



This is a repository copy of *Conformance relations and hyperproperties for doping detection in time and space*.

White Rose Research Online URL for this paper:  
<https://eprints.whiterose.ac.uk/183305/>

Version: Published Version

---

**Article:**

Biewer, S., Dimitrova, R., Fries, M. et al. (4 more authors) (2022) Conformance relations and hyperproperties for doping detection in time and space. *Logical Methods in Computer Science*, 18 (1). 14. ISSN 1860-5974

[https://doi.org/10.46298/lmcs-18\(1:14\)2022](https://doi.org/10.46298/lmcs-18(1:14)2022)

---

**Reuse**

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:  
<https://creativecommons.org/licenses/>

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

## CONFORMANCE RELATIONS AND HYPERPROPERTIES FOR DOPING DETECTION IN TIME AND SPACE

SEBASTIAN BIEWER<sup>a</sup>, RAYNA DIMITROVA<sup>b</sup>, MICHAEL FRIES<sup>c</sup>, MACIEJ GAZDA<sup>d</sup>,  
THOMAS HEINZE<sup>c</sup>, HOLGER HERMANN<sup>a,e</sup>, AND MOHAMMAD REZA MOUSAVI<sup>f</sup>

<sup>a</sup> Saarland University — Computer Science, Saarland Informatics Campus, Saarbrücken, Germany

<sup>b</sup> CISA Helmholtz Center for Information Security, Saarbrücken, Germany

<sup>c</sup> Automotive Powertrain Institute, htw saar — University of Applied Sciences, Saarbrücken, Germany

<sup>d</sup> Department of Computer Science, University of Sheffield, Sheffield, UK

<sup>e</sup> Institute of Intelligent Software, Guangzhou, China

<sup>f</sup> Department of Informatics, King’s College London, London, UK

---

**ABSTRACT.** We present a novel and generalised notion of doping cleanness for cyber-physical systems that allows for perturbing the inputs and observing the perturbed outputs both in the time—and value—domains. We instantiate our definition using existing notions of conformance for cyber-physical systems. As a formal basis for monitoring conformance-based cleanness, we develop the temporal logic **HyperSTL\***, an extension of Signal Temporal Logics with trace quantifiers and a freeze operator. We show that our generalised definitions are essential in a data-driven method for doping detection and apply our definitions to a case study concerning diesel emission tests.

### 1. INTRODUCTION

System doping, in our terminology, is an intentional intervention causing a change in the system’s normal behaviour against the interests of the user or other stakeholders (such as the society at large). Examples of system doping are widespread and range from vendors’ enforcing a monopoly on chargers and spare parts (by checking for and refusing third-party chargers and spare parts, respectively) to tampering with exhaust emission in order to detect and pass emission tests. Doping can be the result of embedding a piece of code or smuggling a piece of electronic circuit into the system and it can be caused by the original developers or by hackers. Software and system doping has been studied in the past couple of years and rigorous theories for it have been developed [BDFH16, DBB<sup>+</sup>17, BDH18]. These theories were subsequently adopted in order to detect doping, or formally, to check system cleanness [HBDK18, BDH19] (corresponding to the absence of doping).

---

*Key words and phrases:* Software Doping, Conformance Testing, Hybrid Systems.






Figure 1: Running Example: Specified (a) and Actual (b) Test Cycles and Emission Footprints obtained from Different (Fictitious) Vehicles (c) and (d).

In the present paper, we extend the theory of doping to the setting of cyber-physical systems (CPS) by exploiting the notions of conformance testing for CPS [AMF14, DMP17, KM15]. The existing theories of software doping define doping in terms of drastic deviations in output as a result of minor deviations in input, where the term “deviation” refers to differences in validity of propositions or values of variables. However, the current notions come short of properly dealing with the issues of retiming and delays, which are commonly present in the signals of CPS. We observe that this is an essential aspect of detecting doping for cyber-physical systems: often the traces to be tested for doping have subtly different timing behaviour, e.g., due to measurement and calibration errors or due to the slight deviations of human actors in acting upon the planned scenarios. The insufficient treatment of retiming and delays can both lead to false negatives, i.e., missing cases of doping, as well as false positives, i.e., reporting spurious doping cases.

To address these issues, we exploit the notion of conformance to devise a general theory of being clean from doping and instantiate that theory with some existing notions of conformance for hybrid systems. We show how these notions can account for retiming and lead to more precise notions of cleanness. Furthermore, we show how the retiming can be synchronized between input and output, leading to a refined notion of cleanness with a rigorous account of the relation between the retiming of input and output.

We illustrate the usefulness of our theory by empirical analysis of diesel engine exhaust emissions in the context of one of the official test cycles, the New European Driving Cycle (NEDC) [Uni13]. In particular, we show that catering for retiming is essential in effectively exploiting the actual driving cycles for performing doping analysis. We thus demonstrate that our new theory remedies a major shortcoming in the existing notions from the literature. To facilitate the presentation, we use throughout the remainder of this paper the following simple running example, which is inspired by our case study.

**Example 1.1.** Fig. 1.(a) shows two test cycles (evolution of speed over time), designed to detect whether the exhaust emission control of a particular vehicle is doped. The test cycle  $i_{st}$ , depicted with a black solid line, is the standard one prescribed by the (fictitious) official regulation, while test cycle  $i_{dev}$ , depicted by a red dotted line, is a slight deviation thereof. If the exhaust emissions measured during the test cycle  $i_{dev}$  turn out to be significantly higher than the ones measured in test cycle  $i_{st}$ , then we can conclude that the exhaust emission system is potentially doped, since it appears tailored to the standard test cycle.

Fig. 1.(b) addresses a notorious problem of testing cars: a human tester is supposed to drive the car as just described, however, she can do this only up to a certain imprecision. Assume her driving of  $i_{dev}$  exhibits a slight time shift  $\tau$  relative to the test cycle, as in  $i_{ddev}$ , while  $i_{st}$  is being driven as intended.

The result of a test is the emission footprint measured at the exhaust pipe of the car. Fig. 1.(c) and Fig. 1.(d) show two different possible test results (obtained from different cars) for the scenario in Fig. 1.(b). Intuitively, the footprints in Fig. 1.(c) provide significant evidence for doping — a slightly different test cycle has resulted in significantly larger footprint. However, due to the time shift on the input side in Fig. 1.(b) the point-wise difference of the two driven test-cycles has grown very large. As we show in the remainder of this paper, the existing theory of doping fails to detect such a clear evidence, due to the minor delay during the execution of the driving cycle. The emission footprint in Fig. 1.(d) is another (synthetic) example of a significant deviation which cannot be detected for the input in Fig. 1.(b) using existing theories; this latter footprint sheds some light on the intricate design decisions in the theory we develop in this paper.

The contributions for this paper can be summarized as follows:

- We define a *general notion of conformance* that can express different ways of comparing execution traces by allowing deviations both in value and in time;
- We define a *general notion of cleanness for hybrid systems*, and show that it subsumes the existing notion of robust cleanness [DBB<sup>+</sup>17];
- We define the notion of synchronized retiming, which provides a rigorous tool for relating the retiming of input and output and use it to produce a refined notion of cleanness;
- We provide a logical account of cleanness (based on the notion of hybrid conformance) by developing a temporal logic, called *HyperSTL\**, that extends Signal Temporal Logic (STL) with a freeze operator and quantifiers over traces; and
- We demonstrate the usefulness of the proposed generic framework by applying it to *software doping tests* in the automotive domain, where we show that the new cleanness definition is able to flag a case of software doping that goes unnoticed when robust cleanness is used.

This paper substantially extends the theoretical material and the experimental results published in the earlier conference publication [DGM<sup>+</sup>20]. In particular, the following contributions of the present paper are new with respect to the earlier conference publication:

- the notion of cleanness with synchronised retiming in Section 5 and its application in Section 7 for doping detection in our case study,
- the logical approach to cleanness in Section 6, the introduction of our logical formalism *HyperSTL\**, its monitoring, and its application to specify hybrid conformance, and
- the redesign of our experimental setup to comply with the NEDC test cycle (with preconditioning and control of ambient temperature), as well as several new experiments to make full use of our theories of retiming. These new experiments led to much more substantial and decisive evidence for our case study.

## 2. RELATED WORK

The term “software doping” was coined around 2015 [HHB15] in media uncovering the diesel exhaust emissions scandal. An informal problem formulation [BDFH16] pointed out the general phenomenon of intentionally added hidden software behaviour, which is not in the interest of the consumer. Shortly after, this observation has been complemented by a set of formal *cleanness* definitions [DBB<sup>+</sup>17] laying the theoretical foundations upon which formal methods to detect such software behaviour can be used. It is possible to detect missing functionality and undesired existing functionality. The definitions support both sequential programs and nondeterministic reactive programs. To check satisfaction of the

definitions, it is necessary to compare two (or more) execution traces of the same system. Such properties are called *hyperproperties* [CS08] (whereas classical properties are *trace properties*). Tool support for analysing hyperproperties typically requires high computational effort [CFK<sup>+</sup>14, FRS15]. There exist several temporal logics for analysing satisfaction of trace properties of various kinds of systems, one of them being *Linear Temporal Logic* (LTL) [Pnu77] for systems producing outputs in discrete time steps and properties that do not consider the time passing between outputs. LTL has been extended to the logic HyperLTL, which can express hyperproperties by allowing explicit quantification of execution traces in front of an LTL formula [CFK<sup>+</sup>14]. Tools for model-checking boolean circuits, satisfiability and monitoring of HyperLTL specifications have been developed [AB16, BSB17, FRS15, FH16, FHS17, FHST17, FHST18, HST19].

*Signal Temporal Logic* (STL) [MN04] is an extension of LTL that adds support for time constraints and real-valued signals. Tools exist that automatically try to falsify STL formulas [Don10, ALFS11]. There has been an extension of STL to HyperSTL in a similar fashion as it was done for HyperLTL [NKJ<sup>+</sup>17]. The syntax of HyperSTL, however, is not able to express the cleanness definitions (for deterministic systems) in a way that allows (efficient) falsification. Freeze Temporal Logic [AH94] introduces a freeze quantifier to “record” the moment of time while evaluating a formula and later use the recorded time to measure time lapse. The concept of freeze quantifier has been used to specify properties of data-dependent systems (using models such as register automata) [DL09] and dynamical systems [BDSV14, DMP17]. In the present paper, we introduce the concept of freeze quantifier [BDSV14] into a logic that is inspired by HyperSTL [NKJ<sup>+</sup>17]. The combination turns out to be expressive enough to specify cleanness with respect to hybrid conformance. Moreover, we use a monitoring procedure inspired by an earlier extension of STL with freeze operator [STB<sup>+</sup>19] to check for cleanness with respect to hybrid conformance. *Robust cleanness* is defined for distance functions on inputs and outputs [DBB<sup>+</sup>17]. When used with temporal logics the distance functions are restricted to those compatible with the logics. To be fully independent, robust cleanness analysis has been embedded into the theory of model-based testing [BDH19] with input-output conformance [Tre92, Tre96].

Notions of conformance for discrete event systems have been discussed for almost a century. The earliest work on this topic dates back to 1960’s when researchers studied model-based testing of digital circuits using Finite State Machine models [Hen64, LY96]. Concurrency theory contributed ideas to this field, such as decoupling (i.e., removing the synchronised assumption between) inputs and outputs and observing failures to engage in a communication (and more specifically quiescence) [DH84, Tre92]. A theory of conformance testing for systems with continuous dynamics was developed by Michiel van Osch [vO06]; this theory did not gain much popularity in practice, partly because of its insufficient treatment of approximation (e.g., differences in values and retiming). Pappas and Girard [GJP08, GP11] proposed the use of Metric Bisimulation for conformance checking in dynamical systems and Pappas and Fainekos [FP09] developed a falsification framework for the same purpose. This research led to two notions of conformance used in the present paper, namely hybrid conformance by Abbas and Fainekos [AMF14] and Skorokhod conformance by Deshmukh, Majumdar, and Prabhu [DMP17].

### 3. PRELIMINARIES

**3.1. Semantic domain.** In this section, we provide definitions regarding semantic domain, conformance, and robust cleanness. We begin with the definition of our semantic domain, called generalised timed traces [GM20]. This definition subsumes both discrete-time state sequences and continuous-time trajectories. A generalised timed trace is a function with a discrete or continuous domain (called time domain) and a co-domain which is a metric space. Intuitively, a generalised timed trace maps each element of its time domain to a state. We require that the set of possible states is a metric space since we study conformance notions that compare traces based on the distance between the states of the traces.

**Definition 3.1.** Let  $(\mathcal{Y}, d_{\mathcal{Y}})$  be a metric space. A  $\mathcal{Y}$ -valued *generalised timed trace* (GTT) is a function  $\mu : \mathcal{T} \rightarrow \mathcal{Y}$  such that  $\mathcal{T} \subseteq \mathbb{R}_{\geq 0}$ . We call  $\mathcal{T}$  the *time domain* of  $\mu$ , denoted  $\text{dom}(\mu)$ .  $\text{GTT}(\mathcal{Y})$  is the set of all  $\mathcal{Y}$ -valued generalised timed traces.

For a GTT  $\mu : \mathcal{T} \rightarrow \mathcal{Y}$  and time  $t_0 \in \mathcal{T}$ , by  $\mu[\dots t_0]$  we denote the prefix of  $\mu$  up to  $t_0$ , i.e., the restriction  $\mu|_{t \in \mathcal{T}: t \leq t_0}$ ; likewise, by  $\mu[t_s \dots t_e]$ , we shall denote the restriction  $\mu|_{t \in \mathcal{T}: t_s \leq t \leq t_e}$ .

A hybrid system is a mapping from generalised (input) timed traces to sets of generalised (output) timed traces.

**Definition 3.2.** A  $\mathcal{Y}$ -valued hybrid system is a function  $H : \text{GTT}(\mathcal{Y}) \rightarrow \mathcal{P}(\text{GTT}(\mathcal{Y}))$  such that for all  $\mu \in \text{GTT}(\mathcal{Y})$  and all  $\mu' \in H(\mu)$  it holds that  $\text{dom}(\mu') = \text{dom}(\mu)$ . We define  $\mathcal{H}(\mathcal{Y})$  to be the set of all  $\mathcal{Y}$ -valued hybrid systems.

In addition, we distinguish deterministic hybrid systems whose output values range over singleton sets only. In what follows, we identify deterministic hybrid systems with functions of the type  $\text{GTT}(\mathcal{Y}) \rightarrow \text{GTT}(\mathcal{Y})$ .

For simplicity, we assume that the input and output domain are defined on the same metric spaces. The generalisation to different spaces is straightforward.

**3.2. Conformance relations.** Recently, a number of notions of conformance for cyber-physical systems have been proposed [ARM17, KM15]. It turns out that these notions, two of which are quoted below, can provide a rigorous basis for doping detection.

Note that throughout the paper, the variables  $\tau$  and  $\epsilon$  (with possible subscripts) always range over non-negative real numbers.

**Definition 3.3.** We say that  $\mathcal{Y}$ -valued GTTs  $\mu_1 : \mathcal{T}_1 \rightarrow \mathcal{Y}$  and  $\mu_2 : \mathcal{T}_2 \rightarrow \mathcal{Y}$  are:

- *trace conformant* with tolerance threshold for signal value  $\epsilon$ , notation  $\text{TraceConf}_{\epsilon}(\mu_1, \mu_2)$ , if  $\mathcal{T}_1 = \mathcal{T}_2$  and for all  $t \in \mathcal{T}_1$ ,  $d_{\mathcal{Y}}(\mu_1(t), \mu_2(t)) \leq \epsilon$
- *hybrid conformant* with thresholds  $\tau$  and  $\epsilon$ , denoted  $\text{HybridConf}_{\tau, \epsilon}(\mu_1, \mu_2)$ , if:
  - $\forall t_1 \in \mathcal{T}_1 \exists t_2 \in \mathcal{T}_2 : |t_2 - t_1| \leq \tau \wedge d_{\mathcal{Y}}(\mu_2(t_2), \mu_1(t_1)) \leq \epsilon$
  - $\forall t_2 \in \mathcal{T}_2 \exists t_1 \in \mathcal{T}_1 : |t_1 - t_2| \leq \tau \wedge d_{\mathcal{Y}}(\mu_1(t_1), \mu_2(t_2)) \leq \epsilon$
- *Skorokhod conformant* with tolerance thresholds  $\tau$  and  $\epsilon$ , notation  $\text{SkorConf}_{\tau, \epsilon}(\mu_1, \mu_2)$ , if  $\mathcal{T}_1$  and  $\mathcal{T}_2$  are intervals and there is a strictly increasing continuous bijection  $r : \mathcal{T}_1 \rightarrow \mathcal{T}_2$  called retiming, such that:
  - for all  $t \in \mathcal{T}_1$ ,  $|r(t) - t| \leq \tau$ , and
  - for all  $t \in \mathcal{T}_1$ ,  $d_{\mathcal{Y}}(\mu_1(t), \mu_2(r(t))) \leq \epsilon$ .

We show in the proposition below and also in our generalisation results in Section 4.2, that these notions are closely related. However, they also have some fundamental differences, that can be illustrated using the example in Fig. 1.

**Example 3.4.** Consider again the example shown in Fig. 1.

We can see that in Fig. 1.(a)  $i_{st}$  and  $i_{dev}$  are trace conformant with value threshold  $\epsilon$ , as they only exhibit point-wise deviations by values less than  $\epsilon$ . In contrast,  $i_{st}$  and  $i_{ddev}$  in Fig. 1.(b) are not trace conformant, yet they are hybrid conformant with time and value margins  $\tau$  and  $\epsilon$ , respectively. The key difference is that the inputs depicted in Fig. 1.(b) are very different if compared point-wise, but if one allows for retiming, they are close enough in value after retiming.

The outputs  $o'(i_{st})$  and  $o'(i_{ddev})$  in Fig. 1.(d) illustrate the fundamental difference between hybrid and Skorokhod conformance: although the order of rising and falling signals are reversed in the two trajectories, they are still hybrid conformant, because hybrid conformance disregards the order. However, Skorokhod conformance requires an order-preserving retiming, and hence distinguishes these two trajectories. On the other hand, such retiming exists, e.g., for  $i_{st}$  and  $i_{ddev}$  in Fig. 1.(b), witnessing their Skorokhod conformance.

We shall use the following notation. We write  $\text{Conf}_1 \sqsubseteq \text{Conf}_2$  whenever for all  $\mu_1 : \mathcal{T}_1 \rightarrow \mathcal{Y}$  and  $\mu_2 : \mathcal{T}_2 \rightarrow \mathcal{Y}$ , we have  $\text{Conf}_1(\mu_1, \mu_2) \implies \text{Conf}_2(\mu_1, \mu_2)$ . We write  $\text{Conf}_1 \sqsubset \text{Conf}_2$  whenever  $\text{Conf}_1 \sqsubseteq \text{Conf}_2$  and  $\neg \text{Conf}_2 \sqsubseteq \text{Conf}_1$ .

**Proposition 3.5.** *For any  $\tau, \epsilon \in \mathbb{R}_{\geq 0}$ , the following relations hold:*

$$\text{TraceConf}_\epsilon \sqsubset \text{SkorConf}_{\tau, \epsilon} \sqsubset \text{HybridConf}_{\tau, \epsilon}$$

**3.3. Robust cleanness.** We shall now state the original definition of robust cleanness from [DBB<sup>+</sup>17], adapted to our framework of hybrid systems. It is based on Definition 7 and Proposition 19 from [DBB<sup>+</sup>17]; the phrasing below abstracts from the so-called parameters of interest and standard inputs. Moreover it is cast in the setting of generalised timed traces rather than discrete-step programs, and stated using trace conformance with different thresholds for inputs and outputs,  $\kappa_I$  and  $\kappa_O$ .

Intuitively, a hybrid system is robustly clean if for every pair of input prefixes on which no difference in the inputs exceeding  $\kappa_I$  has occurred so far (i.e., all sub-prefixes are trace conformant), the corresponding sets of output prefixes are also conformant with respect to  $\kappa_O$ . As we consider nondeterministic systems, Hausdorff distance is used to compare sets of outputs (see [DBB<sup>+</sup>17] for details).

**Definition 3.6.** A hybrid system  $H$  is robustly clean, denoted  $\text{RobustClean}(\kappa_I, \kappa_O)$ , whenever:

$$\begin{aligned} & \forall i_1, i_2 \in \text{GTT}(\mathcal{Y}) : \forall t \in \text{dom}(i_1) \cup \text{dom}(i_2) : \\ & (\forall t' \leq t : \text{TraceConf}_{\kappa_I}(i_1[\dots t'], i_2[\dots t'])) \implies \\ & ((\forall o_1 \in H(i_1) \exists o_2 \in H(i_2) : \text{TraceConf}_{\kappa_O}(o_1[\dots t], o_2[\dots t])) \wedge \\ & (\forall o_2 \in H(i_2) \exists o_1 \in H(i_1) : \text{TraceConf}_{\kappa_O}(o_1[\dots t], o_2[\dots t]))) \end{aligned}$$

Note that in the above definition we do not require that  $\text{dom}(i_1) = \text{dom}(i_2)$ . In practice, robust cleanness is typically applied to pairs of traces that are both defined over  $\mathbb{N}$ . Here, however, for the sake of generality we impose no such restriction. In particular, when the

time domains of two traces are different, for example disjoint, the predicate `RobustClean` will trivially evaluate to *true*.

