



HAL
open science

How Adaptive and Reliable is Your Program?

Valentina Castiglioni, Michele Loreti, Simone Tini

► **To cite this version:**

Valentina Castiglioni, Michele Loreti, Simone Tini. How Adaptive and Reliable is Your Program?. 41th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE), Jun 2021, Valletta, Malta. pp.60-79, 10.1007/978-3-030-78089-0_4 . hal-03740265

HAL Id: hal-03740265

<https://inria.hal.science/hal-03740265>

Submitted on 29 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

How Adaptive and Reliable is your Program?*

Valentina Castiglioni¹[0000-0002-8112-6523], Michele Loreti²[0000-0003-3061-863X],
and Simone Tini³[0000-0002-3991-5123]

¹ Reykjavik University, Reykjavik, Iceland, valentinac@ru.is

² University of Camerino, Camerino, Italy, michele.loreti@unicam.it

³ University of Insubria, Como, Italy, simone.tini@uninsubria.it

Abstract. We consider the problem of *modelling* and *verifying* the behaviour of systems characterised by a close interaction of a *program* with the *environment*. We propose to model the program-environment interplay in terms of the probabilistic modifications they induce on a set of application-relevant data, called *data space*. The behaviour of a system is thus identified with the probabilistic evolution of the initial data space. Then, we introduce a metric, called *evolution metric*, measuring the differences in the *evolution sequences* of systems and that can be used for system verification as it allows for expressing how well the program is fulfilling its tasks. We use the metric to express the properties of *adaptability* and *reliability* of a program, which allow us to identify potential critical issues of it w.r.t. changes in the initial environmental conditions. We also propose an *algorithm*, based on statistical inference, for the evaluation of the evolution metric.

1 Introduction

With the ever-increasing complexity of the digital world and diffusion of IoT systems, cyber-physical systems, and smart devices, we are witnessing the rise of software applications, henceforth *programs*, that must be able to deal with highly changing operational conditions, henceforth *environment*. Examples of such programs are the software components of unmanned vehicles, (on-line) service applications, the devices in a smart house, etc, which have to interact with other programs and heterogeneous devices, and with physical phenomena like wind, temperature, etc. Henceforth, we use the term *system* to denote the combination of the environment and the program acting on it. Hence, the *behaviour* of a system is the result of the program-environment interplay.

The main challenge in the analysis and verification of these systems is then the dynamical and, sometimes, unpredictable behaviour of the environment. The highly dynamical behaviour of physical processes can only be approximated in order to become computationally tractable and it can constitute a safety hazard for the devices in the system (like, e.g., an unexpected gust of wind for a drone that is autonomously setting its trajectory to avoid obstacles); some devices or programs may appear, disappear, or become temporarily unavailable; faults or conflicts may occur (like, e.g., in a smart home the program responsible for the ventilation of a room may open a window causing a conflict with the program that has to limit the noise level); sensors may introduce

* This work has been partially supported by the IRF project “OPEL” (grant No. 196050-051) and by the PRIN project “IT-MaTTerS” (grant No. 2017FTXR7S).

some measurement errors; etc. The introduction of *uncertainties* and *approximations* in these systems is therefore inevitable.

In the literature, we can find a wealth of proposals of stochastic and probabilistic models, as, e.g., *Stochastic Hybrid Systems* [8,28] and *Markov Decision Processes* [32], and *ad hoc* solutions for specific application contexts, as, e.g., establishing safety guarantees for drones flying under particular circumstances [26,41]. Yet, in these studies, either the environment is not explicitly taken into account or it is modelled only deterministically. In addition to that, due to the variety of applications and heterogeneity of systems, no general formal framework to deal with these challenges has been proposed so far. The lack of concise abstractions and of an automatic support makes the analysis and verification of the considered systems difficult, laborious, and error prone.

Our contribution. With this paper we aim at taking a first step towards a solution of the above-mentioned challenges, with a special focus on *verification*, by developing the tools for the verification of the ability of programs to *adjust* their behaviour to the unpredictable environment. Formally, we introduce two *measures* allowing us to assess how well a given program can perform under perturbations in the environmental conditions. We call these two measures *adaptability* and *reliability*. As an example, consider a drone that is autonomously flying along a given trajectory. In this setting, a perturbation can be given by a gust of wind that moves the drone out of its trajectory. We will say that the program controlling the drone is *adaptable* if it can retrieve the initial trajectory within a suitable amount of time. In other words, we say that a program is adaptable if no matter how much its behaviour is affected by the perturbations, it is able to react to them and regain its intended behaviour within a given amount of time. On the other hand, it may be the case that the drone is able to detect the presence of a gust of wind and can oppose to it, being only slightly moved from its initial trajectory. In this case, we say that the program controlling the drone is *reliable*. Hence, *reliability* expresses the ability of a program to maintain its intended behaviour (up-to some reasonable tolerance) despite the presence of perturbations in the environment.

In order to measure the adaptability and reliability of a program, we need to be able to express *how well* it is fulfilling its tasks. The systems that we are considering are strongly characterised by a quantitative behaviour, given by both the presence of uncertainties and the data used by program and environment. It seems then natural, and reasonable, to quantify the differences in the behaviour of systems by means of a metric over data. However, in order to informally discuss our proposal of a *metric semantics* for the kind of systems that we are considering, we need first to explain how the behaviour of these systems is defined, namely to introduce our general formal model for them. Our idea is to favour the modelling of the program-environment interplay over precise specifications of the operational behaviour of a program. In the last decade, many researchers have focused their studies on formal models capturing both the *qualitative* and *quantitative* behaviour of systems: Probabilistic Automata [34], Stochastic Process Algebras [5, 16, 27], Labelled Markov Chains and Stochastic Hybrid Systems [8, 28]. A common feature of these models is that the (quantitative, labelled) transitions expressing the computation steps directly model the behaviour of the system as a whole.

In this paper we take a *different point of view*: we propose to model the behaviour of program and environment separately, and then explicitly represent their *interaction* in

a purely *data-driven* fashion. In fact, while the environmental conditions are (partially) available to the program as a set of data, allowing it to adjust its behaviour to the current situation, the program is also able to use data to (partially) control the environment and fulfil its tasks. It is then natural to model the program-environment interplay in terms of the changes they induce on a set of application-relevant data, henceforth called *data space*. This feature will allow for a significant simplification in modelling the behaviour of the program, which can be *isolated* from that of the environment. Moreover, as common to favour *computational tractability* [1, 2], we adopt a *discrete time* approach.

We can then study the behaviour of the system as a whole by analysing how data evolve in time. In our model, a *system* consists in *three distinct components*: 1. a *process* P describing the behaviour of the program, 2. a *data state* \mathbf{d} describing the current state of the data space, and 3. an *environment evolution* \mathcal{E} describing the effect of the environment on \mathbf{d} . As we focus on the interaction with the environment, we abstract from the internal computation of the program and model only its activity on \mathbf{d} . At each step, a process can *read/update* values in \mathbf{d} and \mathcal{E} applies on the resulting data state, providing a new data state at the next step. To deal with the uncertainties, we introduce *probability* at two levels: (i) we use the *discrete generative probabilistic model* [24] to define processes, and (ii) \mathcal{E} induces a *continuous distribution* over data states. The behaviour of the system is then entirely expressed by its *evolution sequence*, i.e., the sequence of distributions over data states obtained at each step. Given the novelties of our model, as a side contribution we show that this behaviour defines a *Markov process*.

It is now reasonable to define our *metric semantics* in terms of a (time-dependent) distance on the evolution sequences of systems, which we call the *evolution metric*. The *evolution metric* will allow us to: 1. verify how well a program is fulfilling its tasks by comparing it with its specification, 2. compare the activity of different programs in the same environment, 3. compare the behaviour of one program w.r.t. different environments and changes in the initial conditions. The third feature will allow us to measure the *adaptability* and *reliability* of programs. The evolution metric will consist of two components: a *metric on data states* and the *Wasserstein metric* [39]. The former is defined in terms of a (time-dependent) *penalty function* allowing us to compare two data states only on the base of the objectives of the program. The latter lifts the metric on data states to a metric on distributions on data states. We then obtain a metric on evolution sequences by considering the maximal of the Wasserstein distances over time. We provide an *algorithm* for the estimation of the evolution sequences of systems and thus for the evaluation of the evolution metric. Following [37], the Wasserstein metric is evaluated in time $O(N \log N)$, where N is the (maximum) number of samples. We already adopted this approach in [11] in the context of finite-states self-organising collective systems, without any notion of environment or data space.

As an example of application, we use our framework to model a simple smart-room scenario. We consider two programs: the thermostat of a heating system, and an air quality controller. The former has to keep the room temperature within a desired comfort interval. The latter has to keep the quality of the air above a given threshold. We use our algorithm to evaluate the differences between two systems having the same programs but starting from different initial conditions. Finally, we apply it to measure the adaptability and reliability of the considered programs.

2 Background

Measurable spaces. A σ -algebra over a set Ω is a family Σ of subsets of Ω s.t. $\Omega \in \Sigma$, and Σ is closed under complementation and under countable union. The pair (Ω, Σ) is called a *measurable space* and the sets in Σ are called *measurable sets*, ranged over by $\mathbb{A}, \mathbb{B}, \dots$. For an arbitrary family Φ of subsets of Ω , the σ -algebra *generated* by Φ is the smallest σ -algebra over Ω containing Φ . In particular, we recall that given a topology T over Ω , the *Borel σ -algebra* over Ω , denoted $\mathcal{B}(\Omega)$, is the σ -algebra generated by the open sets in T . Given two measurable spaces (Ω_i, Σ_i) , $i = 1, 2$, the *product σ -algebra* $\Sigma_1 \times \Sigma_2$ is the σ -algebra on $\Omega_1 \times \Omega_2$ generated by the sets $\{\mathbb{A}_1 \times \mathbb{A}_2 \mid \mathbb{A}_i \in \Sigma_i\}$.

