

# A Temporal Logic for Asynchronous Hyperproperties

Jan Baumeister<sup>1</sup>, Norine Coenen<sup>1</sup>,  
Borzoo Bonakdarpour<sup>2</sup>, Bernd Finkbeiner<sup>1</sup>, and César Sánchez<sup>3</sup>

<sup>1</sup> CISA Helmholtz Center for Information Security, Germany

<sup>2</sup> Michigan State University, USA

<sup>3</sup> IMDEA Software Institute, Spain

**Abstract.** *Hyperproperties* are properties of computational systems that require more than one trace to evaluate, e.g., many information-flow security and concurrency requirements. Where a trace property defines a set of traces, a hyperproperty defines a set of sets of traces. The temporal logics HyperLTL and HyperCTL\* have been proposed to express hyperproperties. However, their semantics are *synchronous* in the sense that all traces proceed at the same speed and are evaluated at the same position. This precludes the use of these logics to analyze systems whose traces can proceed at different speeds and allow that different traces take stuttering steps independently. To solve this problem in this paper, we propose an *asynchronous* variant of HyperLTL. On the negative side, we show that the model-checking problem for this variant is undecidable. On the positive side, we identify a decidable fragment which covers a rich set of formulas with practical applications. We also propose two model-checking algorithms that reduce our problem to the HyperLTL model-checking problem in the synchronous semantics.

## 1 Introduction

Hyperproperties [7] extend the conventional notion of trace properties [1] from a set of traces to a set of sets of traces. In other words, a hyperproperty stipulates a system property and not the property of just individual traces. Many interesting requirements in computing systems are hyperproperties and cannot be expressed by trace properties. Examples include (1) a wide range of information-flow security policies such as *noninterference* [13] and *observational determinism* [27], (2) sensitivity and robustness requirements in cyber-physical systems [26], and (3) consistency conditions such as *linearizability* in concurrent data structures [4].

HyperLTL [6] is a temporal logic for hyperproperties that enriches LTL with quantifiers allowing explicit and simultaneous quantification over multiple execution traces. For example, the observational determinism security policy [27] stipulates that any two executions that start in two *low-equivalent* states (i.e., states whose value of publicly observable variables are the same), should remain in low-equivalent states. This property can be expressed in HyperLTL as the following formula, called  $\varphi_{\text{OD}}$ ,  $\forall\pi.\forall\pi'.(l_\pi \leftrightarrow l_{\pi'}) \rightarrow \Box(l_\pi \leftrightarrow l_{\pi'})$ . However,

```

 $\ell_1$  int 1 = 0;
 $\ell_2$ 
 $\ell_3$  if (h = 0)
 $\ell_4$      1 := 1 + 1;
 $\ell_5$  else
 $\ell_6$      1 := 1;


```

Fig. 1: Program  $P_1$ 

```

 $\ell_1$  int 1 = 0;
 $\ell_2$ 
 $\ell_3$  if (h = 0)
 $\ell_4$      reg := 1 + 1;
 $\ell_5$      1 := reg;
 $\ell_6$  else
 $\ell_7$      1 := 1;

```

Fig. 2: Program  $P_2$ Fig. 3:  $\mathcal{K}$  with a self-loop

the semantics of HyperLTL (and other formal languages for hyperproperties) is *synchronous*, meaning that they completely abstract away the notion of time passage. In HyperLTL, all traces proceed at the same speed, as all temporal operators move the position on all traces simultaneously. Consider the program  $P_1$  in Fig 1, where input values 0 and 1 are possible for *high-secret* variable  $h$ . This renders two possible traces shown in Fig. 4a that satisfy  $\varphi_{\text{OD}}$ .

