hours. If we wish to monitor a safety property over the simulation, a better use of resources is to abort the simulation whenever a violation (or satisfaction) is conclusive from the observed trace prefix. Such situations demand an *online monitoring algorithm*, which has markedly different requirements. In particular, a good online monitoring algorithm must: (1) be able to generate intermediate estimates of property satisfaction based on *partial signals*, (2) use minimal amount of data storage, and (3) be able to run fast enough in a real-time setting.

A basic online algorithm returns a *true* or *false* satisfaction value when the satisfaction or violation of the property being monitored can be concluded by observing the finite trace prefix. In many cases, the information in the trace prefix is insufficient to produce a conclusive answer. Thus, several kinds of semantics have been proposed to interpret MTL or STL formulae over truncated traces. These semantics typically extend the satisfaction in a Boolean sense as used by offline monitoring algorithms to a richer satisfaction domain, typically having three or four values.

A typical three-valued semantics, for instance, assigns the satisfaction value of *true* if given the truncated signal, it can be decided that the signal will definitely satisfy the property. For example, if the value of the signal w at time 2 is 5, then any trace prefix inclusive of this time-point definitely satisfies the formula $\diamond_{[0,10]}(w > 0)$. Dually, it assigns the satisfaction value of *false* when the trace prefix is enough to establish that the signal definitely violates a given formula. If neither determination can be made, then the semantics assigns a satisfaction value of *unknown* [29]. Four-valued semantics that introduce *presumably true* and *presumably false* in addition to *true* and *false* for signal traces that are likely to respectively satisfy or violate the property have been proposed in [30].

In what follows, we first discuss qualitative monitoring algorithms; these algorithms, given an MTL or STL property, decide from a given prefix of a signal, if the entire signal would satisfy or violate the given property. An orthogonal, but relevant issue for online monitoring is the semantics of the MTL/STL property to be monitored on partial signal traces. We recall that following the tradition of the semantics for LTL on truncated traces [57], there are different notions of satisfaction that can be used to reason over prefixes of signals. We say that a signal-prefix *strongly* satisfies a given property if for *any* suffix the resulting signal would satisfy the property. In other words, the signal-prefix is sufficient to decide the satisfaction of the given property. We say that a signal-prefix *weakly* satisfies a

given property if there is *some* suffix such that the resulting signal would satisfy the property. A third notion of satisfaction is *neutral satisfaction*, which is if the given signal-prefix satisfies the property where the temporal operators are restricted to quantify only over the length of the signal-prefix. Some algorithms for qualitative monitoring make use of such richer notions of satisfaction.

Qualitative Online Monitoring. There are two main flavours of qualitative online monitoring. The first, based on work in [97], uses a modification of an algorithm similar to Algorithm 1. This procedure called *incremental marking*, essentially treats the signal as being available in chunks. The algorithm computes the robustness signal in a bottom-up fashion, starting from the leaves (i.e., atomic formulas) appearing in the STL formula and then for each super-formula, combining the robust satisfaction signals for its subformulae. The algorithm maintains the robust satisfaction signal partitioned as the concatenation of two signals: the first segment containing values that have already been propagated to the super-formula (by virtue of having sufficient information to allow deciding the satisfaction of the super-formula), and the second segment containing values that have not yet been propagated, as they may influence the satisfaction of the super-formula because of a part of the signal not yet available.

The second flavour of qualitative online monitoring makes use of automata-based monitors [71] and the richer notions of strong and weak satisfaction of MTL properties by signal-prefixes. In this approach, the given MTL formula (with future and past modalities) is first rewritten so that all temporal operators bound by finite time intervals appear in the scope of zero or more temporal operators that are unbounded. Essentially, each bounded temporal formula defines a finite time-window into the signal, where signal values would have to be available to evaluate the subformula. The algorithm maintains a time-window for each such bounded temporal subformula as a tableau and updates it using dynamic programming methods. This effectively allows treating bounded temporal subformulae as atomic propositions. The outer unbounded operators are monitored using two-way alternating Büchi automata that accept *informative prefixes* of the signal. An informative prefix is a signal-prefix that allows deciding the satisfaction or violation of the given unbounded temporal formula. The procedure has space-bounds that are linear in the variability of the given signal, and length of the formula, and requires constructing automata that are doubly exponential in the size of the given MTL formula.

Quantitative Online Monitoring. Qualitative monitoring algorithms by their nature are unable to quantify the degree to which the signal corresponding to the given signal-prefix may satisfy or violate the property of interest. Online algorithms for computing robust satisfaction semantics seek to address this gap. We discuss three algorithms below. The operating assumption of each algorithm is the same: given a trace prefix, return a quantitative value that captures a notion of robust satisfaction over the incomplete trace. However, none of the algorithms are probabilistic in nature. The first and the third algorithms

provide a quantitative value based only on the trace prefix observed (and possibly a forecast suffix). The second algorithm, given a formula φ , processes a trace prefix and returns a robust satisfaction interval that always contains the robust satisfaction value w.r.t. φ of any trace with this prefix.

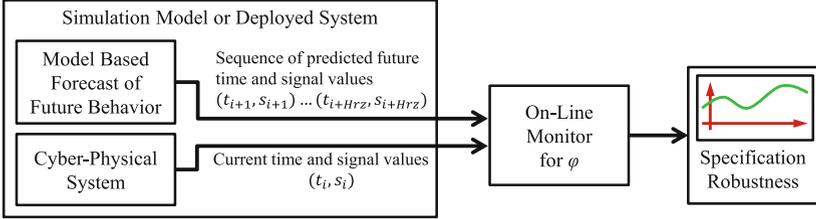


Fig. 3. The online temporal logic monitoring framework for bounded future and unbounded past formulas proposed in [44]. In case the monitored system contains a model which can be used for forecasting future behaviours up to a horizon Hrz , then these behaviours can be used to compute the robustness estimate of the specification φ . If a forecasting mechanism is not available, then it is preferable to monitor past formulas.

Quantitative monitoring of STL with past operators and predictors.

The first algorithm we discuss is for online monitoring of STL formulas with bounded future and unbounded past formulas [44]. For a given STL formula φ , the algorithm computes the horizon, or the number of look-ahead steps that would be required to evaluate the bounded future component of a formula, and the history or the number of samples in the past that would be needed to evaluate a given bounded past formula. The algorithm then maintains a tableau that is updated every time a new signal value becomes available using a dynamic programming based approach. For unbounded past operators, the algorithm exploits the fact that an unbounded past formula can be rewritten such that the computation over an unbounded history can be stored as a *summary* in a variable that is updated each time a new signal value becomes available. For the unobserved parts of the signal, which would be required to compute the satisfaction of a bounded future temporal subformula, the algorithm requires the use of a predictor in order to compute a robustness estimate (see Fig. 3).

In more detail, the algorithm monitors invariants expressed as STL formulas φ . That is, the assumption is that φ should be satisfied at all times of the system execution or model simulation. As an example, consider the specification

$$\varphi \equiv \neg(x \geq 0.5) \rightarrow \Diamond_{[0,1]} \Box_{[0,1]}(x \geq 0.5) \tag{6}$$

which states that if the value of the signal x decreases below 0.5, then within 1s in the past it should have been higher than 0.5 for at least 1s. In Fig. 4, we present a simple example of a signal x and the corresponding robustness value $\rho(\varphi, x, t)$ computed online for Eq. (6). Notice that the signal x decreases below the threshold 0.5 (i.e., the antecedent $\neg(x \geq 0.5)$ is now satisfied) at two points

in time, at 5.5s and at 12.25s. In the first case, no violation of the requirement occurs (barely), while in the second case a violation occurs (the robustness drops below zero). In either case, the robustness of Eq. (6) starts decreasing at time 6 and at time 13, respectively, when the subformula $\Diamond_{[0,1]}\Box_{[0,1]}(x \geq 0.5)$ starts changing robustness value. On the other hand, the robustness increases afterwards because the value of the signal x approaches the threshold 0.5 and, therefore, the antecedent comes closer again to being falsified. The fact that the specification robustness starts decreasing can be used as a warning that the system may be soon violating its requirement; and, thus, remedial action may be required.

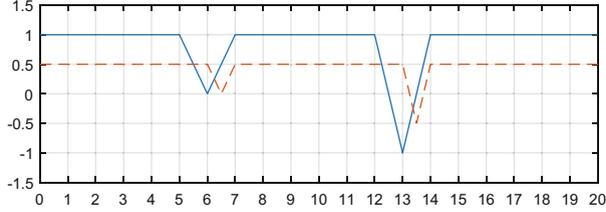


Fig. 4. A signal x (solid blue line) and the corresponding robustness value $\rho(\varphi, x, t)$ (dashed red line) computed for the formula in Eq. (6). (Color figure online)

Quantitative monitoring of robust satisfaction interval for bounded future STL formulas. The second algorithm computes *robust interval semantics* for STL formulas with bounded future formulas [40]. A robust satisfaction interval (ℓ, u) for a given signal-prefix is defined such that ℓ (respectively, u) is the infimum (respectively, supremum) over the robust satisfaction values of the given property for all signals that have the given partial signal as the prefix. For example, consider a constant signal-prefix with magnitude 1, defined over time $t \in [0, 5)$. The robust satisfaction interval for the formula $\Box_{[0,10]}(x > 0)$ over this signal-prefix is $(-\infty, 1]$, while the robust satisfaction interval for the formula $\Diamond_{[0,10]}(x > 0)$ is $[1, \infty)$. The interval semantics generalises the notion of strong and weak semantics. A signal-prefix strongly satisfies (resp. violates) a property if the lower-bound of the robust satisfaction interval is positive (resp. upper-bound of the robust satisfaction interval is negative). A signal-prefix weakly satisfies a property if the upper bound of the robust satisfaction interval is positive.