Given measurable spaces $(\Omega_1, \Sigma_1), (\Omega_2, \Sigma_2)$, a function $f: \Omega_1 \rightarrow \Omega_2$ is said to be Σ_1 -*measurable* if $f^{-1}(\mathbb{A}_2) \in \Sigma_1$ for all $\mathbb{A}_2 \in \Sigma_2$, with $f^{-1}(\mathbb{A}_2) = \{\omega \in \Omega_1 \mid f(\omega) \in \mathbb{A}_2\}$.

Probability spaces. A *probability measure* on a measurable space (Ω, Σ) is a function $\mu: \Sigma \rightarrow [0, 1]$ such that: i) $\mu(\Omega) = 1$, ii) $\mu(\mathbb{A}) \geq 0$ for all $\mathbb{A} \in \Sigma$, and iii) $\mu(\bigcup_{i \in I} \mathbb{A}_i) = \sum_{i \in I} \mu(\mathbb{A}_i)$ for every countable family of pairwise disjoint measurable sets $\{\mathbb{A}_i\}_{i \in I} \subseteq \Sigma$. Then (Ω, Σ, μ) is called a *probability space*.

Notation *With a slight abuse of terminology, we shall henceforth use the term distribution in place of the term probability measure.*

We let $\Delta(\Omega, \Sigma)$ denote the set of all distributions over (Ω, Σ) . For $\omega \in \Omega$, the *Dirac distribution* δ_ω is defined by $\delta_\omega(\mathbb{A}) = 1$, if $\omega \in \mathbb{A}$, and $\delta_\omega(\mathbb{A}) = 0$, otherwise, for all $\mathbb{A} \in \Sigma$. For a countable set of reals $(p_i)_{i \in I}$ with $p_i \geq 0$ and $\sum_{i \in I} p_i = 1$, the *convex combination* of the distributions $\{\mu_i\}_{i \in I} \subseteq \Delta(\Omega, \Sigma)$ is the distribution $\sum_{i \in I} p_i \cdot \mu_i$ in $\Delta(\Omega, \Sigma)$ defined by $(\sum_{i \in I} p_i \cdot \mu_i)(\mathbb{A}) = \sum_{i \in I} p_i \mu_i(\mathbb{A})$, for all $\mathbb{A} \in \Sigma$. A distribution $\mu \in \Delta(\Omega, \Sigma)$ is called *discrete* if $\mu = \sum_{i \in I} p_i \cdot \delta_{\omega_i}$, with $\omega_i \in \Omega$, for some countable set of indexes I . In this case, the *support* of μ is $\text{supp}(\mu) = \{\omega_i \mid i \in I\}$.

The Wasserstein hemimetric. A *metric* on a set Ω is a function $m: \Omega \times \Omega \rightarrow \mathbb{R}^{\geq 0}$ s.t. $m(\omega_1, \omega_2) = 0$ iff $\omega_1 = \omega_2$, $m(\omega_1, \omega_2) = m(\omega_2, \omega_1)$, and $m(\omega_1, \omega_2) \leq m(\omega_1, \omega_3) + m(\omega_3, \omega_2)$, for all $\omega_1, \omega_2, \omega_3 \in \Omega$. We obtain a *hemimetric* by relaxing the first property to $m(\omega_1, \omega_2) = 0$ if $\omega_1 = \omega_2$, and by dropping the requirement on symmetry. A (hemi)metric m is *l-bounded* if $m(\omega_1, \omega_2) \leq l$ for all $\omega_1, \omega_2 \in \Omega$. For a (hemi)metric on Ω , the pair (Ω, m) is a *(hemi)metric space*.

In order to define a *hemimetric on distributions* we use the Wasserstein lifting [39]. We recall that a *Polish space* is a separable completely metrisable topological space.

Definition 1 (Wasserstein hemimetric). *Consider a Polish space Ω and let m be a hemimetric on Ω . For any two distributions μ and ν on $(\Omega, \mathcal{B}(\Omega))$, the Wasserstein lifting of m to a distance between μ and ν is defined by*

$$\mathbf{W}(m)(\mu, \nu) = \inf_{\mathfrak{w} \in \mathfrak{W}(\mu, \nu)} \int_{\Omega \times \Omega} m(\omega, \omega') d\mathfrak{w}(\omega, \omega')$$

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 $(\Omega \times \Omega, \mathcal{B}(\Omega \times \Omega))$ having μ and ν as left and right marginal, respectively, namely $\mathfrak{w}(\mathbb{A} \times \Omega) = \mu(\mathbb{A})$ and $\mathfrak{w}(\Omega \times \mathbb{A}) = \nu(\mathbb{A})$, for all $\mathbb{A} \in \mathcal{B}(\Omega)$.

Despite the Wasserstein distance was originally given on metrics, the Wasserstein hemimetric given above is well-defined. A formal proof of this can be found in [19] and the references therein.

Notation *As elsewhere in the literature, we use the term metric in place of hemimetric.*

3 The model

In this section, we introduce the three components of our systems, namely the *data space*, the *process* describing the behaviour of the program, and the *environment evolution* describing the effects of the environment. The following example perfectly embodies the kind of program-environment interactions we are interested in.

Example 1. We consider a *smart-room scenario* in which the program should guarantee that both the *temperature* and the *air quality* in the room are in a given *comfort zone*. The room is equipped with a *heating system* and an *air filtering system*. Both are equipped with a *sensor* and an *actuator*. In the heating system the sensor is a thermometer that reads the room temperature, while the actuator is used to turn the heater on or off. Similarly, in the air filtering system the sensor perceives the air quality, giving a value in $[0, 1]$, while the actuator activates the air exchangers. The environment models the evolution of temperature and air quality in the room, as described by the following stochastic difference equations, with sample time interval $\Delta\tau = 1$:

$$T(\tau + 1) = T(\tau) + a(e(\tau)) \cdot (T_e - T(\tau)) + h(\tau) \cdot b \cdot (T_h - T(\tau)) \quad (1)$$

$$T_s(\tau) = T(\tau) + n_t(\tau) \quad (2)$$

$$A(\tau + 1) = A(\tau) + e(\tau) \cdot q^+ \cdot (1 - A(\tau)) - (1 - e(\tau)) \cdot q^- \cdot A(\tau) \quad (3)$$

$$A_s(\tau) = A(\tau) + n_a(\tau) \quad (4)$$

Above, $T(\tau)$ and $A(\tau)$ are the room temperature and air quality at time τ , while $T_s(\tau)$ and $A_s(\tau)$ are the respective values read by sensors, which are obtained from the real ones by adding *noises* $n_t(\tau)$ and $n_a(\tau)$, that we assume to be distributed as Gaussian (normal) distributions $\mathcal{N}(0, v_t^2)$ and $\mathcal{N}(0, v_a^2)$, resp., for some suitable v_t^2 and v_a^2 . Then, $h(\tau)$ and $e(\tau)$ represent the state of the actuators of the heating and air filtering system, respectively. Both take value 1 when the actuator is *on*, and 0 otherwise. Following [2, 23, 31], the temperature dynamics depends on two (non negative) values, $a(e(\tau))$ and b , giving the average heat transfer rates normalised w.r.t. the thermal capacity of the room. In detail, $a(e(\tau))$ is the heat loss rate from the room (through walls, windows, etc.) to the external ambient for which we assume a constant temperature T_e . In our case, this value depends on $e(\tau)$, since the loss rate increases when the air exchangers are on. Then, b is the heat transfer rate from the heater, whose temperature is the constant T_h , to the room. The air quality dynamics is similar: when the air exchangers are off, the air quality decreases with a rate q^- , while it increases of a rate q^+ when they are on.

Modelling the data space. We define the data space by means of a *finite* set of *variables* Var representing: i) *environmental conditions* (pressure, temperature, humidity, etc.);

ii) *values perceived by sensors* (unavoidably affected by imprecision and approximations); iii) *state of actuators* (usually elements in a discrete domain). For each $x \in \text{Var}$ we assume a measurable space $(\mathcal{D}_x, \mathcal{B}_x)$, with $\mathcal{D}_x \subseteq \mathbb{R}$ the domain of x and \mathcal{B}_x the Borel σ -algebra on \mathcal{D}_x . Without loosing generality, we can assume that \mathcal{D}_x is either a *finite set* or a *compact* subset of \mathbb{R} . Notably, \mathcal{D}_x is a Polish space. As Var is a finite set, we can always assume it to be ordered, i.e., $\text{Var} = \{x_1, \dots, x_n\}$ for some $n \in \mathbb{N}$.

Definition 2 (Data space). We define the data space over Var , notation \mathcal{D}_{Var} , as the Cartesian product of the variables domains, namely $\mathcal{D}_{\text{Var}} = \times_{i=1}^n \mathcal{D}_{x_i}$. Then, as a σ -algebra on \mathcal{D}_{Var} we consider the the product σ -algebra $\mathcal{B}_{\mathcal{D}_{\text{Var}}} = \times_{i=1}^n \mathcal{B}_{x_i}$.