The synchronous semantics of HyperLTL has a shortcoming which has practical implications as well: formulas are not invariant under *stuttering*. Note that, contrary to LTL, disallowing the use of  $\bigcirc$  does not make the formula invariant under stuttering different, as traces can still stutter independently. This limits the scope of application of HyperLTL to only those settings where different traces can be perfectly aligned. For example, consider program  $P_2$  in Fig. 2, where line  $\ell_4$  in  $P_1$  is refined to its intermediate code using a register that stores the value  $1 + 1$  and then stores this value in memory location 1 in lines  $\ell_4$  and  $\ell_5$ , respectively. Applying the synchronous semantics of HyperLTL results in declaring a violation of  $\varphi_{\text{OD}}$  in the second position. This, however, is not an accurate interpretation of  $\varphi_{\text{OD}}$  (assuming that an attacker only has access to the memory footprint and not the CPU registers or a timing channel), as the two traces are stutter equivalent with respect to the state of variable 1. In fact, the synchronous semantics of HyperLTL may incorrectly identify good programs as bad because it ignores the notion of relative time between traces. This problem is generally amplified in Kripke structures where self-loops correspond to non-deterministic choices that model that the system may remain in a state for some arbitrary time. For instance, consider  $\mathcal{K}$  in Fig. 3 and HyperLTL formula  $\forall\pi.\forall\pi'.((b_\pi \leftrightarrow b_{\pi'}) \mathcal{U} \Box(a_\pi \leftrightarrow a_{\pi'}))$ . Only pairs of traces that take the self-loop the same number of times satisfy this formula. However, since the goal of employing a self-loop is typically to make the duration of staying in a state irrelevant, this semantics is too restrictive.

Besides HyperLTL, other logics have been proposed that allow trace quantification, for example,  $H_\mu$  [14], which adds trace quantifiers to the polyadic  $\mu$ -calculus [2]. For  $H_\mu$ , the model-checking problem is in general undecidable, but two fragments, the  $k$ -synchronous,  $k$ -context bounded fragments, have been identified for which model checking remains decidable [14]. It is not known, however, if any of the commonly used hyperproperties, like observational determinism, noninterference, or linearizability, can be encoded in these fragments.

In this paper, we propose an asynchronous temporal logic for hyperproperties. Our main motivation is to be able to reason about execution traces according




Fig. 4: Synchronous vs. asynchronous semantics for HyperLTL.

to the relative order of the sequences of actions in each trace but not about the duration of each action. We call this temporal logic *Asynchronous HyperLTL* or in short, A-HLTL. The key addition is the notion of *trajectory* that controls the relative speed at which traces progress by choosing at each instant which traces move and which traces stutter. For example, the trajectory shown in Fig. 4c for the two traces of the program in Fig. 2 allows the lower trace to stutter in the first position while the upper trace advances. On the contrary, in the third position, the upper trace stutters while the lower trace moves from the second to the third position. This trajectory enables identification of stutter equivalence of the two traces with respect to state variable  $1$  and, hence, successful verification of observational determinism. In order to reflect the notion of trajectories in our logic, we lift the syntax of HyperLTL by allowing a trajectory modality. This way, the corresponding formula for observational determinism in A-HLTL is the following:

$$\varphi_{OD} \stackrel{\text{def}}{=} \forall \pi. \forall \pi'. E.(li_\pi \leftrightarrow li_{\pi'}) \rightarrow \square(lo_\pi \leftrightarrow lo_{\pi'})$$

where  $E$  denotes the *existence* of a trajectory for temporal operator  $\square$ . The A-HLTL formula for the Kripke structure in Fig. 3 is  $\forall \pi. \forall \pi'. E.((b_\pi \leftrightarrow b_{\pi'}) \mathcal{U} \square(a_\pi \leftrightarrow a_{\pi'}))$ . A-HLTL allows us to reason about relational properties between two different systems that differ on timing, like for example, translation validation [21], which relates executions of the target code with the source code with respect to a (trace or hyper) property.

We show an encoding of the PCP problem into model-checking a formula of the shape  $\forall \pi. \forall \pi'. E.(\square \psi_1(\pi, \pi') \wedge \diamond \psi_2(\pi, \pi'))$ , which implies that model-checking A-HLTL is undecidable, even for the universal fragment. On the positive side, we show two decidable fragments of A-HLTL. The first algorithm is based on a *stuttering construction* in which we modify the Kripke structure to accept all stuttering expansions of the original paths. This algorithm can handle fragment  $\forall \pi_1 \dots \pi_n. E.\psi$ , where the  $\psi$  is a *phase formula*, a class of safety formulas that appear in many hyperproperties and are the building block of expressing trace equivalence. Our second algorithm uses an *acceleration construction* to convert a finite sequence of transitions that do not change phase, into a single transition. This algorithm is able to handle formulas with arbitrary quantification but a simpler kind of phase formulas. A-HLTL is, thus, the first logic for hyperproperties that can express the major asynchronous hyperproperties of interest within

decidable fragments. Moreover, A-HLTL is the first logic for asynchronous hyperproperties with a practical model checking algorithm.

