

DisTL: A Temporal Logic for the Analysis of the Expected Behaviour of Cyber-Physical Systems

Valentina Castiglioni¹, Michele Loreti², and Simone Tini³

¹ Reykjavik University, Reykjavik, Iceland

² University of Camerino, Camerino, Italy

³ University of Insubria, Como, Italy

Abstract. The behaviour of systems characterised by a closed interaction of software components with the environment is inevitably subject to uncertainties. We propose a general framework for the specification and verification of requirements on the behaviour of these systems. We introduce the *Distribution Temporal Logic (DisTL)*, a novel temporal logic allowing us to specify properties on the *expected* behaviour of systems, and to include the presence of uncertainties in the specification. We equip DisTL with a robustness semantics and we prove it sound and complete w.r.t. the semantics induced by the *evolution metric*, i.e., a hemimetric expressing how well a system is fulfilling its tasks with respect to another one. Finally, we give a *statistical model checking algorithm* for DisTL specifications, and we apply our framework to a simple unmanned ground vehicle scenario.

1 Introduction

We have recently proposed a formal framework to model and analyse the behaviour of systems that are subject to *uncertainty*. The most prominent example is that of cyber-physical systems [24] (CPSs), in which software components, or *agents*, must interact with a highly changing and, even, unpredictable environment. To reason on these systems, we have introduced, in [7], the *evolution sequence model*: the behaviour of the system is modelled in terms of the modifications that the interaction of the agents with the environment induce on a set of application-relevant data, called *data state*. As those modifications are subject to uncertainty, induced by the environment and system's approximations, we model them as probability measures, henceforth simply called distributions, on the attainable data states. Hence, the evolution sequence of a system is the sequence of distributions over the data states obtained at each step, and can also be seen as the discrete-time version of the cylinder of all possible trajectories of the system, that takes into account the effects of uncertainty at each step. We have also provided the notion of *evolution metric* between evolution sequences, which allows us to quantify the *behavioural distance* [1, 6, 10] between systems.

Then, we have introduced, in [8], a novel temporal logic, called Robustness Temporal Logic (RobTL), allowing us to specify temporal requirements on the evolution of distances between the nominal behaviour of a system and its *perturbed* version. In particular, we can use RobTL formulae to specify robustness properties [14, 27, 29, 30] against uncertainties of the agents in the system. This is made possible by using atomic propositions of the form $\Delta(\text{exp}, p) \bowtie \eta$, to compare a threshold η with the distance,

specified by an expression exp , between a given evolution sequence and its perturbed version, obtained by applying a perturbation specified by p , starting from a given time step. Then, we combine atomic propositions with classic Boolean and temporal operators, in order to extend these evaluations to the entire evolution sequences.

The expressive power of RobTL comes at a price: besides the behaviour of the agents and the environment, we must be able to specify the perturbation that affects the system, in order to measure its robustness. While our tool STARK, the *Software Tool for the Analysis of Robustness in the unKnown environment* [9] available at <https://github.com/quasylab/jspear>, offers a domain specific language allowing us to do so, it might be the case that we do not know which data are manipulated by a perturbation, nor when and how such manipulations occur. Indeed, it might also be the case that we do not have access to a full specification of the system, but only to a collection of observations on data, following the deployment of the agents in the real world. Although we can still use STARK to perform some comparisons between the *observed* evolution sequence, obtained from the collected observations, and an *ideal* evolution sequence, in this case we cannot use RobTL to specify robustness properties of the system. Our aim, with this paper, is to provide an alternative approach to the study of systems robustness, in order to fill this gap.

Our Contribution: The Distribution Temporal Logic Let us consider a simple scenario: an unmanned ground vehicle is proceeding on a straight path towards a toll booth (henceforth, the objective), where it has to stop to allow the passenger to retrieve the entry ticket to the motorway. With classic probabilistic temporal logics (like, e.g., PCTL [15], CSL [2, 3], probabilistic variants of LTL [23], and probabilistic variants [28, 32] of MTL [18] and STL [20]), we can verify whether the vehicle is going to stop within a certain distance from the objective, and/or whether the probability to do so is above/below a desired threshold. We remark that this is achieved by reasoning in a trace-by-trace fashion: first we check the property over each trajectory of the system, and then we sum up the probability weights of those satisfying it.

Here we want to perform a different analysis: due to the presence of uncertainties, it would be preposterous to require the vehicle to stop precisely at the objective. Our aim is to express that the vehicle is *expected* to stop there, and to allow a certain *variance* on its actual final position, to take uncertainties into account, while guaranteeing that hazardous behaviours are avoided. Technically speaking, we want to specify that, when considering *all possible system behaviours*, the final (i.e., stationary) position of the vehicle agrees with a *desired distribution*, like, e.g., a Gaussian centred over the objective.

To this end, we introduce the *Distribution Temporal Logic (DisTL)*, a novel temporal logic allowing us to express requirements on the *expected* behaviour of the system in the presence of uncertainties and perturbations, by using distributions over data states as atomic propositions. We equip DisTL with a real-valued semantics expressing the *robustness* of the satisfaction of DisTL specifications. The *robustness* of a system s with respect to a formula φ is expressed as a real number $\llbracket \varphi \rrbracket_s \in [-1, 1]$: if it is positive, s satisfies φ . In detail, $\llbracket \varphi \rrbracket_s$ describes how much the behaviour of s has to be modified in order to violate (or satisfy) φ . We can then interpret $\llbracket \varphi \rrbracket_s$ as an indicator of *how well* s behaves with respect to the requirement φ . To formalise “*how well*”, we use the

evolution metric of [7], a (time-dependent) hemimetric on the evolution sequences of systems based on a *hemimetric on data states* and the *Wasserstein metric* [34]. The reason to opt for a hemimetric, instead of a more standard (pseudo)metric, is that it allows us to compare the *relative behaviour* of two systems and thus to express whether one system is *better than* the other. We use the evolution metric to define the robustness of systems with respect to DisTL formulae. As atomic propositions are distributions over data states, by means of the evolution metric we can compare them to the distributions in the evolution sequences of systems. In this way, we obtain useful information on the differences in the behaviour of two systems from the comparison of their robustness. In particular, we prove the robustness to be *sound* and *complete* with respect to our metric semantics: whenever the robustness of s_1 with respect to a formula φ is greater than the distance between s_1 and s_2 , then we can conclude that the robustness of s_2 with respect to φ is positive. We have also implemented in STARK a *statistical algorithm* for the evaluation of systems robustness with respect to DisTL specifications.