Example 2. The data space for the system in Example 1 is defined on the variables T , T_s , h , A , A_s and e . Their domains are $\mathcal{D}_T = \mathcal{D}_{T_s} = [t_m, t_M]$, for suitable values $t_m < t_M$, $\mathcal{D}_A = \mathcal{D}_{A_s} = [0, 1]$, and $\mathcal{D}_h = \mathcal{D}_e = \{0, 1\}$.

When no confusion arises, we will use \mathcal{D} and $\mathcal{B}_{\mathcal{D}}$ in place of \mathcal{D}_{Var} and $\mathcal{B}_{\mathcal{D}_{\text{Var}}}$, respectively. The elements in \mathcal{D} are the n -ples of the form (v_1, \dots, v_n) , with $v_i \in \mathcal{D}_{x_i}$, which can be also identified by means of functions $\mathbf{d}: \text{Var} \rightarrow \mathbb{R}$ from variables to values, with $\mathbf{d}(x) \in \mathcal{D}_x$ for all $x \in \text{Var}$. Each function \mathbf{d} identifies a particular configuration of the data in the data space, and it is thus called a *data state*.

Definition 3 (Data state). A data state is a mapping $\mathbf{d}: \text{Var} \rightarrow \mathbb{R}$ from state variables to values, with $\mathbf{d}(x) \in \mathcal{D}_x$ for all $x \in \text{Var}$.

For simplicity, we shall write $\mathbf{d} \in \mathcal{D}$ in place of $(\mathbf{d}(x_1), \dots, \mathbf{d}(x_n)) \in \mathcal{D}$. Since program and environment interact on the basis of the *current* values of data, we have that at each step there is a data state \mathbf{d} that identifies the *current state of the data space* on which the next computation step is built. Given a data state \mathbf{d} , we let $\mathbf{d}[x = v]$ denote the data state \mathbf{d}' associating v with x , and $\mathbf{d}(y)$ with any $y \neq x$.

Modelling processes. We introduce a simple process calculus allowing us to specify programs that interact with a data state \mathbf{d} in a given environment. We assume that the action performed by a process at a given computation step is determined probabilistically, according to the *generative* probabilistic model [24].

Definition 4 (Syntax of processes). We let \mathcal{P} be the set of processes P defined by:

$$\begin{aligned} P ::= & (\bar{e} \rightarrow \bar{x}).P' \mid \text{if } [e] P_1 \text{ else } P_2 \mid \sum_{i \in I} p_i \cdot P_i \mid P_1 \parallel_p P_2 \mid A \\ e ::= & v \in V \mid x \in \text{Var} \mid \text{op}_k(e_1, \dots, e_k) \end{aligned}$$

where, $V \subseteq \mathbb{R}$ countable, p, p_1, \dots weights in $[0, 1] \cap \mathbb{Q}$, I is finite, A ranges over process variables, op_k indicates a measurable operator $\mathbb{R}^k \rightarrow \mathbb{R}$, and $\bar{\cdot}$ denotes a finite sequence of elements. We assume to have a single definition $A \stackrel{\text{def}}{=} P$ for each process variable A . Moreover, we require that $\sum_{i \in I} p_i = 1$ for any process $\sum_{i \in I} p_i \cdot P_i$.

Process $(\bar{e} \rightarrow \bar{x}).P$ evaluates the sequence of expressions \bar{e} with the current data state \mathbf{d} and assigns the results $\llbracket \bar{e} \rrbracket_{\mathbf{d}}$ to the sequence of variables \bar{x} . We may use \surd to denote the prefix $(\emptyset \rightarrow \emptyset)$. Process $\text{if } [e] P_1 \text{ else } P_2$ behaves either as P_1 when $\llbracket e \rrbracket_{\mathbf{d}} = \top$,

or as P_2 when $\llbracket e \rrbracket_{\mathbf{d}} = \perp$. Then, $\sum_{i=1}^n p_i \cdot P_i$ is the *generative probabilistic choice*: process P_i has probability p_i to move. The *generative probabilistic interleaving* construct $P_1 \parallel_p P_2$ lets the two argument processes to interleave their actions, where at each step P_1 moves with probability p and P_2 with probability $1 - p$. Process variables allow us to specify recursive behaviours by means of equations of the form $A \stackrel{def}{=} P$. To avoid Zeno behaviours we assume that all occurrences of process variables appear *guarded* by prefixing constructs in P . We assume the standard notions of *free* and *bound* process variables. A program is then a *closed* process, i.e., a process without free variables.

Formally, actions performed by a process can be abstracted in terms of the *effects* they have on the data state, i.e., via *substitutions* of the form $\theta = [x_{i_1} \leftarrow v_{i_1}, \dots, x_{i_k} \leftarrow v_{i_k}]$, also denoted $\bar{x} \leftarrow \bar{v}$ if $\bar{x} = x_{i_1}, \dots, x_{i_k}$ and $\bar{v} = v_{i_1}, \dots, v_{i_k}$. Since in Definition 4 operations op_k are assumed to be measurable, we can model the effects as $\mathcal{B}_{\mathcal{D}}$ -measurable functions $\theta: \mathcal{D} \rightarrow \mathcal{D}$ s.t. $\theta(\mathbf{d}) := \mathbf{d}[\bar{x} \leftarrow \bar{v}]$ whenever $\theta = \bar{x} \leftarrow \bar{v}$. We denote by Θ the set of effects. The behaviour of a process can then be defined by means of a function $\text{pstep}: \mathcal{P} \times \mathcal{D} \rightarrow \Delta(\Theta \times \mathcal{P})$ that given a process P and a data state \mathbf{d} yields a *discrete* distribution over $\Theta \times \mathcal{P}$. Function pstep is defined as follows:

$$\begin{aligned}
 \text{(PR1)} \quad & \text{pstep}((\bar{e} \rightarrow \bar{x}).P', \mathbf{d}) = \delta_{(\bar{x} \leftarrow \llbracket \bar{e} \rrbracket_{\mathbf{d}}, P')} \\
 \text{(PR2)} \quad & \text{pstep}(\text{if } [e] P_1 \text{ else } P_2, \mathbf{d}) = \begin{cases} \text{pstep}(P_1, \mathbf{d}) & \text{if } \llbracket e \rrbracket_{\mathbf{d}} = 1 \\ \text{pstep}(P_2, \mathbf{d}) & \text{if } \llbracket e \rrbracket_{\mathbf{d}} = 0 \end{cases} \\
 \text{(PR3)} \quad & \text{pstep}(\sum_i p_i \cdot P_i, \mathbf{d}) = \sum_i p_i \cdot \text{pstep}(P_i, \mathbf{d}) \\
 \text{(PR4)} \quad & \text{pstep}(P_1 \parallel_p P_2, \mathbf{d}) = p \cdot (\text{pstep}(P_1, \mathbf{d}) \parallel_p P_2) + (1 - p) \cdot (P_1 \parallel_p \text{pstep}(P_2, \mathbf{d})) \\
 \text{(PR5)} \quad & \text{pstep}(A, \mathbf{d}) = \text{pstep}(P, \mathbf{d}) \quad (\text{if } A \stackrel{def}{=} P).
 \end{aligned}$$

In rule (PR4), for $\pi \in \Delta(\Theta \times \mathcal{P})$, we let $\pi \parallel_p P$ (resp. $P \parallel_p \pi$) denote the distribution $\pi' \in \Delta(\Theta \times \mathcal{P})$ s.t.: $\pi'(\theta, P') = \pi(\theta, P'')$, whenever $P' = P'' \parallel_p P$ (resp. $P' = P \parallel_p P''$), and 0, otherwise.

Proposition 1 (Properties of process semantics). *Let $P \in \mathcal{P}$ and $\mathbf{d} \in \mathcal{D}$. Then $\text{pstep}(P, \mathbf{d})$ is a discrete distribution with finite support.*

Example 3. We define a program to control the smart-room scenario of Example 1. In detail, we want to guarantee that the temperature in the room is in the interval $Z = [t_{\min}, t_{\max}] \subseteq \mathcal{D}_T$, while the air quality is above a given threshold $q_a \in \mathcal{D}_A$. The following process A_{off}^T (resp. A_{on}^T) turns the heating system *on* (resp. *off*) when the temperature acquired by the sensor goes under $t_{\min} - \varepsilon_t$ (resp. over $t_{\max} + \varepsilon_t$). The use of the tolerance ε_t guarantees that the heating system is not repeatedly turned on/off.

$$\begin{aligned}
 A_{\text{off}}^T & \stackrel{def}{=} \text{if } [T_s < t_{\min} - \varepsilon_t] (1 \rightarrow h).A_{\text{on}}^T \text{ else } \checkmark.A_{\text{off}}^T \\
 A_{\text{on}}^T & \stackrel{def}{=} \text{if } [T_s > t_{\max} + \varepsilon_t] (0 \rightarrow h).A_{\text{off}}^T \text{ else } \checkmark.A_{\text{on}}^T.
 \end{aligned}$$

The behaviour of program components controlling the air filtering system is similar and implemented by the following processes A_{off}^A and A_{on}^A :

$$\begin{aligned}
 A_{\text{off}}^A & \stackrel{def}{=} \text{if } [A_s \leq q_a - \varepsilon_a] (1 \rightarrow e).A_{\text{on}}^A \text{ else } \checkmark.A_{\text{off}}^A \\
 A_{\text{on}}^A & \stackrel{def}{=} \text{if } [A_s > q_a + \varepsilon_a] (0 \rightarrow e).A_{\text{off}}^A \text{ else } \checkmark.A_{\text{on}}^A.
 \end{aligned}$$

The composition of the programs is given by the process $P = A_{off}^T \parallel_{0.5} A_{off}^A$.