We have evaluated the stuttering construction on two sets of cases studies: a range of compiler optimizations and an SPI bus protocol. In both case studies, we were able to prove system correctness using our reduction from A-HLTL to synchronous HyperLTL.

*Organization.* The rest of the paper is structured as follows. Section 2 contains the preliminaries, and Section 3 introduces A-HLTL and presents examples of properties expressible in A-HLTL. Section 4 describes the decidable fragments and present procedures for the model-checking problem. Section 5 shows that the model-checking problem for general A-HLTL formulas is undecidable and present the lower-bound complexity. Experimental results are presented in Section 6. Finally, Section 7 discusses the related work, while Section 8 concludes. Detailed proofs appear in the appendix.

## 2 Preliminaries

Let AP be a set of *atomic propositions* and  $\Sigma = 2^{\text{AP}}$  be the *alphabet*, where we call each element of  $\Sigma$  a *letter*. A *trace* is an infinite sequence  $\sigma = a_0 a_1 \dots$  of letters from  $\Sigma$ . We denote the set of all infinite traces by  $\Sigma^\omega$ . We use  $\sigma(i)$  for  $a_i$  and  $\sigma^i$  for the suffix  $a_i a_{i+1} \dots$ . A *pointed trace* is a pair  $(\sigma, p)$ , where  $p \in \mathbb{N}_0$  is a natural number (called the *pointer*). Pointed traces allow to traverse a trace by moving the pointer. Given a pointed trace  $(\sigma, p)$  and  $n > 0$ , we use  $(\sigma, p) + n$  as a short for  $(\sigma, p + n)$ . We denote the set of all pointed traces by  $\text{PTR} = \{(\sigma, p) \mid \sigma \in \Sigma^\omega \text{ and } p \in \mathbb{N}_0\}$ .

Two pointed traces  $(\sigma, p)$  and  $(\sigma', p')$  are *stuttering equivalent* if there are two infinite sequences of indices  $p = i_0 < i_1 \dots$  and  $p' = j_0 < j_1 \dots$  such that for all  $k \geq 0$  and for all  $l \in [i_k, i_{k+1})$  and  $l' \in [j_k, j_{k+1})$ ,  $\sigma(l) = \sigma'(l')$ . A pointed trace  $(\sigma', p')$  is a *stuttering expansion* of  $(\sigma, p)$  if there is a sequence  $p' = j_0 < j_1 < \dots$  such that for all  $k \geq 0$  and for all  $l \in [j_k, j_{k+1})$ ,  $\sigma(p + k) = \sigma'(l)$ . We say that  $\sigma$  is stuttering equivalent to  $\sigma'$  if  $(\sigma, 0)$  is stuttering equivalent to  $(\sigma', 0)$ , and that  $\sigma'$  is a stuttering expansion of  $\sigma$  if  $(\sigma', 0)$  is a stuttering expansion of  $(\sigma, 0)$ .

A *Kripke structure* is a tuple  $\mathcal{K} = \langle S, S_{\text{init}}, \delta, L \rangle$ , where  $S$  is a set of states,  $S_{\text{init}} \subseteq S$  is the set of initial states,  $\delta \subseteq S \times S$  is a transition relation, and  $L : S \rightarrow \Sigma$  is a labeling function on the states of  $\mathcal{K}$ . We require that for each  $s \in S$ , there exists  $s' \in S$ , such that  $(s, s') \in \delta$ .