In order to show how our techniques can be applied, we consider the unmanned ground vehicle scenario described above as a case study.

2 The Evolution Sequence Model

Evolution Sequences We consider systems consisting of a set of *agents* and an *environment*, whose interaction produces changes on a shared *data space* \mathcal{D} , containing the values assumed by *variables*, representing: (i) physical quantities, (ii) sensors, (iii) actuators, and (iv) internal variables of the agents. Technically, we assume a *finite* set of *variables* Var such that for each $x \in \text{Var}$ the domain $\mathcal{D}_x \subseteq \mathbb{R}$ is either *finite*, or a *compact* subset of \mathbb{R} . Notice that, in particular, this means that \mathcal{D}_x is a Polish space. Moreover, as a σ -algebra over \mathcal{D}_x we assume the Borel σ -algebra, denoted \mathcal{B}_x . The data space \mathcal{D} over Var is then defined as the Cartesian product over the variables domains $\mathcal{D} = \prod_{x \in \text{Var}} \mathcal{D}_x$, and it is equipped with the product σ -algebra $\mathcal{B}_{\mathcal{D}} = \bigotimes_{x \in \text{Var}} \mathcal{B}_x$ [4].

We call *data state* the current state of the data space, and represent it by a mapping $\mathbf{d}: \text{Var} \rightarrow \mathbb{R}$, with $\mathbf{d}(x) \in \mathcal{D}_x$ for all $x \in \text{Var}$. At each step, the agents and the environment induce some changes on the data state, providing a new data state at the next step. Those modifications are also subject to the presence of uncertainties, meaning that it is not always possible to determine exactly the values assumed by data at the next step. Hence, following [7], we model the changes induced at each step as a distribution on the attainable data states. The behaviour of the system is then expressed by its *evolution sequence*, i.e., the sequence of distributions over the data states obtained at each step. In other words, the evolution sequence is the discrete-time version of the cylinder of *all possible trajectories* of the system. In this paper, we do not focus on how evolution sequences are generated: we simply assume a Markov kernel governing the evolution of the system, and the evolution sequence is the *Markov process* generated by it.

Definition 1 (Evolution sequence, [7]). *Given a data space \mathcal{D} , let $\Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ be the set of distributions over the space $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$. Let $\text{step}: \mathcal{D} \rightarrow \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ be the Markov kernel generating the behaviour of a system s having μ as initial distribution. Then, the evolution sequence of s is a countable sequence $\mathcal{S}_{\mu} = \mathcal{S}_{\mu}^0, \mathcal{S}_{\mu}^1, \dots$ of distributions in*

$\Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ such that, for all $\mathbb{D} \in \mathcal{B}_{\mathcal{D}}$:

$$\mathcal{S}_{\mu}^0(\mathbb{D}) = \mu(\mathbb{D}) \quad \mathcal{S}_{\mu}^{i+1}(\mathbb{D}) = \int_{\mathcal{D}} \text{step}(\mathbf{d})(\mathbb{D}) \mathfrak{d}\mathcal{S}_{\mu}^i(\mathbf{d}).$$

We denote by $\mathbf{S}_{\mathcal{D}}$ the set of all possible evolution sequences over \mathcal{D} .

A Distance on Evolution Sequences We introduce a distance measuring the differences in the behaviour of systems that will be used to define the robustness of DisTL specifications. The idea is first to introduce a *distance on distributions over data states* measuring their differences with respect to a *given target*, and then to extend it to the evolution sequences. Following [7], to capture the tasks of the system, we use *penalty functions* $\rho: \mathcal{D} \times \mathbb{N} \rightarrow [0, 1]$ i.e., functions that assign to each data state \mathbf{d} and time step τ a penalty in $[0, 1]$ expressing how far the values of the parameters related to the considered task in \mathbf{d} are from their desired ones at time τ . For brevity, we denote by ρ_{τ} the mapping corresponding to the τ -th element in the list ρ , i.e., $\rho_{\tau}(\mathbf{d}) = \rho(\mathbf{d}, \tau)$.

We have implemented a simple language, PF, to model penalties in STARK:

Definition 2 (Penalties). *Penalties in PF are defined as follows:*

$$\text{pf} ::= \mathbf{0} \mid \mathfrak{f}@_{\tau} \mid \text{pf}_1; \text{pf}_2 \mid \text{pf}^n$$

where pf ranges over PF, n and τ are finite natural numbers, and:

- $\mathbf{0}$ is the null penalty, i.e., at each time step it assigns penalty 0 to any data state;
- $\mathfrak{f}@_{\tau}$ is an atomic penalty, i.e., a function $\mathfrak{f}: \mathcal{D} \rightarrow [0, 1]$ that is applied after τ time steps from the current instant;
- $\text{pf}_1; \text{pf}_2$ is a sequential penalty, i.e., penalty pf_2 is applied at the time step subsequent to the (final) application of pf_1 ;
- pf^n is an iterated penalty, i.e., penalty pf is applied for a total of n times.

This simple language allows us to define some non-trivial penalties that we can use to analyse systems behaviour, like in the following example designed for our case study:

Example 1. As the task of the vehicle is to stop at the objective, it is natural to use the normalisation of its distance from it, stored in variable `p.dist`, as a penalty. However, we notice that since the distance changes in time, in order to have a meaningful penalty value we also need to change the normalisation factor. In fact, if, for instance, we normalise only with respect to the initial distance, the penalty is fated to decrease in time, and thus we may lose important information on the behaviour of the vehicle when it gets closer to the objective. Therefore, assuming that the vehicle is at an initial distance from the objective of 10000 m, we consider the following penalty:

$$\rho_{\text{pos}} = (\mathfrak{f}_{10000}@0)^{100}; (\mathfrak{f}_{7000}@0)^{100}; (\mathfrak{f}_{2500}@0)^{75}; (\mathfrak{f}_{10}@0)^{76} \quad \mathfrak{f}_x(\mathbf{d}) = \frac{\mathbf{d}(\text{p.dist})}{x}.$$

A mapping $\langle \cdot \rangle$ from PF to penalty functions can be defined on the lines of the one proposed in [8] to assign a semantics to perturbation functions.

Proposition 1. For each $\text{pf} \in \text{PF}$, the mapping $\langle \text{pf} \rangle$ is a penalty function.