Modelling the environment. We model the action of the environment by a mapping \mathcal{E} , called *environment evolution*, taking a data state to a distribution over data states.

Definition 5 (Environment evolution). An environment evolution is a map $\mathcal{E}: \mathcal{D} \rightarrow \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ s.t. for each $\mathbb{D} \in \mathcal{B}_{\mathcal{D}}$ the mapping $\mathbf{d} \mapsto \mathcal{E}(\mathbf{d})(\mathbb{D})$ is $\mathcal{B}_{\mathcal{D}}$ -measurable.

Due to the interaction with the program, the probability induced by \mathcal{E} at the next time step *depends only* on the current state of the data space. It is then natural to assume that the behaviour of the environment is modelled as a discrete time Markov process.

Example 4. For our smart-room scenario, the environment evolution \mathcal{E} can be derived directly from Equations (1)–(4). Notice that, in this case, randomness follows from the Gaussian noises associated with the temperature and air quality sensors.

Modelling system's behaviour. We use the notion of *configuration* to model the state of the system at each time step.

Definition 6 (Configuration). A configuration is a triple $c = \langle P, \mathbf{d} \rangle_{\mathcal{E}}$, where P is a process, \mathbf{d} is a data state and \mathcal{E} is an environment evolution. We denote by $\mathcal{C}_{\mathcal{P}, \mathcal{D}, \mathcal{E}}$ the set of configurations defined over \mathcal{P} , \mathcal{D} and \mathcal{E} .

When no confusion arises, we shall write \mathcal{C} in place of $\mathcal{C}_{\mathcal{P}, \mathcal{D}, \mathcal{E}}$.

Let $(\mathcal{P}, \Sigma_{\mathcal{P}})$ be the measurable space of processes, where $\Sigma_{\mathcal{P}}$ is the power set of \mathcal{P} , and $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ be the measurable space of data states. As \mathcal{E} is fixed, we can identify \mathcal{C} with $\mathcal{P} \times \mathcal{D}$ and equip it with the product σ -algebra $\Sigma_{\mathcal{C}} = \Sigma_{\mathcal{P}} \times \mathcal{B}_{\mathcal{D}}$: $\Sigma_{\mathcal{C}}$ is generated by the sets $\{\langle \mathbb{P}, \mathbb{D} \rangle_{\mathcal{E}} \mid \mathbb{P} \in \Sigma_{\mathcal{P}}, \mathbb{D} \in \mathcal{B}_{\mathcal{D}}\}$, where $\langle \mathbb{P}, \mathbb{D} \rangle_{\mathcal{E}} = \{\langle P, \mathbf{d} \rangle_{\mathcal{E}} \mid P \in \mathbb{P}, \mathbf{d} \in \mathbb{D}\}$.

Notation For $\mu_{\mathcal{P}} \in \Delta(\mathcal{P}, \Sigma_{\mathcal{P}})$ and $\mu_{\mathcal{D}} \in \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ we let $\mu = \langle \mu_{\mathcal{P}}, \mu_{\mathcal{D}} \rangle_{\mathcal{E}}$ denote the product distribution on $(\mathcal{C}, \Sigma_{\mathcal{C}})$, i.e., $\mu(\langle \mathbb{P}, \mathbb{D} \rangle_{\mathcal{E}}) = \mu_{\mathcal{P}}(\mathbb{P}) \cdot \mu_{\mathcal{D}}(\mathbb{D})$ for all $\mathbb{P} \in \Sigma_{\mathcal{P}}$ and $\mathbb{D} \in \mathcal{B}_{\mathcal{D}}$. If $\mu_{\mathcal{P}} = \delta_P$ for some $P \in \mathcal{P}$, we shall denote $\langle \delta_P, \mu_{\mathcal{D}} \rangle_{\mathcal{E}}$ simply by $\langle P, \mu_{\mathcal{D}} \rangle_{\mathcal{E}}$.

We aim to express the behaviour of a system in terms of the changes on data. We start with the *one-step* behaviour of a configuration, in which we combine the effects on the data state induced by the activity of the process (given by *pstep*) and the subsequent action by the environment. Formally, we define a function *cstep* that, given a configuration, yields a distribution on $(\mathcal{C}, \Sigma_{\mathcal{C}})$ (Def. 7 below). Then, we use *cstep* to define the *multi-step* behaviour of configuration c as a sequence $\mathcal{S}_{c,0}^{\mathcal{C}}, \mathcal{S}_{c,1}^{\mathcal{C}}, \dots$ of distributions on $(\mathcal{C}, \Sigma_{\mathcal{C}})$. To this end, we show that *cstep* is a Markov kernel (Prop. 3 below). Finally, to abstract from processes and focus only on data, from the sequence $\mathcal{S}_{c,0}^{\mathcal{C}}, \mathcal{S}_{c,1}^{\mathcal{C}}, \dots$, we obtain a sequence of distributions $\mathcal{S}_{c,0}^{\mathcal{D}}, \mathcal{S}_{c,1}^{\mathcal{D}}, \dots$ on $(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ called the *evolution sequence* of the system (Def. 9 below).

Definition 7 (One-step semantics). Function *cstep*: $\mathcal{C} \rightarrow \Delta(\mathcal{C}, \Sigma_{\mathcal{C}})$ is defined for all configurations $\langle P, \mathbf{d} \rangle_{\mathcal{E}} \in \mathcal{C}$ by

$$\text{cstep}(\langle P, \mathbf{d} \rangle_{\mathcal{E}}) = \sum_{(\theta, P') \in \text{supp}(\text{pstep}(P, \mathbf{d}))} \text{pstep}(P, \mathbf{d})(\theta, P') \cdot \langle P', \mathcal{E}(\theta(\mathbf{d})) \rangle_{\mathcal{E}}. \quad (5)$$

The next result follows by $\mathcal{E}(\theta(\mathbf{d})) \in \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ (Def. 5), which ensures that $\langle P', \mathcal{E}(\theta(\mathbf{d})) \rangle_{\mathcal{E}} \in \Delta(\mathcal{C}, \Sigma_{\mathcal{C}})$, and $\text{pstep}(P, \mathbf{d})$ is a discrete distribution in $\Delta(\Theta \times \mathcal{P})$ (Prop. 1).

Proposition 2. *For any configuration $c \in \mathcal{C}$, $\text{cstep}(c)$ is a distribution on $(\mathcal{C}, \Sigma_{\mathcal{C}})$.*

Since $\text{cstep}(c) \in \Delta(\mathcal{C}, \Sigma_{\mathcal{C}})$ for each $c \in \mathcal{C}$, we can rewrite $\text{cstep}: \mathcal{C} \times \Sigma_{\mathcal{C}} \rightarrow [0, 1]$, so that for each configuration $c \in \mathcal{C}$ and measurable set $\mathbb{C} \in \Sigma_{\mathcal{C}}$, $\text{cstep}(c)(\mathbb{C})$ denotes the probability of reaching in one step a configuration in \mathbb{C} starting from c . We can prove that cstep is the Markov kernel of the Markov process modelling our system. This follows by Proposition 2 and by proving that for each $\mathbb{C} \in \Sigma_{\mathcal{C}}$, the mapping $c \mapsto \text{cstep}(c)(\mathbb{C})$ is $\Sigma_{\mathcal{C}}$ -measurable for all $c \in \mathcal{C}$.

Proposition 3. *The function cstep is a Markov kernel.*

Hence, the multi-step behaviour of configuration c can be defined as a time homogeneous Markov process having cstep as Markov kernel and δ_c as initial distribution.

Definition 8 (Multi-step semantics). *Let $c \in \mathcal{C}$ be a configuration. The multi-step behaviour of c is the sequence of distributions $\mathcal{S}_{c,0}^{\mathcal{C}}, \mathcal{S}_{c,1}^{\mathcal{C}}, \dots$ on $(\mathcal{C}, \Sigma_{\mathcal{C}})$ defined inductively as follows:*

$$\begin{aligned} \mathcal{S}_{c,0}^{\mathcal{C}}(\mathbb{C}) &= \delta_c(\mathbb{C}), \text{ for all } \mathbb{C} \in \Sigma_{\mathcal{C}} \\ \mathcal{S}_{c,i+1}^{\mathcal{C}}(\mathbb{C}) &= \int_{\mathcal{C}} \text{cstep}(b)(\mathbb{C}) d(\mathcal{S}_{c,i}^{\mathcal{C}}(b)), \text{ for all } \mathbb{C} \in \Sigma_{\mathcal{C}}. \end{aligned}$$

We can prove that $\mathcal{S}_{c,0}^{\mathcal{C}}, \mathcal{S}_{c,1}^{\mathcal{C}}, \dots$ are well defined, namely they are distributions on $(\mathcal{C}, \Sigma_{\mathcal{C}})$. The proof follows by an easy induction based on Proposition 3.

Proposition 4. *For any $c \in \mathcal{C}$, all $\mathcal{S}_{c,0}^{\mathcal{C}}, \mathcal{S}_{c,1}^{\mathcal{C}}, \dots$ are distributions on $(\mathcal{C}, \Sigma_{\mathcal{C}})$.*

As the program-environment interplay can be observed only in the changes they induce on the data states, we define the *evolution sequence* of a configuration as the sequence of distributions over data states that are reached by it, step-by-step.