**Example 3.7.** Consider the traces depicted in Fig. 1. The input prefixes  $i_{st}$  and  $i_{ddev}$  are given in Fig. 1.(b), and the corresponding pair of outputs is shown in Fig. 1.(c). For  $t \leq t_0$ , we have  $i_{ddev}(t) = 0$  and  $o(i_{ddev})(t) = 0$ . The trace  $i_{st}$  results in output  $o(i_{st})$  and  $i_{ddev}$  results in  $o(i_{ddev})$ . Suppose that  $\epsilon < |i_{st}(t_0) - i_{ddev}(t_0)|$ , and for every  $t < t_0$  it holds that  $|o(i_{st})(t) - o(i_{ddev})(t)| \leq |i_{st}(t) - i_{ddev}(t)|$ . Furthermore,  $\epsilon < |o(i_{st})(t_1) - o(i_{ddev})(t_1)|$  at some time  $t_1 \geq t_0$ . Consider Def. 3.6 instantiated with  $\kappa_I = \kappa_O = \epsilon$ . Clearly, for every  $t < t_0$  the implication in Def. 3.6 is satisfied, since  $|o(i_{st})(t') - o(i_{ddev})(t')| \leq |i_{st}(t') - i_{ddev}(t')|$  for all  $t' \leq t < t_0$ , and  $\kappa_I = \kappa_O$ . For  $t \geq t_0$  the implication is true as well: for  $t' < t_0$  the reasoning is as above, and for all  $t' \geq t_0$  the left-hand side of the implication is false. Hence, regardless of the difference in the output values at  $t_1$ , this pair of inputs satisfies the condition of `RobustClean`( $\epsilon, \epsilon$ ), and, if these are the only traces in a hybrid system  $H$  then we can conclude that  $H$  is `RobustClean`( $\epsilon, \epsilon$ ).

#### 4. CONFORMANCE-BASED CLEANNES

We now define a general notion of conformance-based cleanness and provide two instantiations based on the conformance notions defined in the previous section.

**4.1. Motivation.** The need for considering disturbance in time as well as in value is motivated by our running example from Fig. 1. One of the challenges in performing doping tests for cyber-physical systems is that in such systems timing is rarely perfectly precise, due to imprecision in measurements, or caused by the interaction with the physical world. As illustrated in Example 1.1, for instance, when checking for software doping in a car [BDH19], the input to the system is the value of the car’s speed over time, which is under the control of a driver, and can thus vary from one execution to the other, even if the driver is trying to execute the same input sequence. Clearly, those variations can be in value, as well as in time.

**Example 4.1.** Consider the test setup sketched in Fig. 1. There,  $i_{st}$  and  $i_{ddev}$ , depicted in Fig. 1.(b) define the speed of a car as a function of time. These two input sequences follow a trajectory of values differing by a small margin  $\epsilon$  (the difference in value allowed by the standard defining the doping tests), but also shifted by a small unit of time  $\tau$ . Observe further that  $|i_{st}(t_0) - i_{ddev}(t_0)| \gg \epsilon$ . Thus, without allowing for deviations in time when comparing these input sequences, they will be considered sufficiently different, and as a result their respective exhaust emission outputs will fall out of the comparison when checking for doping according to Def. 3.6, even if the outputs  $H(i_{st}(t))$  and  $H(i_{ddev}(t))$  are vastly different, as depicted in Fig. 1.(c). This results in a false negative, i.e., failing to detect a clearly doped system.

In the above example, we demonstrated that not accounting for timing disturbances when relating input trajectories can result in false negatives in doping detection. Dually, using the traditional comparison for output traces can result in false positives by requiring overly strict matching of outputs.



The above example motivates the need to account for timing deviations in trajectories. Intuitively, for input trajectories this relaxation results in considering more traces as conforming, and thus enforcing more comparisons when checking if a system is clean. For output trajectories this means relaxing the conformance requirement by considering two output sequences as conforming even if their values are not perfectly aligned in time. Furthermore, different types of timing deviations need to be considered in different scenarios, for example, depending on whether the order in which values occur is important or not.

**Example 4.2.** Consider the testing workflow from Example 1.1 and Fig. 1, where inputs  $i_{st}$  and  $i_{ddev}$  are passed to a car. In the second experiment, depicted in Fig. 1.(d), the outputs of the car are  $o'(i_{st})$  and  $o'(i_{ddev})$ , which are hybrid conformant for  $\epsilon$  and  $\tau$ . Hence, this observation of the system is classified as clean under hybrid output conformance. However, the output  $o'(i_{ddev})$  is clearly suspicious, as the values in  $o'(i_{ddev})$  and  $o'(i_{st})$  are reversed. This motivates considering conformance notions that require retimings to be order-preserving. Indeed, using Skorokhod conformance we can detect that the system is doped.

The above examples show that in order to be useful in a diverse set of applications, a software cleanness theory should allow for using a variety of conformance notions. To this end, we next take a more general view on conformance notions, in order to be able to develop a generic conformance-based cleanness framework.

**4.2. Retimings and a more generic view on conformance notions.** So far, we have defined three specific notions of conformance which either coincide, or are closely inspired by ones that have appeared in the literature. In order to define a general framework for cleanness, we also wish to treat notions of conformance in a more generic manner. To this end, we propose an abstract definition of conformance predicates. As conformance predicates admit variations in time, as well as in value, our definition is based on *retimings*, a device that will play a key role in the context of this work. In its general form a retiming is a pair of functions between two time domains. Intuitively, given two GTTs, a retiming will define a mapping from points in each of the traces to points in the other trace. Note that in general the mappings are not required to be injective; this way we can cater for notions of conformance allowing for the so-called local disorder phenomenon (in particular hybrid conformance — see Proposition 4.5).

**Definition 4.3.** A *retiming* is a pair of functions between two time domains, i.e., a pair of the form  $(r_1, r_2)$ , where  $r_1 : \mathcal{T}_1 \rightarrow \mathcal{T}_2$  and  $r_2 : \mathcal{T}_2 \rightarrow \mathcal{T}_1$ , with time domains  $\mathcal{T}_1, \mathcal{T}_2 \subseteq \mathbb{R}_{\geq 0}$ . Given two time domains  $\mathcal{T}_1$  and  $\mathcal{T}_2$ , we denote the set of all retimings between  $\mathcal{T}_1$  and  $\mathcal{T}_2$  with  $\mathcal{RET}(\mathcal{T}_1, \mathcal{T}_2)$ .

Retiming is explicitly present in the definition of Skorokhod conformance; there, each Skorokhod retiming is required to be a strictly increasing continuous bijection. We can express a Skorokhod retiming  $r$  as an instance of our definition as the pair  $(r, r^{-1})$ . In fact, one can also define hybrid conformance, as well as a whole class of conformance notions, using a suitable *family* of retimings.

A family of retimings  $\text{Ret}$  can be further constrained by  $\tau$  to a subset  $\text{Ret}_\tau$  of  $\text{Ret}$  containing only functions that shift time by at most  $\tau$  time units. In order to use a family of retimings for concrete sequences  $\mu_1$  and  $\mu_2$ , it is necessary to consider only functions that match the domains of the sequences. This leads to a generic notion of conformance

associated with a given family of retimings  $\text{Ret}$ , a given time threshold  $\tau$  and a given value threshold  $\epsilon$ .

**Definition 4.4.** Let  $\text{Ret}$  be a family of retimings, and let

$$\begin{aligned} \text{Ret}_\tau &\triangleq \{(r_1, r_2) \in \text{Ret} \mid \forall t \in \text{dom}(r_i) : |r_i(t) - t| \leq \tau \ (i = 1, 2)\}, \\ \text{Ret}_\tau(\mathcal{T}_1, \mathcal{T}_2) &\triangleq \text{Ret}_\tau \cap \mathcal{RET}(\mathcal{T}_1, \mathcal{T}_2). \end{aligned}$$

A *conformance notion* with time threshold  $\tau$  and value threshold  $\epsilon$  induced by  $\text{Ret}$  is a predicate  $\text{Conf}_{\tau, \epsilon}^{\text{Ret}}$  on pairs of GTTs such that, for  $\mu_1 : \mathcal{T}_1 \rightarrow \mathcal{Y}$ ,  $\mu_2 : \mathcal{T}_2 \rightarrow \mathcal{Y}$ :

$$\begin{aligned} \text{Conf}_{\tau, \epsilon}^{\text{Ret}}(\mu_1, \mu_2) \iff \exists (r_1, r_2) \in \text{Ret}_\tau(\mathcal{T}_1, \mathcal{T}_2) : & \forall t \in \mathcal{T}_1 : d_{\mathcal{Y}}(\mu_1(t), \mu_2 \circ r_1(t)) \leq \epsilon \\ & \wedge \forall t \in \mathcal{T}_2 : d_{\mathcal{Y}}(\mu_2(t), \mu_1 \circ r_2(t)) \leq \epsilon. \end{aligned}$$

A conformance notion with *unbounded time deviation* and value threshold  $\epsilon$  induced by  $\text{Ret}$  is a predicate  $\text{Conf}_{\infty, \epsilon}^{\text{Ret}}$  on pairs of GTTs such that, for  $\mu_1 : \mathcal{T}_1 \rightarrow \mathcal{Y}$ ,  $\mu_2 : \mathcal{T}_2 \rightarrow \mathcal{Y}$ :

$$\begin{aligned} \text{Conf}_{\infty, \epsilon}^{\text{Ret}}(\mu_1, \mu_2) \iff \exists (r_1, r_2) \in \text{Ret}(\mathcal{T}_1, \mathcal{T}_2) : & \forall t \in \mathcal{T}_1 : d_{\mathcal{Y}}(\mu_1(t), \mu_2 \circ r_1(t)) \leq \epsilon \\ & \wedge \forall t \in \mathcal{T}_2 : d_{\mathcal{Y}}(\mu_2(t), \mu_1 \circ r_2(t)) \leq \epsilon. \end{aligned}$$

Unless we state explicitly otherwise, we consider conformance notions with finite time threshold  $\tau \in \mathbb{R}_{\geq 0}$ . Using the above definition, we can easily express the specific notions of conformance defined in the previous section by selecting a suitable family of retimings.

**Proposition 4.5.** *The conformance predicates below coincide with the notions of conformance induced by the corresponding families of retimings:*

- *TraceConf $_\epsilon$  is induced by the family of retimings containing only identity functions:  $\text{Ret}_{\text{id}} = \{(\text{id}, \text{id}) \mid \text{id} : \mathcal{T} \rightarrow \mathcal{T} \text{ is the identity on some } \mathcal{T} \subseteq \mathbb{R}_{\geq 0}\}$ .*
- *SkorConf $_{\tau, \epsilon}$  is induced by the family of retimings  $\text{Ret} = \{(r, r^{-1}) \mid r \text{ is a strictly increasing continuous bijection}\}$ .*
- *HybridConf $_{\tau, \epsilon}$  is induced by pairs of arbitrary functions.*

Definition 4.4 also enables us to define other notions of conformance, such as, for instance a “shift conformance”, which, intuitively, shifts all time points by a given constant  $c \in \mathbb{R}$ , i.e.,  $\text{Ret}_c = \{(r, r^{-1}) \mid r(t) = t + c\}$ .

Next, we define a generic notion of cleanness, parametrised by conformance predicates for the input and for the output traces. Instantiating these predicates with existing or new conformance notions, yields different conformance-based notions of cleanness that can capture a variety of cleanness specifications.

**4.3. Definition of Conformance-based Cleanness.** We now extend the notion of robust cleanness [DBB<sup>+</sup>17] to allow for “small” variations in time, in addition to the variations in value. To this end, the new notion makes use of two conformance predicates, one that postulates when two input traces should be considered close enough, and another one that specifies when two output traces are close enough.

Our starting point, the notion of robust cleanness in Definition 3.6, is based on comparison of matching prefixes of a pair of input traces and the corresponding prefixes of the associated output traces. As we now want to accommodate for distance in time, we (1) compare prefixes using a conformance relation, and (2) allow for variation in the length of the compared prefixes that is within the corresponding time-distance threshold. More precisely, when comparing two prefixes, we allow for discarding start and end segments of length at most  $\tau$ .

This intuition is formalized by the predicate  $\text{PrefConf}$  for relaxed comparison of GTT prefixes using a notion of conformance  $\text{Conf}$  with tolerance threshold  $\tau$  for time disturbance. We use cascaded notation to define  $\text{PrefConf}$  as a higher-order function taking  $\text{Conf}$  as its first argument. The predicate  $\text{PrefConf}$  compares two prefixes  $\mu_1$  and  $\mu_2$  by requiring that there exist traces  $\mu_1[t_1^s \dots t_1^e]$  and  $\mu_2[t_2^s \dots t_2^e]$  obtained from them, that are conformant with respect to  $\text{Conf}$ . These traces are obtained by possibly removing a sub-prefix of length at most  $\tau$ , and/or removing extending with a suffix of length at most  $\tau$ .

**Definition 4.6.** Let  $\text{Conf}$  be a notion of conformance on GTTs with tolerance threshold  $\tau \in \mathbb{R}_{\geq 0}$  for time disturbance. For any pair of GTTs  $\mu_1 : \mathcal{T}_1 \rightarrow \mathcal{Y}$ ,  $\mu_2 : \mathcal{T}_2 \rightarrow \mathcal{Y}$ , and  $t \in \mathcal{T} = \mathcal{T}_1 \cup \mathcal{T}_2$ , the predicate  $\text{PrefConf}$  is defined as:

$$\begin{aligned} \text{PrefConf}(\mu_1, \mu_2, t) \iff & \exists t_1^s \in [0, \tau] \cap \mathcal{T}_1, \exists t_1^e \in [t - \tau, t + \tau] \cap \mathcal{T}_1, \\ & \exists t_2^s \in [0, \tau] \cap \mathcal{T}_2, \exists t_2^e \in [t - \tau, t + \tau] \cap \mathcal{T}_2: \\ & \text{Conf}(\mu_1[t_1^s \dots t_1^e], \mu_2[t_2^s \dots t_2^e]). \end{aligned}$$

For conformance notions with unbounded timing deviation  $\text{PrefConf}$  coincides with  $\text{Conf}$ .

The predicate  $\text{PrefConf}$  provides a generic notion of prefix-conformance. By instantiating it with conformance relations  $\text{Conf}_I$  and  $\text{Conf}_O$  for input and output traces respectively, we define the notion of  $(\text{Conf}_I, \text{Conf}_O)$ -cleanness.

For deterministic systems  $(\text{Conf}_I, \text{Conf}_O)$ -cleanness requires that for all pairs of input prefixes for which all sub-prefixes are prefix-conformant w.r.t.  $\text{Conf}_I$ , the corresponding pair of output prefixes are prefix-conformant w.r.t.  $\text{Conf}_O$ .

**Definition 4.7.** A deterministic system  $H$  is  $(\text{Conf}_I, \text{Conf}_O)$ -clean if

$$\begin{aligned} \forall i_1, i_2 \in \text{GTT}(\mathcal{Y}) : \forall t \in \text{dom}(i_1) \cup \text{dom}(i_2) : \\ (\forall t' \leq t : \text{PrefConf}_I(i_1, i_2, t')) \implies \text{PrefConf}_O(H(i_1), H(i_2), t). \end{aligned}$$

The above definition naturally generalises to nondeterministic hybrid systems, by comparing sets of possible output prefixes using Hausdorff distance as in [DBB<sup>+</sup>17].

**Definition 4.8.** A system  $H$  is  $(\text{Conf}_I, \text{Conf}_O)$ -clean if

$$\begin{aligned} \forall i_1, i_2 \in \text{GTT}(\mathcal{Y}) : \forall t \in \text{dom}(i_1) \cup \text{dom}(i_2) : \\ (\forall t' \leq t : \text{PrefConf}_I(i_1, i_2, t')) \implies \\ ((\forall o_1 \in H(i_1) \exists o_2 \in H(i_2) : \text{PrefConf}_O(o_1, o_2, t)) \wedge \\ (\forall o_2 \in H(i_2) \exists o_1 \in H(i_1) : \text{PrefConf}_O(o_1, o_2, t))). \end{aligned}$$

Robust cleanness [DBB<sup>+</sup>17] can be now formulated as conformance-based cleanness, which establishes that  $(\text{Conf}_I, \text{Conf}_O)$ -cleanness is a generalisation. Using hybrid conformance, we define hybrid-conformance cleanness, and similarly, plugging in Skorokhod conformance, we define Skorokhod-conformance cleanness. Formally:

- A hybrid system  $H$  is robustly clean, denoted  $\text{RobustClean}(\kappa_I, \kappa_O)$ , if and only if  $H$  is  $(\text{TraceConf}_{\kappa_I}, \text{TraceConf}_{\kappa_O})$ -clean.
- A hybrid system  $H$  is *hybrid-conformance clean with conformance thresholds*  $(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ , which we denote by  $\text{HybridClean}(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ , if and only if  $H$  is  $(\text{HybridConf}_{\tau_I, \epsilon_I}, \text{HybridConf}_{\tau_O, \epsilon_O})$ -clean.
- A hybrid system  $H$  is *Skorokhod-conformance clean with conformance thresholds*  $(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ , denoted  $\text{SkorClean}(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ , if and only if  $H$  is  $(\text{SkorConf}_{\tau_I, \epsilon_I}, \text{SkorConf}_{\tau_O, \epsilon_O})$ -clean.

**4.4. Properties.** We will now establish some key relations between the cleanness notions defined previously. We begin by lifting the implication between conformance relations to implication between cleanness notions defined using those relations.

**Proposition 4.9.** *Suppose that  $\text{Conf}_I^1 \sqsupseteq \text{Conf}_I^2$  and  $\text{Conf}_O^1 \sqsubseteq \text{Conf}_O^2$ . Then for any system  $H$ :  $H$  is  $(\text{Conf}_I^1, \text{Conf}_O^1)$ -clean  $\implies H$  is  $(\text{Conf}_I^2, \text{Conf}_O^2)$ -clean.*

The proposition above has two important corollaries. The first one explains the relationships between the original robust cleanness, and notions of cleanness based on Skorokhod conformance and hybrid conformance, in particular stating the conservative generalisation property for the latter notions. The second corollary compares cleanness notions with different conformance thresholds.

**Corollary 4.10.** *For all  $\tau_I, \tau_O, \epsilon_I, \epsilon_O \in \mathbb{R}_{\geq 0}$ , the following implications hold:*

- (1)  $\text{RobustClean}(\epsilon_I, \epsilon_O) \implies \text{SkorClean}(0, \epsilon_I, \tau_O, \epsilon_O) \implies \text{HybridClean}(0, \epsilon_I, \tau_O, \epsilon_O)$ ,
- (2)  $\text{HybridClean}(\tau_I, \epsilon_I, 0, \epsilon_O) \implies \text{SkorClean}(\tau_I, \epsilon_I, 0, \epsilon_O) \implies \text{RobustClean}(\epsilon_I, \epsilon_O)$ .

*Also,  $\text{RobustClean}(\epsilon_I, \epsilon_O) = \text{SkorClean}(0, \epsilon_I, 0, \epsilon_O) = \text{HybridClean}(0, \epsilon_I, 0, \epsilon_O)$  and hence  $\text{SkorClean}$  and  $\text{HybridClean}$  are conservative extensions of robust cleanness.*

**Corollary 4.11.** *For all  $\epsilon_I, \epsilon'_I, \epsilon_O, \epsilon'_O, \tau_I, \tau'_I, \tau_O, \tau'_O$  that satisfy the inequalities  $\epsilon'_I \leq \epsilon_I$ ,  $\tau'_I \leq \tau_I$ ,  $\epsilon'_O \geq \epsilon_O$ ,  $\tau'_O \geq \tau_O$  the following implications hold:*

- (1)  $\text{RobustClean}(\epsilon_I, \epsilon_O) \implies \text{RobustClean}(\epsilon'_I, \epsilon'_O)$ ;
- (2)  $\text{HybridClean}(\epsilon_I, \tau_I, \epsilon_O, \tau_O) \implies \text{HybridClean}(\epsilon'_I, \tau'_I, \epsilon'_O, \tau'_O)$ ;
- (3)  $\text{SkorClean}(\epsilon_I, \tau_I, \epsilon_O, \tau_O) \implies \text{SkorClean}(\epsilon'_I, \tau'_I, \epsilon'_O, \tau'_O)$ .