Then we use penalty functions to obtain a *distance on data states*:

Definition 3 (Metric on data states). Let $\rho: \mathcal{D} \times \mathbb{N} \rightarrow [0, 1]$ be a penalty function on \mathcal{D} , and let $\tau \in \mathbb{N}$ be a time step. The metric on data states in \mathcal{D} , $\mathfrak{m}_\tau^\rho: \mathcal{D} \times \mathcal{D} \rightarrow [0, 1]$, is defined, for all $\mathbf{d}_1, \mathbf{d}_2 \in \mathcal{D}$, by: $\mathfrak{m}_\tau^\rho(\mathbf{d}_1, \mathbf{d}_2) = \max\{\rho_\tau(\mathbf{d}_2) - \rho_\tau(\mathbf{d}_1), 0\}$.

Note that $\mathfrak{m}_\tau^\rho(\mathbf{d}_1, \mathbf{d}_2)$ is a hemimetric expressing how much \mathbf{d}_2 is *worse* than \mathbf{d}_1 according to ρ_τ . Then, we need to lift the hemimetric \mathfrak{m}_τ^ρ to a hemimetric over $\Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$. To this end, we make use of the *Wasserstein lifting* [34]: for any two distributions μ, ν on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, the Wasserstein lifting of \mathfrak{m}_τ^ρ to a distance between μ and ν is defined by

$$\mathbf{W}(\mathfrak{m}_\tau^\rho)(\mu, \nu) = \inf_{\mathfrak{w} \in \mathfrak{W}(\mu, \nu)} \int_{\mathcal{D} \times \mathcal{D}} \mathfrak{m}_\tau^\rho(\mathbf{d}, \mathbf{d}') \, \mathfrak{w}(\mathbf{d}, \mathbf{d}')$$

where $\mathfrak{W}(\mu, \nu)$ is the set of the couplings of μ and ν , namely the set of joint distributions \mathfrak{w} over the product space $(\mathcal{D} \times \mathcal{D}, \mathcal{B}(\mathcal{D} \times \mathcal{D}))$ having μ and ν as left and right marginal, respectively, i.e., $\mathfrak{w}(\mathbb{D} \times \mathcal{D}) = \mu(\mathbb{D})$ and $\mathfrak{w}(\mathcal{D} \times \mathbb{D}) = \nu(\mathbb{D})$, for all $\mathbb{D} \in \mathcal{B}(\mathcal{D})$. (See [7, 13] for a discussion on the definition of the Wasserstein lifting over hemimetrics.)

The evolution hemimetric of [7] is then obtained as the *infinity norm* of the tuple of the Wasserstein distances between the distributions in the evolution sequences. Since in most applications the changes on data induced by the system can be appreciated only along wider time intervals than a computation step by the agents, we consider a *discrete, finite* set OT of time steps at which the modifications on data give us useful information on the evolution of the system.

Definition 4 (Evolution metric). Assume a finite set OT of time steps, a penalty function ρ and the metrics on data states \mathfrak{m}_τ^ρ . Then, the evolution metric over ρ and OT, is the mapping $\mathfrak{E}\mathfrak{m}_{\text{OT}}^\rho: \mathbf{S}_{\mathcal{D}} \times \mathbf{S}_{\mathcal{D}} \rightarrow [0, 1]$ defined, for all $\mathcal{S}_1, \mathcal{S}_2$, by

$$\mathfrak{E}\mathfrak{m}_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2) = \max_{\tau \in \text{OT}} \mathbf{W}(\mathfrak{m}_\tau^\rho)(\mathcal{S}_{1,\tau}, \mathcal{S}_{2,\tau}).$$

We remark that since we are interested in verifying requirements over a finite horizon (and we use only *bounded* temporal operators in the logic), the choice of having OT finite is not too restrictive.

Case Study: Unmanned Ground Vehicle As a running example, we consider the unmanned ground vehicle scenario described in the Introduction. Since our objective is merely to showcase the features of the new logic, we work under the following simplifying assumptions: 1. All the objects on the scene, including the vehicle, are one-dimensional; 2. When we start the simulation, the (program controlling the) vehicle already knows its distance from the point at which it has to stop (i.e., its distance from the button on the toll booth, henceforth referred to as *the objective*); 3. The acceleration of the vehicle can assume two values: a positive one $A \text{ m/s}^2$, and a negative one $-B \text{ m/s}^2$ (with $B > 0$); 4. The vehicle is equipped with a speed sensor, s_speed , that is subject to uncertainty (related to instrument accuracy); 5. The speed can never be negative (i.e., we do not allow the vehicle to shift into reverse). Every TIMER steps, the vehicle

decides whether to accelerate or brake, according to the sensor readings and its distance from the objective.

To specify the details of the behaviour of the vehicle we use our tool STARK. The vehicle is modelled as a STARK component, consisting of a set of local variables, used to allocate sensor readings and the value of the acceleration actuator, and a STARK controller, i.e., the process rendering the behaviour of the vehicle. Due to space limitations, we give only an informal presentation. The interested reader can find the full STARK specification at <https://github.com/quasylab/jspear/tree/Tony>. The behaviour of the controller is specified by means of four processes, or *states*: Ctrl, Accelerate, Decelerate and Stop. The computation starts from Ctrl that checks, every TIMER steps, whether the vehicle can accelerate or if it has to brake, and sets the acceleration actuator accordingly. The decision is taken on the basis of the sensed speed and the distance from the objective. States Accelerate and Decelerate manage, respectively, the acceleration and braking phases: the vehicle maintains a constant acceleration (of $A \text{ m/s}^2$ in the case of Accelerate, and of $-B \text{ m/s}^2$ for Decelerate), for TIMER steps; then Ctrl is woken up for a new check. When the speed becomes zero, and it is not possible to get closer to the objective, process Stop sets the acceleration actuator to 0 m/s^2 , and the vehicle becomes stationary.

To model the evolution of the scenario we use a STARK environment, that allows us to set the position of the objective, and to model the movement of the vehicle towards it. In this simple scenario, the uncertainty in the model is given by the precision error in the readings of sensor `s.speed`. Hence, we include a random noise in the updates of that value made by the environment.