Definition 9 (Evolution sequence). *The evolution sequence of $c = \langle P, \mathbf{d} \rangle_{\mathcal{E}}$ is a sequence $\mathcal{S}_c^{\mathcal{D}} \in \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})^{\omega}$ of distributions over \mathcal{D} such that $\mathcal{S}_c^{\mathcal{D}} = \mathcal{S}_{c,0}^{\mathcal{D}} \dots \mathcal{S}_{c,n}^{\mathcal{D}} \dots$ if and only if for all $i \geq 0$ and for all $\mathbb{D} \in \mathcal{B}_{\mathcal{D}}$, $\mathcal{S}_{c,i}^{\mathcal{D}}(\mathbb{D}) = \mathcal{S}_{c,i}^{\mathcal{C}}(\langle \mathcal{P}, \mathbb{D} \rangle_{\mathcal{E}})$.*

4 Towards a metric for systems

We aim at defining a *distance* over the systems described in the previous section, called the *evolution metric*, allowing us to do the following:

1. Verify how well a program is fulfilling its tasks.
2. Establish whether one program behaves better than another one in an environment.
3. Compare the interactions of a program with different environments.

These three objectives can be naturally obtained thanks to the possibility of modelling the program in isolation from the environment typical of our model, and to our purely data-driven system semantics. Intuitively, since the behaviour of a system is entirely described by its evolution sequence, the evolution metric m will indeed be defined as a distance on the evolution sequences of systems. However, in order to obtain the proper technical definition of m , some considerations are due.

Firstly, we notice that in most applications the tasks of the program can be expressed in a purely data-driven fashion. We can identify a set of *parameters of interest* such that, at any time step, any difference between them and the data actually obtained can be interpreted as a flaw in system behaviour. We use a *penalty function* ρ to quantify these differences. From the penalty function we can obtain a *distance on data states*, namely a 1-bounded *hemimetric* $m^{\mathcal{D}}$ expressing how much a data state \mathbf{d}_2 is worse than a data state \mathbf{d}_1 according to parameters of interests. Secondly, we recall that the evolution sequence of a system consists in a sequence of *distributions* over data states. Hence, we use the *Wasserstein metric* to lift $m^{\mathcal{D}}$ to a distance $\mathbf{W}(m^{\mathcal{D}})$ over distributions over data states. Informally, with the Wasserstein metric we can express how much worse a configuration is expected to behave w.r.t. another one at a given time. Finally, we need to lift $\mathbf{W}(m^{\mathcal{D}})$ to a distance on the entire evolution sequences of systems. For our purposes, a reasonable choice is to take the maximum over time of the pointwise (w.r.t. time) Wasserstein distances (see Remark 1 below for further details on this choice).

A metric on data states. We start by proposing a metric on data states, seen as *static components* in isolation from processes and environment. To this end, we introduce a *penalty function* $\rho: \mathcal{D} \rightarrow [0, 1]$, a continuous function that assigns to each data state \mathbf{d} a penalty in $[0, 1]$ expressing how far the values of the parameters of interest in \mathbf{d} are from their desired ones (hence $\rho(\mathbf{d}) = 0$ if \mathbf{d} respects all the parameters). Since some parameters can be time-dependent, so is ρ : at any time step τ , the τ -penalty function ρ_τ compares the data states w.r.t. the values of the parameters expected at time τ .

Example 5. We recall, from Example 2, that $\mathcal{D}_T = [t_m, t_M]$ and $\mathcal{D}_A = [0, 1]$. The task of our program is to keep the value of T within the comfort zone $Z = [t_{\min}, t_{\max}]$, for some t_{\min}, t_{\max} , and that of A above a threshold $q_a \in \mathcal{D}_A$ (cf. Example 3). Hence, we define a penalty function that assigns the penalty 0 if the value of T is in Z and that of A is greater or equal to q_a , otherwise it is proportional to how much T and A are far from Z and q_a , respectively. We let $\rho_\tau(\mathbf{d}) = \max\{\rho^T(\mathbf{d}(T)), \rho^A(\mathbf{d}(A))\}$, where $\rho^T(t)$ is 0 if $t \in [t_{\min}, t_{\max}]$ and $\frac{\max\{t - t_{\max}, t_{\min} - t\}}{\max\{t_M - t_{\max}, t_{\min} - t_m\}}$ otherwise, while $\rho^A(q) = \max\{0, q_a - q\}$.

A formal definition of the penalty function is beyond the purpose of this paper, also due to its context-dependent nature. Besides, notice that we can assume that ρ already includes some tolerances w.r.t. the exact values of the parameters in its evaluation, and thus we do not consider them. The *(timed) metric on data states* is then defined as the asymmetric difference between the penalties assigned to them by the penalty function.

Definition 10 (Metric on data states). For any time step τ , let $\rho_\tau: \mathcal{D} \rightarrow [0, 1]$ be the τ -penalty function on \mathcal{D} . The τ -metric on data states in \mathcal{D} , $m_{\rho, \tau}^{\mathcal{D}}: \mathcal{D} \times \mathcal{D} \rightarrow [0, 1]$, is defined, for all $\mathbf{d}_1, \mathbf{d}_2 \in \mathcal{D}$, by $m_{\rho, \tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_2) = \max\{\rho_\tau(\mathbf{d}_2) - \rho_\tau(\mathbf{d}_1), 0\}$.

Notice that $m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_2) > 0$ iff $\rho_\tau(\mathbf{d}_2) > \rho_\tau(\mathbf{d}_1)$, i.e., the penalty assigned to \mathbf{d}_2 is higher than that assigned to \mathbf{d}_1 . For this reason, we say that $m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_2)$ expresses *how worse* \mathbf{d}_2 is than \mathbf{d}_1 w.r.t. the objectives of the system. It is not hard to see that for all $\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3 \in \mathcal{D}$ we have $m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_2) \leq 1$, $m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_1) = 0$, and $m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_2) \leq m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_1, \mathbf{d}_3) + m_{\rho,\tau}^{\mathcal{D}}(\mathbf{d}_3, \mathbf{d}_2)$, thus ensuring that $m_{\rho,\tau}^{\mathcal{D}}$ is a 1-bounded hemimetric.

Proposition 5. *Function $m_{\rho,\tau}^{\mathcal{D}}$ is a 1-bounded hemimetric on \mathcal{D} .*

Lifting $m_{\rho,\tau}^{\mathcal{D}}$ to distributions. The second step to obtain the evolution metric consists in lifting $m_{\rho,\tau}^{\mathcal{D}}$ to a metric on distributions on data states. Among the several notions of lifting in the literature (see [33] for a survey), we opt for that of Wasserstein, since: i) it preserves the properties of the ground metric; ii) it allows us to deal with discrete and continuous measures; iii) it is computationally tractable via statistical inference. According to Def. 1, the Wasserstein lifting of $m_{\rho,\tau}^{\mathcal{D}}$ to a distance between two distributions $\mu, \nu \in \Delta(\mathcal{D}, \mathcal{B}_{\mathcal{D}})$ is defined by

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

The evolution metric. We now need to lift $\mathbf{W}(m_{\rho,\tau}^{\mathcal{D}})$ to a distance on evolution sequences. To this end, we observe that the evolution sequence of a configuration includes the distributions over data states induced after *each* computation step. Thus, the time step between two distributions is determined by the program. However, it could be the case that the changes on data induced by the environment can be appreciated only along wider time intervals. Our running example is a clear instance of this situation: while we can reasonably assume that the duration of the computation steps of the thermostat is of the order of a millisecond, the variations in the temperature that can be detected in the same time interval are indeed negligible w.r.t. the program's task. A significant temperature rise or drop can be observed only in longer time. To deal with this kind of situations, we introduce the notion of *observation times*, namely a *discrete* set OT of time steps at which the modifications induced by the program-environment interplay give us useful information on the evolution of the system. Hence, a comparison of the evolution sequences based on the differences in the distributions reached at the times in OT can be considered meaningful. Moreover, considering only the differences at the observation times will favour the computational tractability of the evolution metric.

We define the evolution metric as a sort of *weighted infinity norm* of the tuple of the Wasserstein distances between the distributions in the evolution sequences. As weight we consider a non-increasing function $\lambda: \text{OT} \rightarrow (0, 1]$ expressing how much the distance at time τ affects the overall distance between configurations c_1 and c_2 . We refer to λ as to the *discount function*, and to $\lambda(\tau)$ as to the *discount factor at time τ* .

Definition 11 (Evolution metric). *Assume a set OT of observation times and a discount function λ . Let ρ be a penalty function and let $m_{\rho,\tau}^{\mathcal{D}}$ be the metric on data states defined on it. Then, the λ -evolution metric over ρ and OT is the mapping $\mathbf{m}_{\rho,\text{OT}}^\lambda: \mathcal{C} \times \mathcal{C} \rightarrow [0, 1]$ defined, for all configurations $c_1, c_2 \in \mathcal{C}$, by*

$$\mathbf{m}_{\rho,\text{OT}}^\lambda(c_1, c_2) = \sup_{\tau \in \text{OT}} \lambda(\tau) \cdot \mathbf{W}(m_{\rho,\tau}^{\mathcal{D}})(\mathcal{S}_{c_1,\tau}^{\mathcal{D}}, \mathcal{S}_{c_2,\tau}^{\mathcal{D}}).$$

Since $m_{\rho,\tau}^{\mathcal{D}}$ is a 1-bounded hemimetric (Proposition 5) and lifting \mathbf{W} preserves such a property, we can easily derive the same property for m_{OT}^{λ} .