In lieu of a detailed exposition of the algorithm, we demonstrate it with an example in Fig. 6 for the formula in Eq. (7).

$$\varphi \equiv \Box_{[0,a]}(\neg(y > 0) \vee \Diamond_{[b,c]}(x > 0)) \quad (7)$$

We assume that the algorithm starts in a state where it has processed the partial signal till time t_2 , and show the effect of receiving data at time-points t_3 , t_4 and t_5 . The algorithm maintains lists of robust satisfaction intervals at each node in the syntax tree of φ , and updates these lists when new time-points in the trace become available. Each row in the table adjacent to a node shows the

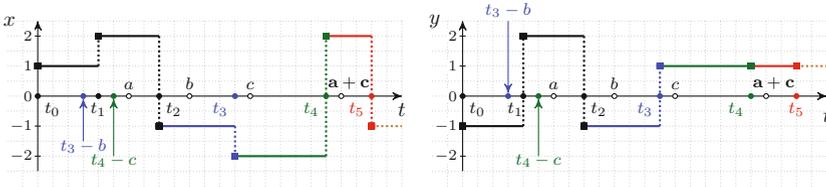


Fig. 5. These plots show the signals $x(t)$ and $y(t)$. Each signal begins at time $t_0 = 0$, and we consider three partial signals: till time t_3 (black + blue), and till time t_4 (black + blue + green), and till time t_5 (black + blue + green + red) (Color figure online)

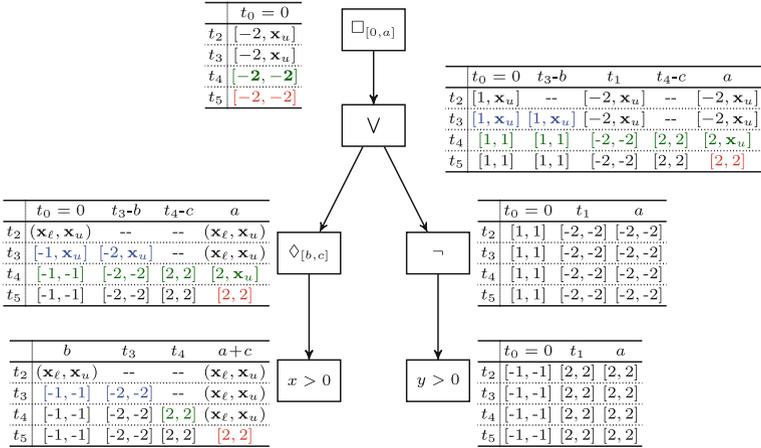


Fig. 6. We show a snapshot of the list at each node in the syntax tree of formula given in Eq. (7) maintained by the algorithm for four different (incremental) partial traces of the signals $x(t)$ and $y(t)$ shown in Fig. 5. Each row indicates the state of each list at the time indicated in the first column. An entry marked -- indicates that the corresponding element did not exist in the list at that time. Each colored entry indicates that the entry was affected by availability of a signal fragment of the corresponding color. (Color figure online)

state of the list after the algorithm processes the value at the time indicated in the first column.

The first row of the table shows the snapshot of the lists at time t_2 . Observe that in the lists for the subformula $y > 0$, $\neg y > 0$, because $a < b$, the data required to compute the interval at t_0 , t_1 and the time a , is available, and hence each of the intervals is singular. On the other hand, for the subformula $x > 0$, the time horizon is $[b, a + c]$, and no signal value is available at any time in this interval. Thus, at time t_2 , all elements of the list at the node corresponding to $x > 0$ are (x_ℓ, x_u) corresponding to the user-provided greatest lower bound and lowest upper bound on the signal $\mathbf{x} = (x, y)$.

To compute the values of $\diamond_{[b,c]}(x > 0)$ at any time t , we take the supremum over values from times $t + b$ to $t + c$. As the time horizon for the node corresponding to $\diamond_{[b,c]}(x > 0)$ is $[0, a]$, t ranges over $[0, a]$. In other words, we wish to perform the sliding maximum over the interval $[0 + b, a + c]$, with a window of length $c - b$. We can use the algorithm for computing the sliding window maximum as discussed in the earlier section on offline monitoring. One caveat is that we need to store separate monotonic edges for the upper and lower bounds of the robust satisfaction intervals. The algorithm then proceeds upward on the syntax tree, only updating the list of a node only when there is an update to the lists of its children.

The second row in each table is the effect of obtaining a new time point (at time t_3) for both signals. Note that this does not affect the lists at the node $y > 0$ or the node $\neg y > 0$, as the robust satisfaction intervals are already singular, but does update the intervals for the node $x > 0$. The algorithm then invokes the sliding window computation on the list for $x > 0$ to update the list for $\diamond_{[b,c]}(x > 0)$. Finally, we remark that the run of this algorithm shows that at time t_4 , the interval for the formula φ is $[-2, -2]$, which yields a negative upper bound, showing that the formula is not satisfied irrespective of the suffixes of x and y . In other words, the satisfaction of φ is known before we have all the data required to fully determine the robust satisfaction value for the formula.

Quantitative monitoring for unbounded horizon STL formulas. Robust satisfaction intervals are meaningless for STL formulas with unbounded future operators. Hence, for such formulas, nominal quantitative semantics have been proposed in [41]. These semantics essentially compute the robust satisfaction value of the formula restricted to the trace prefix that is available. However, a key challenge is to avoid storing the entire signal or performing repeated computation whenever a new time-point becomes available. This can be circumvented by storing the *summary* of computation of the robust satisfaction value and then incrementally updating the summary variables. The number of summary variables required is finite and independent of the trace length, although exponential in the length of the formula. Nominal quantitative satisfaction values for arbitrary unbounded horizon STL formula can be monitored in this fashion. The rules for rewriting the nominal robustness computation of an unbounded horizon STL formula in terms of summary variables is involved, and we omit the details for brevity. Interested readers can find them in [41].

4 Extensions

It is out of doubt that STL has gained in the last decade an increased popularity among engineers for its conciseness and expressive power enabling to specify complex behavioural properties related to the order and the temporal distance among discrete *events* such as the satisfaction of predicates (e.g., threshold crossing) over the real variables.

However, the pure time-domain nature of this specification language sometimes has revealed to be a technical impediment to overcome for an

immediate applicability to cyber-physical systems. In particular, abnormal signal behaviours such as undesirable oscillations and complex topological requirements (i.e., the spatial distribution of the entities generating signals) are very challenging to capture using only the time-domain.

For this reason, in the last decade STL has inspired a number of extensions that have been successfully applied in many applications ranging from identifying oscillatory behaviours in analog circuits [104], biological systems [32] and music melodies [52] to specifying spatiotemporal requirements in reaction-diffusion systems [12, 13, 16, 102], smart grids [69] and mobile ad-hoc sensor networks [15]. In the following we aim to provide an overview of some of the STL extensions recently proposed.

4.1 Monitoring Complex Oscillatory Properties

The dynamical behaviour of a physical system often exhibits complex oscillatory patterns representing an infinite periodically behaviour. Real-life analog signals are characterised by omnipresent noise, i.e., random perturbations of the desired signal. Damped oscillations or oscillations with increasing amplitude are peculiar aspects in many biological systems [21, 22, 32]. Abnormal oscillations due to the presence of *spikes* or *hunting* oscillations [104] are considered undesired behaviours in analog circuit design. Signal-processing tasks such as *peak* detection [5, 19, 34] is common in many medical cyber-physical systems whose correctness impacts the performance and the sensitivity of the computational devices involved. Providing a concise formal specification language expressive

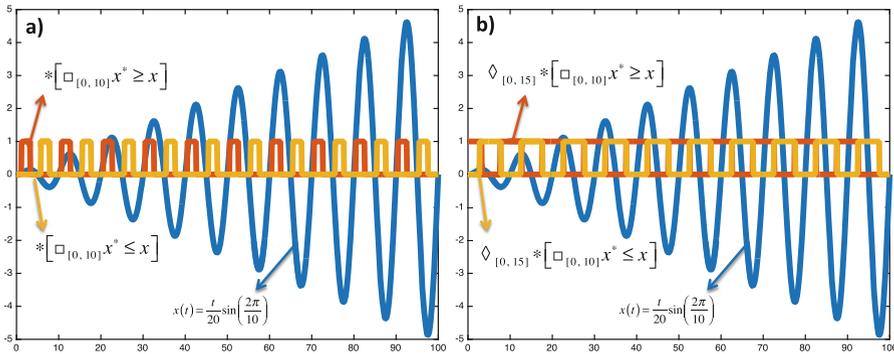


Fig. 7. The freezing operator of STL-* enables in (a) the specification of the notions of local maxima ($*[\square_{[0,10]}x^* \geq x]$) and local minima ($*[\square_{[0,10]}x^* \leq x]$) by comparing the value x with all the *frozen* values (represented by the expression x^*) within a time interval $[0, 10]$. This feature is suitable to capture oscillatory behaviours without knowing a-priori specific signal thresholds. In (b) freezing operators are nested within the classical *eventually* STL operators to detect the possible occurrence of a local maxima and a local minima within $[0, 15]$ time units.

enough to characterise such patterns and to efficiently monitor them is a very challenging task.

For example, the classic STL is not expressive enough to distinguish classes of oscillatory patterns such as damped oscillations or oscillations with increasing amplitude, because it is not able to globally reference and compare local properties (i.e., local minima/maxima) of a signal. Motivated by this necessity the authors in [32] have proposed an extension of STL (named STL-*), augmenting STL with a freezing operator that allows to record the signal values during the evaluation of a sub-property, and to reuse it for comparison in the other parts of the formula. This operator increases the expressive power of STL and for instance it enables to express and to capture various dynamic aspects of oscillations (see Fig. 7). A quantitative semantics for STL-* is also proposed in [33]. However, the price to pay for this enhanced expressiveness is an higher computational complexity for the monitoring procedure. In [32] the authors show that the monitoring algorithm for STL with nesting freeze quantifiers is polynomial in the number of intervals on which the signals is defined and the size of the syntactic tree of the formula, but it is exponential in the number of the nested freeze operators in the formula.