3 The Distribution Temporal Logic

We now introduce the *Distribution Temporal Logic* (DisTL) which allows us to specify requirements on the *expected* behaviour of the system, in the presence of uncertainties and perturbations. The logic bases on two atomic properties, $\mathbf{target}(\mu)_q^\rho$ and $\mathbf{brink}(\mu)_q^\rho$, where μ is a distribution over data states in $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, ρ is a penalty function and q is a real in $[0, 1]$. We will use $\mathbf{target}(\mu)_q^\rho$ to express a desirable behaviour, whereas $\mathbf{brink}(\mu)_q^\rho$ can be used for unwanted, or hazardous, behaviours. These formulae are evaluated over a evolution sequence \mathcal{S} and a time step τ . Let us analyse the formula $\mathbf{target}(\mu)_q^\rho$. To establish whether the system exhibits the desired behaviour, we compare the given distribution μ with the distribution \mathcal{S}_τ : our means of comparison is the Wasserstein lifting of the hemimetric between data states evaluated with respect to the penalty ρ . (Notice that ρ is a parameter of the formula $\mathbf{target}(\mu)_q^\rho$. This is due to the fact that the penalty is not a property of the system but part of the requirements imposed on its behaviour.) As μ is our target distribution, it is natural to check whether \mathcal{S}_τ is *worse than* μ , i.e., to evaluate the distance $\mathbf{W}(\mathfrak{m}_\tau^\rho)(\mu, \mathcal{S}_\tau)$. Given the presence of uncertainties, it would not be feasible to say that the system satisfies the considered formula if and only if $\mathbf{W}(\mathfrak{m}_\tau^\rho)(\mu, \mathcal{S}_\tau) = 0$. Instead, we use the parameter q as a tolerance on the distance: if \mathcal{S}_τ is such that $\mathbf{W}(\mathfrak{m}_\tau^\rho)(\mu, \mathcal{S}_\tau) \leq q$, then the behaviour of the system can be considered acceptable. In other words, q is the *maximal acceptable distance* between the desired behaviour μ and the current behaviour \mathcal{S}_τ .

Conversely, in the formula $\mathbf{brink}(\mu)_q^\rho$ the distribution μ expresses some *unwanted, hazardous*, behaviour. Hence, the distribution \mathcal{S}_τ reached by the system must be better than μ , i.e., $\mathbf{W}(\mathbf{m}_\tau^\rho)(\mathcal{S}_\tau, \mu) > 0$. Also in this case, due to the presence of uncertainties, we need to make use of a threshold parameter q : assuming a distribution \mathcal{S}_τ acceptable when it is only *slightly* better than μ can still lead to an unwanted behaviour (because, in this case, the difference between the two distributions may be due only to some noise). Hence, we let q be the *minimal required distance* between \mathcal{S}_τ and μ , so that \mathcal{S}_τ is an acceptable behaviour if and only if $\mathbf{W}(\mathbf{m}_\tau^\rho)(\mathcal{S}_\tau, \mu) \geq q$.

Let $\text{var}(\mu) \subseteq \text{Var}$ be the set of data variables over which the distribution μ is defined. Similarly, for a penalty function ρ , we can consider the set $\text{var}(\rho) \subseteq \text{Var}$.

Definition 5 (DisTL). *The modal logic DisTL consists in the set of formulae \mathcal{L} defined by the following syntax:*

$$\varphi ::= \top \mid \mathbf{target}(\mu)_q^\rho \mid \mathbf{brink}(\mu)_q^\rho \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi_1 \mathcal{U}^{[a,b]} \varphi_2$$

with φ ranging over \mathcal{L} , $\mu \in \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ a distribution over data states, ρ a penalty function such that $\text{var}(\rho) \subseteq \text{var}(\mu)$, $q \in [0, 1]$, and $[a, b]$ an interval in OT.

Disjunction and negation are the standard Boolean connectives, and $\varphi_1 \mathcal{U}^{[a,b]} \varphi_2$ is the *bounded until* operator stating that φ_1 is satisfied until, at a time in $[a, b]$, φ_2 is. As expected, other standard operators can be defined as macros in our logic:

$$\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2) \quad \diamond^{[a,b]} \varphi \equiv \top \mathcal{U}^{[a,b]} \varphi \quad \square^{[a,b]} \varphi \equiv \neg \diamond^{[a,b]} \neg\varphi.$$

Formulae are evaluated over evolution sequences and time steps. In a *quantitative semantics* approach, for a formula φ , an evolution sequence \mathcal{S} , and a time step τ , the value $\llbracket \varphi \rrbracket_{\mathcal{S}, \tau} \in [-1, 1]$ expresses the *robustness* of \mathcal{S} with respect to φ at time τ , i.e., how much the behaviour \mathcal{S}_τ can be modified, either while preserving the validity of property φ (if φ is already satisfied), or in order to obtain it.

Definition 6 (DisTL: quantitative semantics). *For any evolution sequence \mathcal{S} , time step τ , and DisTL formula φ , the robustness of \mathcal{S} with respect to φ at τ , notation $\llbracket \varphi \rrbracket_{\mathcal{S}, \tau} \in [-1, 1]$, is defined inductively with respect to the structure of φ as follows:*

$$\begin{aligned} \llbracket \top \rrbracket_{\mathcal{S}, \tau} &= 1 \\ \llbracket \mathbf{target}(\mu)_q^\rho \rrbracket_{\mathcal{S}, \tau} &= q - \mathbf{W}(\mathbf{m}_\tau^\rho)(\mu, \mathcal{S}_\tau) \\ \llbracket \mathbf{brink}(\mu)_q^\rho \rrbracket_{\mathcal{S}, \tau} &= \mathbf{W}(\mathbf{m}_\tau^\rho)(\mathcal{S}_\tau, \mu) - q \\ \llbracket \neg\varphi \rrbracket_{\mathcal{S}, \tau} &= -\llbracket \varphi \rrbracket_{\mathcal{S}, \tau} \\ \llbracket \varphi_1 \vee \varphi_2 \rrbracket_{\mathcal{S}, \tau} &= \max\{\llbracket \varphi_1 \rrbracket_{\mathcal{S}, \tau}, \llbracket \varphi_2 \rrbracket_{\mathcal{S}, \tau}\} \\ \llbracket \varphi_1 \mathcal{U}^{[a,b]} \varphi_2 \rrbracket_{\mathcal{S}, \tau} &= \max_{\tau' \in [\tau+a, \tau+b]} \min\left\{ \llbracket \varphi_2 \rrbracket_{\mathcal{S}, \tau'}, \min_{\tau'' \in [\tau+a, \tau')} \llbracket \varphi_1 \rrbracket_{\mathcal{S}, \tau''} \right\}. \end{aligned}$$

Intuitively, the value $\mathbf{W}(\mathbf{m}_\tau^\rho)(\mu, \mathcal{S}_\tau)$ quantifies the difference between the distribution \mathcal{S}_τ reached by the system at time τ and μ . Hence, on the one hand the robustness $\llbracket \mathbf{target}(\mu)_q^\rho \rrbracket_{\mathcal{S}, \tau}$ expresses whether the distribution in the evolution sequence of \mathcal{S} is within the maximal acceptable distance q from μ . On the other hand, it also expresses

how much \mathcal{S}_τ can be modified while guaranteeing that the behaviour of the system remains within the specified parameters. Clearly, the closer μ and \mathcal{S}_τ , the higher the robustness. Similarly, $\llbracket \mathbf{brink}(\mu)_q^\rho \rrbracket_{\mathcal{S}, \tau}$ quantifies the robustness with respect to μ (and ρ) in terms of how much \mathcal{S}_τ may *get close* to μ while keeping the minimal required distance q . Hence, the farther \mathcal{S}_τ and μ , the higher the robustness. The semantics of boolean connectives and bounded until is standard. Notice that due to the potential asymmetry of our distances, it is not true in general that $\mathbf{brink}(\mu)_q^\rho = \neg \mathbf{target}(\mu)_{1-q}^\rho$.