**Example 4.12.** Consider the testing workflow in Fig. 1. The inputs passed to a car are  $i_{st}$  and  $i_{ddev}$ , depicted in Fig. 1.(b). One of the test results is presented in Fig. 1.(c), where  $i_{st}$  reveals output  $o(i_{st})$  and  $i_{ddev}$  reveals  $o(i_{ddev})$ . We assume that  $\epsilon < |i_{st}(t_0) - i_{ddev}(t_0)|$  and  $\epsilon < |o(i_{st})(t_1) - o(i_{ddev})(t_1)|$  at some time  $t_1 \geq t_0$ .

- As we saw in Example 3.7, for inputs  $i_{st}$  and  $i_{ddev}$ , the car that emits the outputs depicted in Fig. 1.(c) is deemed  $\text{RobustClean}(\epsilon, \epsilon)$ . Note, that in the presence of other inputs the car used for testing might not be  $\text{RobustClean}(\epsilon, \epsilon)$ .
- As explained in Example 3.4,  $i_{st}$  and  $i_{ddev}$  are hybrid conformant for  $\epsilon$  and  $\tau$ , i.e., the predicate  $\text{PrefConf}_I$  on the left-hand side of the implication in Def. 4.7 holds.  $\text{PrefConf}_O$ , however, fails at time  $t_1$  for signals  $o(i_{st})$  and  $o(i_{ddev})$ . Hence, the system tested in Fig. 1.(c) is not  $\text{HybridClean}(\epsilon, \tau, \epsilon, \tau)$ .

We now discuss testing and falsification of conformance-based cleanness. For systems with discrete time domains the existing methods for verifying [DBB<sup>+</sup>17] or testing [BDH19] robust cleanness can be readily applied.

In the case of hybrid cleanness, existing methods for testing hybrid conformance, such as [AHF<sup>+</sup>14] and [ACM<sup>+</sup>18] can be extended to testing and falsification of hybrid cleanness of hybrid systems consisting of traces with finite time domains. Methods for checking Skorokhod conformance were presented in [DMP17]. Due to the quantification over all time-points  $t'$  in our Definition 4.7 and Definition 4.8, it is not clear how to directly extend them to testing Skorokhod cleanness.

## 5. CLEANNES WITH SYNCHRONIZED RETIMINGS

**5.1. Practical motivation.** An intuitive and useful notion of doping cleanness should capture precisely what we expect from a clean system subject to disturbances in time and value. In this regard, one can observe that even the more discriminating SkorClean predicate has certain drawbacks. The following example motivates why one may want to resort to the finer definition to be proposed in this section.

**Example 5.1.** Consider the scenario of particle emission cleanness presented in Example 3.4 and the input (velocity)- and output trajectories depicted in Fig. 2. Assume that for some input trajectory  $i_1$ , the vehicle shows the output (emission) profile  $o_1$ ; for a second input  $i_2(t) = i_1(t - \tau)$ , consider two possible output trajectories: one output is  $o_1(t - \tau)$ , i.e., it is shifted in the same manner as input; this is assumed to be the best response to  $i_2$ . The other output is of the form  $o_1(t - \tau - \delta)$ , where  $\delta > 0$  can be arbitrarily small, i.e., it is the optimal output with an arbitrary small shift to the right. Skorokhod-conformance cleanness with  $\tau_I = \tau_O$  would accept the first output, but it would reject the second one. A potential solution could be to increase the value of  $\tau_O$  so that it is significantly larger than  $\tau_I$ , but this increases the imprecision by accepting too many trajectories shifted far to the left from  $o_1(t - \tau)$ .

Intuitively, when the input shifts by some  $\tau$ , we would like to compare the corresponding output trajectory with the one that is shifted accordingly. In the above-mentioned case, one would therefore ideally like to perform conformance check of output against  $o_1(t - \tau)$ , rather than  $o_1(t)$ .




Figure 2: Imprecision problem in Skorokhod-conformance cleanness without synchronisation of retimings. While  $o_1(t - \tau)$  is the best expected response to  $i_2$ , no trajectory to the right of  $o_1(t - \tau)$  is accepted when  $\tau_I = \tau_O = \tau$ .

**5.2. Formal theory of synchronized retiming.** In order to alleviate this imprecision, we propose a definition of conformance-based cleanness with synchronised retimings, in which we do not check the conformance of the resulting outputs directly, but rather check conformance of each of the outputs against the transformation of the other output with the retiming that is expected, based on the retiming of the corresponding input. Note that the expected retiming of the output is not always precisely the same as that of the input. Instead, we assume that the set of expected output retimings to a given input retiming is available through a synchronisation function.

As mentioned earlier, we can include in the set of conforming output trajectories the best expected response  $o_1(t - \tau)$  by allowing a sufficiently large  $\tau_O$ , but this comes at the price of introducing imprecision in the conformance relation. By shifting the reference point of conformance comparison, our cleanness with synchronised retimings avoids this imprecision. What is more important, by performing the synchronisation independently of  $\tau_O$ , we introduce the opportunity to constrain the set of conforming output traces to those traces that are as close as desired to the ideal expected output behaviour.

We proceed to formalise the enhanced notion of cleanness with synchronisation outlined above. We start with two auxiliary definitions.

Definition 4.4 entails that whenever the conformance predicate  $\text{Conf}$  holds for certain pair of timed traces, it is witnessed by at least one relevant retiming  $(r_1, r_2)$ . The following operator “extracts” all such witness retimings:

$$\text{Wit Conf}_{\tau, \epsilon}^{\text{Ret}}(\mu_1, \mu_2) \triangleq \{(r_1, r_2) \in \text{Ret}_{\tau}(\mathcal{T}_1, \mathcal{T}_2) \mid \\ \forall t \in \mathcal{T}_1 : d_{\mathcal{Y}}(\mu_1(t) - \mu_2 \circ r_1(t)) \leq \epsilon \wedge \\ \forall t \in \mathcal{T}_2 : d_{\mathcal{Y}}(\mu_1 \circ r_2(t) - \mu_2(t)) \leq \epsilon\}.$$

Similarly, we define the collection of all retimings that witness prefix conformance ( $\text{PrefixConf}$  predicate, Definition 4.6):

$$\text{PrefixWit Conf}(\mu_1, \mu_2, t) \triangleq \{(r_1, r_2) \in \text{Wit Conf}(\mu_1[t_1^s \dots t_1^e], \mu_2[t_2^s \dots t_2^e]) \mid \\ t_1^s \in [0, \tau] \cap \mathcal{T}_1, t_1^e \in [t - \tau, t + \tau] \cap \mathcal{T}_1, \\ t_2^s \in [0, \tau] \cap \mathcal{T}_2, t_2^e \in [t - \tau, t + \tau] \cap \mathcal{T}_2\}.$$

Note that the domains of the retimings in  $\text{PrefixWit Conf}(\mu_1, \mu_2, t)$  can be smaller than the domains of  $\mu_1$  and  $\mu_2$ .

Synchronisation is realised through a function  $\text{Sync}$  specifying all allowed pairs of output retimings for a given pair of input retimings. With this, we can extend the definition of  $(\text{Conf}_I, \text{Conf}_O)$  cleanness as follows. Given two inputs that are conformant, i.e.,  $\text{PrefixConf}_I(i_1, i_2)$ , we may pick any pair  $(r_1, r_2)$  from the set  $\text{PrefixWit Conf}_I(i_1, i_2)$  of pairs of retiming functions for which the input conformance holds. This pair induces another pair  $(r'_1, r'_2) \in \text{Sync}(r_1, r_2)$  of retiming functions for the output timeline. For those, the prefix conformance predicates  $\text{Conf}_O(o_1 \circ r'_2, o_2)$  and  $\text{Conf}_O(o_1, o_2 \circ r'_1)$  must hold. This is formally expressed in the following definition.

**Definition 5.2.** A deterministic system  $H$  is  $(\text{Conf}_I, \text{Conf}_O)$ -clean with synchronised re-timing through  $\text{Sync} : \mathcal{RET}(\mathcal{T}_1, \mathcal{T}_2) \rightarrow \mathcal{P}(\mathcal{RET}(\mathcal{T}_1, \mathcal{T}_2))$  if the following holds:

$$\forall i_1, i_2 \in \text{GTT}(\mathcal{Y}) : \forall t \in \text{dom}(i_1) \cup \text{dom}(i_2) \\ (\forall t' \leq t : \text{PrefixConf}_I(i_1, i_2, t')) \implies \\ \exists (r_1, r_2) \in \text{PrefixWit Conf}_I(i_1, i_2, t) : \\ \exists (r'_1, r'_2) \in \text{Sync}(r_1, r_2) : \\ \text{PrefixConf}_O(H(i_1) \circ r'_2, H(i_2), t) \wedge \\ \text{PrefixConf}_O(H(i_1), H(i_2) \circ r'_1, t)$$

Through the function  $\text{Sync}$ , which is a parameter to the above definition, we can specify the allowed retimings for the output, such as, for example a scaling of the input retiming when the timelines of the input and the output have different scales. It is the responsibility of the cleanness tester or verifier to accurately specify the expected behaviour, as an inappropriately chosen  $\text{Sync}$  function can result in declaring doped systems to be clean.

One important aspect of this definition is that by selecting a suitable synchronisation function  $\text{Sync}$  we can incorporate in the cleanness check any available knowledge regarding the expected output behaviours for conforming input trajectories. The following proposition states how the conformance-based notions of cleanness can be recovered by choosing appropriate retimings.

**Proposition 5.3.** *For a given class of retimings  $\text{Ret}$ , an arbitrary induced conformance relation on inputs  $\text{Conf}_I = \text{Conf}_{\tau, \epsilon}^{\text{Ret}}$ , and a conformance relation on outputs  $\text{Conf}_O$ , there exists a synchronised retiming  $\text{Sync}$  such that  $(\text{Conf}_I, \text{Conf}_O)$ -cleanness (as per Definition 4.7) is a special instance of  $(\text{Conf}_I, \text{Conf}_O)$ -cleanness through  $\text{Sync}$ .*

By setting  $\text{Sync}(r_1, r_2) = \{(\text{id}, \text{id})\}$  we obtain the corresponding notion of  $(\text{Conf}_I, \text{Conf}_O)$ -cleanness, which, in particular, means that cleanness with synchronised retimings is also a conservative generalization of robust cleanness.

**Example 5.4.** Consider the behaviour introduced in Example 5.1. As for the retiming witnessing conformance between  $i_1$  and  $i_2$ , let us take the most obvious one i.e.  $(r_1, r_2) = (t - \tau, t + \tau)$ . If  $(r_1, r_2)$  covers the whole domain of the output trace, then we can use  $\text{Sync}(r_1, r_2) = \{(r_1, r_2)\}$ , according to which the output should be retimed in the same way as the input is. By reusing the same retiming of input for output,  $o_1(t - \tau - \delta)$  conforms to the retimed output  $o_1$  with respect to the margin  $\delta$ .

We use this theory in our experimental setup in Section 7 and show how it can lead to a more accurate analysis of emission data in practice.

## 6. EXPRESSING CLEANNES IN TIMED HYPER LOGICS

In this section we introduce a logic that is capable of characterizing the notions of robust and hybrid cleanness. Since robust cleanness can be characterized in the logic  $\text{HyperLTL}$  [CFK<sup>+</sup>14], the logic we propose is a temporal logic for hyperproperties. Our semantic domain consists of generalized timed traces, and thus, our logic extends Signal Temporal Logic (STL) [MN04] with quantifiers over traces. In order to be able to express deviations in time, our logic uses freeze quantifiers as the mechanism for comparing values at different time points. More precisely, the proposed logic is obtained by extending  $\text{STL}^*$  [BDSV14] with trace quantifiers. In the remainder of the section we provide the formal definition of the logic and discuss its applicability in the context of specifying and monitoring cleanness of hybrid systems.

**6.1. Preliminaries.** For the presentation in this section it will be convenient to consider hybrid systems as sets of GTTs, where each GTT represents a pair of input and output GTTs. The reason for this is that we will define a logic whose formulas refer to both the inputs and the outputs of a hybrid system over time, and are therefore interpreted over sets of such combined traces that contain both the input to the system and the system's output.

Formally, we will represent a  $\mathcal{Y}'$ -valued hybrid system  $H : \text{GTT}(\mathcal{Y}') \rightarrow \mathcal{P}(\text{GTT}(\mathcal{Y}'))$  as a subset of  $\text{GTT}(\mathcal{Y})$  where  $\mathcal{Y} = \mathcal{Y}' \times \mathcal{Y}'$ , defined as  $\{\mu \in \text{GTT}(\mathcal{Y}' \times \mathcal{Y}') \mid \exists \mu_I, \mu_O \in \text{GTT}(\mathcal{Y}'), \mu_O \in H(\mu_I), \mu(t) = (\mu_I(t), \mu_O(t)) \text{ for all } t \in \text{dom}(\mu_I)\}$ . The definition of the GTTs  $\mu$  in this set is possible since according to Definition 3.2 we have that for all  $\mu_I \in \text{GTT}(\mathcal{Y}')$  and all  $\mu_O \in H(\mu_I)$  it holds that  $\text{dom}(\mu_I) = \text{dom}(\mu_O)$ .

For the rest of this subsection, whenever we refer to a GTT  $\mu \in H$  of a  $\mathcal{Y}'$ -valued hybrid system  $H$ , we mean a function  $\mu : \mathcal{T} \rightarrow \mathcal{Y}$  defined as above, with  $\mathcal{Y} = \mathcal{Y}' \times \mathcal{Y}'$ . Given  $\mu \in H$

such that  $\mu(t) = (\mu^I(t), \mu^O(t))$  for  $t \in \text{dom}(\mu)$ , we denote with  $\mu^I : \mathcal{T} \rightarrow \mathcal{Y}'$  the projection of  $\mu$  on the input component and with  $\mu^O : \mathcal{T} \rightarrow \mathcal{Y}'$  its projection on the output component.

Let  $\mu : \mathcal{T} \rightarrow \mathcal{Y}$  be a GTT. If  $\mathcal{T}$  is an interval of the form  $[0, b] \subseteq \mathbb{R}_{\geq 0}$  with  $b \in \mathbb{Q}_{>0}$ , or of the form  $[0, \infty)$ , and  $\mathcal{Y} = \mathbb{R}^n$  for some  $n \in \mathbb{N}$ , we say that  $\mu$  is a *real-valued signal*, and define  $\text{length}(\mu) = b$ , respectively  $\text{length}(\mu) = \infty$ , to be the *time length* of  $\mu$ . If  $\mathcal{T}$  is instead a strictly increasing sequence  $t_0, t_1, t_2, \dots$  of rational numbers such that  $t_0 = 0$  we say that  $\mu$  is a *timed word* and similarly define its *time length* as  $\text{length}(\mu) = \max \mathcal{T}$  if  $\mathcal{T}$  is finite and  $\text{length}(\mu) = \infty$  otherwise.

Let  $X$  be a finite set of *real-valued variables*. We denote with  $\mathbb{R}^X$  the set of possible valuations of  $X$ . In the rest of this section we assume that the range of all GTTs that we consider is  $\mathbb{R}^X$  for a given finite set of real-valued variables. We will assume that the variables in  $X$  are indexed, i.e.,  $X = \{x_1, x_2, \dots, x_n\}$  for some  $n \in \mathbb{N}$ , and use  $\mathbb{R}^n$  instead of  $\mathbb{R}^X$  with the expected interpretation. An *atomic predicate* over  $X$  is a function  $\alpha : \mathbb{R}^X \rightarrow \mathbb{B}$ . Recall from Definition 3.1 that a GTT  $\mu : \mathcal{T} \rightarrow \mathcal{Y}$  is defined for a metric space  $(\mathcal{Y}, d_{\mathcal{Y}})$ . When  $\mathcal{Y} = \mathbb{R}^X$  for some set of variables  $X$ , we assume that the metric  $d$  can be expressed as an arithmetic expression  $d_{\mathcal{Y}}(X, X')$  over the variables  $X \cup X'$ , where  $X' = \{x' \mid x \in X\}$ . More precisely, we have that the expression  $d_{\mathcal{Y}}(X, X')$  evaluates to  $u \in \mathbb{R}$  for valuations  $v \in \mathbb{R}^X$  and  $v' \in \mathbb{R}^{X'}$  of  $X$  and  $X'$ , respectively, if and only if  $d_{\mathcal{Y}}(v, v') = u$ . For a  $\mathcal{Y}$ -valued hybrid system  $H$  we denote with  $d_I$  and  $d_O$  the arithmetic expressions that define the metrics associated with the underlying metric spaces for the input and output values of  $H$ .

**6.2. The logic HyperSTL\*.** We now define the logic HyperSTL\*, which extends the logic STL\* [BDSV14] with *quantifiers over traces*, that are used to relate multiple GTTs in a hybrid system. To this end, let  $V_{\text{trace}}$  be a countably infinite set of *trace variables*. For a set  $X$  of real-valued variables and a given trace variable  $\pi \in V_{\text{trace}}$ , let  $X_{\pi} = \{x_{\pi} \mid x \in X\}$  be the set of variables indexed with  $\pi$ .

Let  $\mathcal{I} = \{1, \dots, m\}$  for some  $m \in \mathbb{N}$  be a finite index set. As in the logic STL\*, the index set  $\mathcal{I}$  consists of the indices of the positions in the *frozen time vector*. Intuitively, at each position of the frozen time vector a time point can be stored. For a trace variable  $\pi \in V_{\text{trace}}$  and an index  $i \in \mathcal{I}$ , let  $X_{\pi}^{*i} = \{x_{\pi}^{*i} \mid x \in X\}$  be the set of variables indexed with  $\pi$  and  $*_i$ .

**6.2.1. Syntax.** Let  $X$  be a finite set of real-valued variables, and  $AP$  be a set of atomic predicates over the set of indexed variables  $\bigcup_{\pi \in V_{\text{trace}}} X_{\pi} \cup \bigcup_{\pi \in V_{\text{trace}}, i \in \mathcal{I}} X_{\pi}^{*i}$ .

HyperSTL\* formulas are defined according to the following grammar.

$$\begin{aligned} \Phi &::= \exists \pi. \Phi \mid \forall \pi. \Phi \mid \varphi, \\ \varphi &::= \alpha \mid \top \mid \neg \varphi \mid \varphi \vee \varphi \mid \varphi \mathcal{U}_J \varphi \mid \varphi \mathcal{S}_J \varphi \mid *_i \varphi, \end{aligned}$$

where  $\pi \in V_{\text{trace}}$  is a trace variable,  $\alpha$  is an atomic predicate from  $AP$ ,  $J \subseteq \mathbb{R}_{\geq 0}$  is an interval with endpoints in  $\mathbb{Q}_{\geq 0} \cup \{\infty\}$ , and  $i \in \mathcal{I}$  is an index.

The operators  $\mathcal{U}$  and  $\mathcal{S}$  are the temporal operators *Until* and *Since*. The  $*_i$  operator, for  $i \in \mathcal{I}$  is the *signal-value freeze* operator. Their semantics is formally defined below. When the interval  $J$  is of the form  $[0, \infty)$  we often omit it for convenience.

The Boolean constant  $\perp$  (false), additional Boolean operators, as well as additional temporal operators are defined in the usual way. More concretely, we define  $\diamond_J \varphi \stackrel{\Delta}{=} \top \mathcal{U}_J \varphi$ ,  $\square_J \varphi \stackrel{\Delta}{=} \neg \diamond_J \neg \varphi$ ,  $\heartsuit_J \varphi \stackrel{\Delta}{=} \top \mathcal{S}_J \varphi$ , and  $\blacklozenge_J \varphi \stackrel{\Delta}{=} \neg \heartsuit_J \neg \varphi$ .



A HyperSTL\* formula is *well-formed* if each occurrence of a trace quantifier introduces a unique variable name, and it is *closed* if every occurrence of a variable in  $X_\pi$  is in the scope of a quantifier for  $\pi$ . We will consider only well-formed HyperSTL\* formulas.

Note that in HyperSTL\*, unlike [BDSV14], we also allow the past operator  $\mathcal{S}$ , as well as arbitrary intervals  $J$  in the operators  $\mathcal{U}$  and  $\mathcal{S}$ . We define HyperSTL\*<sub>fin</sub> to be the fragment of HyperSTL\* such that every interval  $J$  is of the form  $[a, b]$ , where  $a, b \in \mathbb{Q}_{\geq 0}$  and  $a < b$ .

**6.2.2. Semantics.** HyperSTL\* formulas are interpreted over trace assignments and register valuations. A *trace assignment* is a partial function with finite domain from  $V_{trace}$  to the set of GTTs in a given hybrid system. Formally, given a hybrid system  $H$  represented as a set of input-output GTTs, a trace assignment  $\Pi$  is a partial function  $\Pi : V_{trace} \rightarrow H$ . *Register valuations* are  $|\mathcal{I}|$ -dimensional vectors over  $\mathbb{R}_{\geq 0}$ .