Although STL-* is more expressive than STL, its analysis is still limited to the time-domain representation of a signal. However, oscillatory patterns (i.e., chirp signals, hunting behaviours, noise filtering) are in general very challenging and tedious to investigate using only time and a time-frequency analysis is essential sometime to efficiently detect them. Time-frequency analysis is an important branch of *signal processing* and it is based on the study of the *spectrogram* (see Fig. 8), a representation of the frequencies' magnitude in a signal as they vary with time. Spectrograms can be calculated from digitally sampled data in the time-domain representation of the signal using extensions of the classic Fourier transforms such as the Short Time Fourier Transform (STFT) or Wavelet Transforms (WTs). Although a preliminary work on combining time and frequency domain specifications for periodic signals is reported in [37], the first attempt to provide a unified formalism to express time-frequency properties of a signal is the *time-frequency logic* (TFL) introduced in [52]. TFL extends STL with predicates evaluating the magnitude of a particular frequency range in a point in time. The semantics of TFL operates over a spectrogram generated using STFT. In [52] TFL was applied to detect musical patterns, but it can be easily used in other application domains. More recently, TFL was extended in [104] to operate over spectrograms generated using WTs. These spectrograms generally provide a better trade-off between the resolution in the time domain and the resolution in the frequency domain w.r.t. STFT.

4.2 Monitoring Spatio-Temporal Behaviours

The components in CPS are generally distributed across space and connected via a communication infrastructure. The complex behaviour of each individual component due to a fully-integrated hybridisation of computational (logical) and physical action and the interactions between these components via the network

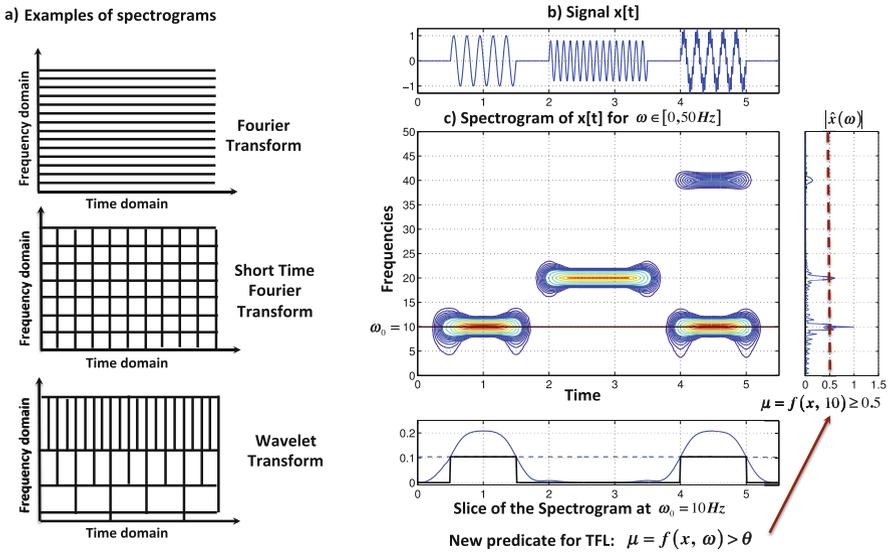


Fig. 8. In (a) we show three different examples of spectrogram obtained using different time to frequency domain signal transformations: Fourier Transform (FT), Short Time Fourier Transform (STFT) and Wavelet Transform (WT). FT provides the average magnitude of a signal for a certain interval of frequencies along the entire duration of the signal. FT are indeed not suitable to localise the time when a change of frequencies occurs. STFT consists instead in dividing a longer time signal into shorter segments of the same length and in computing the FT for each of these segments. This provides a better time resolution than FT. However, STFT has a fixed resolution that depends on the time length chosen for the shorter segments. A wide window gives better frequency resolution but poor time resolution and vice-versa. WT generally provides the best combination of good time resolution for high-frequency events and good frequency resolution for low-frequency events. In (b) we show an example of a signal with in (c) its STFT spectrogram for the frequencies in the range 0–50 Hz. On the right of the spectrogram we provide the projection of the magnitude associated with each frequency, while on the bottom we show a slice of the spectrogram for the frequency 10 Hz. TFL provides a special predicate that enables to compute such slicing and to compare the magnitude w.r.t. a user-defined threshold.

enable them to produce very rich and complicated emergent spatiotemporal behaviours, often impossible to predict at design time. Examples include smart grids, robotics teams or collections of genetically engineered living cells. In such examples, temporal logics may be not sufficient to capture also topological spatial requirements. For example, the notion of being *surrounded* or *spatial superpositioning* (averaging resources in a space) are not available in the standard STL and encoding them with specific functions may result cumbersome.

Recently, three different spatiotemporal extensions of STL, *Spatial-Temporal Logic* (SpaTeL) [69], the *Signal Spatio-Temporal Logic* (SSTL) [16, 102] and the *Spatio-Temporal Reach and Escape Logic* (STREL) [15] have been proposed to

accommodate the growing need of expressing not just temporal but also spatiotemporal requirements in CPS.

SpaTeL [69] is the unification of STL and Tree-Spatial-Superposition-Logic (TSSL) introduced in [12,13] to classify and detect spatial patterns. TSSL reasons over quadtrees, spatial data structures that are constructed by recursively partitioning the space into uniform quadrants. TSSL uses the notion of spatial superposition (introduced in [68]) that provides a way to describe statistically the distribution of discrete states in a particular partition of the space and that enable to specify self-similar and fractal-like structures that generally characterise the patterns emerging in nature. SpaTeL is equipped with both a qualitative and a quantitative semantics that provide a measure or robustness of how much the property is satisfied or violated. In [69] this measure of robustness is used as a fitness function to guide the parameter synthesis process for the neighbourhood prices in a demand-side management system model of a smart grid using particle swarm optimisation (PSO) algorithms.

SSTL [16,102] instead extends STL with three spatial modalities *somewhere*, *everywhere* and *surround*, which can be nested arbitrarily with the original STL temporal operator. SSTL is interpreted over a discrete model of space represented as a finite undirected graph. Each edge of the graph is labeled with a positive weight that can be used to represent the distance between two nodes. This provides a metric structure to the space, in terms of shortest path distances. However, the weight can be used to encode also other kind of information (i.e., the average travelling time between two cities). In [102], the authors provide a qualitative and quantitative semantics of SSTL and efficient monitoring algorithms for both semantics.

STREL [15] generalises SSTL by introducing two new spatial operators: *reach* and *escape*. These operators enable to express the same spatial modalities in SSTL and to compute the monitoring procedure locally at each location using the information of its neighbours. While SSTL can operate only on static spatial-temporal models (the position of the locations remains always fixed), STREL can handle also dynamic/mobile networks. Moreover, in [15] the authors show that for a certain class of models (called *euclidean spatial models*), all the spatial properties expressed in STREL that satisfy a model will also satisfy all the model transformations using rotation, translation and reflection.

4.3 Matching and Measuring Temporal Patterns over CPS Behaviours

In the introduction of this document, we mentioned that declarative specification languages are typically based on temporal logics or regular expressions. Both formalisms have their merits and weaknesses. Temporal logic are often good for describing global behaviours and the expected relations between the events and the states that evolve over time. In contrast, regular expressions are convenient for expressing local temporal patterns (consecutive sequences of events at states) that happen in a behaviour. In the digital hardware community, it has been observed that the necessary expressiveness and succinctness of the specification

language is truly achieved when temporal logic is combined with regular expressions. In fact, both IEEE formal specification language standards, SystemVerilog Assertions (SVA) [126] and Property Specification Language (PSL) [55] adopt this combined approach.

In the context of continuous-time applications, Timed Regular Expressions (TRE) were proposed in [11] as a real-time extension of regular expressions. For a long time, this formalism was subject to theoretical studies, but without any real practical relevance. More recently, a novel algorithm for matching and extracting TRE patterns from hybrid behaviours was developed in [124]. The original offline pattern matching procedure was extended with an online version in [125]. These results on TRE pattern matching enabled the combination of STL with regular expressions also in the continuous-time setting, as it was shown in [64]. Finally, automated extractions of quantitative measurements from CPS behaviours based on TRE patterns was proposed in [65].

5 Applications to Cyber-Physical Systems (CPS)

In this section, we provide an overview of the important applications of temporal logic-based monitoring techniques presented thus far and conclude with a presentation of the practical challenges that need to be addressed through future research.

5.1 Practical Considerations for CPS Monitoring

CPS integrate computation and control of physical processes to enable safety critical applications in many domains including medical devices, automotive systems, avionics and power systems [88]. The problem of monitoring CPS has been a productive area for the runtime verification community as a whole, leading to many important considerations such as monitoring timed properties, quantitative semantics, simulation-guided falsification and online monitoring challenges.

Both offline and online monitoring setups present unique challenges for CPS applications. As described in Sect. 3, the offline monitoring setup analyses trace data collected from running a system, after the execution has terminated. A key challenge includes that of monitoring large volumes of data efficiently [26]. At the same time, richer specification languages with higher computational worst-case complexities can be accommodated in an offline monitoring setup, exacerbating the challenge of efficient offline monitoring.

Online monitoring, on the other hand, is constrained by the limited ability to store the trace as the system being monitored executes and the hard real time demands on the computation time. Furthermore, in practice, monitoring is often restricted to perform a single pass through the trace in the forward direction. This naturally restricts us to specification languages that can be monitored efficiently in an online fashion. For instance, the presence of unbounded until operators in the specification can potentially require a large lookahead to resolve the truth of the formula appropriately [39, 66, 70]. Finally, if the monitoring shares

the same platform as the deployed system, it should be *non-intrusive* as much as possible: in other words, its consumption of resources such as CPU time, memory and I/O must not interfere with that of the running application [127]. The latter concern is especially acute for CPS, wherein instrumentation can potentially interfere with critical timing properties of the system being observed.

Finally, physical measurements are well known to be noisy and they require specialised sensors. As a result, the problem of monitoring under incomplete and noisy measurements is especially relevant for CPS applications [24, 83].