Example 2. Let us now use DisTL to formalise the requirement on the final position of the vehicle discussed at the beginning of this section: we want to express that it is distributed like a Gaussian centred on the objective and with variance σ^2 , for some σ . In the STARK implementation of the system, the position of the vehicle is modelled in terms of its physical distance from the objective: when variable `p.dist` equals 0, the position of the vehicle corresponds to the objective. Hence, we can use `p.dist` $\sim \mathcal{G}(0, \sigma^2)$, henceforth denoted μ_{pos} , as the target distribution over the position. Given the penalty ρ_{pos} defined in Example 1, and a desired tolerance ε , we can use the formula

$$\varphi_1 = \square^{[\tau_1, \tau_2]} \mathbf{target}(\mu_{\text{pos}})_\varepsilon^{\rho_{\text{pos}}}$$

to capture the requirement on the final position, where the time interval $[\tau_1, \tau_2]$ is chosen according to the systems parameters.

Clearly, we can use DisTL also to express strict requirements (i.e., without approximations and tolerances): for instance, we must require that the vehicle is stationary in the final position, i.e., that its speed equals 0. This can be done by means of a Dirac (or point) distribution $\delta_0(\text{p.speed})$, henceforth denoted μ_{sp} , where `p.speed` is the variable storing the value of the physical speed of the vehicle. Consider the penalty function $\rho_{\text{sp}}(\mathbf{d}) = (\mathbf{d}(\text{p.speed})/\text{MAX_SPEED}@0)^{\mathfrak{h}}$, where `MAX_SPEED` is the parameter storing the maximal speed of the vehicle. We can then use the atomic formula $\mathbf{target}(\mu_{\text{sp}})_0^{\rho_{\text{sp}}}$ to express that the speed of the vehicle *must* be 0 (notice the tolerance 0). Then, we can combine the two formulae to express that the vehicle is expected to stop in an ε -neighbourhood of the objective, within a time horizon \mathfrak{h} (determined according to the other parameters of the system, like `TIMER`, `A`, and `B`):

$$\varphi_2 = \diamond^{[0, \mathfrak{h}]} (\mathbf{target}(\mu_{\text{sp}})_0^{\rho_{\text{sp}}} \wedge \mathbf{target}(\mu_{\text{pos}})_\varepsilon^{\rho_{\text{pos}}}).$$

Soundness and Completeness of the Robustness Semantics We can show that DisTL characterises the distance between evolution sequences. More precisely, the quantitative semantics of DisTL induces a distance between evolution sequences that coincides with the symmetrisation of the hemimetric $\mathfrak{E}m_{\text{OT}}^\rho$ and is therefore a pseudometric. Clearly, since the evolution metric is defined in terms of a given penalty function ρ , it will be characterised by the distance over formulae in \mathcal{L}_ρ , which is the sub-class of \mathcal{L} with atomic propositions of the form $(\cdot)_{(\cdot)}^\rho$.

Definition 7 (DisTL distance). *The DisTL distance between $\mathcal{S}_1, \mathcal{S}_2 \in \mathbf{S}_{\mathcal{D}}$ with respect to a penalty function ρ and OT is defined as*

$$\mathcal{L}_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2) = \sup_{\varphi \in \mathcal{L}_\rho, \tau \in \text{OT}} |\llbracket \varphi \rrbracket_{\mathcal{S}_1, \tau} - \llbracket \varphi \rrbracket_{\mathcal{S}_2, \tau}|.$$

Firstly, we show that the symmetrisation of $\mathfrak{E}m_{\text{OT}}^\rho$ is an upper bound to $\mathfrak{L}_{\text{OT}}^\rho$.

Lemma 1. *For any penalty function ρ , and $\mathcal{S}_1, \mathcal{S}_2 \in \mathbf{S}_{\mathcal{D}}$ we have:*

$$\mathfrak{L}_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2) \leq \max \{ \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2), \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_2, \mathcal{S}_1) \}.$$

Then, we show that for all evolution sequences $\mathcal{S}_1, \mathcal{S}_2$ there exists a formula φ in \mathcal{L}_ρ such that the symmetrisation of $\mathfrak{E}m_{\text{OT}}^\rho$ coincides with the difference in the evaluations of φ over $\mathcal{S}_1, \mathcal{S}_2$.

Lemma 2. *For all $\mathcal{S}_1, \mathcal{S}_2 \in \mathbf{S}_{\mathcal{D}}$ and penalty functions ρ , there is a formula $\varphi \in \mathcal{L}_\rho$ with $|\llbracket \varphi \rrbracket_{\mathcal{S}_1, \tau} - \llbracket \varphi \rrbracket_{\mathcal{S}_2, \tau}| = \max \{ \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2), \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_2, \mathcal{S}_1) \}$, for some $\tau \in \text{OT}$.*

From Lemma 1 and Lemma 2 we infer that the DisTL distance $\mathfrak{L}_{\text{OT}}^\rho$ and the symmetrisation of $\mathfrak{E}m_{\text{OT}}^\rho$ coincide.

Theorem 1. *For all evolution sequences \mathcal{S}_1 and \mathcal{S}_2 we have that:*

$$\mathfrak{L}_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2) = \max \{ \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2), \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_2, \mathcal{S}_1) \}.$$

Theorem 1 entails the *soundness* (Lemma 1) and *completeness* (Lemma 2) of our notion of robustness. In particular, as a direct consequence of Theorem 1, we can obtain the following classic result (see, e.g., [11]): whenever the robustness of a evolution sequence \mathcal{S} with respect to a formula φ is greater than the distance between \mathcal{S} and \mathcal{S}' , then the robustness of \mathcal{S}' with respect to φ is positive as well.