Proposition 6. *Function $m_{\rho,\text{OT}}^{\lambda}$ is a 1-bounded hemimetric on \mathcal{C} .*

Notice that if λ is a *strictly* non-increasing function, then it specifies how much the distance of *future events* is mitigated and, moreover, it guarantees that to obtain upper bounds on the evolution metric only a *finite* number of observations is needed.

Remark 1. Usually, due to the presence of uncertainties, the behaviour of a system can be considered acceptable even if it differs from its intended one *up-to a certain tolerance*. Similarly, the properties of adaptability and reliability that we aim to study will check whether a program is able to perform well in a perturbed environment *up-to a given tolerance*. In this setting, the choice of defining the evolution metric as the point-wise maximal distance in time between the evolution sequences of systems is natural and reasonable: if in the worst case (the maximal distance) the program keeps the parameters of interest within the given tolerance, then its entire behaviour can be considered acceptable. However, with this approach we have that a program is only as good as its worst performance, and one could argue that there are application contexts in which our evolution metric would be less meaningful. For these reasons, we remark that we could have given a *parametric* version of Definition 11 and defining the evolution metric in terms of a generic *aggregation function* f over the tuple of Wasserstein distances. Then, one could choose the best instantiation for f according to the chosen application context. The use of a parametric definition would have not affected the technical development of our paper. However, to keep the notation and presentation as simple as possible, we opted to define $m_{\rho,\text{OT}}^{\lambda}$ directly in the weighted infinity norm form. A similar reasoning applies to the definition of the penalty function that we gave in Example 5.

5 Estimating the evolution metric

In this section we show how the evolution metric can be estimated via statistical techniques. Firstly, we show how we can estimate the evolution sequence of a given configuration c . Then, we evaluate the distance between two configurations c_1 and c_2 on their estimated evolution sequences.

Computing empirical evolution sequences. To compute the empirical evolution sequence of a configuration c the following function EST can be used.

```

1: function EST( $c, k, N$ )
2:    $\forall i : (0 \leq i \leq k) : E_i \leftarrow \emptyset$ 
3:    $counter \leftarrow 0$ 
4:   while  $counter < N$  do
5:      $(c_0, \dots, c_k) \leftarrow \text{SIM}(c, k)$ 
6:      $\forall i : E_i \leftarrow E_i, c_i$ 
7:      $counter \leftarrow counter + 1$ 
8:   end while
9:   return  $E_0, \dots, E_k$ 
10: end function

```

Function EST(c, k, N) invokes N times function SIM, i.e., any simulation algorithm sampling a sequence of configurations c_0, \dots, c_k , modelling k steps of a computation from $c = c_0$. Then, the sequence E_0, \dots, E_k is computed, where E_i is the tuple c_i^1, \dots, c_i^N of configurations observed at time i in each of the N sampled computations.




Fig. 1: Estimated distribution of the temperature after 50 steps with $N = 10^2$, $N = 10^3$ and $N = 10^4$. As comfort zone we consider the interval $[15, 20]$.

Each E_i can be used to estimate the distribution $\mathcal{S}_{c,i}^C$. For any i , with $0 \leq i \leq k$, we let $\hat{\mathcal{S}}_{c,i}^{C,N}$ be the distribution s.t. for any $\mathbb{C} \in \Sigma_C$ we have $\hat{\mathcal{S}}_{c,i}^{C,N}(\mathbb{C}) = \frac{|E_i \cap \mathbb{C}|}{N}$. Finally, we let $\hat{\mathcal{S}}_c^{\mathcal{D},N} = \hat{\mathcal{S}}_{c,0}^{\mathcal{D},N} \dots \hat{\mathcal{S}}_{c,k}^{\mathcal{D},N}$ be the *empirical evolution sequence* s.t. for any measurable set of data states $\mathbb{D} \in \mathcal{B}_{\mathcal{D}}$ we have $\hat{\mathcal{S}}_{c,i}^{\mathcal{D},N}(\mathbb{D}) = \hat{\mathcal{S}}_{c,i}^{C,N}(\langle \mathcal{P}, \mathbb{D} \rangle_{\mathcal{E}})$. Then, by applying the weak law of large numbers to the i.i.d samples, we get that when N goes to infinite both $\hat{\mathcal{S}}_{c,i}^{C,N}$ and $\hat{\mathcal{S}}_{c,i}^{\mathcal{D},N}$ converge weakly to $\mathcal{S}_{c,i}^C$ and $\mathcal{S}_{c,i}^{\mathcal{D}}$ respectively:

$$\lim_{N \rightarrow \infty} \hat{\mathcal{S}}_{c,i}^{C,N} = \mathcal{S}_{c,i}^C \quad \lim_{N \rightarrow \infty} \hat{\mathcal{S}}_{c,i}^{\mathcal{D},N} = \mathcal{S}_{c,i}^{\mathcal{D}}. \quad (6)$$

The tool and the scripts of the examples are available (in Python) at <https://github.com/quasylab/spear>.

Example 6. We apply our simulation to the heating system from Sect. 3, with initial configuration $c_1 = \langle P, \{T = 5.0, T_s = 5.0, h = 0, A = 0.5, A_s = 0.5, e = 0\} \rangle_{\mathcal{E}}$, where P is the process in Ex. 3, and $[t_{\min}, t_{\max}] = [15, 20]$. In Fig. 1 the probability distribution of the temperature after 50 steps is reported.

Computing distance between two configurations. Function EST allows us to collect independent samples at each time step i from 0 to a deadline k . These samples can be used to estimate the distance between two configurations c_1 and c_2 . Following a similar approach to [37], to estimate the Wasserstein distance $\mathbf{W}(m_{\rho,i}^{\mathcal{D}})$ between two (unknown) distributions $\mathcal{S}_{c_1,i}^{\mathcal{D}}$ and $\mathcal{S}_{c_2,i}^{\mathcal{D}}$ we can use N independent samples $\{c_1^1, \dots, c_1^N\}$ taken from $\mathcal{S}_{c_1,i}^{\mathcal{D}}$ and $\ell \cdot N$ independent samples $\{c_2^1, \dots, c_2^{\ell \cdot N}\}$ taken from $\mathcal{S}_{c_2,i}^{\mathcal{D}}$. After that, we exploit the i -penalty function ρ and we consider the two sequences of values: $\{\omega_j = \rho_i(\mathbf{d}_1^j) | \langle P_1^j, \mathbf{d}_1^j \rangle_{\mathcal{E}_1} = c_1^j\}$ and $\{\nu_h = \rho_i(\mathbf{d}_2^h) | \langle P_2^h, \mathbf{d}_2^h \rangle_{\mathcal{E}_2} = c_2^h\}$. We can assume, without loss of generality, that these sequences are ordered, i.e., $\omega_j \leq \omega_{j+1}$ and $\nu_h \leq \nu_{h+1}$. The value $\mathbf{W}(m_{\rho,i}^{\mathcal{D}})(\mathcal{S}_{c_1,i}^{\mathcal{D}}, \mathcal{S}_{c_2,i}^{\mathcal{D}})$ can be approximated as $\frac{1}{\ell N} \sum_{h=1}^{\ell N} \max\{\nu_h - \omega_{\lceil \frac{h}{\ell} \rceil}, 0\}$. The next theorem, based on results in [37, 40], ensures that the larger the number of samplings the closer the gap between the estimated value and the exact one.

Theorem 1. Let $\mathcal{S}_{c_1,i}^C, \mathcal{S}_{c_2,i}^C \in \Delta(\mathcal{C}, \Sigma_C)$ be unknown, and ρ be a penalty function. Let $\{\omega_j = \rho_i(\mathbf{d}_1^j)\}$ and $\{\nu_h = \rho_i(\mathbf{d}_2^h)\}$ be the ordered sequences obtained from independent samples taken from $\mathcal{S}_{c_1,i}^C$ and $\mathcal{S}_{c_2,i}^C$, respectively. Then, it holds, a.s., that $\mathbf{W}(m_{\rho,i}^{\mathcal{D}})(\mathcal{S}_{c_1,i}^{\mathcal{D}}, \mathcal{S}_{c_2,i}^{\mathcal{D}}) = \lim_{N \rightarrow \infty} \frac{1}{\ell N} \sum_{h=1}^{\ell N} \max\{\nu_h - \omega_{\lceil \frac{h}{\ell} \rceil}, 0\}$.

```

1: function DIST( $c_1, c_2, \rho, \lambda, OT, N, \ell$ )
2:    $k \leftarrow \max_{OT}$ 
3:    $E_{1,1}, \dots, E_{1,k} \leftarrow \text{EST}(c_1, k, N)$ 
4:    $E_{2,1}, \dots, E_{2,k} \leftarrow \text{EST}(c_2, k, \ell N)$ 
5:    $m \leftarrow \infty$ 
6:   for all  $i \in OT$  do
7:      $m_i \leftarrow \text{COMPW}(E_{1,i}, E_{2,i}, \rho_i)$ 
8:      $m \leftarrow \min\{m, \lambda(i) \cdot m_i\}$ 
9:   end for
10:  return  $m$ 
11: end function

1: function COMPW( $E_1, E_2, \rho$ )
2:   $((P_1^1, \mathbf{d}_1^1)_{\mathcal{E}_1}, \dots, (P_1^N, \mathbf{d}_1^N)_{\mathcal{E}_1}) \leftarrow E_1$ 
3:   $((P_2^1, \mathbf{d}_2^1)_{\mathcal{E}_2}, \dots, (P_2^{\ell N}, \mathbf{d}_2^{\ell N})_{\mathcal{E}_2}) \leftarrow E_2$ 
4:   $\forall j : (1 \leq j \leq N) : \omega_j \leftarrow \rho(\mathbf{d}_1^j)$ 
5:   $\forall h : (1 \leq h \leq \ell N) : \nu_h \leftarrow \rho(\mathbf{d}_2^h)$ 
6:  re index  $\{\omega_j\}$  s.t.  $\omega_j \leq \omega_{j+1}$ 
7:  re index  $\{\nu_h\}$  s.t.  $\nu_h \leq \nu_{h+1}$ 
8:  return  $\frac{1}{\ell N} \sum_{h=1}^{\ell N} |\omega_{\lceil \frac{h}{\ell} \rceil} - \nu_h|$ 
9: end function

```

Fig. 2: Functions used to estimate the evolution metric on systems.