A *path* of a Kripke structure is an infinite sequence of states  $s(0)s(1)\dots \in S^\omega$ , such that  $s(0) \in S_{\text{init}}$  and  $(s(i), s(i+1)) \in \delta$ , for all  $i \geq 0$ . A *trace* of a Kripke structure is a trace  $\sigma(0)\sigma(1)\sigma(2)\dots \in \Sigma^\omega$ , such that there exists a path  $s(0)s(1)\dots \in S^\omega$  with  $\sigma(i) = L(s(i))$  for all  $i \geq 0$ . Abusing notation we use  $\sigma = L(\rho)$  to denote that  $\sigma$  is the trace corresponding to path  $\rho$ . We denote by  $\text{Traces}(\mathcal{K}, s)$  the set of all traces of  $\mathcal{K}$  with paths that start in state  $s \in S$ . We denote by  $\text{Traces}(\mathcal{K}, A)$  the set of all traces that start from some state in  $A \subseteq S$  and  $\text{Traces}(\mathcal{K})$  as a short for  $\text{Traces}(\mathcal{K}, S_{\text{init}})$ .

**HyperLTL.** HyperLTL [6] is a temporal logic that extends LTL [18, 20] for hyperproperties, which allows reasoning about multiple execution traces simultaneously. The syntax of HyperLTL is:

$$\begin{aligned} \varphi &::= \exists\pi.\varphi \mid \forall\pi.\varphi \mid \psi \\ \psi &::= a_\pi \mid \psi \vee \psi \mid \neg\psi \mid \bigcirc\psi \mid \psi \mathcal{U} \psi \end{aligned}$$

where  $\pi$  is a *trace variable* from an infinite supply of trace variables. The intended meaning of  $a_\pi$  is that proposition  $a \in \Sigma$  holds in the current time in trace  $\pi$ . Trace quantifiers  $\exists\pi$  and  $\forall\pi$  allow reasoning simultaneously about different traces of the computation. Atomic predicates  $a_\pi$  refer to a single trace  $\pi$ . Given a HyperLTL formula  $\varphi$ , we use  $\text{Vars}(\varphi)$  for the set of trace variables quantified in  $\varphi$ . A formula  $\varphi$  is well-formed if for all atoms  $a_\pi$  in  $\varphi$ ,  $\pi$  is quantified in  $\varphi$  (i.e.,  $\pi \in \text{Vars}(\varphi)$ ) and if no trace variable is quantified twice in  $\varphi$ . Given a set of traces  $T$ , the semantics of a HyperLTL formula  $\varphi$  is defined in terms of trace assignments, which is a (partial) map from trace variables to indexed traces  $\Pi : \text{Vars}(\varphi) \rightarrow \text{PTR}$ . The trace assignment with empty domain is denoted by  $\Pi_\emptyset$ . We use  $\text{Dom}(\Pi)$  for the subset of  $\text{Vars}(\varphi)$  for which  $\Pi$  is defined. Given a trace assignment  $\Pi$ , a trace variable  $\pi$ , a trace  $\sigma$  and a pointer  $p$ , we denote by  $\Pi[\pi \mapsto (\sigma, p)]$  the assignment that coincides with  $\Pi$  for every trace variable except for  $\pi$ , which is mapped to  $(\sigma, p)$ . Also, we use  $\Pi + n$  to denote the trace assignment  $\Pi'$  such that  $\Pi'(\pi) = \Pi(\pi) + n$  for all  $\pi \in \text{Dom}(\Pi) = \text{Dom}(\Pi')$ . The semantics of HyperLTL is:

$$\begin{aligned} \Pi \models_T \exists\pi.\varphi &\quad \text{iff} \quad \text{for some } \sigma \in T, \Pi[\pi \mapsto (\sigma, 0)] \models_T \varphi \\ \Pi \models_T \forall\pi.\varphi &\quad \text{iff} \quad \text{for all } \sigma \in T, \Pi[\pi \mapsto (\sigma, 0)] \models_T \varphi \\ \Pi \models_T \psi &\quad \text{iff} \quad \Pi \models \psi \\ \Pi \models a_\pi &\quad \text{iff} \quad a \in \sigma(p), \text{ where } (\sigma, p) = \Pi(\pi) \\ \Pi \models_T \psi_1 \vee \psi_2 &\quad \text{iff} \quad \Pi \models_T \psi_1 \text{ or } \Pi \models_T \psi_2 \\ \Pi \models \neg\psi &\quad \text{iff} \quad \Pi \not\models \psi \\ \Pi \models \bigcirc\psi &\quad \text{iff} \quad (\Pi + 1) \models \psi \\ \Pi \models \psi_1 \mathcal{U} \psi_2 &\quad \text{iff} \quad \text{for some } j \geq 0 \ (\Pi + j) \models \psi_2 \\ &\quad \text{and for all } 0 \leq i < j, (\Pi + i) \models \psi_1 \end{aligned}$$