Another important classification, especially for CPS applications, involves the problem of modelling the physical environment surrounding the closed loop being monitored [79]. Herein, different monitoring setups are distinguished by the complexity and fidelity of the physical plant models and the software runtime setup used [84]. Software-in-the-loop monitors properties using the control system being designed and a mathematical model of the plant. Hardware-in-the-Loop monitoring (also known as Processor in the Loop) executes the controller on the runtime platform used during deployment, while using a mathematical model of the plant [120]. The monitoring challenges include the problem of mapping the “simulation time” of the plant model to the real-time elapsed for the real-time software. Model-based development environments such as Matlab™, Simulink/Stateflow™, Modelica™, Scade™ and DSpace™ support software/hardware in the loop testing and runtime monitoring for complex CPS [76].

Real-time Monitoring of CPS. Monitoring of real systems in real time during their execution requires adapting the techniques and the algorithms presented so far in this document. In the case of real-time monitoring, the property observers are implemented on a physical device that is connected to the system-under-test (SUT) [105, 118, 119]. Several considerations must be taken into account - the real-time sensing of the SUT signals and the environment, the frequency of the monitor operation that must be at least as high as the frequency at which the SUT works, the limited availability of resources that are available on the monitor device, etc. When considering real-time monitors implemented in embedded software or hardware, it is often the case that both the computations are done at the periodic intervals over the quantised input signals. In that case, STL is interpreted over signals that are defined in discrete time and over finite domains. While in this setting, STL is not more expressive than LTL, keeping explicit real-time operators and numerical inputs allows efficient implementation of monitors and allows giving them both qualitative and quantitative semantics. A translation of STL specifications into real-time monitors implemented in FPGA is proposed in [77]. A quantitative semantics for such STL properties based on the *weighted edit distance*, together with the algorithms for computing the robustness degree of a trace with respect to a property are developed in [78].

5.2 Property Falsification and Parameter Synthesis

In this section, we summarise work on the application of specification robustness of traces (Sect. 2.2) to two very related problems: property falsification and parameter synthesis.

Falsification techniques attempt to find a counterexample to a given property and to a given model of a system (or even to the actual system). Falsification is a rather valuable approach, generalising manual testing to an automatic search for violating counterexamples. However, the core challenge in falsification is the question of where to search for violations. This is a very challenging problem for CPS due to the continuous nature of the input space which consists of initial conditions and input signals to the model.

The problem of parameter synthesis for property satisfaction can be thought of as a dual problem to the falsification problem. In the case of parameter synthesis, the search is typically not over the space of initial conditions and inputs to the model, but over the model's parameters. In addition, the goal is not to find a falsifying behaviour, but rather either a set of satisfying behaviours for a range of parameters, or a parameter which induces a model which robustly satisfies the property.

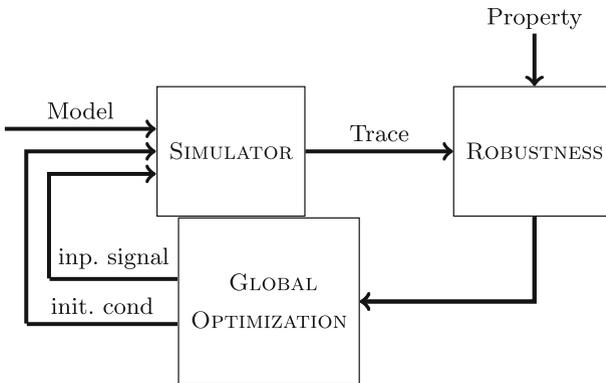


Fig. 9. Setup for the falsification of properties using robustness metrics.

To automate the search, we may use specification robustness to provide guidance on what to search for. Figure 9 illustrates the overall setup for falsification and/or parameter synthesis. As pointed out in [103] and in [2], robustness is a natural measure of a distance between a signal and a property. A tool that tries to minimise robustness by searching over inputs and initial conditions is in essence a tool for finding counterexamples. Similarly, a tool that attempts to maximise robustness by searching over a range of parameters is a parameter synthesis tool. To this end, a global optimisation engine is used to systematically

guide the search for inputs (parameters) that minimise (maximise) the overall robustness. The approach does not need to find the globally minimum (maximum) robustness. Rather, it stops whenever the specification robustness crosses a given threshold. For falsification methods, this threshold is typically the value zero indicating a change from positive robustness values (correct behaviour) to negative robustness values (incorrect behaviour). Also, in the case of counterexample search, even if a falsification is not obtained, then the minimum robustness obtained can serve as useful information to provide the engineer.

The idea of minimising robustness to search for falsification or maximising robustness for parameter synthesis has been implemented in tools such as Breach [46] and S-Taliro [10]. However, the challenges in these problems lie in the choice of a global optimisation solver that can efficiently (quickly) converge to local optima. In theory, any such solver will have to fundamentally grapple with the underlying undecidability of finding falsifying inputs for programs, in general. On the other hand, approaches such as Nelder-Meade algorithm [101], Simulated Annealing [103], Ant-Colony Optimization [9], Gaussian Process Optimization with Upper Confidence Bound [18], and the Cross-Entropy method [115] have been used to report success on large examples. Path-planning based methods like RRTs (rapidly exploring random trees) combined with online monitoring of STL robustness have shown promise in systems with hybrid dynamics [53].

As the model fidelity and complexity increases, so does the model simulation time. Long computation times are acceptable in optimal design applications, e.g., [58], since the system must be designed once; however, they can be problematic in system testing applications. Typically, the developers may be willing to wait overnight for test results, but most probably they will not be willing to wait for a week, for example. To improve the performance of stochastic optimisation methods, e.g., [9, 18, 103, 115], in [1, 6] they proposed hybrid techniques where robustness descent directions are analytically computed and interleaved with stochastic optimization methods. Such a process guarantees fast convergence to local optimiser points. On the other hand, it requires a white-box model, i.e., the mathematical model must be known to the falsification algorithm. Recently, the aforementioned restriction was relaxed in [130] where it was shown that the descent directions can be approximated as long as the system simulator can provide linearizations of the model along the simulation trace of the system. Finally, we should remark that some of the techniques which were used in [6] for falsification were initially used in [50] for parameter synthesis using simulations for computing reachable sets.

Automotive Systems. Automotive systems present an important application for many aspects of runtime monitoring and property falsification, discussed thus far. As automobiles become ever more autonomous, it is important to check the functional correctness of their core components.

The work in [62] describes the application of S-Taliro to automotive system models. Therein, they show the presence of unexpected behaviours in an automatic transmission model that were not revealed by previous testing approaches.

Falsification methods for stochastic systems were applied to stochastic models of automotive systems in [3]. The presented framework for robustness guided falsification in this section is also used as an intermediate step in specification mining methods [74, 129]. The methods presented in [74, 129] are primarily applied to automotive applications. In [131], the authors applied Breach as part of a compositional verification scheme for complex automotive system with many sub-modules. The Breach requirement mining feature was used at the system level to induce pre-conditions for sub-modules, making it easier to apply successfully model-checking analysis at the module level. The contribution in [85] formulated a library of control-theoretic specifications that can be expressed in STL and showed its application to an automotive powertrain control benchmark. Both STL and TRE have been recently used to specify, monitor and measure hybrid properties of the DS13 standard [65, 106].

CPS Engineering Education. STL monitoring was used in the context of a MOOC³ (Massively Open Online Class) teaching basic concepts of CPS design [82]. A key assignment was for the student to design, simulate and execute on real hardware a control algorithm driving a robot in an environment with obstacles. In order to evaluate hundreds of students contributions, a simulator was designed and equipped with STL monitoring capabilities. Grading was then done by evaluating a set of test cases and STL properties implementing fault monitors, i.e., each STL property evaluated to true would indicate a specific type of fault. The system would then return either some feedback if the user were a student or a partial grade if the user were an instructor.

Systems and Synthetic Biology. The growing need of computational models and methods [25] to investigate and to design complex biological systems with a predictable behaviour has also benefited greatly from the use of the aforementioned monitoring techniques. STL has become popular also among bioscientists to specify in a concise and unambiguous way the behaviour of several cellular and molecular mechanisms. The quantitative semantics of STL and its extensions has triggered the development of several *parameter synthesis* techniques and invaluable tools [12, 17, 18, 47, 69] to automatically characterise the parameter region of a biological model responsible for a behaviour of interest. Similarly to falsification analysis, parameter synthesis leverages an optimisation process using a particular heuristic. The only difference is that the objective function is to maximise (instead of minimising) the robustness with respect to an STL requirement. This approach has been successfully employed to study several biological case studies. Examples include the study of the onset of new blood vessel sprouting [48], the programmed cell death (apoptosis) [123], the effect of iron metabolism on blood cell specialisation [100] and the logical characterisation of an oscillator of the circadian clock in the *Ostreococcus Tauri* [19].

³ <https://www.edx.org/course/cyber-physical-systems-uc-berkeleyx-cecs149-1x>.

Medical Devices. The growing area of closed-loop medical devices has led to devices such as implantable pacemakers and artificial pancreas that provide life sustaining treatments in real-time. As a result, the problem of monitoring and verifying their operation takes on great significance. The broader area of closed loop medical devices has received a lot of recent interest from the formal verification community. This started with work on pacemakers and implantable cardiac defibrillators (ICDs) that includes hybrid automata models for excitable cells in the heart [20, 67], leading to approaches that employ these models to test closed loop systems [80, 81, 109].

Other examples of safety critical medical devices include the ones used in intensive care. In [34], the authors propose a method to automatically detect ineffective breathing efforts in patients in intensive care subject to assisted ventilation. Their approach is based on learning and monitoring STL specification discriminating between normal and ineffective breaths.

More recently, also the artificial pancreas concept has emerged as an important approach to treat type-1 diabetes, approaching a *de facto* cure [86].

Artificial Pancreas Control Systems. The artificial pancreas concept refers to a series of increasingly sophisticated devices (see Fig. 10) that automate the delivery of insulin to patients with type-1 diabetes in a closed loop, automatically responding to changes in the patient’s blood glucose levels and activities such as meals and exercise [38, 72, 86]. However, such systems can pose risks to the patient arising from defects and malfunctions. Short-term risks include extremely low blood glucose levels called *hypoglycemia*, that can lead to seizures, loss of consciousness, coma or even death in extreme cases.