The outlined procedure is realised by functions DIST and COMPW in Fig. 2. The former takes as input the two configurations to compare, the penalty function (seen as the sequence of the i -penalty functions), the discount function λ , the bounded set OT if observation times, and the parameters N and ℓ used to obtain the samplings of computation. Function DIST collects the samples E_i of possible computations during the observation period $[0, \max_{OT}]$, where \max_{OT} denotes the last observation time. Then, for each $i \in OT$, the distance at time i is computed via the function $\text{COMPW}(E_{1,i}, E_{2,i}, \rho_i)$. As the penalty function allows us to reduce the evaluation of the Wasserstein distance in \mathbb{R}^n to its evaluation on \mathbb{R} , due to the sorting of $\{\nu_h \mid h \in [1, \dots, \ell N]\}$ the complexity of function COMPW is $O(\ell N \log(\ell N))$ (cf. [37]).


Example 7. We change the initial value of the air quality in the configuration c_1 in Example 6, and consider $c_2 = \langle P, \{T = 5.0, T_s = 5.0, h = 0, A = 0.3, A_s = 0.3, e = 0\} \rangle_{\mathcal{E}}$. Figure 3a shows the variation in time of the distance between c_1 and c_2 .

6 Adaptability and reliability of programs

In this section we exploit the evolution metric to study some dependability properties of programs, which we call *adaptability* and *reliability*, w.r.t. a data state and an environment. Both notions entail the ability of the process to induce a *similar* behaviour in systems that start from *similar* initial conditions. They differ in how time is considered.

The notion of adaptability imposes some constraints on the *long term* behaviour of systems, disregarding their possible initial dissimilarities. Given the thresholds $\eta_1, \eta_2 \in [0, 1)$ and an observable time $\tilde{\tau}$, we say that a program P is adaptable w.r.t. a data state \mathbf{d} and an environment evolution \mathcal{E} if whenever P starts its computation from a data state \mathbf{d}' that differs from \mathbf{d} for at most η_1 , then we are guaranteed that the distance between the evolution sequences of the two systems after time $\tilde{\tau}$ is bounded by η_2 .

Definition 12 (Adaptability). *Let $\tilde{\tau} \in OT$ and $\eta_1, \eta_2 \in [0, 1)$. We say that P is $(\tilde{\tau}, \eta_1, \eta_2)$ -adaptable w.r.t. the data state \mathbf{d} and the environment evolution \mathcal{E} if $\forall \mathbf{d}' \in \mathcal{D}$ with $m_{\rho, 0}^{\mathcal{D}}(\mathbf{d}, \mathbf{d}') \leq \eta_1$ it holds $m_{\{\tau \in OT \mid \tau \geq \tilde{\tau}\}}^{\lambda}(\langle P, \mathbf{d} \rangle_{\mathcal{E}}, \langle P, \mathbf{d}' \rangle_{\mathcal{E}}) \leq \eta_2$.*



(a) In blue: pointwise distance between c_1 and c_2 . In red: $m_{\rho, \{\tau' \geq \tau\}}^\lambda(c_1, c_2)$, for each τ (Ex. 7). (b) Adaptability of P in c_1 (from Ex. 6) for $M = 100$, $\eta_1 = 0.2$ (Ex. 8).

Fig. 3: Examples of the evaluation of the evolution metric (assuming λ being the constant 1).

We remark that one can always consider the data state \mathbf{d} as the ideal model of the world used for the specification of P , and the data state \mathbf{d}' as the real world in which P has to execute. Hence, the idea behind adaptability is that even if the initial behaviour of the two systems is quite different, P is able to reduce the gap between the real evolution and the desired one within the time threshold $\tilde{\tau}$. Notice that being $(\tilde{\tau}, \eta_1, \eta_2)$ -adaptable for $\tilde{\tau} = \min\{\tau \mid \tau \in \text{OT}\}$ is equivalent to being (τ, η_1, η_2) -adaptable for all $\tau \in \text{OT}$.

The notion of reliability strengthens that of adaptability by bounding the distance on the evolution sequences from the beginning. A program is reliable if it guarantees that small variations in the initial conditions cause only bounded variations in its evolution.

Definition 13 (Reliability). Let $\eta_1, \eta_2 \in [0, 1)$. We say that P is (η_1, η_2) -reliable w.r.t. the data state \mathbf{d} and the environment evolution \mathcal{E} if $\forall \mathbf{d}' \in \mathcal{D}$ with $m_{\rho, 0}^{\mathcal{D}}(\mathbf{d}, \mathbf{d}') \leq \eta_1$ it holds $m_{\text{OT}}^\lambda(\langle P, \mathbf{d} \rangle_{\mathcal{E}}, \langle P, \mathbf{d}' \rangle_{\mathcal{E}}) \leq \eta_2$.

We can use our algorithm to verify adaptability and reliability of a given program. Given a configuration $\langle P, \mathbf{d} \rangle_{\mathcal{E}}$, a set OT of observation times and a given threshold $\eta_1 \geq 0$, we can sample M variations $\{\mathbf{d}_1, \dots, \mathbf{d}_M\}$ of \mathbf{d} , s.t. for any i , $m_{\rho, 0}^{\mathcal{D}}(\mathbf{d}, \mathbf{d}_i) \leq \eta_1$. Then, for each sampled data state we can estimate the distance between $c = \langle P, \mathbf{d} \rangle_{\mathcal{E}}$ and $c_i = \langle P, \mathbf{d}_i \rangle_{\mathcal{E}}$ at the different time steps in OT, namely $m_{\{\tau \in \text{OT} \mid \tau \geq \tilde{\tau}\}}^\lambda(c, c_i)$ for any $\tilde{\tau} \in \text{OT}$. Finally, for each $\tilde{\tau} \in \text{OT}$, we let $l_{\tilde{\tau}} = \max_i \{m_{\{\tau \in \text{OT} \mid \tau \geq \tilde{\tau}\}}^\lambda(c, c_i)\}$. We can observe that, for the chosen η_1 , each $l_{\tilde{\tau}}$ gives us a lower bound to the $\tilde{\tau}$ -adaptability of the program. Similarly, for $\tau_{\min} = \min_{\text{OT}} \tau$, $l_{\tau_{\min}}$ gives a lower bound for its reliability.

Example 8. Figure 3b shows the evaluation of $l_{\tilde{\tau}}$ for the program P in the configuration c_1 from Example 6 with parameters $M = 100$ and $\eta_1 = 0.2$. Observe that the initial perturbation is not amplified and after 12 steps it is almost absorbed. In particular, our program is $(12, 0.2, \eta_2)$ -adaptable w.r.t. the data state and the environment evolution in Example 6, for any $\eta_2 \geq 0.05 + e_{\mathbf{W}}^{12}$, where $e_{\mathbf{W}}^{12}$ is the approximation error $e_{\mathbf{W}}^{12} = |\mathbf{W}(m_{\rho, 12}^{\mathcal{D}})(\hat{\mathcal{S}}_{c_1, 12}^{\mathcal{D}, 1000}, \hat{\mathcal{S}}_{c_2, 12}^{\mathcal{D}, 10000}) - \mathbf{W}(m_{\rho, 12}^{\mathcal{D}})(\mathcal{S}_{c_1, 12}^{\mathcal{D}}, \mathcal{S}_{c_2, 12}^{\mathcal{D}})|$. We refer the interested reader to [36, Corollary 3.5, Equation (3.10)] for an estimation of $e_{\mathbf{W}}^{12}$.

7 Concluding remarks

As a first step for future research we will provide a simple logic, defined in the vein of *Signal Temporal Logic* (STL) [30], that can be used to specify requirements on the evolution sequences of a system. Our intuition is that we can exploit the evolution metric, and the algorithm we have proposed, to develop a quantitative model checking tool for this type of systems. Moreover, we would like to enhance the modelling of time. Firstly we could relax the timing constraints on the evolution metric by introducing a *time tolerance* and defining a *stretched evolution metric* as a Skorokhod-like metric [35], as those used for conformance testing [18]. Then, we could provide an extension of our techniques to the case in which also the program shows a continuous time behaviour.