Let  $\Pi$  be a trace assignment with domain  $U_{trace} \subseteq V_{trace}$ , and let  $\alpha$  be an atomic predicate defined over the variables in  $\bigcup_{\pi \in U_{trace}} X_\pi \cup \bigcup_{\pi \in U_{trace}, i \in \mathcal{I}} X_\pi^{*i}$ . Consider a time point  $t \in \mathbb{R}_{\geq 0}$  such that for each  $\pi \in U_{trace}$  for which a variable from  $X_\pi$  occurs in  $\alpha$  it holds that  $t \in \text{dom}(\Pi(\pi))$ , and a register valuation  $T$  such that for each pair  $\pi \in U_{trace}$  and  $i \in \mathcal{I}$  such that a variable from  $X_\pi^{*i}$  occurs in  $\alpha$  it holds that  $T(i) \in \text{dom}(\Pi(\pi))$ . Then, the value of the atomic predicate  $\alpha$  at the tuple  $(\Pi, T, t)$  is defined as:

$$\alpha(\Pi, T, t) \triangleq \alpha((\Pi(\pi)(t))_{\pi \in U_{trace}}, (\Pi(\pi)(T(i)))_{\pi \in U_{trace}, i \in \mathcal{I}}).$$

Intuitively, the atomic predicate is evaluated using the signal values at time point  $t$  and at the time points stored in the frozen time vector  $T$ . If for some of the indexed variables that occur in  $\alpha$  the corresponding time point is not in the time domain of the corresponding trace, then the value of the atomic predicate is *undefined*.

To define the semantics of HyperSTL\*, we define the function  $\text{Value}$  that maps a formula  $\Psi$ , a trace assignment  $\Pi$ , a register assignment  $T$  and a time point  $t$  to a value in the set  $\{\top, \text{F}, \text{U}\}$ , which indicates whether  $\Psi$  is true ( $\top$ ), false ( $\text{F}$ ) or undefined ( $\text{U}$ ) at  $(\Pi, T, t)$ . Formally, for a hybrid system  $H$ , a trace assignment  $\Pi$ , a register valuation  $T$ , and  $t \in \mathbb{R}_{\geq 0}$ , the value  $\text{Value}(\Psi, H, \Pi, T, t)$  is defined by induction on the structure of HyperSTL\* formulas.

- If  $\Psi = \alpha$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \alpha(\Pi, T, t) & \text{if the value of } \alpha \text{ is defined at } (\Pi, T, t), \\ \text{U} & \text{otherwise.} \end{cases}$$

- If  $\Psi = \top$ , then  $\text{Value}(\Psi, H, \Pi, T, t) = \top$ .
- If  $\Psi = \neg\varphi$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \top & \text{if } \text{Value}(\varphi, H, \Pi, T, t) = \text{F}, \\ \text{F} & \text{if } \text{Value}(\varphi, H, \Pi, T, t) = \top, \\ \text{U} & \text{if } \text{Value}(\varphi, H, \Pi, T, t) = \text{U}. \end{cases}$$

- If  $\Psi = \varphi_1 \vee \varphi_2$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \top & \text{if } \text{Value}(\varphi_1, H, \Pi, T, t) = \top \text{ or } \text{Value}(\varphi_2, H, \Pi, T, t) = \top, \\ \text{F} & \text{if } \text{Value}(\varphi_1, H, \Pi, T, t) = \text{F} \text{ and } \text{Value}(\varphi_2, H, \Pi, T, t) = \text{F}, \\ \text{U} & \text{otherwise.} \end{cases}$$

- If  $\Psi = \varphi_1 \mathcal{U}_J \varphi_2$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \top & \text{if for some } t' \geq t \text{ such that } t' - t \in J : \text{Value}(\varphi_2, H, \Pi, T, t') = \top \text{ and} \\ & \text{for all } t'' \in [t, t') : \text{Value}(\varphi_1, H, \Pi, T, t'') \in \{\top, \text{U}\} \text{ or} \\ & t'' - t \in J \text{ and } \text{Value}(\varphi_2, H, \Pi, T, t'') = \top, \\ \text{F} & \text{otherwise.} \end{cases}$$

- If  $\Psi = \varphi_1 \mathcal{S}_J \varphi_2$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \top & \text{if for some } t' \in [0, t] \text{ such that } t - t' \in J : \text{Value}(\varphi_2, H, \Pi, T, t') = \top \text{ and} \\ & \text{for all } t'' \in (t', t] : \text{Value}(\varphi_1, H, \Pi, T, t'') \in \{\top, \text{U}\} \text{ or} \\ & t'' - t \in J \text{ and } \text{Value}(\varphi_2, H, \Pi, T, t'') = \top, \\ \text{F} & \text{otherwise.} \end{cases}$$

- If  $\Psi = *_i \varphi$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \text{Value}(\varphi, H, \Pi, T[i \mapsto t], t) & \text{if } t \in \text{dom}(\Pi(\pi)) \text{ for all } x_{\pi}^{*i} \text{ that appear in } \varphi, \\ \text{U} & \text{otherwise.} \end{cases}$$

where  $T[i \mapsto t](j) \triangleq t$  if  $j = i$  and  $T[i \mapsto t](j) \triangleq T(j)$  if  $j \neq i$ .

- If  $\Psi = \exists \pi. \Phi$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \top & \text{if for some } \mu \in H : \text{Value}(\Phi, H, \Pi[\pi \mapsto \mu], T, t) = \top, \\ \text{F} & \text{otherwise.} \end{cases}$$

- If  $\Psi = \forall \pi. \Phi$ , then

$$\text{Value}(\Psi, H, \Pi, T, t) = \begin{cases} \top & \text{if for all } \mu \in H : \text{Value}(\Phi, H, \Pi[\pi \mapsto \mu], T, t) \in \{\top, \text{U}\}, \\ \text{F} & \text{otherwise.} \end{cases}$$

Note that if a formula is closed, then its value is always either  $\top$  or  $\text{F}$ .

For a hybrid system  $H$ , a trace assignment  $\Pi$ , a register valuation  $T$ ,  $t \in \mathbb{R}_{\geq 0}$ , and a HyperSTL\* formula  $\Phi$  we can define the satisfaction relation  $\models$  where

$$(H, \Pi, T, t) \models \Phi \text{ if and only if } \text{Value}(\Phi, H, \Pi, T, t) = \top.$$

We say that a hybrid system  $H$  *satisfies a closed formula*  $\Phi$ , denoted  $H \models \Phi$ , if and only if it holds that  $(H, \Pi_{\emptyset}, T_0, 0) \models \Phi$ , where  $\Pi_{\emptyset}$  is the empty trace assignment and  $T_0$  is the register valuation in which 0 is stored at every index.

**Example 6.1.** Let  $H = \{\mu_1, \mu_2\}$  be a hybrid system that consists of two generalized timed traces,  $\mu_1$  with  $\text{dom}(\mu_1) = \{0, 2, 4\}$  and  $\mu_2$  with  $\text{dom}(\mu_2) = \{0, 1, 3\}$ , where  $\mu_1(t) = 0$  for all  $t \in \{0, 2, 4\}$ , and  $\mu_2(0) = \mu_2(1) = 1$  and  $\mu_2(3) = 0$ .

Consider the HyperSTL\* formula  $\Phi_1 = \forall \pi_1. \forall \pi_2. \square_{[0,4]}(x_{\pi_1} = x_{\pi_2})$  that states that for every pair of timed traces and every time point in the interval  $[0, 4]$  the value of the two traces must be equal (i.e., they agree on the value of variable  $x$ ). We have that  $H \not\models \Phi_1$  since the two traces differ at time point  $t = 0$ . If, on the other hand we consider the formula  $\Phi_2 = \forall \pi_1. \forall \pi_2. \square_{[1,4]}(x_{\pi_1} = x_{\pi_2})$  obtained from  $\Phi_1$  by replacing the interval  $[0, 4]$  by  $[1, 4]$ , we have that  $H \models \Phi_2$ . The justification behind this is that there is no time point in the interval  $[1, 4]$  where we witness a violation of the atomic predicate  $x_{\pi_1} = x_{\pi_2}$ . In particular, in the time interval  $[1, 4]$  there is no point at which both traces are defined.

Now, consider the HyperSTL\* formula  $\Phi_3 = \forall \pi_1. \forall \pi_2. \diamond_{[0,4]}(x_{\pi_1} = x_{\pi_2})$  that states that for every pair of traces, in the interval  $[0, 4]$  there *exists* a time point where the values of the two traces are the same. We have that  $H \not\models \Phi_3$ , as expected, since there is no point in this interval where both traces are defined and have the same value. Note that we also have  $H \not\models \forall \pi_1. \forall \pi_2. \diamond_{[1,4]}(x_{\pi_1} = x_{\pi_2})$  for the interval where the value of  $x_{\pi_1} = x_{\pi_2}$  is undefined.

Finally, let  $\Phi_4 = \forall \pi_1. \forall \pi_2. \diamond_{[0,3]} * \diamond_{[0,1]} x_{\pi_1}^* = x_{\pi_2}$ . The formula  $\Phi_4$  states that for every pair of traces there exists a time point  $t_1$  in  $[0, 3]$  such that there is a time point  $t_2$  at most 1 time unit later, such that the value of the first trace at  $t_1$  is equal to the value of the second trace at time  $t_2$ . Here  $t_1$  is the frozen time point per the semantics of the freeze operator  $*$ . We have that  $H \models \Phi_4$ . To see this, when  $\pi_1$  is  $\mu_1$  and  $\pi_2$  is  $\mu_2$  let  $t_1 = 2$  and  $t_2 = 3$ , and when  $\pi_1$  is  $\mu_2$  and  $\pi_2$  is  $\mu_1$  let  $t_1 = 3$  and  $t_2 = 4$ .

**Remark 6.2.** Due to the generality of our semantic domain, which generalizes both continuous signals and timed words, we have to address the issue of having to define the interpretation of HyperSTL\* formulas over all time points in  $\mathbb{R}_{\geq 0}$  while the considered traces might not be defined at all points. Furthermore, the semantics of the logic has to account for the fact that formulas, even atomic propositions, refer to different traces which are possibly defined over different time domains. To this end, we defined the function *Value* that assigns values in the set  $\{\top, \text{F}, \text{U}\}$ . For instance, if  $\text{Value}(\alpha, H, \Pi, T, t) = \text{U}$ , then  $\text{Value}(\alpha \vee \neg \alpha, H, \Pi, T, t) = \text{U}$ . For the temporal operators, our semantics is reminiscent of that in [GM20], in the sense that for evaluating  $\varphi_1 \mathcal{U} \varphi_2$  the subformula  $\varphi_1$  is evaluated only in time points where its value is defined, and the time point where the obligation  $\varphi_2$  must hold is one where its value is defined. The treatment in  $\mathcal{S}$  is analogous. Our semantics interprets trace quantifiers over the traces for which the formula has a defined value.

Other temporal logics for timed hyperproperties face similar issues, which we discuss in Remark 6.4. The logic HyperSTL [NKJ<sup>+</sup>17], on the other hand is not affected by such difficulties, since its semantics is defined over continuous signals defined over a whole interval.

**Remark 6.3.** In our definition of the semantics of  $\varphi_1 \mathcal{U}_J \varphi_2$ , similarly to [DMP17] and [GM20], we account for the fact that in a dense time domain there might not exist a *first* time point where  $\varphi_2$  is satisfied. Therefore we allow for  $\varphi_1$  to be violated at intermediate time points as long as at those points the value of  $\varphi_2$  is  $\top$  and the constraint imposed by  $J$  is satisfied. More precisely,  $\varphi_1 \mathcal{U}_J \varphi_2$  is  $\top$  at time point  $t$  if there exists a time point  $t' \geq t$  such that  $t' - t \in J$ ,  $\varphi_2$  is  $\top$  at  $t'$ , and for all intermediate points  $t'' \in [t, t')$  it holds that if  $\varphi_1$  is  $\text{F}$  at  $t''$ , then  $t'' - t \in J$  and  $\varphi_2$  is  $\top$  at  $t''$ . The analogous holds for  $\mathcal{S}$ .

**Remark 6.4.** Existing temporal logics for timed hyperproperties have also faced the challenge of dealing with timed traces that are defined over different sets of time points.

In [HZJ19] this leads to the consideration of two different semantics of their logic HyperMTL: an asynchronous semantics that does not require the time stamps in two timed traces to match, and a synchronous semantics in which the range of quantifiers is restricted to the traces that synchronize with the current trace assignment. The logic HyperMTL includes for each trace variable a Boolean constant  $\top$  (true) indexed with that variable, which allows for expressing syntactically in formulas the requirement that the current time point is in the domain of the corresponding trace. In contrast, in our logic HyperSTL\* we account for undefined values on the semantic level in the definition of the value function, and values at different points in time on different traces can be related via the freeze operator.

The authors of [BPS20] provide an alternative logic HyperMTL by extending the logic MTL with quantifiers over traces in the point-wise semantics. The semantics of their logic

has both a synchronous and an asynchronous layer. At the synchronous layer, traces are compared at the same points in time, and if a trace is undefined at a given point, the value at the closest previous event is used. At the asynchronous layer, an asynchronous version of the  $\mathcal{U}$  operator allows for a bounded difference in the time points when the obligation of the *Until* formula is fulfilled in different traces. Our logic, on the other hand, allows for a general and flexible way of relating time points on different traces via the freeze operator.

**6.3. Expressing robust and hybrid cleanness.** Using the logic HyperSTL\* we can express trace and hybrid conformance and robust and hybrid cleanness. We begin by first formalizing the conformance notions, and then provide the characterization of cleanness.

Let  $\pi_1$  and  $\pi_2$  be two trace variables, and let  $\tau$  and  $\epsilon$  be non-negative rational constants. We can express hybrid conformance with thresholds  $\tau$  and  $\epsilon$ , i.e.,  $\text{HybridConf}_{\tau,\epsilon}$ , as follows:

$$\varphi_{\tau,\epsilon}^{\text{HybridConf}} = \left( \Box * _1 \left( \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \right) \right) \wedge \left( \Box * _2 \left( \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon \right) \right)$$

where  $d_{\mathcal{Y}}(X, X')$  is an arithmetic expression characterizing the metric  $d_{\mathcal{Y}}$ . Note that in the above formula the trace variables  $\pi_1$  and  $\pi_2$  are not quantified, and hence it is not closed.

Intuitively, the formula states that for every time point on the trace described by  $\pi_1$  it holds that within  $\tau$  time units in the past or in the future, there exists a point on the trace described by  $\pi_2$  where the value is  $\epsilon$ -close to the value of  $\pi_1$  at the current time point, and symmetrically for the other direction with traces  $\pi_1$  and  $\pi_2$  swapped.

**Proposition 6.5.** *Let  $H$  be a deterministic hybrid system defined over a set of real-valued variables  $X$  such that  $0 \in \text{dom}(\mu)$  for each  $\mu \in H$ . Let  $\tau, \epsilon \geq 0$  be rational constants, and  $\mu_1, \mu_2 \in H$ . Let  $\pi_1$  and  $\pi_2$  be trace variables and  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$ . Then*

$$\text{HybridConf}_{\tau,\epsilon}(\mu_1, \mu_2) \text{ if and only if } (H, \Pi, T_0, 0) \models \varphi_{\tau,\epsilon}^{\text{HybridConf}}.$$

*Proof.* ( $\implies$ ) First, suppose that  $\text{HybridConf}_{\tau,\epsilon}(\mu_1, \mu_2)$ .

By Definition 3.3, we have that for all  $t_1 \in \text{dom}(\mu_1)$  there exists  $t_2 \in \text{dom}(\mu_2)$  such that  $|t_1 - t_2| \leq \tau$  and  $d_{\mathcal{Y}}(\mu_2(t_2), \mu_1(t_1)) \leq \epsilon$ . Hence, when  $t_1 \in \text{dom}(\mu_1)$ , we have that  $\text{Value}(*_1(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t_1) = \top$ .

If  $t_1 \notin \text{dom}(\mu_1)$ , then, from the definition of the semantics of the operator  $*_1$  we have that  $\text{Value}(*_1(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t_1) = \top$ .

Thus,  $\text{Value}(\Box * _1(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, 0) = \top$ .

$\text{Value}(\Box * _2(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon), H, \Pi, T_0, 0) = \top$  can be shown by applying the same reasoning as above, this time for  $\mu_2$ .

From the two facts we showed, we can conclude that  $(H, \Pi, T_0, 0) \models \varphi_{\tau,\epsilon}^{\text{HybridConf}}$ .

( $\impliedby$ ) For the other direction, assume that  $(H, \Pi, T_0, 0) \models \varphi_{\tau,\epsilon}^{\text{HybridConf}}$ .

Let  $t_1 \in \text{dom}(\mu_1)$ . Since  $(H, \Pi, T_0, 0) \models \varphi_{\tau,\epsilon}^{\text{HybridConf}}$ , by the semantics of  $\Box$  we have that  $\text{Value}(*_1(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \vee \Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t_1) \in \{\top, \text{U}\}$ . Taking into account that we are considering the case when  $t_1 \in \text{dom}(\mu_1)$ , and that by definition

- $\text{Value}(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon, H, \Pi, T_0, t_1) \neq \text{U}$  and
- $\text{Value}(\Diamond_{[0,\tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon, H, \Pi, T_0, t_1) \neq \text{U}$ ,

we conclude that  $\text{Value}(*_1(\diamond_{[0,\tau]} d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon) \vee \diamond_{[0,\tau]} d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t_1) = \top$ . Therefore, there exists  $t_2 \in \text{dom}(\mu_2)$  such that  $|t_1 - t_2| \leq \tau$  and  $d\mathcal{Y}(\mu_2(t_2), \mu_1(t_1)) \leq \epsilon$ .

The reasoning for the other conjunct when  $t_2 \in \text{dom}(\mu_2)$  is symmetric. We hence obtain  $\text{HybridConf}_{\tau,\epsilon}(\mu_1, \mu_2)$ , which concludes the proof.  $\square$

As a special case, we can express  $\text{TraceConf}_\epsilon$  as

$$\varphi_\epsilon^{\text{TraceConf}} = \left( \Box *_1(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon) \right) \wedge \left( \Box *_2(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon) \right).$$

Note that the formula  $\varphi_\epsilon = \Box d\mathcal{Y}(X_{\pi_1}, X_{\pi_2}) \leq \epsilon$  does not characterize trace conformance as it does not assert the requirement that the time domains of the two traces must be the same. The formula  $\varphi_\epsilon^{\text{TraceConf}}$ , on the other hand, requires that each time point where one of the traces is defined, must be matched by a value of the other trace at the same time point.

**Proposition 6.6.** *Let  $H$  be a deterministic hybrid system defined over a set of real-valued variables  $X$  such that  $0 \in \text{dom}(\mu)$  for each  $\mu \in H$ . Let  $\epsilon \geq 0$  be a rational constant, and  $\mu_1, \mu_2 \in H$ . Let  $\pi_1$  and  $\pi_2$  be trace variables and  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$ . Then*

$$\text{TraceConf}_\epsilon(\mu_1, \mu_2) \text{ if and only if } (H, \Pi, T_0, 0) \models \varphi_\epsilon^{\text{TraceConf}}.$$

*Proof.* ( $\implies$ ) First, suppose that  $\text{TraceConf}_\epsilon(\mu_1, \mu_2)$ .

By Definition 3.3, we have that  $\text{dom}(\mu_1) = \text{dom}(\mu_2)$  and for all  $t_1 \in \text{dom}(\mu_1)$  it holds that  $d\mathcal{Y}(\mu_2(t_1), \mu_1(t_1)) \leq \epsilon$ . This is equivalent to the conjunction of the following statements:

- for all  $t_1 \in \text{dom}(\mu_1)$ , it holds that  $t_1 \in \text{dom}(\mu_2)$  and  $d\mathcal{Y}(\mu_2(t_1), \mu_1(t_1)) \leq \epsilon$ , and
- for all  $t_2 \in \text{dom}(\mu_2)$ , it holds that  $t_2 \in \text{dom}(\mu_1)$  and  $d\mathcal{Y}(\mu_1(t_2), \mu_2(t_2)) \leq \epsilon$ .

Therefore, if  $t \in \text{dom}(\mu_1) = \text{dom}(\mu_2)$  it holds that

$$\begin{aligned} \text{Value}(*_1(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t) &= \top \text{ and} \\ \text{Value}(*_2(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon), H, \Pi, T_0, t) &= \top. \end{aligned}$$

When  $t \notin \text{dom}(\mu_1) = \text{dom}(\mu_2)$  we have that that

$$\begin{aligned} \text{Value}(*_1(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t) &= \text{U} \text{ and} \\ \text{Value}(*_2(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon), H, \Pi, T_0, t) &= \text{U}. \end{aligned}$$

Thus,  $\text{Value}(\varphi_\epsilon^{\text{TraceConf}}, H, \Pi, T_0, 0) = \top$ , which is what we had to prove.