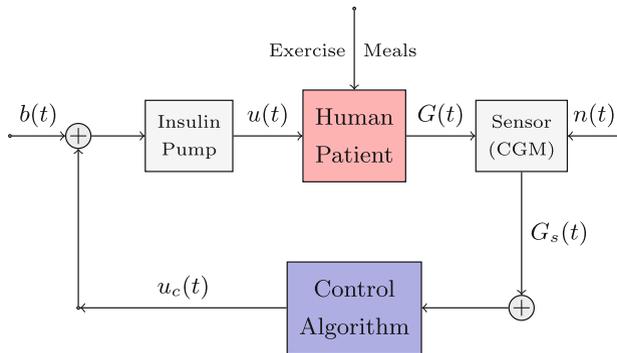


Fig. 10. Overview of the key components of an artificial pancreas control system. $b(t)$: external user commanded insulin, $u(t)$: insulin infused to patient, $G(t)$: blood glucose level of the patient, $n(t)$: sensor measurement error (noise), $G_s(t)$: glucose level estimated/reported by sensor, $u_c(t)$: insulin infusion commanded by the algorithm.

Cameron et al. use robustness-guided falsification techniques for checking properties of closed loop control systems for the artificial pancreas [36]. Their

work investigates a PID controller proposed in [121, 122, 128] based on published descriptions of the control system available. The simulation environment incorporates this controller in a closed loop with models of the patient [99], the sensors and actuators. Their work formulated nearly six different temporal properties of the closed loop and obtained falsification for three of them. However, they could not falsify the remaining three properties that governed the absence of prolonged hypoglycemia and hyperglycemia in the patient.

Another recent study [116] was performed to test a predictive pump shutoff controller designed in [35] that has undergone outpatient clinical trials [92]. This study involved the entire controller software *as is*, without any modifications. At the same time, the closed loop simulation permits us to pose a rich set of questions that compare the closed loop performance with a corresponding open loop under the same meal inputs and physiological model conditions. The falsification discovered adverse noise patterns in the CGM sensor that could trick the Kalman filter into predicting inaccurate forecasts for the future glucose value, and thus prevent appropriate pump shutoff/resumption. At the same time, critical properties such as not commanding excess insulin when the patient is in hypoglycemia could not be violated. The study concluded the need to investigate these violations under more realistic patterns of CGM noise.

6 Tools

Due to relatively low computational complexity of the online and offline monitoring algorithms, many software tools have been developed over the last two decades. Among the first tools that were developed for monitoring (a subset or superset) of Boolean-valued temporal logic specifications were the Temporal Rover [54], MaC [89], Java PathExplorer [70] and LOLA [39]. Since then, there has been a wealth of research on on- and off-line monitoring of requirements expressed in some form of temporal logic (see for example the competition at the *Runtime Verification* conference series [14, 23, 63, 112]) and several publicly available tools have resulted from this effort, for example:

1. RV-Monitor [91]: is available at <https://runtimeverification.com/monitor/>
2. MonPoly [27]: is available at http://www.infsec.ethz.ch/research/projects/mon_enf.html
3. LTLFO2Mon [28] is available at <https://github.com/jckuester/ltlfo2mon>

The review in this section focuses only on tools that can reason about real-time properties of traces (output signals) since this is a necessity for testing and monitoring for Embedded and Cyber-Physical Systems. In addition, the focus is on publicly available software tools for off- and on-line monitoring that can be readily downloaded and utilized in testing and monitoring applications. Some of these tools are open source with licenses that allow extensions and redistribution. In the following, we group the software tools survey into two main broader categories: Boolean semantics and multi-valued semantics.

6.1 Software Tools for Boolean Semantics

In the first category, i.e., Boolean semantics, the tool AMT [108] available at:

<http://www-verimag.imag.fr/DIST-TOOLS/TEMPO/AMT/content.html>

analyzes STL properties over analog system output signals. In particular, the properties analyzed by AMT are an extension of the industrial specification language PSL with STL requirements. The software tool AMT is a stand alone executable with a graphical interface where the user provides the STL/PSL properties, the signals and whether the analysis is going to be offline or incremental. In return, the tool plots the Boolean satisfaction of each property over time.

6.2 Software Tools for Quantitative Semantics

When considering software tools for evaluating quantitative semantics of STL over signals, there are several options. The following tools are publicly available and they can monitor both real valued and Boolean signals:

1. Breach [46]: available at <https://github.com/decyphir/breach>
2. S-Taliro [10]: available at <https://sites.google.com/a/asu.edu/s-taliro/>
3. U-Check [31]: available at <https://github.com/dmilios/U-check>

Breach and S-Taliro are add-on toolboxes for the Matlab environment while U-Check is a stand-alone program written in Java. Breach and S-Taliro provide analysis tools for black box testing of models and hardware-in-the-loop systems while U-Check deals with stochastic models (Continuous-Time Markov Chains).

The efficient evaluation of STL requirements over real-valued and Boolean signals gave raise to a number of semi-formal verification methods from testing based verification [59] to parameter mining [129, 132] to falsification [2] to

Table 2. Tools for reasoning with multi-valued temporal logics and their functionality.

Analysis	Functionality	Breach	S-Taliro	U-Check
Signal	Offline testing	[46, 49]	[60–62]	
	Online monitoring	[40]	[43]	
System (best effort)	Falsification	[46]	[2]	[31]
	Coverage testing		[45]	
	Requirement mining	[129]	[74, 132]	
	Parameter synthesis	[46]	[117]	
	Conformance		[4]	
Specification	Visual specifications		[73, 75]	
	Debugging		[44]	

synthesis [111]. As reviewed in Sect. 5, the aforementioned methods have been applied to a wide range of practical applications. Table 2 provides an overview with references of the various analysis methods that each tool supports.

7 Conclusion

Cyber-Physical Systems (CPS) combine heterogeneous and networked computational entities with physical components interacting with them through sensors and actuators. Continuous and hybrid behaviours naturally arise from such dynamical systems. Here, we have provided an in-depth overview of the state-of-the-art techniques for CPS monitoring.

The common denominator of all these methods is the possibility to express in a very powerful, concise and unambiguous way the properties of interest using a formal specification language. In this work, we have mainly focused our attention on Signal Temporal Logic (STL), a formalism enabling the designer to reason about real-time properties over real-valued signals. In the recent years, there has been a great effort to provide efficient algorithms to support online and offline monitoring of STL formulas over (system output) signals.

The introduction of novel quantitative semantics has considerably widened the spectrum of applications from just monitoring qualitatively real-time signals to providing novel falsification analysis and parameter synthesis techniques in model-based testing as well as hardware-in-the-loop testing. As a consequence, the application domains have also grown dramatically, ranging now from automotive systems to synthetic biology and medical devices.

We believe that these techniques will play more and more a key role in industry in the design and engineering safe and resilient CPS and/or to equip them with real-time hardware-based monitors enabling CPS self-awareness and adaptation.

Acknowledgment. E. Bartocci and D. Ničković acknowledge the partial support of the EU ICT COST Action IC1402 on Runtime Verification beyond Monitoring (ARVI) and of the HARMONIA (845631) project, funded by a national Austrian grant from Austrian FFG under the program IKT der Zukunft. E. Bartocci acknowledges the partial support of the Austrian National Research Network S 11405-N23 (RiSE/SHiNE) of the Austrian Science Fund (FWF). G. Fainekos acknowledges the support of the NSF CAREER award 1350420.

References

1. Abbas, H., Fainekos, G.: Computing descent direction of MTL robustness for non-linear systems. In: Proceedings of ACC 2013: The 2013 American Control Conference, pp. 4405–4410 (2013)
2. Abbas, H., Fainekos, G.E., Sankaranarayanan, S., Ivancic, F., Gupta, A.: Probabilistic temporal logic falsification of cyber-physical systems. *ACM Trans. Embed. Comput. Syst.* **12**(s2), 95:1–95:30 (2013)

3. Abbas, H., Hoxha, B., Fainekos, G., Ueda, K.: Robustness-guided temporal logic testing and verification for stochastic cyber-physical systems. In: Proceedings of the 4th Annual IEEE International Conference on Cyber Technology in Automation, Control and Intelligent, pp. 1–6. IEEE (2014)
4. Abbas, H., Mittelman, H., Fainekos, G.E.: Formal property verification in a conformance testing framework. In: Proceedings of MEMOCODE 2014: The 12th ACM-IEEE International Conference on Formal Methods and Models for System Design, pp. 155–164. IEEE (2014)
5. Abbas, H., Rodionova, A., Bartocci, E., Smolka, S.A., Grosu, R.: Quantitative regular expressions for Arrhythmia detection algorithms. In: Feret, J., Koepl, H. (eds.) CMSB 2017. LNCS, vol. 10545, pp. 23–39. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67471-1_2
6. Abbas, H., Winn, A., Fainekos, G.E., Julius, A.A.: Functional gradient descent method for metric temporal logic specifications. In: Proceedings of ACC 2014: The American Control Conference, pp. 2312–2317. IEEE (2014)
7. Akazaki, T., Hasuo, I.: Time robustness in MTL and expressivity in hybrid system falsification. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9207, pp. 356–374. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21668-3_21
8. Alur, R., Feder, T., Henzinger, T.A.: The benefits of relaxing punctuality. *J. ACM* **43**(1), 116–146 (1996)
9. Annapureddy, Y.S.R., Fainekos, G.E.: Ant colonies for temporal logic falsification of hybrid systems. In: Proceedings of IECON 2010: The 36th Annual Conference on IEEE Industrial Electronics Society, pp. 91–96 (2010)
10. Annapureddy, Y., Liu, C., Fainekos, G., Sankaranarayanan, S.: S-TALiRO: a tool for temporal logic falsification for hybrid systems. In: Abdulla, P.A., Leino, K.R.M. (eds.) TACAS 2011. LNCS, vol. 6605, pp. 254–257. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19835-9_21
11. Asarin, E., Caspi, P., Maler, O.: Timed regular expressions. *J. ACM* **49**(2), 172–206 (2002)
12. Aydin-Gol, E., Bartocci, E., Belta, C.: A formal methods approach to pattern synthesis in reaction diffusion systems. In: Proceedings of CDC 2014: The 53rd IEEE Conference on Decision and Control, pp. 108–113. IEEE (2014)
13. Bartocci, E., Aydin-Gol, E., Haghghi, I., Belta, C.: A formal methods approach to pattern recognition and synthesis in reaction diffusion networks. *IEEE Trans. Control Netw. Syst.* **PP**(99), 1–12 (2016)
14. Bartocci, E., Bonakdarpour, B., Falcone, Y.: First international competition on software for runtime verification. In: Bonakdarpour, B., Smolka, S.A. (eds.) RV 2014. LNCS, vol. 8734, pp. 1–9. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11164-3_1
15. Bartocci, E., Bortolussi, L., Loreti, M., Nenzi, L.: Monitoring mobile and spatially distributed cyber-physical systems. In: Proceedings of MEMOCODE 2017: The 15th ACM-IEEE International Conference on Formal Methods and Models for System Design, pp. 146–155. ACM (2017)
16. Bartocci, E., Bortolussi, L., Milios, D., Nenzi, L., Sanguinetti, G.: Studying emergent behaviours in morphogenesis using signal spatio-temporal logic. In: Abate, A., Šafránek, D. (eds.) HSB 2015. LNCS, vol. 9271, pp. 156–172. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-26916-0_9