The use of metrics for the analysis of systems stems from [17, 22, 29] where, in a process algebraic setting, it is argued that metrics are indeed more informative than behavioural equivalences when quantitative information on the behaviour is taken into account. The Wasserstein lifting has then found several successful applications: from the definition of *behavioural metrics* (e.g., [7, 12, 20]), to privacy [9, 10, 15] and machine learning (e.g., [4, 25, 38]). Usually, one can use behavioural metrics to quantify how well an implementation (I) meets its specification (S). In [14] the authors do so by setting a two players game with weighted choices, and the cost of the game is interpreted as the distance between I and S . Hence the authors propose three distance functions: *correctness*, *coverage*, and *robustness*. Correctness expresses how often I violates S , coverage is its dual, and robustness measures how often I can make an unexpected error with the resulting behaviour still meeting S . A similar game-based approach is used in [13] to define a *masking fault-tolerant* distance. Briefly, a system is masking fault-tolerant if faults do not induce any observable behaviour in the system. Hence, the proposed distance measures how many faults are tolerated by I while being masked by the states of the system. Notice that the notion of robustness from [14] and the masking fault-tolerant distance from [13] are quite different from our reliability. In fact, we are not interested in counting how many times an error occurs, but in checking whether the system is able to regain the desired behaviour after the occurrence of an error.

Systems showing an highly dynamic behaviour are usually modelled as Stochastic Hybrid Systems (SHSs) (see, e.g., [6, 8]), which allow for combining in a single model the discrete, continuous and probabilistic features of systems. Our model clearly belongs to a subclass of SHSs. However, as previously outlined, our approach differs from that of SHSs since we model the program, the environment and their interaction (the data state) as three distinct components. This choice allows us to study the behaviour of the program by means of the evolution metric. It would be interesting to investigate if, and how, our method can be extended to the general class of SHSs.

We remark here that our objective in this paper was to provide some tools for the *analysis* of the interaction of a *given* program with the environment, and not for the *synthesis* of a program. However, for programs that are *controllers*, some metric-based approaches have been proposed for their synthesis [3, 21]. We will then study whether our approach can be combined with some learning techniques in order to design and synthesise robust controllers.

References

1. Abate, A., D’Innocenzo, A., Benedetto, M.D.D.: Approximate abstractions of stochastic hybrid systems. *IEEE Trans. Automat. Contr.* **56**(11), 2688–2694 (2011)
2. Abate, A., Katoen, J., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. *Eur. J. Control* **16**(6), 624–641 (2010)
3. Abate, A., Prandini, M.: Approximate abstractions of stochastic systems: A randomized method. In: *Proceedings of CDC-ECC 2011*. pp. 4861–4866 (2011)
4. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: *Proceedings of ICML 2017*. pp. 214–223 (2017)
5. Bernardo, M., Nicola, R.D., Loreti, M.: A uniform framework for modeling nondeterministic, probabilistic, stochastic, or mixed processes and their behavioral equivalences. *Inf. Comput.* **225**, 29–82 (2013)
6. Bloom, H.A.P., (eds.), J.L.: *Stochastic Hybrid Systems: Theory and Safety Critical Applications*. No. 337 in *Lecture Notes in Control and Information Sciences*, Springer-Verlag, Berlin/Heidelberg (2006)
7. van Breugel, F.: A behavioural pseudometric for metric labelled transition systems. In: *Proceedings of CONCUR 2005*. pp. 141–155 (2005)
8. Cassandras, C.G., Lygeros, J., (eds.): *Stochastic Hybrid Systems*. No. 24 in *Control Engineering*, CRC Press, Boca Raton, 1st edn. (2007)
9. Castiglioni, V., Chatzikokolakis, K., Palamidessi, C.: A logical characterization of differential privacy via behavioral metrics. In: *Proceedings of FACS 2018*. *Lecture Notes in Computer Science*, vol. 11222, pp. 75–96 (2018)
10. Castiglioni, V., Chatzikokolakis, K., Palamidessi, C.: A logical characterization of differential privacy. *Sci. Comput. Program.* **188**, 102388 (2020)
11. Castiglioni, V., Loreti, M., Tini, S.: Measuring adaptability and reliability of large scale systems. In: *Proceedings of ISoLA 2020*. *Lecture Notes in Computer Science*, vol. 12477, pp. 380–396 (2020)
12. Castiglioni, V., Loreti, M., Tini, S.: The metric linear-time branching-time spectrum on non-deterministic probabilistic processes. *Theor. Comput. Sci.* **813**, 20–69 (2020)
13. Castro, P.F., D’Argenio, P.R., Demasi, R., Putruele, L.: Measuring masking fault-tolerance. In: *Proceedings of TACAS 2019, Part II*. *Lecture Notes in Computer Science*, vol. 11428, pp. 375–392 (2019)
14. Cerný, P., Henzinger, T.A., Radhakrishna, A.: Simulation distances. *Theor. Comput. Sci.* **413**(1), 21–35 (2012)
15. Chatzikokolakis, K., Gebler, D., Palamidessi, C., Xu, L.: Generalized bisimulation metrics. In: *Proceedings of CONCUR 2014*. pp. 32–46 (2014)
16. Ciocchetta, F., Hillston, J.: Bio-pepa: An extension of the process algebra PEPA for biochemical networks. *Electron. Notes Theor. Comput. Sci.* **194**(3), 103–117 (2008)
17. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theor. Comput. Sci.* **318**(3), 323–354 (2004)
18. Deshmukh, J.V., Majumdar, R., Prabhu, V.S.: Quantifying conformance using the skorokhod metric. *Formal Methods in System Design* **50**(2-3), 168–206 (2017)
19. Faugeras, O.P., Rüschemdorf, L.: Risk excess measures induced by hemi-metrics. *Probability, Uncertainty and Quantitative Risk* **3**:6 (2018)
20. Gebler, D., Larsen, K.G., Tini, S.: Compositional bisimulation metric reasoning with probabilistic process calculi. *Log. Methods Comput. Sci.* **12**(4) (2016)
21. Ghosh, S., Bansal, S., Sangiovanni-Vincentelli, A.L., Seshia, S.A., Tomlin, C.: A new simulation metric to determine safe environments and controllers for systems with unknown dynamics. In: *Proceedings of HSCC 2019*. pp. 185–196 (2019)

22. Giacalone, A., Jou, C.C., Smolka, S.A.: Algebraic reasoning for probabilistic concurrent systems. In: Proceedings of IFIP Work, Conf. on Programming, Concepts and Methods. pp. 443–458 (1990)
23. Girard, A., Gößler, G., Mouelhi, S.: Safety controller synthesis for incrementally stable switched systems using multiscale symbolic models. *IEEE Trans. Automat. Contr.* **61**(6), 1537–1549 (2016)
24. van Glabbeek, R.J., Smolka, S.A., Steffen, B.: Reactive, generative and stratified models of probabilistic processes. *Inf. Comput.* **121**(1), 59–80 (1995)
25. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: Proceedings of Advances in Neural Information Processing Systems. pp. 5767–5777 (2017)
26. Heredia, G., Jimenez-Cano, A.E., Sánchez, I., Llorente, D., Vega, V., Braga, J., Acosta, J.Á., Ollero, A.: Control of a multicopter outdoor aerial manipulator. In: Proceedings of IROS 2014. pp. 3417–3422. IEEE (2014)
27. Hillston, J., Hermans, H., Herzog, U., Mertsiotakis, V., Rettelbach, M.: Stochastic process algebras: integrating qualitative and quantitative modelling. In: Proceedings of International Conference on Formal Description Techniques 1994. IFIP, vol. 6, pp. 449–451 (1994)
28. Hu, J., Lygeros, J., Sastry, S.: Towards a theory of stochastic hybrid systems. In: Proceedings of HSCC 2000. Lecture Notes in Computer Science, vol. 1790, pp. 160–173 (2000)
29. Kwiatkowska, M.Z., Norman, G.: Probabilistic metric semantics for a simple language with recursion. In: Proceedings of MFCS’96. Lecture Notes in Computer Science, vol. 1113, pp. 419–430 (1996)
30. 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)
31. Malhame, R., Yee Chong, C.: Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Trans. Automat. Contr.* **30**(9), 854–660 (1985)
32. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley Series in Probability and Statistics, Wiley, United States (2005)
33. Rachev, S.T., Klebanov, L.B., Stoyanov, S.V., Fabozzi, F.J.: The Methods of Distances in the Theory of Probability and Statistics. Springer (2013)
34. Segala, R.: Modeling and verification of randomized distributed real-time systems. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA (1995)
35. Skorokhod, A.V.: Limit theorems for stochastic processes. *Theory Probab. Appl.* **1**, 261–290 (1956)
36. Sriperumbudur, B.K., Fukumizu, K., Gretton, A., Schölkopf, B., Lanckriet, G.R.G.: On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics* **6**, 1550–1599 (2021)
37. Thorsley, D., Klavins, E.: Approximating stochastic biochemical processes with Wasserstein pseudometrics. *IET Syst. Biol.* **4**(3), 193–211 (2010)
38. Tolstikhin, I.O., Bousquet, O., Gelly, S., Schölkopf, B.: Wasserstein auto-encoders. In: Proceedings of ICLR 2018 (2018)
39. Vaserstein, L.N.: Markovian processes on countable space product describing large systems of automata. *Probl. Peredachi Inf.* **5**(3), 64–72 (1969)
40. Villani, C.: Optimal transport: old and new, vol. 338. Springer (2008)
41. Virágh, C., Nagy, M., Gershenson, C., Vásárhelyi, G.: Self-organized UAV traffic in realistic environments. In: Proceedings of IROS 2016. pp. 1645–1652. IEEE (2016)