( $\impliedby$ ) For the other direction, assume that  $(H, \Pi, T_0, 0) \models \varphi_\epsilon^{\text{TraceConf}}$ .

Let  $t_1 \in \text{dom}(\mu_1)$ . We have that  $\text{Value}(*_1(\diamond_{[0,0]} d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon), H, \Pi, T_0, t_1) = \top$ . This implies that  $t_1 \in \text{dom}(\mu_2)$  and  $\text{Value}(d\mathcal{Y}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon, H, \Pi, \{1 \mapsto t_1\}, t_1) = \top$ .

Thus, we can conclude that for all  $t_1 \in \text{dom}(\mu_1)$ , it holds that  $t_1 \in \text{dom}(\mu_2)$  and  $d\mathcal{Y}(\mu_2(t_1), \mu_1(t_1)) \leq \epsilon$ . Analogously, we can show that for all  $t_2 \in \text{dom}(\mu_2)$ , it holds that  $t_2 \in \text{dom}(\mu_1)$  and  $d\mathcal{Y}(\mu_1(t_2), \mu_2(t_2)) \leq \epsilon$ . Hence, by Definition 3.3,  $\text{TraceConf}_\epsilon(\mu_1, \mu_2)$ .  $\square$

We use the idea of the above encoding to define a closed **HyperSTL\*** formula that characterizes hybrid cleanliness,  $\text{HybridClean}(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ , for deterministic hybrid systems.

For the rest of the section we consider hybrid systems  $H$  such that for every  $\mu \in H$  it holds that  $0 \in \text{dom}(\mu)$ , that is, we assume that all traces are defined at time point 0.

Furthermore, we assume that the set of variables  $X$  defining the states of the hybrid system  $H$  contains an *explicit clock variable*  $c$  representing the current time, that is never reset. That is,  $c$  simply captures the time-stamps of the values of the GTTs in  $H$ . Formally, for every GTT  $\mu \in H$ , and every  $t \in \mathbb{R}_{\geq 0}$ , it holds that  $\mu(t)(c) = t$ . With that, the atomic propositions in  $AP$  can refer to the current time point, and the freeze operator captures the

current time-stamp together with the current values of the other variables. Let  $\alpha \in AP$  be an atomic proposition and  $r, r_i \in \mathbb{R}_{\geq 0}$  for  $i \in \mathcal{I}$  be non-negative real constants. We denote by  $\alpha[r, r_1, \dots, r_{|\mathcal{I}|}]$  the atomic predicate obtained from  $\alpha$  by replacing each variable  $c_\pi$  by  $r$ , and each variable  $c_\pi^{*i}$  by  $r_i$ , for all  $\pi \in V_{trace}$  and  $i \in \mathcal{I}$ . By the definition of the clock variable  $c$ , for every trace assignment  $\Pi$ , register valuation  $T$ , and  $t \in \mathbb{R}_{\geq 0}$  we have that

$$\text{Value}(\alpha, H, \Pi, T, t) = \text{Value}(\alpha[t, T(1), \dots, T(|\mathcal{I}|)], H, \pi, T, t),$$

when for every  $\pi \in V_{trace}$  for which  $c_\pi$  occurs in  $\alpha$  it holds that  $t \in \text{dom}(\Pi(\pi))$  and for every  $\pi \in V_{trace}$  and  $i \in \mathcal{I}$  for which  $c_\pi^{*i}$  occurs in  $\alpha$  it holds that  $T(i) \in \text{dom}(\Pi(\pi))$ .

Let  $\tau_I$  and  $\epsilon_I$  be non-negative rational constants defining the threshold values for the input conformance relation, and  $\tau_O$  and  $\epsilon_O$  be the ones for the output conformance.

Let  $\pi$  and  $\pi'$  be trace variables and  $i, s, e \in \mathcal{I}$ . First, we define the formulas

$$\begin{aligned} \varphi_{\tau_I, \epsilon_I}^{\text{matchI}}(\pi, i, \pi', s, e) &= *_i \left( \diamond_{[0, \tau_I]}(d_I(X_{\pi'}, X_{\pi}^{*i}) \leq \epsilon_I \wedge c_{\pi'} \geq c_{\pi'}^{*s} \wedge c_{\pi'} \leq c_{\pi'}^{*e}) \vee \right. \\ &\quad \left. \diamond_{[0, \tau_I]}(d_I(X_{\pi'}, X_{\pi}^{*i}) \leq \epsilon_I \wedge c_{\pi'} \geq c_{\pi'}^{*s} \wedge c_{\pi'} \leq c_{\pi'}^{*e}) \right), \\ \varphi_{\tau_O, \epsilon_O}^{\text{matchO}}(\pi, i, \pi', s, e) &= *_i \left( \diamond_{[0, \tau_O]}(d_O(X_{\pi'}, X_{\pi}^{*i}) \leq \epsilon_O \wedge c_{\pi'} \geq c_{\pi'}^{*s} \wedge c_{\pi'} \leq c_{\pi'}^{*e}) \vee \right. \\ &\quad \left. \diamond_{[0, \tau_O]}(d_O(X_{\pi'}, X_{\pi}^{*i}) \leq \epsilon_O \wedge c_{\pi'} \geq c_{\pi'}^{*s} \wedge c_{\pi'} \leq c_{\pi'}^{*e}) \right), \end{aligned}$$

where  $d_I(X, X')$  and  $d_O(X, X')$  are the arithmetic expressions characterizing the metrics on the sets of input and output values of the considered hybrid system.

Intuitively, the formula  $\varphi_{\tau_I, \epsilon_I}^{\text{matchI}}(\pi, i, \pi', s, e)$  evaluated at time point  $t$  and register valuation  $T$  is true if and only if there exists a time point  $t' \in [t - \tau_I, t + \tau_I] \cap [T(s), T(e)]$  such that the input value at time  $t$  on the trace represented by  $\pi$  and the input value at time  $t'$  on the trace represented by  $\pi'$  are  $\epsilon_I$ -close. The formula  $\varphi_{\tau_O, \epsilon_O}^{\text{matchO}}(\pi, i, \pi', s, e)$  states the same for the output values. The need to constrain the time  $t'$  where the match of the values at  $t$  must be found comes from the fact that in the definition of cleanness in Section 4, prefixes are compared using the predicate  $\text{PrefConf}$ . Recall that  $\text{PrefConf}$  compares two prefixes  $\mu_1$  and  $\mu_2$  by requiring that there exist segments  $\mu_1[t_1^s \dots t_1^e]$  and  $\mu_2[t_2^s \dots t_2^e]$  obtained from them, that are conformant. In the above formulas, the frozen values of the clock variable  $c$  represent the end points of the interval for the trace assigned to  $\pi'$ .

Using the formula  $\varphi_{\tau_I, \epsilon_I}^{\text{matchI}}(\pi, i, \pi', s, e)$  we define the formula  $\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfI}}(\pi_1, \pi_2)$  which is true if and only if the current time point defines a pair of prefixes of the traces represented by  $\pi_1$  and  $\pi_2$  for which there exist hybrid conforming segments obtained from the prefixes by possibly removing a prefix/suffix of length at most  $\tau_I$  or adding a suffix of length at most  $\tau_I$ .

The formula is defined as a disjunction over the possible ways in which the end-points of the two segments are ordered on the time line. Let  $P$  be a set of 4-tuples of the indices  $\{3, 4, 5, 6\} \subseteq \mathcal{I}$  such that  $(s_1, e_1, s_2, e_2) \in P$  if and only if  $s_1, e_1, s_2, e_2$  are pairwise different, and  $s_1, s_2 \in \{3, 4\}$  and  $e_1, e_2 \in \{5, 6\}$ . For  $\bar{p} = (s_1, e_1, s_2, e_2) \in P$ , abusing notation we define  $\bar{p}(s_1) = \bar{p}(e_1) = \pi_1$  and  $\bar{p}(s_2) = \bar{p}(e_2) = \pi_2$ . That is, the function  $\bar{p}$  maps each endpoint

index to the trace variable with which it is associated. We now define

$$\begin{aligned}
\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfI}}(\pi_1, \pi_2) &= \bigvee_{\bar{p}=(s_1, e_1, s_2, e_2) \in P} \left( *7 \varphi_{\pi_1} \wedge *7 \varphi_{\pi_2} \right) \\
\varphi_{\pi_1} &= \diamond \left( c_{\bar{p}(3)} \leq \tau_I \wedge *3 \diamond \left( c_{\bar{p}(4)} \leq \tau_I \wedge \right. \right. \\
&\quad \left. *4 \diamond \left( c_{\bar{p}(5)} \geq (c_{\pi_1}^{*7} - \tau_I) \wedge c_{\bar{p}(5)} \leq (c_{\pi_1}^{*7} + \tau_I) \wedge \right. \right. \\
&\quad \left. \left. *5 \diamond \left( c_{\bar{p}(6)} \geq (c_{\pi_1}^{*7} - \tau_I) \wedge c_{\bar{p}(6)} \leq (c_{\pi_1}^{*7} + \tau_I) \wedge *6 \varphi^{\text{ConfI}} \right) \right) \right) \\
\varphi_{\pi_2} &= \diamond \left( c_{\bar{p}(3)} \leq \tau_I \wedge *3 \diamond \left( c_{\bar{p}(4)} \leq \tau_I \wedge \right. \right. \\
&\quad \left. *4 \diamond \left( c_{\bar{p}(5)} \geq (c_{\pi_2}^{*7} - \tau_I) \wedge c_{\bar{p}(5)} \leq (c_{\pi_2}^{*7} + \tau_I) \wedge \right. \right. \\
&\quad \left. \left. *5 \diamond \left( c_{\bar{p}(6)} \geq (c_{\pi_2}^{*7} - \tau_I) \wedge c_{\bar{p}(6)} \leq (c_{\pi_2}^{*7} + \tau_I) \wedge *6 \varphi^{\text{ConfI}} \right) \right) \right) \\
\varphi^{\text{ConfI}} &= \square \left( \left( c_{\pi_1} \geq c_{\pi_1}^{*s_1} \wedge c_{\pi_1} \leq c_{\pi_1}^{*e_1} \right) \rightarrow \varphi_{\tau_I, \epsilon_I}^{\text{matchI}}(\pi_1, 1, \pi_2, s_2, e_2) \right) \wedge \\
&\quad \square \left( \left( c_{\pi_2} \geq c_{\pi_2}^{*s_2} \wedge c_{\pi_2} \leq c_{\pi_2}^{*e_2} \right) \rightarrow \varphi_{\tau_I, \epsilon_I}^{\text{matchI}}(\pi_2, 2, \pi_1, s_1, e_1) \right).
\end{aligned}$$

The conjunct  $*7\varphi_{\pi_1}$  handles the case when the current time point  $t$  (i.e., the last time point for the considered prefixes) is in the domain of  $\pi_1$ . The second conjunct handles the case when  $t$  is in the domain of  $\pi_2$ . If neither is the case, then the value of the conjunction is  $\mathbf{U}$ . Since the formulas  $*7\varphi_{\pi_1}$  and  $*7\varphi_{\pi_2}$  differ only in the trace on which the time-point  $t$  is frozen, if both of them have a defined value, than these values are necessarily the same.

Each of  $\varphi_{\pi_1}$  and  $\varphi_{\pi_2}$  asserts the existence of a sequence of time points defining the compared segments of the two traces and their input conformance (formula  $\varphi^{\text{ConfI}}$ ). The formula  $\varphi^{\text{ConfI}}$  captures the requirement that the two input prefixes ending at the current time point have segments (defined by the pairs of time points  $c_{\pi_1}^{*s_1}$  and  $c_{\pi_1}^{*e_1}$ , and  $c_{\pi_2}^{*s_2}$  and  $c_{\pi_2}^{*e_2}$ , respectively) that are hybrid conformant, as in Definition 4.6.

The formula  $\varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2)$  is defined analogously using  $\varphi_{\tau_O, \epsilon_O}^{\text{matchO}}(\pi, i, \pi', s, e)$ .

**Proposition 6.7.** *Let  $H$  be a deterministic hybrid system defined over a set of real-valued variables  $X$  that includes an explicit clock variable  $c$ , and such that  $0 \in \text{dom}(\mu)$  for each  $\mu \in H$ . Let  $\tau_I, \tau_O, \epsilon_I, \epsilon_O \geq 0$  be rational constants, and the predicates  $\text{PrefConfI}$  and  $\text{PrefConfO}$  be instantiated using  $\text{HybridConf}_{\tau_I, \epsilon_I}$  and  $\text{HybridConf}_{\tau_O, \epsilon_O}$  respectively. That is, let  $\text{ConfI} = \text{HybridConf}_{\tau_I, \epsilon_I}$  and  $\text{ConfO} = \text{HybridConf}_{\tau_O, \epsilon_O}$ . Let  $\mu_1, \mu_2 \in H$ , let  $\pi_1$  and  $\pi_2$  be trace variables and  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$  a trace assignment.*

(1) *If  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , then*

*$\text{PrefConfI}(\mu_1^I, \mu_2^I, t)$  is true if and only if  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfI}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\mathbf{T}, \mathbf{U}\}$ .*

(2) *If  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , then*

*$\text{PrefConfO}(\mu_1^O, \mu_2^O, t)$  is true if and only if  $\text{Value}(\varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\mathbf{T}, \mathbf{U}\}$ .*

(3) *If  $t \notin \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  then  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfI}}(\pi_1, \pi_2), H, \Pi, T_0, t) = \mathbf{U}$ .*

(4) *If  $t \notin \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  then  $\text{Value}(\varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) = \mathbf{U}$ .*

*Proof.* We show (1), the proof for (2) is analogous.

( $\implies$ ) Suppose that  $\text{PrefConfI}(\mu_1^I, \mu_2^I, t)$  is true. By Definition 4.6, there exist  $t_1^s, t_2^s \in [0, \tau_I]$  and  $t_1^e, t_2^e \in [t - \tau_I, t + \tau_I]$  such that  $\text{ConfI}(\mu_1^I[t_1^s \dots t_1^e], \mu_2^I[t_2^s \dots t_2^e])$  is true.

Let  $t_3, t_4, t_5, t_6$  be a permutation of  $t_1^s, t_2^s, t_1^e, t_2^e$  such that  $t_3 \leq t_4 \leq t_5 \leq t_6$ . Then,  $t_3, t_4 \in [0, \tau_I]$  and  $t_5, t_6 \in [t - \tau_I, t + \tau_I]$ . Let  $\bar{p} = (s_1, e_1, s_2, e_2) \in P$  be defined according to the permutation  $t_3, t_4, t_5, t_6$ , that is,  $s_1 = i$  where  $t_1^s$  is  $t_i$ , and so on. Let  $t_7 = t$ . We define the register valuation  $T = \{i \mapsto t_i \mid i \in \{3, 4, 5, 6, 7\}\}$ .

Since  $\text{ConfI}(\mu_1^I[t_1^s \dots t_1^e], \mu_2^I[t_2^s \dots t_2^e])$ , we have that  $\text{Value}(\varphi^{\text{ConfI}}, H, \Pi, T, t_6) = \mathbf{T}$ .

Since  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , assume, without loss of generality that  $t \in \text{dom}(\mu_1)$ . By the choice of  $t_3, t_4, t_5, t_6$  we have that  $t_3 \leq t_4 \leq t_5 \leq t_6$ , and  $t_3, t_4 \in [0, \tau]$  and  $t_5, t_6 \in [t - \tau_I, t + \tau_I]$  which, together with the definition of  $\varphi_{\pi_1}$  implies that  $\text{Value}(\varphi_{\pi_1}, H, \Pi, \{7 \mapsto t_7\}, t_7) = \top$ .

If  $t \in \text{dom}(\mu_2)$  we can show as above that the value of  $*_7\varphi_{\pi_2}$  is  $\top$ . Otherwise, the value of  $*_7\varphi_{\pi_2}$  is  $\perp$ . Hence, we conclude that  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\top, \perp\}$ .

( $\Leftarrow$ ) Now, suppose that  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\top, \perp\}$ . Hence, there exists  $\bar{p} = (s_1, e_1, s_2, e_2) \in P$  for which we have  $\text{Value}((*_7\varphi_{\pi_1} \wedge *_7\varphi_{\pi_2}), H, \Pi, T_0, t) \in \{\top, \perp\}$ . Since  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , assume, without loss of generality, that  $t \in \text{dom}(\mu_1)$ . Then we must have that  $\text{Value}(*_7\varphi_{\pi_1}, H, \Pi, T_0, t) = \top$ . Let  $t_7 = t$ . Then,  $\text{Value}(\varphi_{\pi_1}, H, \Pi, \{7 \mapsto t_7\}, t) = \top$ . Therefore, according to the definition of  $\varphi_{\pi_1}$  we have that there exist  $t_3 \leq t_4 \leq t_5 \leq t_6$  such that  $t_3, t_4 \in [0, \tau]$  and  $t_5, t_6 \in [t - \tau_I, t + \tau_I]$ , and, furthermore,  $\text{Value}(\varphi^{\text{Confl}}, H, \Pi, T, t_6) = \top$  for the valuation  $T = \{i \mapsto t_i \mid i \in \{3, 4, 5, 6, 7\}\}$ .

We define  $t_1^s, t_2^s, t_1^e, t_2^e$  to be the permutation of  $t_3, t_4, t_5, t_6$  determined by  $\bar{p}$ , that is  $t_1^s = t_{s_1}$ , and so on. Then, since  $\text{Value}(\varphi_{\pi_1}, H, \Pi, \{7 \mapsto t_7\}, t) = \top$ , we have that  $t_1^s \leq t_1^e$ ,  $t_2^s \leq t_2^e$ , and  $t_1^s, t_2^s \in [0, \tau_I]$  and  $t_1^e, t_2^e \in [t - \tau_I, t + \tau_I]$ . Hence,  $\text{Value}(\varphi^{\text{Confl}}, H, \Pi, T, t_6) = \top$  implies that  $\text{Conf}_I(\mu_1^I[t_1^s \dots t_1^e], \mu_2^I[t_2^s \dots t_2^e])$ , which allows us to conclude that  $\text{PrefConf}_I(\mu_1^I, \mu_2^I, t)$ .

(3) and (4) follow directly from the semantics of the operators  $\wedge$  and  $*$ .  $\square$

We now define the formula characterizing hybrid cleanness as

$$\Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{HybridClean}} = \forall \pi_1. \forall \pi_2. \square \left( (\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2)) \rightarrow \varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2) \right).$$

**Proposition 6.8.** *Let  $H$  be a deterministic hybrid system defined over a set of real-valued variables  $X$  that includes an explicit clock variable  $c$ , and such that  $0 \in \text{dom}(\mu)$  for each  $\mu \in H$ . Let  $\tau_I, \tau_O, \epsilon_I, \epsilon_O \geq 0$  be rational constants. It holds that*

$$H \text{ is HybridClean}(\tau_I, \epsilon_I, \tau_O, \epsilon_O) \text{ if and only if } H \models \Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{HybridClean}}.$$

*Proof.* ( $\Rightarrow$ ) First, suppose that  $H$  is  $\text{HybridClean}(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ . Let  $\mu_1, \mu_2 \in H$  be two arbitrarily chosen traces. Let  $\pi_1$  and  $\pi_2$  be trace variables, and let  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$ .

Let  $t \geq 0$  be an arbitrary time point. If  $\text{Value}(\varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\top, \perp\}$ , then it holds that  $\text{Value}((\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2)) \rightarrow \varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\top, \perp\}$ . If, on the other hand we have that  $\text{Value}(\varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) = \text{F}$ , then it holds that  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  and by Proposition 6.7 we have that  $\text{PrefConf}_O(\mu_1^O, \mu_2^O, t)$  is false. According to Definition 4.7, this means that there exists  $t' \leq t$  such that  $\text{PrefConf}_I(\mu_1^I, \mu_2^I, t')$  is false. Applying Proposition 6.7 we obtain for the time point  $t'$  that  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), H, \Pi, T_0, t') = \text{F}$ , and hence  $\text{Value}(\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), H, \Pi, T_0, t) = \text{F}$ . This implies that  $\text{Value}((\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2)) \rightarrow \varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) = \top$ .

Therefore, since  $t \geq 0$  was chosen arbitrarily,  $(H, \Pi, T_0, 0) \models \square \left( (\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2)) \rightarrow \varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2) \right)$ . Since  $\mu_1$  and  $\mu_2$  were arbitrary, we conclude  $(H, \Pi_\emptyset, T_0, 0) \models \Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{HybridClean}}$ .