17. Bartocci, E., Bortolussi, L., Nenzi, L.: A temporal logic approach to modular design of synthetic biological circuits. In: Gupta, A., Henzinger, T.A. (eds.) CMSB 2013. LNCS, vol. 8130, pp. 164–177. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40708-6_13
18. Bartocci, E., Bortolussi, L., Nenzi, L., Sanguinetti, G.: System design of stochastic models using robustness of temporal properties. *Theor. Comput. Sci.* **587**, 3–25 (2015)
19. Bartocci, E., Bortolussi, L., Sanguinetti, G.: Data-driven statistical learning of temporal logic properties. In: Legay, A., Bozga, M. (eds.) FORMATS 2014. LNCS, vol. 8711, pp. 23–37. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10512-3_3
20. Bartocci, E., Corradini, F., Berardini, M.R.D., Entcheva, E., Smolka, S.A., Grosu, R.: Modeling and simulation of cardiac tissue using hybrid I/O automata. *Theor. Comput. Sci.* **410**(33–34), 3149–3165 (2009)
21. Bartocci, E., Corradini, F., Merelli, E., Tesei, L.: Model checking biological oscillators. *Electr. Notes Theor. Comput. Sci.* **229**(1), 41–58 (2009)
22. Bartocci, E., Corradini, F., Merelli, E., Tesei, L.: Detecting synchronisation of biological oscillators by model checking. *Theor. Comput. Sci.* **411**(20), 1999–2018 (2010)
23. Bartocci, E., Falcone, Y., Bonakdarpour, B., Colombo, C., Decker, N., Havelund, K., Joshi, Y., Klaedtke, F., Milewicz, R., Reger, G., Rosu, G., Signoles, J., Thoma, D., Zălinescu, E., Zhang, Y.: First international competition on runtime verification: rules, benchmarks, tools, and final results of CRV 2014. *Int. J. Softw. Tools Technol. Transf.*, 1–40, April 2017
24. Bartocci, E., Grosu, R., Karmarkar, A., Smolka, S.A., Stoller, S.D., Zadok, E., Seyster, J.: Adaptive runtime verification. In: Qadeer, S., Tasiran, S. (eds.) RV 2012. LNCS, vol. 7687, pp. 168–182. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-35632-2_18
25. Bartocci, E., Liò, P.: Computational modeling, formal analysis, and tools for systems biology. *PLoS Comput. Biol.* **12**(1), 1–22 (2016)
26. Basin, D., Caronni, G., Ereth, S., Harvan, M., Klaedtke, F., Mantel, H.: Scalable offline monitoring. In: Bonakdarpour, B., Smolka, S.A. (eds.) RV 2014. LNCS, vol. 8734, pp. 31–47. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11164-3_4
27. Basin, D., Harvan, M., Klaedtke, F., Zălinescu, E.: MONPOLY: monitoring usage-control policies. In: Khurshid, S., Sen, K. (eds.) RV 2011. LNCS, vol. 7186, pp. 360–364. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29860-8_27
28. Bauer, A., Küster, J.-C., Vegliach, G.: From propositional to first-order monitoring. In: Legay, A., Bensalem, S. (eds.) RV 2013. LNCS, vol. 8174, pp. 59–75. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40787-1_4
29. Bauer, A., Leucker, M., Schallhart, C.: Monitoring of real-time properties. In: Arun-Kumar, S., Garg, N. (eds.) FSTTCS 2006. LNCS, vol. 4337, pp. 260–272. Springer, Heidelberg (2006). https://doi.org/10.1007/11944836_25
30. Bauer, A., Leucker, M., Schallhart, C.: Comparing LTL semantics for runtime verification. *J. Logic Comput.* **20**(3), 651–674 (2010)
31. Bortolussi, L., Milios, D., Sanguinetti, G.: U-check: model checking and parameter synthesis under uncertainty. In: Campos, J., Haverkort, B.R. (eds.) QEST 2015. LNCS, vol. 9259, pp. 89–104. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-22264-6_6

32. Brim, L., Dluhos, P., Safránek, D., Vejpustek, T.: STL*: Extending signal temporal logic with signal-value freezing operator. *Inf. Comput.* **236**, 52–67 (2014)
33. Brim, L., Vejpustek, T., Safránek, D., Fabriková, J.: Robustness analysis for value-freezing signal temporal logic. In: *Proceedings of HSB 2013: The Second International Workshop on Hybrid Systems and Biology*. EPTCS, vol. 125, pp. 20–36 (2013)
34. Bufo, S., Bartocci, E., Sanguinetti, G., Borelli, M., Lucangelo, U., Bortolussi, L.: Temporal logic based monitoring of assisted ventilation in intensive care patients. In: Margaria, T., Steffen, B. (eds.) *ISoLA 2014*. LNCS, vol. 8803, pp. 391–403. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45231-8_30
35. Cameron, F., Wilson, D.M., Buckingham, B.A., Arzumanyan, H., Clinton, P., Chase, H.P., Lum, J., Maahs, D.M., Calhoun, P.M., Bequette, B.W.: Inpatient studies of a Kalman-filter-based predictive pump shutoff algorithm. *J. Diabetes Sci. Technol.* **6**(5), 1142–1147 (2012)
36. Cameron, F., Fainekos, G., Maahs, D.M., Sankaranarayanan, S.: Towards a verified artificial pancreas: challenges and solutions for runtime verification. In: Bartocci, E., Majumdar, R. (eds.) *RV 2015*. LNCS, vol. 9333, pp. 3–17. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23820-3_1
37. Chakarov, A., Sankaranarayanan, S., Fainekos, G.: Combining time and frequency domain specifications for periodic signals. In: Khurshid, S., Sen, K. (eds.) *RV 2011*. LNCS, vol. 7186, pp. 294–309. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29860-8_22
38. Cobelli, C., Man, C.D., Sparacino, G., Magni, L., Nicolao, G.D., Kovatchev, B.P.: Diabetes: Models, signals and control (methodological review). *IEEE Rev. Biomed. Eng.* **2**, 54–95 (2009)
39. D’Angelo, B., Sankaranarayanan, S., Sanchez, C., Robinson, W., Finkbeiner, B., Sipma, H., Mehrotra, S., Manna, Z.: LOLA: runtime monitoring of synchronous systems. In: *Proceedings of TIME 2005: The 12th International Symposium on Temporal Representation and Reasoning*, pp. 166–174. IEEE (2005)
40. Deshmukh, J.V., Donzé, A., Ghosh, S., Jin, X., Juniwal, G., Seshia, S.A.: Robust online monitoring of signal temporal logic. In: Bartocci, E., Majumdar, R. (eds.) *RV 2015*. LNCS, vol. 9333, pp. 55–70. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23820-3_4
41. Deshmukh, J.V., Donzé, A., Ghosh, S., Jin, X., Garvit, J., Seshia, S.A.: Robust online monitoring of signal temporal logic. *Formal Methods Syst. Des.* **51**(1), 5–30 (2017)
42. Deshmukh, J.V., Majumdar, R., Prabhu, V.S.: Quantifying conformance using the Skorokhod metric. In: Kroening, D., Păsăreanu, C.S. (eds.) *CAV 2015*. LNCS, vol. 9207, pp. 234–250. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21668-3_14
43. Dokhanchi, A., Hoxha, B., Fainekos, G.: On-line monitoring for temporal logic robustness. In: Bonakdarpour, B., Smolka, S.A. (eds.) *RV 2014*. LNCS, vol. 8734, pp. 231–246. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11164-3_19
44. Dokhanchi, A., Hoxha, B., Fainekos, G.E.: Metric interval temporal logic specification elicitation and debugging. In: *Proceedings of MEMOCODE 2015: The 13th ACM/IEEE International Conference on Formal Methods and Models for Codesign*, pp. 70–79. IEEE (2015)