Corollary 1. *Let φ be any formula in \mathcal{L}_ρ , $\tau \in \text{OT}$ and let $i \in \{1, 2\}$. Whenever $\llbracket \varphi \rrbracket_{\mathcal{S}_i, \tau} \geq \max \{ \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_1, \mathcal{S}_2), \mathfrak{E}m_{\text{OT}}^\rho(\mathcal{S}_2, \mathcal{S}_1) \}$, then $\llbracket \varphi \rrbracket_{\mathcal{S}_{3-i}, \tau} \geq 0$.*

4 Statistical Model Checking

In this section we present an algorithm, based on statistical techniques and simulation, that allows us to estimate the robustness of a system \mathbf{s} with respect to a DisTL formula φ . This algorithm consists in three basic steps:

- (i) A randomised procedure, based on simulation, that allows us to estimate the evolution sequence of system \mathbf{s} , assuming an initial data state $\mathbf{d}_\mathbf{s}$: Starting from $\mathbf{d}_\mathbf{s}$ we sample N sequences of data states $\mathbf{d}_0^j, \dots, \mathbf{d}_k^j$, for $j = 1, \dots, N$; then all the data states collected at time i are used to estimate the distribution $\mathcal{S}_{\delta_{\mathbf{d}_\mathbf{s}}}^i$.
- (ii) A mechanism to estimate the Wasserstein distance between two (unknown) distributions μ and ν on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$, similar to the one presented in [31]: To estimate $\mathbf{W}(\mathbf{m}_i^\rho)(\mu, \nu)$ we use N independent samples $\{\mathbf{d}_1^1, \dots, \mathbf{d}_1^N\}$ taken from μ and ℓN independent samples $\{\mathbf{d}_2^1, \dots, \mathbf{d}_2^{\ell N}\}$ taken from ν .
- (iii) A procedure that computes the robustness by inspecting the syntax of φ and by using the first two components.

The proposed approach has been implemented in Java, as part of STARK, and is available at <https://github.com/quasylab/jspear/tree/Tony>. We omit the presentation of the first two steps (that have already been discussed at length in [7, 8]), and we give an overview of the third step. We limit ourselves to recall the following result, from [7], on the estimation of the Wasserstein distance:

```

1: function SAT( $\mathbf{d}_s, \tau, \varphi, \ell, N$ )
2:    $h \leftarrow \text{HORIZON}(\varphi)$ 
3:    $\bar{E} \leftarrow \text{ESTIMATE}(\mathbf{d}_s, h, N)$ 
4:   return EVAL( $\bar{E}, \tau, \varphi, \ell, N$ )
5: end function

```

Fig. 1: Function used to evaluate system robustness with respect to a formula.

Theorem 2 ([7]). Let $\mu, \nu \in \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ be unknown. Let $\{\mathbf{d}_1^1, \dots, \mathbf{d}_1^N\}$ be independent samples taken from μ , and $\{\mathbf{d}_2^1, \dots, \mathbf{d}_2^{\ell \cdot N}\}$ independent samples taken from ν . Let $\{\omega_j = \rho_i(\mathbf{d}_1^j)\}$ and $\{\nu_h = \rho_i(\mathbf{d}_2^h)\}$ be the ordered sequences obtained by applying the penalty ρ_i to the samples. Then, it holds, almost surely, that $\mathbf{W}(\mathfrak{m}_i^\rho)(\mu, \nu) = \lim_{N \rightarrow \infty} \frac{1}{\ell N} \sum_{h=1}^{\ell N} \max \left\{ \nu_h - \omega_{\lceil \frac{h}{\ell} \rceil}, 0 \right\}$.

Statistical Estimation of Robustness The computation of the robustness of a system \mathbf{s} with respect to a formula φ , at a given time step τ and starting from the data state \mathbf{d}_s , is performed via the function SAT given in Figure 1. Together with the data state \mathbf{d}_s , the step τ , and the formula φ , function SAT takes as parameters the two integers ℓ and N identifying the number of samplings that will be used to estimate the Wasserstein metric. This function consists of three steps. First the *time horizon* h of the formula φ is computed (by induction on the structure of φ) to identify the number of steps needed to evaluate the robustness. In the second step, function ESTIMATE is used to simulate the evolution sequence of \mathbf{s} from \mathbf{d}_s by collecting the sets of samplings $\bar{E} = E_0, \dots, E_h$. Then, in the third step, function EVAL, presented in Figure 2, is used for the evaluation of the robustness. The structure of EVAL is similar to the monitoring function for STL defined in [20]. Function EVAL is defined recursively on the syntax of φ . In the cases of the atomic formulae $\mathbf{target}(\mu)_q^\rho$ and $\mathbf{brink}(\mu)_q^\rho$, firstly we use function SAMPLE to obtain $\ell \cdot N$ independent samples of the distribution μ . Then, we use function WASS to compute the Wasserstein distance between the sampling of μ and the sampling E_τ of the distribution reached at step τ by \mathbf{s} .

The following theorem guarantees that when N goes to infinite, the robustness computed by function SAT converges, almost surely, to the exact value.

Theorem 3. For any formula φ , system \mathbf{s} , data state \mathbf{d}_s , time step τ , and integer $\ell > 0$

$$\lim_{N \rightarrow \infty} \text{SAT}(\mathbf{d}_s, \tau, \varphi, \ell, N) = \llbracket \varphi \rrbracket_{\mathcal{S}_{\mathbf{d}_s, \tau}}.$$

Proof. The proof follows by induction on the structure of the formula φ , using Theorem 2 to deal with the base cases of $\varphi = \mathbf{target}(\mu)_p^\rho$ and $\varphi = \mathbf{brink}(\mu)_p^\rho$.