( $\Leftarrow$ ) Now, suppose that  $(H, \Pi_\emptyset, T_0, 0) \models \Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{HybridClean}}$ . Let  $\mu_1, \mu_2 \in H$  be two arbitrarily chosen traces. Let  $\pi_1$  and  $\pi_2$  be trace variables, and define  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$ .

Let  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  be such that for every  $t' \leq t$  with  $t' \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  it holds that  $\text{PrefConf}_I(\mu_1^I, \mu_2^I, t')$  is true. By Proposition 6.7, for every such  $t'$  we have  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), H, \Pi, T_0, t') \in \{\top, \perp\}$ . If  $t' \notin \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , we have  $\text{Value}(\varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), H, \Pi, T_0, t') = \perp$ . Thus,  $\text{Value}(\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfl}}(\pi_1, \pi_2), \Pi, T_0, t) = \top$ .



By assumption,  $\text{Value}(\Box \varphi_{\tau_I, \epsilon_I}^{\text{PrefConfI}}(\pi_1, \pi_2)) \rightarrow \varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\mathbb{T}, \mathbb{U}\}$ . Therefore,  $\text{Value}(\varphi_{\tau_O, \epsilon_O}^{\text{PrefConfO}}(\pi_1, \pi_2), H, \Pi, T_0, t) \in \{\mathbb{T}, \mathbb{U}\}$ . Since  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , by Proposition 6.7, this implies that  $\text{PrefConf}_O(\mu_1^O, \mu_2^O, t)$  is true. Since  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  was chosen arbitrarily, we conclude that  $\text{HybridClean}(\tau_I, \epsilon_I, \tau_O, \epsilon_O)$ .  $\square$

In the special case when  $\tau_I = \tau_O = 0$ , i.e., when we consider robust cleanness, the characterization of cleanness is simpler, since in the definition of  $\text{PrefConf}$  we only need to consider the actual compared prefixes (and not their truncated or extended versions) because the timing deviation is 0. As per the definition of robust cleanness, we consider  $\text{PrefConf}$  instantiated with  $\text{TraceConf}$ .

Thus, we can characterize robust cleanness as

$$\begin{aligned} \Phi_{\epsilon_I, \epsilon_O}^{\text{RobustClean}} &= \forall \pi_1. \forall \pi_2. \Box \left( (\Box \varphi_{\epsilon_I}^I) \rightarrow \varphi_{\epsilon_O}^O \right), \text{ where} \\ \varphi_{\epsilon_I}^I &= \left( *_1 \Diamond_{[0,0]} d_I(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon_I \right) \wedge \left( *_2 \Diamond_{[0,0]} d_I(X_{\pi_1}, X_{\pi_2}^{*2}) \leq \epsilon_I \right) \text{ and} \\ \varphi_{\epsilon_O}^O &= \left( *_3 \Diamond_{[0,0]} d_O(X_{\pi_2}, X_{\pi_1}^{*3}) \leq \epsilon_O \right) \wedge \left( *_4 \Diamond_{[0,0]} d_O(X_{\pi_1}, X_{\pi_2}^{*4}) \leq \epsilon_O \right). \end{aligned}$$

**Proposition 6.9.** *Let  $H$  be a deterministic hybrid system defined over a set of real-valued variables  $X$  such that  $0 \in \text{dom}(\mu)$  for each  $\mu \in H$ . Let  $\epsilon_I, \epsilon_O \geq 0$  be rational constants. It holds that*

$$H \text{ is RobustClean}(\epsilon_I, \epsilon_O) \text{ if and only if } H \models \Phi_{\epsilon_I, \epsilon_O}^{\text{RobustClean}}.$$

*Proof.* ( $\implies$ ) First, suppose that  $H$  is  $\text{RobustClean}(\epsilon_I, \epsilon_O)$ . Let  $\mu_1, \mu_2 \in H$  be two arbitrarily chosen traces. Let  $\pi_1$  and  $\pi_2$  be trace variables, and let  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$ .

Let  $t \geq 0$  be an arbitrary time point. If  $\text{Value}(\varphi_{\epsilon_O}^O, H, \Pi, T_0, t) \in \{\mathbb{T}, \mathbb{U}\}$ , then it holds that  $\text{Value}(\Box((\Box \varphi_{\epsilon_I}^I) \rightarrow \varphi_{\epsilon_O}^O), H, \Pi, T_0, t) \in \{\mathbb{T}, \mathbb{U}\}$ . If, on the other hand we have that  $\text{Value}(\varphi_{\epsilon_O}^O, H, \Pi, T_0, t) = \mathbb{F}$ , then it holds that at least one of the conjuncts of  $\varphi_{\epsilon_O}^O$  is false. Suppose that, without loss of generality,  $\text{Value}(*_3 \Diamond_{[0,0]} d_O(X_{\pi_2}, X_{\pi_1}^{*3}) \leq \epsilon_O, H, \Pi, T_0, t) = \mathbb{F}$ . Hence,  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  and according to Definition 4.7, this means that there exists  $t' \leq t$  such that  $\text{Value}(\varphi_{\epsilon_I}^I, H, \Pi, T_0, t') = \mathbb{F}$ . Therefore,  $\text{Value}(\Box \varphi_{\epsilon_I}^I, H, \Pi, T_0, t) = \mathbb{F}$ . Thus, in this case we obtain that  $\text{Value}(\Box((\Box \varphi_{\epsilon_I}^I) \rightarrow \varphi_{\epsilon_O}^O), H, \Pi, T_0, t) = \mathbb{T}$ .

Therefore, since  $t \geq 0$  was chosen arbitrarily, we have  $(H, \Pi, T_0, 0) \models \Box((\Box \varphi_{\epsilon_I}^I) \rightarrow \varphi_{\epsilon_O}^O)$ . Since  $\mu_1$  and  $\mu_2$  were arbitrary, we can conclude  $(H, \Pi_\emptyset, T_0, 0) \models \Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{RobustClean}}$ .

( $\impliedby$ ) Now, suppose that  $(H, \Pi_\emptyset, T_0, 0) \models \Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{RobustClean}}$ . Let  $\mu_1, \mu_2 \in H$  be two arbitrarily chosen traces. Let  $\pi_1$  and  $\pi_2$  be trace variables, and define  $\Pi = \{\pi_1 \mapsto \mu_1, \pi_2 \mapsto \mu_2\}$ .

Let  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  be such that for every  $t' \leq t$  with  $t' \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$  it holds that  $\text{Value}(\varphi_{\epsilon_I}^I, H, \Pi, T_0, t') = \mathbb{T}$ . If  $t' \notin \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ , we have  $\text{Value}(\varphi_{\epsilon_I}^I, H, \Pi, T_0, t') = \mathbb{U}$ . Therefore, we have that  $\text{Value}(\Box \varphi_{\epsilon_I}^I, \Pi, T_0, t) = \mathbb{T}$ .

By assumption,  $\text{Value}(\Box((\Box \varphi_{\epsilon_I}^I) \rightarrow \varphi_{\epsilon_O}^O), H, \Pi, T_0, t) \in \{\mathbb{T}, \mathbb{U}\}$ . Therefore, we obtain that  $\text{Value}(\varphi_{\epsilon_O}^O, H, \Pi, T_0, t) = \mathbb{T}$ , since  $t \in \text{dom}(\mu_1) \cup \text{dom}(\mu_2)$ . Since  $t$  was chosen arbitrarily, we can conclude that  $\text{RobustClean}(\epsilon_I, \epsilon_O)$ .  $\square$

**6.4. Monitoring HyperSTL\*<sub>fin</sub> over finite-length real-valued signals.** We now consider the fragment  $\text{HyperSTL}^*_{fin}$  and describe a method for offline monitoring of  $\text{HyperSTL}^*_{fin}$  properties on finite sets of finite-length signals.

If the given traces are finite timed words, we can obtain from them piecewise linear signals by linear interpolation, or piecewise constant signals by fixing the value for each half-closed interval between time points to be the value at the starting point of this interval.

If the signals in the set are of different time length, we take the minimum length across the set, and consider the traces up to that length. Thus, we ensure that for some  $B$ , all traces are defined over the interval  $[0, B]$ . Furthermore, we only consider formulas in  $\text{HyperSTL}^*_{fin}$  for which the bounds of the temporal operators are such that every subformula has a defined value when the formula is evaluated over traces with time domain  $[0, B]$ .

Our method handles the trace quantifiers similarly to the algorithm for offline monitoring of  $\text{HyperLTL}$  formulas on finite traces given in [FHST17, FHST19]. The method iterates over tuples of generalized timed traces. The arity of the tuples is determined by the quantifier prefix of the formula. For instance, for monitoring a formula of the form  $\Phi = \forall \pi_1 \dots \forall \pi_n \exists \pi'_1 \dots \exists \pi'_m \cdot \varphi$  we will evaluate  $\varphi$  on tuples of GTTs of arity  $n + m$ , to either determine that  $\Phi$  is satisfied over the given set of traces, or return an  $n$ -tuple witnessing a violation. In order to evaluate the trace-quantifier-free formula  $\varphi$  on an  $(n + m)$ -tuple of traces we compute a satisfaction evidence by using the method proposed in [BDSV14]. Note that unlike [FHST17, FHST19] we consider finite traces and formulas with bounded temporal operators. Since we assume that the length of the traces is sufficient for all subformulas of a given  $\text{HyperSTL}^*_{fin}$  formula of interest to have a defined value, the truth value of the quantifier-free part of the formula is defined. As we consider a fixed set of recorded traces, we can check offline the satisfaction of  $\text{HyperSTL}^*_{fin}$  formulas with arbitrary quantifier alternations similarly to [FHST17, FHST19], as outlined above.

Let  $\varphi$  be a trace-quantifier-free formula. Let  $H$  be a finite set of  $\mathbb{R}^l$ -valued finite-length signals, and let  $K \in H^k$  be a tuple of traces of arity  $k$ . Then, we can interpret  $K$  as a real-valued signal  $\kappa$  of order  $k \times l$ . The *satisfaction set* of the formula  $\varphi$  over the signal  $\kappa$  is defined analogously to [BDSV14]. Similarly to [BDSV14], we assume that the traces in  $H$  are piecewise linear and that the atomic predicates are linear, in order to make the computation of the satisfaction set tractable. Note that the atomic predicates used in the characterization of hybrid cleanness in Section 6.3 are linear when the expressions  $d_I(X, X')$  and  $d_O(X, X')$  are linear. The explicit clock variable defines a linear signal. Clearly, if the traces in  $H$  are piecewise linear, then so is the signal  $\kappa$ . Thus, the satisfaction set for  $\varphi$  can be calculated effectively, represented as convex polytopes.

The satisfaction set for the formula  $\varphi$  given the signal  $\kappa$  is a subset of  $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}}$ , defined inductively with respect to structure of  $\varphi$ . Here, a signal  $\kappa$  of order  $k \times l$  is interpreted as a trace assignment for  $k$  trace variables, in which each trace variable is assigned a GTT in the form of a real-valued signal of order  $l$ .

$$\begin{aligned}
\text{Sat}(\alpha, \kappa) &= \{(t, T) \in \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}} \mid \text{Value}(\alpha, H, \kappa, T, t) = \top\}; \\
\text{Sat}(\top, \kappa) &= \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}}; \\
\text{Sat}(\neg \varphi', \kappa) &= (\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}}) \setminus \text{Sat}(\varphi', \kappa); \\
\text{Sat}(\varphi_1 \vee \varphi_2, \kappa) &= \text{Sat}(\varphi_1, \kappa) \cup \text{Sat}(\varphi_2, \kappa); \\
\text{Sat}(\varphi_1 \mathcal{U}_{[a,b]} \varphi_2, \kappa) &= \{(t, T) \in \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}} \mid \exists t' \in [t + a, t + b] : (t', T) \in \text{Sat}(\varphi_2, \kappa) \wedge \\
&\quad \forall t'' \in [t, t'] : (t'', T) \in \text{Sat}(\varphi_1, \kappa) \cup \text{Sat}(\varphi_2, \kappa) \cap [t + a, t + b]\}; \\
\text{Sat}(\varphi_1 \mathcal{S}_{[a,b]} \varphi_2, \kappa) &= \{(t, T) \in \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}} \mid \exists t' \in [t - b, t - a] : (t', T) \in \text{Sat}(\varphi_2, \kappa) \wedge \\
&\quad \forall t'' \in (t', t] : (t'', T) \in \text{Sat}(\varphi_1, \kappa) \cup \text{Sat}(\varphi_2, \kappa) \cap [t - b, t - a]\}; \\
\text{Sat}(*_i \varphi', \kappa) &= \{(t, T) \in \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^{\mathcal{I}} \mid (t, T[i \mapsto t]) \in \text{Sat}(\varphi', \kappa)\}.
\end{aligned}$$

Once the satisfaction sets for the atomic predicates appearing in the given formula have been computed, the satisfaction sets for the composite formulas can be constructed by following the inductive definition above. Under the assumptions we made above about the signals and the atomic predicates, the satisfaction sets for the atomic predicates can be computed directly using the method described in [BDSV14]. For further details, we refer to [BDSV14].

In order to perform monitoring for the formula  $\Phi_{\tau_I, \epsilon_I, \tau_O, \epsilon_O}^{\text{HybridClean}}$  defined in Section 6.3, we bound the temporal operators based on the signal length  $B$ , and consider the case when  $\tau_I > 0$  and  $\tau_O > 0$  which ensures that the intervals in the operators are non-singular.

## 7. CASE STUDY

In this section we evaluate the proposed notion of conformance-based cleanness in a real application context, known as the Diesel Emissions Scandal [BDH19, HBDK18, KHB18, BDH18, DBB<sup>+</sup>17, CLP<sup>+</sup>17, BDFH16]: Starting in fall 2015, millions of diesel cars were found being equipped with defeat devices reducing the effectiveness of emission cleaning systems during real-world usage — in contrast to the regulator-defined driving scenarios on a chassis dynamometer, where the amount of emitted pollutants stay well below the applicable limits.

It was soon suspected, that the singularities of the testing procedure were straightforward to identify and hence made cheating easy; in particular, there was only a single test cycle for testing emission cleaning systems, to be executed under very particular conditions. In the European Union, this was the *New European Driving Cycle* (NEDC) [Uni13], the speed profile of which is shown in Fig. 3. The NEDC consists of four repetitions of an elementary *urban driving cycle* (UDC) followed by one *extra urban driving cycle* (EUDC). Each test run is preceded by a preconditioning phase (PRECON), in which three EUDCs are driven consecutively. Between PRECON and the test, the vehicle has to cool down for 6 to 36 hours at an ambient temperature between 20 and 30 degrees Celsius.

Robust cleanness gives us a way of deriving additional test cycles that are reasonable w.r.t. the official NEDC. For a concrete context, “reasonable” is defined by an accompanying formally defined *contract*. The contributions of this paper enable us to go beyond previous work [BDH19] where a contract allowed inputs and outputs to deviate only in the value domains, but not in the time domain.

**7.1. Experimental Setup.** Our empirical studies apply the theory developed in the previous sections in a very specific setting. The system under test is a Nissan NV200 Evalia equipped with a Renault 1.5 dci (110hp) diesel engine and approved w.r.t. regulation *Euro 6b*. All tests were conducted in November 2020. As shown in Fig. 3, the car is fixed on a chassis dynamometer and attached to a portable emissions measurement system (PEMS) in preparation for the test. The PEMS is connected to the onboard diagnostics (OBD) interface of the vehicle. During a test, the PEMS measures the amount of several gases at the end of the car’s exhaust pipe and logs the data received from OBD. The PEMS is able to internally synchronise the times of gas measurements and OBD data. We will not consider the internal PEMS retiming and instead analyse the final data set. As input, we consider the OBD speed data, as output, the sum of emitted NO and NO<sub>2</sub> (abbreviated as NO<sub>x</sub>). The input and output is sampled by 1 Hz. The amount of NO<sub>x</sub> emitted along different runs is comparable only to a limited extend. This is because the emission cleaning system used can have internal regeneration phases, which — from an external observer perspective — are




Figure 3: Left: New European Driving Cycle (NEDC); Right: Test Setup with Nissan NV200 Evalia on a chassis dynamometer attached to a PEMS.




Figure 4: PERMNEDC (solid, black line) compared to NEDC (dashed, blue line)

triggered nondeterministically. Hence, for the formal analysis we consider the accumulated output over 1180 seconds (the length of a complete NEDC). This is also the value decisive for type approval according to the official regulations.

Formally, a contract specifying *clean* behaviour for a system is given as a tuple. Previous work [BDH19] proposes the concrete contract  $\mathcal{C}_t = \langle d_I, d_O, \epsilon_I, \epsilon_O \rangle$  for diesel doping tests, where  $d_I(i_1, i_2) \triangleq |i_1 - i_2|$  and  $d_O(o_1, o_2) \triangleq |o_1 - o_2|$ , and  $\epsilon_I = 15$  km/h and  $\epsilon_O = 180$  mg/km.  $\mathcal{C}_t$  is based on robust cleanliness, i.e.,  $\text{TraceConf}_{\epsilon_I}$  and  $\text{TraceConf}_{\epsilon_O}$  conformance for inputs and outputs. In the following, we will add the designated conformance notions to the tuple, i.e., instead of  $\mathcal{C}_t$  we write  $\mathcal{C} = \langle d_I, d_O, \text{TraceConf}_{\epsilon_I}, \text{TraceConf}_{\epsilon_O} \rangle$  (implicitly encoding  $\epsilon_I$  and  $\epsilon_O$ ). In  $\mathcal{C}$  the degree of deviation is constrained by a threshold on the pointwise difference of the new and the reference test input. Comparisons of data at different time points are not possible. The theory of conformance-based doping tests developed in this paper improves upon the previous testing methodology by enabling variation in the time domain. This gives us the possibility of reordering NEDC segments, lengthening a test beyond the time limits of the original NEDC, and of increased tolerance regarding human-caused input distortions during test execution. With the parameters  $d_I, d_O, \epsilon_I, \epsilon_O$  and  $\text{TraceConf}_{\epsilon_O}$  in  $\mathcal{C}$  fixed, we show how adaptations of  $\mathcal{C}$  with different input conformance relations are of benefit for software doping analysis.

**NEDC Permutations:** Based on the conformance notions in this work, we propose a new test cycle PERMNEDC in which NEDC segments are permuted on the time axis. Fig. 4 shows the test cycle. In each of the four UDC segments the three non-zero speed-phases are permuted. The transformation from NEDC to PERMNEDC can be




Figure 5: DOUBLENEDC (solid, black line) compared to NEDC (dashed, blue line)

described by a retiming function  $r_p$ . An explicit definition of  $r_p$  is space consuming, hence we omit it. Along with the new cycle, we propose two suitable variants of contract  $\mathcal{C}$  with different input conformances. Neither input conformance is constrained by a time threshold; in other words,  $\tau = \infty$ , so we omit  $\tau$  in the index.

- Contract  $\mathcal{C}_a$  is as  $\mathcal{C}$ , but entails input conformance  $\text{Conf}_{\epsilon_I}^{\text{Ret}_a}$ , where  $\text{Ret}_a = \{(r, r^{-1}) \mid r \in \mathcal{T} \rightarrow \mathcal{T} \text{ and } r \text{ is total and bijective}\}$  is the family of retimings that allows any reordering of the NEDC inputs. Notably, no inputs can be added or removed.
- Contract  $\mathcal{C}_p$  adjusts  $\mathcal{C}$  by enforcing input conformance  $\text{Conf}_{\epsilon_I}^{\text{Ret}_p}$ , where  $\text{Ret}_p = \{(r_p, r_p^{-1})\}$  is the family of retimings that only allows the particular retiming  $r_p$  used to design the test cycle as discussed above. This input conformance is stricter than  $\text{Conf}_{\epsilon_I}^{\text{Ret}_a}$  above; it enforces that PERMNEDC is not permuted any further by the driver.

**NEDC Lengthening:** Conformance-based doping tests can run longer than the NEDC; this is not possible with robust cleanness. We propose the test cycle DOUBLENEDC, which consists of two consecutive NEDCs. In contrast to all other test cycles in this paper, DOUBLENEDC produces two outputs: the first after 1180 seconds, and the second after 2360 seconds. The first half of this cycle is a classical “cold” NEDC (i.e., the engine cooled down before the test execution). The second half is a “hot” NEDC, since the cool-down phase was implicitly skipped. Also, the PRECON phase is skipped implicitly; there is only a single EUDC (instead of three) prior to the second NEDC. The inputs of both NEDCs can be compared by using the retiming functions  $\text{id}$  and  $r_d = \lambda t. t \bmod 1180$ .  $r_d$  maps time points of the global “test case clock” to the local time point in the NEDC time domain. DOUBLENEDC requires to adapt contract  $\mathcal{C}$  to  $\mathcal{C}_d$  by replacing robust cleanness by cleanness with synchronised retiming: it entails input conformance  $\text{Conf}_{\epsilon_I}^{\text{Ret}_d}$ , which allows only  $\text{id}$  and  $r_d$  as retiming functions, i.e.,  $\text{Ret}_d = \{(\text{id}, r_d)\}$ . Similarly,  $\mathcal{C}_d$  must include the synchronisation retiming function  $\text{Sync}_d(r_1, r_2) = (\text{id}, r_d)$ , which enforces that both DOUBLENEDC outputs are compared to the single NEDC output (independent of the input retimings  $r_1$  and  $r_2$ ).