45. Dokhanchi, A., Zutshi, A., Sriniva, R.T., Sankaranarayanan, S., Fainekos, G.: Requirements driven falsification with coverage metrics. In: Proceedings of EMSOFT: The 12th International Conference on Embedded Software, pp. 31–40. IEEE (2015)
46. Donzé, A.: Breach, a toolbox for verification and parameter synthesis of hybrid systems. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 167–170. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_17
47. Donzé, A., Clermont, G., Legay, A., Langmead, C.J.: Parameter synthesis in nonlinear dynamical systems: application to systems biology. In: Batzoglu, S. (ed.) RECOMB 2009. LNCS, vol. 5541, pp. 155–169. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02008-7_11
48. Donzé, A., Fanchon, E., Gattepaille, L.M., Maler, O., Tracqui, P.: Robustness analysis and behavior discrimination in enzymatic reaction networks. PLoS ONE **6**(9), e24246 (2011)
49. Donzé, A., Ferrère, T., Maler, O.: Efficient robust monitoring for STL. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 264–279. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_19
50. Donzé, A., Krogh, B., Rajhans, A.: Parameter synthesis for hybrid systems with an application to simulink models. In: Majumdar, R., Tabuada, P. (eds.) HSCC 2009. LNCS, vol. 5469, pp. 165–179. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00602-9_12
51. Donzé, A., Maler, O.: Robust satisfaction of temporal logic over real-valued signals. In: Chatterjee, K., Henzinger, T.A. (eds.) FORMATS 2010. LNCS, vol. 6246, pp. 92–106. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15297-9_9
52. Donzé, A., Maler, O., Bartocci, E., Nickovic, D., Grosu, R., Smolka, S.: On temporal logic and signal processing. In: Chakraborty, S., Mukund, M. (eds.) ATVA 2012. LNCS, pp. 92–106. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33386-6_9
53. Dreossi, T., Dang, T., Donzé, A., Kapinski, J., Jin, X., Deshmukh, J.V.: Efficient guiding strategies for testing of temporal properties of hybrid systems. In: Havelund, K., Holzmann, G., Joshi, R. (eds.) NFM 2015. LNCS, vol. 9058, pp. 127–142. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-17524-9_10
54. Drusinsky, D.: Monitoring temporal rules combined with time series. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 114–117. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45069-6_11
55. Eisner, C., Fisman, D.: A Practical Introduction to PSL. Springer, Heidelberg (2006). <https://doi.org/10.1007/978-0-387-36123-9>
56. Eisner, C., Fisman, D., Havlicek, J.: A topological characterization of weakness. In: Proceedings of PODC 2005: The 24th Annual ACM Symposium on Principles of Distributed Computing, pp. 1–8. ACM (2005)
57. Eisner, C., Fisman, D., Havlicek, J., Lustig, Y., McIsaac, A., Van Campenhout, D.: Reasoning with temporal logic on truncated paths. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 27–39. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45069-6_3
58. Fainekos, G.E., Giannakoglou, K.C.: Inverse design of airfoils based on a novel formulation of the ant colony optimization method. Inverse Prob. Eng. **11**(1), 21–38 (2003)

59. Fainekos, G.E., Girard, A., Pappas, G.J.: Temporal logic verification using simulation. In: Asarin, E., Bouyer, P. (eds.) FORMATS 2006. LNCS, vol. 4202, pp. 171–186. Springer, Heidelberg (2006). https://doi.org/10.1007/11867340_13
60. Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications. In: Havelund, K., Núñez, M., Roşu, G., Wolff, B. (eds.) FATES/RV 2006. LNCS, vol. 4262, pp. 178–192. Springer, Heidelberg (2006). https://doi.org/10.1007/11940197_12
61. Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications for continuous-time signals. *Theor. Comput. Sci.* **410**(42), 4262–4291 (2009)
62. Fainekos, G.E., Sankaranarayanan, S., Ueda, K., Yazarel, H.: Verification of automotive control applications using S-TaLiRo. In: Proceedings of ACC 2012: The 2012 American Control Conference, pp. 3567–3572. IEEE (2012)
63. Falcone, Y., Ničković, D., Reger, G., Thoma, D.: Second international competition on runtime verification. In: Bartocci, E., Majumdar, R. (eds.) RV 2015. LNCS, vol. 9333, pp. 405–422. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23820-3_27
64. Ferrère, T.: Assertions and measurements for mixed-signal simulation. Ph.D. thesis. Université Grenoble-Alpes, France (2016)
65. Ferrère, T., Maler, O., Ničković, D., Ulus, D.: Measuring with timed patterns. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9207, pp. 322–337. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21668-3_19
66. Finkbeiner, B., Sipma, H.B.: Checking finite traces using alternating automata. *Formal Methods Syst. Des.* **24**(2), 101–127 (2004)
67. Grosu, R., Batt, G., Fenton, F.H., Glimm, J., Le Guernic, C., Smolka, S.A., Bartocci, E.: From cardiac cells to genetic regulatory networks. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 396–411. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_31
68. Grosu, R., Smolka, S.A., Corradini, F., Wasilewska, A., Entcheva, E., Bartocci, E.: Learning and detecting emergent behavior in networks of cardiac myocytes. *Commun. ACM* **52**(3), 97–105 (2009)
69. Haghghi, I., Jones, A., Kong, Z., Bartocci, E., Grosu, R., Belta, C.: SpaTeL: a novel spatial-temporal logic and its applications to networked systems. In: Proceedings of HSCC 2015: The 18th International Conference on Hybrid Systems: Computation and Control, pp. 189–198. IEEE (2015)
70. Havelund, K., Rosu, G.: Monitoring Java programs with Java pathexplorer. *Electron. Not. Theoret. Comput. Sci.* **55**(2), 200–217 (2001)
71. Ho, H.-M., Ouaknine, J., Worrell, J.: Online monitoring of metric temporal logic. In: Bonakdarpour, B., Smolka, S.A. (eds.) RV 2014. LNCS, vol. 8734, pp. 178–192. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11164-3_15
72. Hovorka, R.: Continuous glucose monitoring and closed-loop systems. *Diabet. Med.* **23**(1), 1–12 (2005)
73. Hoxha, B., Bach, H., Abbas, H., Dokhanci, A., Kobayashi, Y., Fainekos, G.: Towards formal specification visualization for testing and monitoring of cyber-physical systems. In: International Workshop on Design and Implementation of Formal Tools and Systems, DIFTS 2014 (2014)
74. Hoxha, B., Dokhanchi, A., Fainekos, G.: Mining parametric temporal logic properties in model based design for cyber-physical systems. *Int. J. Softw. Tools Technol. Transf.* (2017). (in press)

75. Hoxha, B., Mavridis, N., Fainekos, G.E.: VISPEC: a graphical tool for elicitation of MTL requirements. In: Proceedings of IROS 2015: The 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3486–3492. IEEE (2015)
76. MathWorks, Inc.: Test generated code with SIL and PIL simulations, cf. <https://www.mathworks.com/help/ecoder/examples/software-and-processor-in-the-loop-sil-and-pil-simulation.html>
77. Jaksic, S., Bartocci, E., Grosu, R., Kloibhofer, R., Nguyen, T., Ničković, D.: From signal temporal logic to FPGA monitors. In: Proceedings of MEMOCODE 2015: The 13th ACM/IEEE International Conference on Formal Methods and Models for Codesign, pp. 218–227. IEEE (2015)
78. Jakić, S., Bartocci, E., Grosu, R., Ničković, D.: Quantitative monitoring of STL with edit distance. In: Falcone, Y., Sánchez, C. (eds.) RV 2016. LNCS, vol. 10012, pp. 201–218. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46982-9_13
79. Jensen, J.C., Chang, D.H., Lee, E.A.: A model-based design methodology for cyber-physical systems. In: Proceedings of IEEE Workshop on Design, Modeling, and Evaluation of Cyber-Physical Systems (CyPhy), pp. 1666–1671. IEEE (2011)
80. Jiang, Z., Pajic, M., Alur, R., Mangharam, R.: Closed-loop verification of medical devices with model abstraction and refinement. *Int. J. Softw. Tools Technol. Transfer* **16**(2), 191–213 (2014)
81. Jiang, Z., Pajic, M., Moarref, S., Alur, R., Mangharam, R.: Modeling and verification of a dual chamber implantable pacemaker. In: Flanagan, C., König, B. (eds.) TACAS 2012. LNCS, vol. 7214, pp. 188–203. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28756-5_14
82. Juniwal, G., Donzé, A., Jensen, J.C., Seshia, S.A.: CPSGrader: synthesizing temporal logic testers for auto-grading an embedded systems laboratory. In: Proceedings of EMSOFT 2014: The 2014 International Conference on Embedded Software, pp. 24:1–24:10. IEEE (2014)
83. Kalajdzic, K., Bartocci, E., Smolka, S.A., Stoller, S.D., Grosu, R.: Runtime verification with particle filtering. In: Legay, A., Bensalem, S. (eds.) RV 2013. LNCS, vol. 8174, pp. 149–166. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40787-1_9
84. Kane, A.: Runtime monitoring for safety-critical embedded systems. Ph.D. thesis, Carnegie Mellon University, College of Engineering (2015)
85. Kapinski, J., Jin, X., Deshmukh, J., Donzé, A., Yamaguchi, T., Ito, H., Kaga, T., Kobuna, S., Seshia, S.: ST-Lib: a library for specifying and classifying model behaviors. In: SAE Technical Paper. SAE International (2016)
86. Kowalski, A.: Pathway to artificial pancreas revisited: moving downstream. *Diabetes Care* **38**, 1036–1043 (2015)
87. Koymans, R.: Specifying real-time properties with metric temporal logic. *Real-Time Syst.* **2**(4), 255–299 (1990)
88. Lee, E.A.: Cyber physical systems: design challenges. In: Proceedings of ISORC 2011: The 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing, pp. 363–369, May 2008
89. Lee, I., Kannan, S., Kim, M., Sokolsky, O., Viswanathan, M.: Runtime assurance based on formal specifications. In: Proceedings of PDPTA 1999: The International Conference on Parallel and Distributed Processing Techniques and Applications, pp. 279–287. CSREA Press (1999)
90. Lemire, D.: Streaming maximum-minimum filter using no more than three comparisons per element. *Nord. J. Comput.* **13**(4), 328–339 (2006)