Example 3. The procedure outlined above can be used for the evaluation of the requirements on the vehicle scenario presented in Example 2. Let \mathcal{V} be the evolution sequence of the system obtained from the following initial parameters: $\text{p_dist} = 10000$; $\text{p_speed} = 25.0$; $\text{MAX_SPEED} = 40.0$; $\text{A} = 0.25$; $\text{B} = 2.0$; $\text{TIMER} = 1.0$. In Figure 3a we report the evaluations, in \mathcal{V} and time step 0, of various instances of the formula φ_1 .

```

1: function EVAL( $\bar{E}, \tau, \varphi, \ell, N$ )
2:   match  $\varphi$ 
3:     with  $\top$  :
4:       return 1.0
5:     with target( $\mu$ ) $_q^\rho$  :
6:        $E' \leftarrow \text{SAMPLE}(\mu, \ell \cdot N)$ 
7:       return  $q - \text{WASS}(E', E_\tau, \rho)$ 
8:     with brink( $\mu$ ) $_p^\rho$  :
9:        $E' \leftarrow \text{SAMPLE}(\mu, \ell \cdot N)$ 
10:      return  $\text{WASS}(E_\tau, E', \rho) - q$ 
11:     with  $\neg\varphi_1$  :
12:       return  $-\text{EVAL}(\bar{E}, \tau, \varphi_1, \ell, N)$ 
13:     with  $\varphi_1 \vee \varphi_2$  :
14:       return  $\max\{\text{EVAL}(\bar{E}, \tau, \varphi_1, \ell, N), \text{EVAL}(\bar{E}, \tau, \varphi_2, \ell, N)\}$ 
15:     with  $\varphi_1 \mathcal{U}^{[a,b]} \varphi_2$  :
16:        $res \leftarrow 0.0$ 
17:       for  $i \in [\tau + a, \tau + b]$  do
18:          $res_1 \leftarrow 0.0$ 
19:         for  $j \in [\tau + a, i]$  do
20:            $res_1 \leftarrow \min\{res_1, \text{EVAL}(\bar{E}, j, \varphi_1, \ell, N)\}$ 
21:         end for
22:          $res \leftarrow \max\{res, \min\{res_1, \text{EVAL}(\bar{E}, i, \varphi_2, \ell, N)\}\}$ 
23:       end for
24:       return  $res$ 
25: end function

```

Fig. 2: Evaluation of robustness.

We consider a different instance for each $\tau_1 \in [250, 280]$, while we fix the other parameters: $\tau_2 = 350$; $N = 100$ and $\ell = 10$. For each instance, we consider three variations, according to the threshold of the atomic formula: $\varepsilon \in \{0.1, 0.3, 0.5\}$.

Finally, we remark that $\llbracket \varphi_2 \rrbracket_{\mathcal{V}, \tau} = 0.0$ for all $\tau \in [0, 350]$, as shown in Figure 3b. In terms of robustness semantics, an evaluation to 0.0 is usually considered as non-informative, as it gives us no information on how much the behaviour of the system can be modified in order to validate (or invalidate) the considered property. However, in this specific case, 0.0 is the best value that we can obtain. In fact, the formula φ_2 contains a strict requirement on the speed, which is required to be exactly 0.0. Hence, $\llbracket \varphi_2 \rrbracket_{\mathcal{V}, \tau} = 0.0$ means that this requirement is actually met and, at the same time, even the tiniest modification in the behaviour might cause the formula to be invalidated.

5 Concluding Remarks

Differently from the other probabilistic temporal logics usually considered in the literature, DisTL can be used to express the properties of the distributions expressing the transient and expected behaviour of the system. Up to our knowledge, [33] is the only




Fig. 3: Evaluations of various instances of formulae φ_1 and φ_2 .

other paper proposing to substitute probabilistic guarantees on the temporal properties with a richer description of the probabilistic events. In detail, [33] introduces *ProbSTL*, a stochastic variant of STL tailored to the incremental runtime verification of safe behaviour of robotic systems under uncertainties, based on a predictive stream reasoning tool: their stochastic signal is given by the prediction on the possible future trajectories of a system, taking uncertainties into account. Yet, ProbSTL specifications are tested only on the current trajectory of the system. This is the main difference with our work, since our logic has been built to express the overall uncertain behaviour of the system. This disparity is also a consequence of the different application context: off-line verification for us, runtime verification in [33]. However, as future work, we plan to develop a predictive model for the runtime monitoring of DisTL specifications. In particular, inspired by [5, 22] where deep neural networks are used as reachability predictors for predictive monitoring, we intend to integrate our work with learning techniques, to favour the computation and evaluation of the predictions.

In Markov processes as *transformers of distributions* [17, 19], state-to-state transition probabilities are interpreted as a single distribution over the state space. We remark that the state space in [17, 19] is finite and discrete, whereas our evolution sequences are defined in the continuous setting, which means that we are not introducing any limiting assumption on the behaviour of the environment. Moreover, the temporal logics used to model check properties of transformers of distributions, respectively *iLTL* in [19] and the *almost acyclic* Büchi automata in [17], have a boolean semantics, and are thus not comparable to DisTL, in which formulae are interpreted in terms of robustness.

A statistical model checking algorithm for PCTL specifications over Markov chains has been proposed in [35], using stratified sampling. This allows for the generation of negatively correlated samples, thus reducing the number of samples needed to obtain confident results at the price of restricting the form of the PCTL formulae to be checked. We plan to study the use of stratified sampling in our model checking algorithm.

We also plan to investigate the application of our framework to the analysis of biological systems. Some quantitative extensions of temporal logics have already been proposed in that setting (e.g. [12, 25, 26]) to capture the notion of robustness from [16] or similar proposals [21]. It would be interesting to see whether the use of DisTL and evolution sequences can lead to new results in this setting.