**Human Time Imprecision Tolerance:** Diesel doping tests are executed by humans driving a car. Humans tend to make mistakes when driving. Mistakes can be the over- or undershooting of the targeted speed (the error is on the value axis), or accelerations or decelerations happening too early or too late (the error is a shift on the time axis), or superpositions thereof. To compensate for both value and time errors, we use hybrid conformance. As a formal contract, this would be expressed by a variant of  $\mathcal{C}$ , in which the input conformance is replaced by  $\text{HybridConf}_{\tau_I, \epsilon_I}$  for some  $\tau_I > 0$ . For the purpose of demonstration, we will later analyse several such variants of  $\mathcal{C}$ , each variant with a




Figure 6: SINENEDC (solid, black line) compared to NEDC (dashed, blue line)

unique value for  $\tau_I$  and  $\epsilon_I$ , i.e., we consider the contract  $\mathcal{C}(\tau_I, \epsilon_I)$  parametrised in  $\tau_I$  and  $\epsilon_I$ . Concrete values for  $\tau_I$  and  $\epsilon_I$  must be specified when using the contract.

A test cycle that reflects drivings rich of acceleration and deceleration phases—and is hence particularly prone to human driving errors—is SINENEDC [BDH19]. SINENEDC is defined as the NEDC superimposed by a sine curve, formally

$$\text{SINENEDC}(t) = \max\{0, \text{NEDC}(t) + 5 \sin(0.5t)\},$$

with a maximum input deviation from NEDC of 5km/h, compare Fig. 6. We will evaluate SINENEDC under several variants of  $\mathcal{C}(\tau_I, \epsilon_I)$ .

Human time imprecision is as yet not considered in test cycles PERMNEDC and DOUBLENEDC, both cycles require a cycle-specific conformance predicate. However, tolerance for human imprecision can be added to these predicates by means of conformance and retiming composition. Let  $\text{Ret}^{(1)}$  and  $\text{Ret}^{(2)}$  be two families of retimings. Then  $\text{Ret}^{(2)} \circ \text{Ret}^{(1)} \triangleq \{(r_1^{(2)} \circ r_1^{(1)}, r_2^{(2)} \circ r_2^{(1)} \mid (r_1^{(2)}, r_2^{(2)}) \in \text{Ret}^{(2)} \text{ and } (r_1^{(1)}, r_2^{(1)}) \in \text{Ret}^{(1)}\}$  is the component-wise function composition. The definition for conformance composition is  $\text{Conf}_{\tau_2, \epsilon}^{\text{Ret}^{(2)}} \circ \text{Conf}_{\tau_1, \epsilon}^{\text{Ret}^{(1)}} \triangleq \text{Conf}_{\infty, \epsilon}^{\text{Ret}^{(3)}}$ , where  $\text{Ret}^{(3)} = \text{Ret}^{(2)} \circ \text{Ret}^{(1)}$  composes the individual retimings. The  $\tau_1$ - and  $\tau_2$ -constraints on  $\text{Ret}^{(1)}$  and  $\text{Ret}^{(2)}$  are applied before the composition. It is not necessary to apply further timing constraints to the resulting retiming, hence we allow infinite  $\tau$ . To overcome the human imprecisions for PERMNEDC and DOUBLENEDC, we use the parametrised contracts  $\mathcal{C}_p(\tau_I, \epsilon_I)$  and  $\mathcal{C}_d(\tau_I, \epsilon_I)$ , adaptations of  $\mathcal{C}_p$  and  $\mathcal{C}_d$ , with input conformances  $\text{HybridPermConf}_{\tau_I, \epsilon_I} = \text{HybridConf}_{\tau_I, \epsilon_I} \circ \text{Conf}_{\epsilon_I}^{\text{Ret}_p}$  and  $\text{HybridDoubleConf}_{\tau_I, \epsilon_I} = \text{HybridConf}_{\tau_I, \epsilon_I} \circ \text{Conf}_{\epsilon_I}^{\text{Ret}_d}$ , respectively. As for hybrid conformance in  $\mathcal{C}(\tau_I, \epsilon_I)$ , we will specify concrete  $\tau_I$  and  $\epsilon_I$  upon usage of the contracts. Notably, for DOUBLENEDC, this does not have effects on the output conformance, because  $\text{Sync}_d$  does not consider the input retiming. This is important, because outputs are available only at time points 1180 and 2360 and must not be moved to time points different than that. We do not compose  $\text{Conf}_{\epsilon_I}^{\text{Ret}_a}$  and hybrid conformance, because  $\text{Ret}_a$  allows any possible NEDC permutation, which naturally reduces the effect of timing imprecisions.

Table 1 summarises the contracts presented above.

$\mathcal{C}$	$= \langle d_I, d_O, \text{TraceConf}_{\epsilon_I},$	$\text{TraceConf}_{\epsilon_O}$	$\rangle$
$\mathcal{C}_a$	$= \langle d_I, d_O, \text{Conf}_{\epsilon_I}^{\text{Ret}_a},$	$\text{TraceConf}_{\epsilon_O}$	$\rangle$
$\mathcal{C}_p$	$= \langle d_I, d_O, \text{Conf}_{\epsilon_I}^{\text{Ret}_p},$	$\text{TraceConf}_{\epsilon_O}$	$\rangle$
$\mathcal{C}_d$	$= \langle d_I, d_O, \text{Conf}_{\epsilon_I}^{\text{Ret}_d},$	$\text{TraceConf}_{\epsilon_O}, \text{Sync}_d$	$\rangle$
$\mathcal{C}(\tau_I, \epsilon_I)$	$= \langle d_I, d_O, \text{HybridConf}_{\tau_I, \epsilon_I},$	$\text{TraceConf}_{\epsilon_O}$	$\rangle$
$\mathcal{C}_p(\tau_I, \epsilon_I)$	$= \langle d_I, d_O, \text{HybridConf}_{\tau_I, \epsilon_I} \circ \text{Conf}_{\epsilon_I}^{\text{Ret}_p},$	$\text{TraceConf}_{\epsilon_O}$	$\rangle$
$\mathcal{C}_d(\tau_I, \epsilon_I)$	$= \langle d_I, d_O, \text{HybridConf}_{\tau_I, \epsilon_I} \circ \text{Conf}_{\epsilon_I}^{\text{Ret}_d},$	$\text{TraceConf}_{\epsilon_O}, \text{Sync}_d$	$\rangle$

Table 1: Overview of the evaluated contracts. For  $\mathcal{C}$ ,  $\mathcal{C}_a$ ,  $\mathcal{C}_p$  and  $\mathcal{C}_d$ ,  $\epsilon_I = 15$  km/h. For all contracts,  $\epsilon_O = 180$  mg/km,  $d_I(i_1, i_2) = |i_1 - i_2|$  and  $d_O(o_1, o_2) = |o_1 - o_2|$ .  $\text{Ret}_a$ ,  $\text{Ret}_p$  and  $\text{Ret}_d$  are defined as explained in the text above.

**7.2. Computing Parameters of Hybrid Conformance.** In some experiments we compute, for a fixed time threshold  $\tau$  and two test cycles, the minimal value error  $\epsilon$  such that hybrid conformance holds for the input. The implementation of this computation is inspired by the the HyperSTL\* formula  $\varphi_{\tau, \epsilon}^{\text{HybridConf}}$ , i.e., it computes  $\min_{\epsilon} \varphi_{\tau, \epsilon}^{\text{HybridConf}}$  for two test executions  $\pi_1$  and  $\pi_2$ . The algorithm is sketched below.

$$\begin{aligned} \text{localMin}(t_1, \mu_1, \mu_2, \tau) &= \min \{d_{\mathcal{Y}}(\mu_2[t_2], \mu_1[t_1]) \mid t_2 \in [t_1 - \tau; t_1 + \tau] \cap \text{dom}(\mu_2)\} \\ \text{globalMin}(\mu_1, \mu_2, \tau) &= \max \{\text{localMin}(t_1, \mu_1, \mu_2, \tau) \mid t_1 \in \text{dom}(\mu_1)\} \\ \epsilon_{\min}(\mu_1, \mu_2, \tau) &= \max \{\text{globalMin}(\mu_1, \mu_2, \tau), \text{globalMin}(\mu_2, \mu_1, \tau)\} \end{aligned}$$

Here,  $\text{localMin}(t_1, x_1, x_2, \tau)$  computes the minimal  $\epsilon$  for subformula  $\diamond_{[0, \tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon \vee \diamond_{[0, \tau]} d_{\mathcal{Y}}(X_{\pi_2}, X_{\pi_1}^{*1}) \leq \epsilon$ , where the value of  $X_{\pi_1}^{*1}$  is frozen at time  $t_1$ .  $\text{globalMin}(x_1, x_2, \tau)$  reflects the *Globally* and *Freeze* operator: it finds the maximum by quantifying over all  $t_1 \in \text{dom}(x_1)$  and by calling  $\text{localMin}$  with the frozen time value  $t_1$ . To reflect the complete formula,  $\epsilon_{\min}(x_1, x_2, \tau)$  returns the maximum of the conjuncts, which are the results of  $\text{globalMin}$  for both combinations of  $x_1$  and  $x_2$ .

The computations for  $\text{HybridPermConf}$  and  $\text{HybridDoubleConf}$  proceed in two steps. Both  $\text{Ret}_p$  and  $\text{Ret}_d$  are singleton sets; it is known which retiming must be applied first. For two traces  $\mu_1$  and  $\mu_2$  and retiming  $(r_1, r_2)$ , there are shifted traces  $\mu'_1 = \mu_1 \circ r_2$  and  $\mu'_2 = \mu_2 \circ r_1$ . The minimal  $\epsilon$  for hybrid conformance is given by  $\epsilon_{\min}^{(r_1, r_2)}(\mu_1, \mu_2, \tau) = \max \{\text{globalMin}(\mu_1, \mu'_2, \tau), \text{globalMin}(\mu_2, \mu'_1, \tau)\}$ .

**7.3. Test Results & Verdicts.** We executed each of NEDC, PERM NEDC, DOUBLE NEDC and SIN NEDC two times. We identify a concrete test execution by a suffix -1 or -2 to test cycle identifier (e.g., NEDC-1 is the first and NEDC-2 the second execution of NEDC). Raw data and the implementation of the analysis is available online [BFH20]. For NEDC, we combined the result of both executions to an average value of 182 mg/km of  $\text{NO}_x$ . Notably, the Euro 6b regulation (to which our car is supposed to conform to) allows at most 80 mg/km, and the car under test is certified with 60.8 mg/km according to its documentation. The car is 3 years old.




Figure 7: PERMNEDC-1 speed (black) and NEDC speed (blue) in km/h, and accumulated  $\text{NO}_x$  for PERMNEDC-1 (red) and NEDC (orange) in mg/km.

For doping detection, a test verdict is only meaningful if its input trace is conformant to that of the average NEDC execution; otherwise, the test is trivially passed. We will first evaluate PERMNEDC w.r.t.  $\mathcal{C}_a$  and  $\mathcal{C}_p$ , DOUBLENEDC w.r.t.  $\mathcal{C}_d$ , and SINENEDC w.r.t.  $\mathcal{C}$ . To demonstrate the effects of hybrid conformance, we then analyse the experiments w.r.t. the parametrised variants of the contracts  $\mathcal{C}$ ,  $\mathcal{C}_p$  and  $\mathcal{C}_d$ , respectively. By definition of the test cycles, the nominal value difference for PERMNEDC and DOUBLENEDC after retiming is zero, and for SINENEDC it is 5 km/h. Though, due to human imprecisions, the actual differences are significantly higher.

- The executions of PERMNEDC are shown in Fig. 7 and 8. The amount of emitted  $\text{NO}_x$  were 392 mg/km for PERMNEDC-1 and 316 mg/km for PERMNEDC-2.  $\text{Conf}_{\epsilon_I}^{\text{Ret}_a}$  does hold for  $\epsilon_I \geq 3$  km/h for both executions; with contract  $\mathcal{C}_a$ , which defines  $\epsilon_I = 15$  km/h, drastic deviations of  $\text{NO}_x$  can be detected as doping. It is detected for PERMNEDC-1, i.e., the cleanliness test fails, as the difference of  $\text{NO}_x$  (compared to NEDC) is 210 mg/km and hence greater than  $\epsilon_O = 180$  mg/km defined by  $\mathcal{C}_a$ . Test PERMNEDC-2 passes with an  $\text{NO}_x$  difference of 134 mg/km which is within the contract.

With contract  $\mathcal{C}_p$  and input conformance  $\text{Conf}_{\epsilon_I}^{\text{Ret}_p}$ , the test verdict for PERMNEDC-1 is different.  $\text{Conf}_{\epsilon_I}^{\text{Ret}_p}$  would only hold for  $\epsilon_I \geq 16$  km/h, which is above the contract defined threshold of 15 km/h. Hence, PERMNEDC-1 is not adduced and the test trivially passed.

- DOUBLENEDC-1 and 2, shown in Fig. 9 and 10, lead to an average emission of 305 mg/km, respectively 308 mg/km of  $\text{NO}_x$ . Executions of DOUBLENEDC are twice as long as regular NEDC tests and produce two outputs. The measurements for DOUBLENEDC-1 report (229, 382) mg/km, for DOUBLENEDC-2 (207, 408) mg/km. To determine the verdicts for contract  $\mathcal{C}_d$ , we first check if  $\text{Conf}_{\epsilon_I}^{\text{Ret}_d}$  holds. This turns out not to hold for DOUBLENEDC-2, because we observed value deviations of up to 25 km/h. This test is therefore trivially passed. For DOUBLENEDC-1 all value deviations remain below the 15 km/h threshold; this test run is thus to be considered relevant for output comparison. According to the retiming synchronisation in  $\mathcal{C}_d$ , each of the outputs 229 and 382 must be






Figure 8: PERMNEDC-2 speed (black) and NEDC speed (blue) in km/h, and accumulated  $\text{NO}_x$  for PERMNEDC-2 (red) and NEDC (orange) in mg/km.




Figure 9: DOUBLENEDC-1 speed (black) and NEDC speed (blue) in km/h, and accumulated  $\text{NO}_x$  for DOUBLENEDC-1 (red) and NEDC (orange) in mg/km.

compared to the NEDC output 182. The output conformance is violated for the second output, with a difference of 200 mg/km, exceeding the allowed  $\epsilon_O = 180$  mg/km threshold. Hence, DOUBLENEDC-1 fails — doping is detected.

- During the test executions of SINENEDC, we measured 483 mg/km and 632 mg/km. The test progression is shown in Fig. 11 and 12. In SINENEDC-1, speed values deviate by up to 18 km/h, which exceeds the  $\epsilon_I$  threshold in  $\mathcal{C}$ , so this test run is trivially passed. SINENEDC-2 respects the  $\epsilon_I$  threshold because inputs never deviate by more than 13 km/h. Consequently, SINENEDC-2 convicts our test car of doping, as the output difference of 450 mg/km is 2.5 times the allowed threshold  $\epsilon_O$ .
- As discussed, we use hybrid conformance to compensate for human driving imprecisions. In this context, Table 2 details the effect of a choice of  $\tau$  on the maximal value error. We fix a maximum value that we allow for the time offset  $\tau_I$ . For this  $\tau_I$  we analyse our dataset to find the minimal  $\epsilon_I$  such that for the combination of  $\tau_I$  and  $\epsilon_I$  the input




Figure 10: DOUBLENEDC-2 speed (black) and NEDC speed (blue) in km/h, and accumulated  $\text{NO}_x$  for DOUBLENEDC-2 (red) and NEDC (orange) in mg/km.




Figure 11: SINENEDC-1 speed (black) and NEDC speed (blue) in km/h, and accumulated  $\text{NO}_x$  for SINENEDC-1 (red) and NEDC (orange) in mg/km.

Test Name	Contract	Input	Conformance	$\tau_I = 0$	$\tau_I = 1$	$\tau_I = 2$	$\tau_I = 3$	$\tau_I = 5$	$\tau_I = 10$	$\tau_I = 15$	$\tau_I = 20$
PERMNEDC-1	$C_p(\tau_I, \epsilon_I)$	HybridPermConf		$\epsilon_I = 16$	$\epsilon_I = 16$	$\epsilon_I = 16$	$\epsilon_I = 11$	$\epsilon_I = 8$	$\epsilon_I = 8$	$\epsilon_I = 8$	$\epsilon_I = 8$
PERMNEDC-2	$C_p(\tau_I, \epsilon_I)$	HybridPermConf		$\epsilon_I = 11$	$\epsilon_I = 10$	$\epsilon_I = 7$	$\epsilon_I = 7$	$\epsilon_I = 7$	$\epsilon_I = 7$	$\epsilon_I = 7$	$\epsilon_I = 7$
DOUBLENEDC-1	$C_d(\tau_I, \epsilon_I)$	HybridDoubleConf		$\epsilon_I = 15$	$\epsilon_I = 12$	$\epsilon_I = 11$	$\epsilon_I = 9$	$\epsilon_I = 6$	$\epsilon_I = 6$	$\epsilon_I = 6$	$\epsilon_I = 6$
DOUBLENEDC-2	$C_d(\tau_I, \epsilon_I)$	HybridDoubleConf		$\epsilon_I = 25$	$\epsilon_I = 18$	$\epsilon_I = 10$	$\epsilon_I = 8$	$\epsilon_I = 8$	$\epsilon_I = 8$	$\epsilon_I = 8$	$\epsilon_I = 8$
SINENEDC-1	$C(\tau_I, \epsilon_I)$	HybridConf		$\epsilon_I = 18$	$\epsilon_I = 16$	$\epsilon_I = 15$	$\epsilon_I = 12$	$\epsilon_I = 9$	$\epsilon_I = 7$	$\epsilon_I = 6$	$\epsilon_I = 6$
SINENEDC-2	$C(\tau_I, \epsilon_I)$	HybridConf		$\epsilon_I = 13$	$\epsilon_I = 11$	$\epsilon_I = 9$	$\epsilon_I = 9$	$\epsilon_I = 7$	$\epsilon_I = 7$	$\epsilon_I = 7$	$\epsilon_I = 7$

Table 2: Comparison of minimal value thresholds  $\epsilon_I$  for fixed  $\tau_I$ . Values are given as km/h and time in seconds.

traces under consideration satisfy the cycle-specific hybrid conformance. For  $\tau_I = 0$  we get exactly the  $\epsilon_I$  for which the two traces satisfy  $\text{Conf}_{\epsilon_I}^{\text{Ret}^p}$  (for PERMNEDC),  $\text{Conf}_{\epsilon_I}^{\text{Ret}^d}$  (for DOUBLENEDC), and  $\text{TraceConf}_{\epsilon_O}$  (for SINENEDC). Table 2 shows the computed  $\epsilon_I$




Figure 12: SINENEDC-2 speed (black) and NEDC speed (blue) in km/h, and accumulated  $\text{NO}_x$  for SINENEDC-2 (red) and NEDC (orange) in mg/km.

values for  $\tau_I = 0, 1, 2, 5, 10, 15$  and 20 seconds. As expected, an increasing  $\tau_I$  induces the minimal  $\epsilon_I$  to decrease. At  $\tau_I = 5$  the decrease in the value error reduces notably. This happens because the error is only partially caused by the incorrect timing of the driver. From the values reported in Table 2 we see that if we allow deviation for the input  $\tau_I = 2$ , and keep  $\epsilon_I = 15$ , then we have that  $\text{HybridDoubleConf}_{\tau_I, \epsilon_I}(\text{NEDC}, \text{DOUBLENEDC-2})$  and  $\text{HybridConf}_{\tau_I, \epsilon_I}(\text{NEDC}, \text{SINENEDC-1})$  hold. For time threshold  $\tau_I = 3$  seconds  $\text{HybridPermConf}_{\tau_I, \epsilon_I}(\text{NEDC}, \text{PERMNEDC-1})$  also holds. Thus, under hybrid conformance these pairs of traces will be considered in the cleanness test for contracts  $\mathcal{C}_d(2, 15)$ ,  $\mathcal{C}(2, 15)$  and  $\mathcal{C}_p(3, 15)$ , respectively, while under their original contract and input conformance they are to be dismissed.

**7.4. Evaluation and Discussion.** The amounts of emitted  $\text{NO}_x$  observed during our experiments provide clear indications of software doping regarding the car’s emission cleaning system. The conformance-based contracts provide the formal basis for this verdict, as discussed above. We here complement this fact with a more intuitive explanation of the behaviour observed.