91. Luo, Q., Zhang, Y., Lee, C., Jin, D., Meredith, P.O.N., Șerbănuță, T.F., Roșu, G.: RV-Monitor: efficient parametric runtime verification with simultaneous properties. In: Bonakdarpour, B., Smolka, S.A. (eds.) RV 2014. LNCS, vol. 8734, pp. 285–300. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11164-3_24
92. Maahs, D.M., Calhoun, P., Buckingham, B.A., et al.: A randomized trial of a home system to reduce nocturnal hypoglycemia in type 1 diabetes. *Diabetes Care* **37**(7), 1885–1891 (2014)
93. Majumdar, R., Prabhu, V.S.: Computing the Skorokhod distance between polygonal traces. In: Proceedings of HSCC 2015: The 18th International Conference on Hybrid Systems: Computation and Control, pp. 199–208. ACM (2015)
94. Majumdar, R., Prabhu, V.S.: Computing distances between reach flowpipes. In: Proceedings of HSCC 2016: The 19th International Conference on Hybrid Systems: Computation and Control, pp. 267–276. ACM (2016)
95. Maler, O.: Some thoughts on runtime verification. In: Falcone, Y., Sánchez, C. (eds.) RV 2016. LNCS, vol. 10012, pp. 3–14. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46982-9_1
96. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Lakhnech, Y., Yovine, S. (eds.) FORMATS/FTRTFT 2004. LNCS, vol. 3253, pp. 152–166. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30206-3_12
97. Maler, O., Ničković, D.: Monitoring properties of analog and mixed-signal circuits. *STTT* **15**(3), 247–268 (2013)
98. Maler, O., Nickovic, D., Pnueli, A.: Checking temporal properties of discrete, timed and continuous behaviors. In: Avron, A., Dershowitz, N., Rabinovich, A. (eds.) Pillars of Computer Science. LNCS, vol. 4800, pp. 475–505. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78127-1_26
99. Man, C.D., Raimondo, D.M., Rizza, R.A., Cobelli, C.: GIM, simulation software of meal glucose-insulin model. *J. Diabetes Sci. Tech.* **1**(3), 323–330 (2007)
100. Mobilia, N., Donzé, A., Marc Moulis, J., Fanchon, E.: Producing a set of models for the iron homeostasis network. In: Proceedings of HSB 2013: The Second International Workshop on Hybrid Systems and Biology. EPTCS, vol. 125, pp. 92–98 (2013)
101. Nelder, J.A., Mead, R.: A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965)
102. Nenzi, L., Bortolussi, L., Ciancia, V., Loreti, M., Massink, M.: Qualitative and quantitative monitoring of spatio-temporal properties. In: Bartocci, E., Majumdar, R. (eds.) RV 2015. LNCS, vol. 9333, pp. 21–37. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23820-3_2
103. Nghiem, T., Sankaranarayanan, S., Fainekos, G.E., Ivancic, F., Gupta, A., Pappas, G.J.: Monte-carlo techniques for falsification of temporal properties of non-linear hybrid systems. In: Proceedings of HSCC 2010: The 13th ACM International Conference on Hybrid Systems: Computation and Control, pp. 211–220. ACM (2010)
104. Nguyen, L., Kapinski, J., Jin, X., Deshmukh, J., Butts, K., Johnson, T.: Abnormal data classification using time-frequency temporal logic. In: Proceedings of HSCC 2017: The 20th ACM International Conference on Hybrid Systems: Computation and Control, pp. 237–242. ACM (2017)
105. Nguyen, T., Bartocci, E., Ničković, D., Grosu, R., Jaksic, S., Selyunin, K.: The HARMONIA project: hardware monitoring for automotive systems-of-systems. In: Margaria, T., Steffen, B. (eds.) ISoLA 2016. LNCS, vol. 9953, pp. 371–379. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47169-3_28

106. Nguyen, T., Ničković, D.: Assertion-based monitoring in practice – checking correctness of an automotive sensor interface. In: Lang, F., Flammini, F. (eds.) FMICS 2014. LNCS, vol. 8718, pp. 16–32. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10702-8_2
107. Nickovic, D.: Checking timed and hybrid properties: theory and applications. Ph.D. thesis. Université Joseph Fourier, Grenoble, France (2008)
108. Nickovic, D., Maler, O.: AMT: a property-based monitoring tool for analog systems. In: Raskin, J.-F., Thiagarajan, P.S. (eds.) FORMATS 2007. LNCS, vol. 4763, pp. 304–319. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75454-1_22
109. Pajic, M., Mangharam, R., Sokolsky, O., Arney, D., Goldman, J., Lee, I.: Model-driven safety analysis of closed-loop medical systems. *IEEE Trans. Ind. Inform.* **10**(1), 3–16 (2014)
110. Pnueli, A.: The temporal logic of programs. In: Proceedings of the 18th Annual Symposium on Foundations of Computer Science, pp. 46–57. IEEE (1977)
111. Raman, V., Donzé, A., Sadigh, D., M. Murray, R., Seshia, S.A.: Reactive synthesis from signal temporal logic specifications. In: Proceedings of the HSCC 2015: The 18th International Conference on Hybrid Systems: Computation and Control, pp. 239–248. ACM (2015)
112. Reger, G., Hallé, S., Falcone, Y.: Third international competition on runtime verification. In: Falcone, Y., Sánchez, C. (eds.) RV 2016. LNCS, vol. 10012, pp. 21–37. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46982-9_3
113. Rizk, A., Batt, G., Fages, F., Soliman, S.: On a continuous degree of satisfaction of temporal logic formulae with applications to systems biology. In: Heiner, M., Uhrmacher, A.M. (eds.) CMSB 2008. LNCS (LNAI), vol. 5307, pp. 251–268. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88562-7_19
114. Rodionova, A., Bartocci, E., Ničković, D., Grosu, R.: Temporal logic as filtering. In: Proceedings of HSCC 2016: The 19th International Conference on Hybrid Systems: Computation and Control, pp. 11–20. ACM (2016)
115. Sankaranarayanan, S., Fainekos, G.: Falsification of temporal properties of hybrid systems using the cross-entropy method. In: Proceedings of HSCC 2012: The 15th ACM International Conference on Hybrid Systems: Computation and Control, pp. 125–134. ACM (2012)
116. Sankaranarayanan, S., Kumar, S.A., Cameron, F., Bequette, B.W., Fainekos, G.E., Maahs, D.M.: Model-based falsification of an artificial pancreas control system. *SIGBED Rev.* **14**(2), 24–33 (2017)
117. Sankaranarayanan, S., Miller, C., Raghunathan, R., Ravanbakhsh, H., Fainekos, G.E.: A model-based approach to synthesizing insulin infusion pump usage parameters for diabetic patients. In: Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing, pp. 1610–1617. IEEE (2012)
118. Selyunin, K., Jaksic, S., Nguyen, T., Reidl, C., Hafner, U., Bartocci, E., Nickovic, D., Grosu, R.: Runtime monitoring with recovery of the SENT communication protocol. In: Majumdar, R., Kunčak, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 336–355. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63387-9_17
119. Selyunin, K., Nguyen, T., Bartocci, E., Grosu, R.: Applying runtime monitoring for automotive electronic development. In: Falcone, Y., Sánchez, C. (eds.) RV 2016. LNCS, vol. 10012, pp. 462–469. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46982-9_30
120. Short, M., Pont, M.J.: Hardware in the loop simulation of embedded automotive control system. In: Proceedings of 2005 IEEE Intelligent Transportation Systems, pp. 426–431. IEEE, September 2005

121. Steil, G.M.: Algorithms for a closed-loop artificial pancreas: the case for proportional-integral-derivative control. *J. Diabetes Sci. Technol.* **7**, 1621–1631 (2013)
122. Steil, G., Pantoleon, A., Rebrin, K.: Closed-loop insulin delivery - the path to physiological glucose control. *Adv. Drug Deliv. Rev.* **56**(2), 125–144 (2004)
123. Stoma, S., Donzé, A., Bertaux, F., Maler, O., Batt, G.: STL-based analysis of TRAIL-induced apoptosis challenges the notion of type I/type II cell line classification. *PLoS Comput. Biol.* **9**(5), e1003056 (2013)
124. Ulus, D., Ferrère, T., Asarin, E., Maler, O.: Timed pattern matching. In: Legay, A., Bozga, M. (eds.) *FORMATS 2014*. LNCS, vol. 8711, pp. 222–236. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10512-3_16
125. Ulus, D., Ferrère, T., Asarin, E., Maler, O.: Online timed pattern matching using derivatives. In: Chechik, M., Raskin, J.-F. (eds.) *TACAS 2016*. LNCS, vol. 9636, pp. 736–751. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49674-9_47
126. Vijayaraghavan, S., Ramanathan, M.: *A Practical Guide for SystemVerilog Assertions*. Springer, New York (2006). <https://doi.org/10.1007/b137011>
127. Watterson, C., Heffernan, D.: Runtime verification and monitoring of embedded systems. *IET Softw.* **1**(5), 172–179 (2007)
128. Weinzimer, S., Steil, G., Swan, K., Dziura, J., Kurtz, N., Tamborlane, W.: Fully automated closed-loop insulin delivery versus semiautomated hybrid control in pediatric patients with type 1 diabetes using an artificial pancreas. *Diabetes Care* **31**, 934–939 (2008)
129. Xiaoqing, J., Donzé, A., Deshmukh, J.V., Seshia, S.A.: Mining requirements from closed-loop control models. In: *Proceedings of HSCC 2013: The ACM International Conference on Hybrid Systems: Computation and Control*, pp. 43–52. ACM (2013)
130. Yaghoubi, S., Fainekos, G.: Hybrid approximate gradient and stochastic descent for falsification of nonlinear systems. In: *Proceedings of ACC 2017: The 2017 American Control Conference*, pp. 529–534. IEEE (2017)
131. Yamaguchi, T., Kaga, T., Donzé, A., Seshia, S.A.: Combining requirement mining, software model checking, and simulation-based verification for industrial automotive systems. In: *Proceedings of FMCAD 2016: The 16th International Conference on Formal Methods in Computer-Aided Design*, pp. 201–204 (2016)
132. Yang, H., Hoxha, B., Fainekos, G.: Querying parametric temporal logic properties on embedded systems. In: Nielsen, B., Weise, C. (eds.) *ICTSS 2012*. LNCS, vol. 7641, pp. 136–151. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34691-0_11