References

1. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the Future in Systems Theory. In: Proceedings of ICALP'03, pp. 1022–1037. ICALP '03, Springer (2003). https://doi.org/10.1007/3-540-45061-0_79
2. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.K.: Verifying continuous time Markov chains. In: Proceedings of CAV '96. Lecture Notes in Computer Science, vol. 1102, pp. 269–276 (1996). https://doi.org/10.1007/3-540-61474-5_75
3. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.K.: Model-checking continuous-time Markov chains. *ACM Trans. Comput. Log.* **1**(1), 162–170 (2000). <https://doi.org/10.1145/343369.343402>
4. Bogachev, V.I.: Measure Theory. No. v. 1 in Measure Theory, Springer-Verlag, Berlin/Heidelberg (2007). <https://doi.org/10.1007/978-3-540-34514-5>
5. Bortolussi, L., Cairoli, F., Paoletti, N., Smolka, S.A., Stoller, S.D.: Neural predictive monitoring. In: Proceedings of RV 2019. Lecture Notes in Computer Science, vol. 11757, pp. 129–147 (2019). https://doi.org/10.1007/978-3-030-32079-9_8
6. Castiglioni, V., Loreti, M., Tini, S.: The metric linear-time branching-time spectrum on nondeterministic probabilistic processes. *Theor. Comput. Sci.* **813**, 20–69 (2020). <https://doi.org/10.1016/j.tcs.2019.09.019>
7. Castiglioni, V., Loreti, M., Tini, S.: How adaptive and reliable is your program? In: Proceedings of FORTE 2021. Lecture Notes in Computer Science, vol. 12719, pp. 60–79 (2021). https://doi.org/10.1007/978-3-030-78089-0_4
8. Castiglioni, V., Loreti, M., Tini, S.: RobTL: A temporal logic for the robustness of cyber-physical systems. *CoRR abs/2212.11158* (2022). <https://doi.org/10.48550/arXiv.2212.11158>
9. Castiglioni, V., Loreti, M., Tini, S.: STARK: A software tool for the analysis of robustness in the unknown environment. In: Proceedings of COORDINATION 2023. Lecture Notes in Computer Science, vol. 13908 (2023). https://doi.org/10.1007/978-3-031-35361-1_6, to appear.
10. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theor. Comput. Sci.* **318**(3), 323–354 (2004). <https://doi.org/10.1016/j.tcs.2003.09.013>
11. Donzé, A., Maler, O.: Robust satisfaction of temporal logic over real-valued signals. In: Proceedings of FORMATS 2010. Lecture Notes in Computer Science, vol. 6246, pp. 92–106 (2010). https://doi.org/10.1007/978-3-642-15297-9_9
12. Fages, F., Rizk, A.: On temporal logic constraint solving for analyzing numerical data time series. *Theor. Comput. Sci.* **408**(1), 55–65 (2008). <https://doi.org/10.1016/j.tcs.2008.07.004>
13. Faugeras, O.P., Rüschemdorf, L.: Risk excess measures induced by hemi-metrics. *Probability, Uncertainty and Quantitative Risk* **3:6** (2018). <https://doi.org/10.1186/s41546-018-0032-0>
14. Fränzle, M., Kapinski, J., Prabhakar, P.: Robustness in cyber-physical systems. *Dagstuhl Reports* **6**(9), 29–45 (2016). <https://doi.org/10.4230/DagRep.6.9.29>
15. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Asp. Comput.* **6**(5), 512–535 (1994). <https://doi.org/10.1007/BF01211866>
16. Kitano, H.: Towards a theory of biological robustness. *Molecular Systems Biology* **3**(1), 137 (2007). <https://doi.org/https://doi.org/10.1038/msb4100179>
17. Korthikanti, V.A., Viswanathan, M., Agha, G., Kwon, Y.: Reasoning about mdps as transformers of probability distributions. In: Proceedings of QEST 2010. pp. 199–208. IEEE Computer Society (2010). <https://doi.org/10.1109/QEST.2010.35>
18. Koymans, R.: Specifying real-time properties with metric temporal logic. *Real Time Syst.* **2**(4), 255–299 (1990). <https://doi.org/10.1007/BF01995674>

19. Kwon, Y., Agha, G.: Linear inequality LTL (iltl): A model checker for discrete time markov chains. In: Proceedings of ICFEM 2004. Lecture Notes in Computer Science, vol. 3308, pp. 194–208 (2004). https://doi.org/10.1007/978-3-540-30482-1_21
20. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Proceedings of FORMATS and FTRTFT 2004. Lecture Notes in Computer Science, vol. 3253, pp. 152–166 (2004). https://doi.org/10.1007/978-3-540-30206-3_12
21. Nasti, L., Gori, R., Milazzo, P.: Formalizing a notion of concentration robustness for biochemical networks. In: Proceedings of STAF 2018. Lecture Notes in Computer Science, vol. 11176, pp. 81–97 (2018). https://doi.org/10.1007/978-3-030-04771-9_8
22. Phan, D., Paoletti, N., Zhang, T., Grosu, R., Smolka, S.A., Stoller, S.D.: Neural state classification for hybrid systems. In: Proceedings of ATVA 2018. Lecture Notes in Computer Science, vol. 11138, pp. 422–440 (2018). https://doi.org/10.1007/978-3-030-01090-4_25
23. Pnueli, A.: The temporal logic of programs. In: Proceedings of FOCS 1977. pp. 46–57. IEEE Computer Society (1977). <https://doi.org/10.1109/SFCS.1977.32>
24. Rajkumar, R., Lee, I., Sha, L., Stankovic, J.A.: Cyber-physical systems: the next computing revolution. In: DAC. pp. 731–736. ACM (2010)
25. Rizk, A., Batt, G., Fages, F., Soliman, S.: A general computational method for robustness analysis with applications to synthetic gene networks. *Bioinform.* **25**(12) (2009). <https://doi.org/10.1093/bioinformatics/btp200>
26. Rizk, A., Batt, G., Fages, F., Soliman, S.: Continuous valuations of temporal logic specifications with applications to parameter optimization and robustness measures. *Theor. Comput. Sci.* **412**(26), 2827–2839 (2011). <https://doi.org/10.1016/j.tcs.2010.05.008>
27. Runger, M., Tabuada, P.: A notion of robustness for cyber-physical systems. *IEEE Trans. Autom. Control.* **61**(8), 2108–2123 (2016)
28. Sadigh, D., Kapoor, A.: Safe control under uncertainty with probabilistic signal temporal logic. In: Proceedings of Robotics: Science and Systems XII 2016 (2016). <https://doi.org/10.15607/RSS.2016.XII.017>
29. Shahrokni, A., Feldt, R.: A systematic review of software robustness. *Information and Software Technology* **55**(1), 1–17 (2013). <https://doi.org/https://doi.org/10.1016/j.infsof.2012.06.002>
30. Sontag, E.D.: *Input to State Stability: Basic Concepts and Results*, pp. 163–220. Springer Berlin Heidelberg (2008). https://doi.org/10.1007/978-3-540-77653-6_3
31. Thorsley, D., Klavins, E.: Approximating stochastic biochemical processes with Wasserstein pseudometrics. *IET Syst. Biol.* **4**(3), 193–211 (2010). <https://doi.org/10.1049/iet-syb.2009.0039>
32. Tiger, M., Heintz, F.: Stream reasoning using temporal logic and predictive probabilistic state models. In: Proceedings of TIME 2016. pp. 196–205 (2016). <https://doi.org/10.1109/TIME.2016.28>
33. Tiger, M., Heintz, F.: Incremental reasoning in probabilistic signal temporal logic. *Int. J. Approx. Reason.* **119**, 325–352 (2020). <https://doi.org/10.1016/j.ijar.2020.01.009>
34. Vaserstein, L.N.: Markovian processes on countable space product describing large systems of automata. *Probl. Peredachi Inf.* **5**(3), 64–72 (1969)
35. Wang, Y., Roohi, N., West, M., Viswanathan, M., Dullerud, G.E.: Statistical verification of PCTL using antithetic and stratified samples. *Formal Methods Syst. Des.* **54**(2), 145–163 (2019). <https://doi.org/10.1007/s10703-019-00339-8>