- **PERMNEDC** slightly reorders NEDC segments in the UDC part of the test cycle. During this part, the measured  $\text{NO}_x$  does not significantly differ from the NEDC reference. However, during the (unmodified) EUDC part, the amount of emissions grows significantly. It is very unlikely to find a physical explanation for the  $\text{NO}_x$  increase; and very likely, that the cleaning system is optimised specifically for the NEDC.
- The **DOUBLENEDC** executions appear to reveal that the emission cleaning system optimisation can also rely on engine temperature or execution time instead of speed data. Physically, many of the common emission cleaning techniques require a hot engine to work properly (and none of them requires a cold engine). Therefore, a lower  $\text{NO}_x$  value can be expected if the NEDC is run with a hot engine. In our experiments, however, the  $\text{NO}_x$  emissions in the hot half are almost two times higher than in the initial cold part. In

other words, the emission cleaning performance is reduced after the first NEDC execution. There is no physical explanation for this behaviour. Inside the software, detecting the end of an NEDC trip can be implemented very easily, for instance with a timer counting from 1180 — the length of NEDC — to zero.

- With SINENEDC, we test the cleaning system during driving behaviour which is rich in accelerations and decelerations. An increased amount of  $\text{NO}_x$  can possibly be explained by physical phenomena. However, we measured an increase of factors 2.7 and 3.5; these numbers can be safely considered as too high for a trustworthy emission cleaning system.

Software doping theory provides the basis for detecting software behaviour violating a formal contract. In this, physical aspects of the emission cleaning system should be considered during the construction of test cases, and test cycles for which drastically higher emissions can be explained physically, should not be considered. The test cycles we used for our experiments were picked with automotive expertise to avoid physically stressful cycles. If test cases are generated automatically from a contract, the physical constraints could be captured by the contract.

The contracts we use for our experiments can be interpreted as very generous in favour of the manufacturers. Input thresholds such as 15 km/h and 2 seconds appear as reasonable values, keeping all tests close enough to the original NEDC. For the output threshold, we use a very large deviation value of 180 mg/km, which allows  $\text{NO}_x$  emissions to almost double compared to the original NEDC value. Despite the generosity of the contracts, our experiments have been able to reveal doping for all experiments except PERMNEDC-2.

The analysis of the data shows that it is indeed necessary to not only consider a deviation of value, but to also allow for timing deviations. Considering value and timing deviations offers a rich set of potential test cycles for doping tests and allows to realistically verify conformance of a test cycle and a reference cycle; especially when the quality of the studied driving tests suffers from the human-caused input distortions. In this regard, cleanness notions entailing hybrid conformance are more adequate than conformance notions demanding punctual test executions, such as robust cleanness. Without hybrid conformance, more of the doping cases we have detected would slip through.

Finally, while hybrid conformance is central to the case study considered here, our generic theory of conformance-based cleanness allows for using other conformance notions as appropriate for the CPS under test.

#### ACKNOWLEDGMENTS

The work of Sebastian Biewer and Holger Hermanns is supported by the ERC Grant 695614 (POWVER) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) grant 389792660 as part of TRR 248, see <https://perspicuous-computing.science>. The work of Sebastian Biewer is supported by the Saarbrücken Graduate School of Computer Science. The work of Holger Hermanns is supported by the Key-Area Research and Development Program Grant 2018B010107004 of Guangdong Province. The work of Mohammad Reza Mousavi has been partially supported by the UKRI Trustworthy Autonomous Systems Node in Verifiability, with Grant Award Reference EP/V026801/1.

## REFERENCES

- [AB16] Shreya Agrawal and Borzoo Bonakdarpour. Runtime verification of k-safety hyperproperties in hyperltl. In *IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016*, pages 239–252. IEEE Computer Society, 2016. doi:10.1109/CSF.2016.24.
- [ACM<sup>+</sup>18] Hugo L. S. Araujo, Gustavo Carvalho, Morteza Mohaqeqi, Mohammad Reza Mousavi, and Augusto Sampaio. Sound conformance testing for cyber-physical systems: Theory and implementation. *Sci. Comput. Program.*, 162:35–54, 2018. doi:10.1016/j.scico.2017.07.002.
- [AH94] Rajeev Alur and Thomas A. Henzinger. A really temporal logic. *J. ACM*, 41(1):181–204, 1994. doi:10.1145/174644.174651.
- [AHF<sup>+</sup>14] Houssam Abbas, Bardh Hoxha, Georgios E. Fainekos, Jyotirmoy V. Deshmukh, James Kapinski, and Koichi Ueda. Wip abstract: Conformance testing as falsification for cyber-physical systems. In *ACM/IEEE International Conference on Cyber-Physical Systems, ICCPS, Berlin, Germany, April 14-17, 2014*, page 211. IEEE Computer Society, 2014. doi:10.1109/ICCPS.2014.6843724.
- [ALFS11] Yashwanth Annpureddy, Che Liu, Georgios E. Fainekos, and Sriram Sankaranarayanan. S-taliro: A tool for temporal logic falsification for hybrid systems. In Parosh Aziz Abdulla and K. Rustan M. Leino, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 17th International Conference, TACAS 2011, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011, Saarbrücken, Germany, March 26-April 3, 2011. Proceedings*, volume 6605 of *Lecture Notes in Computer Science*, pages 254–257. Springer, 2011. doi:10.1007/978-3-642-19835-9\_21.
- [AMF14] Houssam Abbas, Hans D. Mittelmann, and Georgios E. Fainekos. Formal property verification in a conformance testing framework. In *Twelfth ACM/IEEE International Conference on Formal Methods and Models for Codesign, MEMOCODE 2014, Lausanne, Switzerland, October 19-21, 2014*, pages 155–164. IEEE, 2014. doi:10.1109/MEMCOD.2014.6961854.
- [ARM17] A. Aerts, M. Reniers, and M.R. Mousavi. Chapter 19 - model-based testing of cyber-physical systems. In Houbing Song, Danda B. Rawat, Sabina Jeschke, and Christian Brecher, editors, *Cyber-Physical Systems, Intelligent Data-Centric Systems*, pages 287 – 304. Academic Press, Boston, 2017. URL: <http://www.sciencedirect.com/science/article/pii/B9780128038017000195>, doi:<https://doi.org/10.1016/B978-0-12-803801-7.00019-5>.
- [BDFH16] Gilles Barthe, Pedro R. D’Argenio, Bernd Finkbeiner, and Holger Hermanns. Facets of software doping. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation: Discussion, Dissemination, Applications - 7th International Symposium, ISOFA 2016, Part II*, volume 9953 of *LNCS*, pages 601–608, 2016. URL: [http://dx.doi.org/10.1007/978-3-319-47169-3\\_46](http://dx.doi.org/10.1007/978-3-319-47169-3_46), doi:10.1007/978-3-319-47169-3\_46.
- [BDH18] Sebastian Biewer, Pedro R. D’Argenio, and Holger Hermanns. Cyber-physical doping tests. In *3rd Workshop on Monitoring and Testing of Cyber-Physical Systems, MT@CPSWeek 2018, Porto, Portugal, April 10, 2018*, pages 18–19. IEEE, 2018. doi:10.1109/MT-CPS.2018.00016.
- [BDH19] Sebastian Biewer, Pedro R. D’Argenio, and Holger Hermanns. Doping tests for cyber-physical systems. In David Parker and Verena Wolf, editors, *Quantitative Evaluation of Systems, 16th International Conference, QEST 2019, Glasgow, UK, September 10-12, 2019, Proceedings*, volume 11785 of *Lecture Notes in Computer Science*, pages 313–331. Springer, 2019. doi:10.1007/978-3-030-30281-8\_18.
- [BDSV14] Lubos Brim, Petr Dluhos, David Safranek, and Tomas Vejpustek. STL\*: Extending signal temporal logic with signal-value freezing operator. *Inf. Comput.*, 236:52–67, 2014. doi:10.1016/j.ic.2014.01.012.
- [BFH20] Sebastian Biewer, Michael Fries, and Thomas Heinze. Conformance relations and hyperproperties for doping detection in time and space (supplementary material). <https://www.powver.org/publications/conformance-based-doping-detection>, 2020.
- [BPS20] Borzoo Bonakdarpour, Pavithra Prabhakar, and César Sánchez. Model checking timed hyperproperties in discrete-time systems. In Ritchie Lee, Susmit Jha, and Anastasia Mavridou, editors, *NASA Formal Methods - 12th International Symposium, NFM 2020, Moffett Field, CA, USA, May 11-15, 2020, Proceedings*, volume 12229 of *Lecture Notes in Computer Science*, pages 311–328. Springer, 2020. doi:10.1007/978-3-030-55754-6\_18.
- [BSB17] Noel Brett, Umair Siddique, and Borzoo Bonakdarpour. Rewriting-based runtime verification for alternation-free hyperltl. In Axel Legay and Tiziana Margaria, editors, *Tools and Algorithms*

- for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Part II, volume 10206 of *Lecture Notes in Computer Science*, pages 77–93, 2017. doi:10.1007/978-3-662-54580-5\_5.
- [CFK<sup>+</sup>14] Michael R. Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K. Micinski, Markus N. Rabe, and César Sánchez. Temporal logics for hyperproperties. In Martín Abadi and Steve Kremer, editors, *POST 2014*, volume 8414 of *LNCS*, pages 265–284. Springer, 2014. URL: [http://dx.doi.org/10.1007/978-3-642-54792-8\\_15](http://dx.doi.org/10.1007/978-3-642-54792-8_15), doi:10.1007/978-3-642-54792-8\_15.
- [CLP<sup>+</sup>17] Moritz Contag, Guo Li, Andre Pawlowski, Felix Domke, Kirill Levchenko, Thorsten Holz, and Stefan Savage. How they did it: An analysis of emission defeat devices in modern automobiles. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 231–250. IEEE Computer Society, 2017. doi:10.1109/SP.2017.66.
- [CS08] Michael R. Clarkson and Fred B. Schneider. Hyperproperties. In *CSF'08*, pages 51–65, 2008. URL: <http://dx.doi.org/10.1109/CSF.2008.7>, doi:10.1109/CSF.2008.7.
- [DBB<sup>+</sup>17] Pedro R. D'Argenio, Gilles Barthe, Sebastian Biewer, Bernd Finkbeiner, and Holger Hermanns. Is your software on dope? - formal analysis of surreptitiously "enhanced" programs. In Hongseok Yang, editor, *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10201 of *Lecture Notes in Computer Science*, pages 83–110. Springer, 2017. doi:10.1007/978-3-662-54434-1\_4.
- [DGM<sup>+</sup>20] Rayna Dimitrova, Maciej Gazda, Mohammad Reza Mousavi, Sebastian Biewer, and Holger Hermanns. Conformance-based doping detection for cyber-physical systems. In Alexey Gotsman and Ana Sokolova, editors, *Formal Techniques for Distributed Objects, Components, and Systems - 40th IFIP WG 6.1 International Conference, FORTE 2020, Held as Part of the 15th International Federated Conference on Distributed Computing Techniques, DisCoTec 2020, Valletta, Malta, June 15-19, 2020, Proceedings*, volume 12136 of *Lecture Notes in Computer Science*, pages 59–77. Springer, 2020. doi:10.1007/978-3-030-50086-3\_4.
- [DH84] Rocco De Nicola and Matthew Hennessy. Testing equivalences for processes. *Theor. Comput. Sci.*, 34:83–133, 1984. doi:10.1016/0304-3975(84)90113-0.
- [DL09] Stéphane Demri and Ranko Lazic. LTL with the freeze quantifier and register automata. *ACM Trans. Comput. Log.*, 10(3):16:1–16:30, 2009. doi:10.1145/1507244.1507246.
- [DMP17] Jyotirmoy V. Deshmukh, Rupak Majumdar, and Vinayak S. Prabhu. Quantifying conformance using the skorokhod metric. *Formal Methods in System Design*, 50(2-3):168–206, 2017.
- [Don10] Alexandre Donzé. Breach, A toolbox for verification and parameter synthesis of hybrid systems. In Tayssir Touili, Byron Cook, and Paul B. Jackson, editors, *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, volume 6174 of *Lecture Notes in Computer Science*, pages 167–170. Springer, 2010. doi:10.1007/978-3-642-14295-6\_17.
- [FH16] Bernd Finkbeiner and Christopher Hahn. Deciding hyperproperties. In Josée Desharnais and Radha Jagadeesan, editors, *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, volume 59 of *LIPICs*, pages 13:1–13:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.CONCUR.2016.13.
- [FHS17] Bernd Finkbeiner, Christopher Hahn, and Marvin Stenger. EAHyper: Satisfiability, implication, and equivalence checking of hyperproperties. In Rupak Majumdar and Viktor Kuncak, editors, *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part II*, volume 10427 of *Lecture Notes in Computer Science*, pages 564–570. Springer, 2017. doi:10.1007/978-3-319-63390-9\_29.
- [FHST17] Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. Monitoring hyperproperties. In Shuvendu K. Lahiri and Giles Regeer, editors, *Runtime Verification - 17th International Conference, RV 2017, Seattle, WA, USA, September 13-16, 2017, Proceedings*, volume 10548 of *Lecture Notes in Computer Science*, pages 190–207. Springer, 2017. doi:10.1007/978-3-319-67531-2\_12.
- [FHST18] Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. RVHyper: A runtime verification tool for temporal hyperproperties. In Dirk Beyer and Marieke Huisman, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 24th International*

- Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part II*, volume 10806 of *Lecture Notes in Computer Science*, pages 194–200. Springer, 2018. doi:10.1007/978-3-319-89963-3\_11.
- [FHST19] Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. Monitoring hyperproperties. *Formal Methods Syst. Des.*, 54(3):336–363, 2019. doi:10.1007/s10703-019-00334-z.
- [FP09] Georgios E. Fainekos and George J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theor. Comput. Sci.*, 410(42):4262–4291, 2009. doi:10.1016/j.tcs.2009.06.021.
- [FRS15] Bernd Finkbeiner, Markus N. Rabe, and César Sánchez. Algorithms for model checking HyperLTL and HyperCTL\*. In Daniel Kroening and Corina S. Pasareanu, editors, *CAV 2015*, volume 9206 of *LNCS*, pages 30–48. Springer, 2015. URL: [http://dx.doi.org/10.1007/978-3-319-21690-4\\_3](http://dx.doi.org/10.1007/978-3-319-21690-4_3), doi:10.1007/978-3-319-21690-4\_3.
- [GJP08] Antoine Girard, A. Agung Julius, and George J. Pappas. Approximate simulation relations for hybrid systems. *Discrete Event Dynamic Systems*, 18(2):163–179, 2008. doi:10.1007/s10626-007-0029-9.
- [GM20] Maciej Gazda and Mohammad Reza Mousavi. Logical characterisation of hybrid conformance. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPICs*, pages 130:1–130:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ICALP.2020.130.
- [GP11] Antoine Girard and George J. Pappas. Approximate bisimulation: A bridge between computer science and control theory. *Eur. J. Control*, 17(5-6):568–578, 2011. doi:10.3166/ejc.17.568-578.
- [HBDK18] Holger Hermanns, Sebastian Biewer, Pedro R. D’Argenio, and Maximilian A. Köhl. Verification, testing, and runtime monitoring of automotive exhaust emissions. In Gilles Barthe, Geoff Sutcliffe, and Margus Veanes, editors, *LPAR-22. 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Awassa, Ethiopia, 16-21 November 2018*, volume 57 of *EPiC Series in Computing*, pages 1–17. EasyChair, 2018. URL: <http://www.easychair.org/publications/paper/xgqz>.
- [Hen64] F. C. Hennie. Fault detecting experiments for sequential circuits. In *5th Annual Symposium on Switching Circuit Theory and Logical Design, Princeton, New Jersey, USA, November 11-13, 1964*, pages 95–110. IEEE Computer Society, 1964. doi:10.1109/SWCT.1964.8.
- [HHB15] Thorsten Hapke, Peter Hornung, and Joachim Becker. Schummeln auch in Europa. ARD / Norddeutscher Rundfunk, <https://www.tagesschau.de/wirtschaft/vw-schummelsoftware-101.html>, 2015. Online; accessed: 2019-04-19.
- [HST19] Christopher Hahn, Marvin Stenger, and Leander Tentrup. Constraint-based monitoring of hyperproperties. In Tomás Vojnar and Lijun Zhang, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part II*, volume 11428 of *Lecture Notes in Computer Science*, pages 115–131. Springer, 2019. doi:10.1007/978-3-030-17465-1\_7.
- [HZJ19] Hsi-Ming Ho, Ruoyu Zhou, and Timothy M. Jones. On verifying timed hyperproperties. In Johann Gamper, Sophie Pinchinat, and Guido Sciavicco, editors, *26th International Symposium on Temporal Representation and Reasoning, TIME 2019, October 16-19, 2019, Málaga, Spain*, volume 147 of *LIPICs*, pages 20:1–20:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.TIME.2019.20.
- [KHB18] Maximilian A. Köhl, Holger Hermanns, and Sebastian Biewer. Efficient monitoring of real driving emissions. In Christian Colombo and Martin Leucker, editors, *Runtime Verification - 18th International Conference, RV 2018, Limassol, Cyprus, November 10-13, 2018, Proceedings*, volume 11237 of *Lecture Notes in Computer Science*, pages 299–315. Springer, 2018. doi:10.1007/978-3-030-03769-7\_17.
- [KM15] Narges Khakpour and Mohammad Reza Mousavi. Notions of conformance testing for cyber-physical systems: Overview and roadmap (invited paper). In Luca Aceto and David de Frutos-Escrig, editors, *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid*,

- Spain, September 1-4, 2015*, volume 42 of *LIPICs*, pages 18–40. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. doi:10.4230/LIPICs.CONCUR.2015.18.
- [LY96] D. Lee and M. Yannakakis. Principles and methods of testing finite-state machines - a survey. *Proceedings of the IEEE*, 84(8):1089–1123, 1996. doi:10.1109/5.533956.
- [MN04] Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In Yassine Lakhnech and Sergio Yovine, editors, *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems, Joint International Conferences on Formal Modelling and Analysis of Timed Systems, FORMATS 2004 and Formal Techniques in Real-Time and Fault-Tolerant Systems, FTRTFT 2004, Grenoble, France, September 22-24, 2004, Proceedings*, volume 3253 of *Lecture Notes in Computer Science*, pages 152–166. Springer, 2004. doi:10.1007/978-3-540-30206-3\_12.
- [NKJ<sup>+</sup>17] Luan Viet Nguyen, James Kapinski, Xiaoqing Jin, Jyotirmoy V. Deshmukh, and Taylor T. Johnson. Hyperproperties of real-valued signals. In Jean-Pierre Talpin, Patricia Derler, and Klaus Schneider, editors, *Proceedings of the 15th ACM-IEEE International Conference on Formal Methods and Models for System Design, MEMOCODE 2017, Vienna, Austria, September 29 - October 02, 2017*, pages 104–113. ACM, 2017. doi:10.1145/3127041.3127058.
- [Pnu77] Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977. doi:10.1109/SFCS.1977.32.
- [STB<sup>+</sup>19] David Safránek, Matej Troják, Vojtech Bruza, Tomáš Vejpustek, Jan Papoušek, Martin Demko, Samuel Pastva, Ales Pejznoch, and Lubos Brim. Barbaric robustness monitoring revisited for stl\* in parasim. In Luca Bortolussi and Guido Sanguinetti, editors, *Computational Methods in Systems Biology - 17th International Conference, CMSB 2019, Trieste, Italy, September 18-20, 2019, Proceedings*, volume 11773 of *Lecture Notes in Computer Science*, pages 356–359. Springer, 2019. doi:10.1007/978-3-030-31304-3\_26.
- [Tre92] Jan Tretmans. *A formal Approach to conformance testing*. PhD thesis, University of Twente, The Netherlands, 1992.
- [Tre96] Jan Tretmans. Conformance testing with labelled transition systems: Implementation relations and test generation. *Computer Networks and ISDN Systems*, 29(1):49–79, 1996. doi:10.1016/S0169-7552(96)00017-7.
- [Uni13] United Nations. UN Vehicle Regulations - 1958 Agreement, Revision 2, Addendum 100, Regulation No. 101, Revision 3 — E/ECE/324/Rev.2/Add.100/Rev.3, 2013. URL: <http://www.unece.org/trans/main/wp29/wp29regs101-120.html>.
- [vO06] Michiel van Osch. Hybrid input-output conformance and test generation. In Klaus Havelund, Manuel Núñez, Grigore Rosu, and Burkhart Wolff, editors, *Formal Approaches to Software Testing and Runtime Verification, First Combined International Workshops, FATES 2006 and RV 2006, Seattle, WA, USA, August 15-16, 2006, Revised Selected Papers*, volume 4262 of *Lecture Notes in Computer Science*, pages 70–84. Springer, 2006. doi:10.1007/11940197\_5.