Note that quantifiers assign traces to trace variables and set the pointer to the initial position 0. We say that a set of traces  $T$  is a model of a HyperLTL formula  $\varphi$ , denoted  $T \models \varphi$  whenever  $\Pi_\emptyset \models_T \varphi$ . A Kripke structure  $\mathcal{K}$  is a model of a HyperLTL formula  $\varphi$ , denoted by  $\mathcal{K} \models \varphi$ , whenever  $\text{Traces}(\mathcal{K}) \models \varphi$ .

### 3 Asynchronous HyperLTL

We introduce a temporal logic A-HLTL as an extension of HyperLTL to express asynchronous hyperproperties.

**Trajectories.** To model the asynchronous passage of time, we now introduce the notion of a trajectory, which chooses when traces move and when they stutter. Let  $\mathcal{V}$  be a set of trace variables and let  $I \subseteq \mathcal{V}$ . The  $I$ -successor of a trace assignment  $\Pi$ , denoted by  $\Pi + I$ , is the trace assignment  $\Pi'$  such that  $\Pi'(\pi) = \Pi(\pi) + 1$  if  $\pi \in I$  and  $\Pi'(\pi) = \Pi(\pi)$  otherwise. That is, the pointers of indices in  $I$  advance by one step, while the others remain the same. A *trajectory*  $t : t(0)t(1)t(2)\cdots$  for a formula  $\varphi$  is an infinite sequence of non-empty subsets of  $\text{Vars}(\varphi)$ . Essentially, in each step of the trajectory one or more of the traces make progress. A trajectory is fair for a trace variable  $\pi \in \text{Vars}(\varphi)$  if there are infinitely many positions  $j$  such that  $\pi \in t(j)$ . A trajectory is fair if it is fair for all trace variables in  $\text{Vars}(\varphi)$ . Given a trajectory  $t$ , by  $t^i$ , we mean the suffix  $t(i)t(i+1)\cdots$ . Furthermore, for a set of trace variables  $\mathcal{V}$ , we use  $\text{TRJ}_{\mathcal{V}}$  for set of all trajectories for indices from  $\mathcal{V}$ .

### 3.1 Syntax and Semantics of Asynchronous HyperLTL

The syntax of Asynchronous HyperLTL is:

$$\begin{aligned} \varphi &::= \exists\pi.\varphi \mid \forall\pi.\varphi \mid \mathbf{E}\psi \mid \mathbf{A}\psi \\ \psi &::= a_{\pi} \quad \neg\psi \quad \psi_1 \vee \psi_2 \mid \psi_1 \mathbf{U} \psi_2 \mid \bigcirc\psi \end{aligned}$$

where  $a \in \text{AP}$ ,  $\pi$  is a trace variable from an infinite supply  $\mathcal{V}$  of trace variables,  $\mathbf{E}$  is the existential trajectory modality and  $\mathbf{A}$  is the universal trajectory modality. The intended meaning of  $\mathbf{E}$  is that there is a trajectory that gives an interpretation of the relative passage of time between the traces for which the temporal formula that relates the traces is satisfied. Dually,  $\mathbf{A}$  means that for all trajectories, the resulting alignment makes the inner formula true. It is important to note that there is no nesting of trajectory modalities and that all temporal operators in a formula are interpreted with respect to a single modality.

We use the usual syntactic sugar for Boolean operators  $\text{true} \stackrel{\text{def}}{=} a_{\pi} \vee \neg a_{\pi}$ ,  $\text{false} \stackrel{\text{def}}{=} \neg \text{true}$ ,  $\varphi_1 \wedge \varphi_2 \stackrel{\text{def}}{=} \neg(\neg\varphi_1 \vee \neg\varphi_2)$ , and the syntactic sugar for temporal operators  $\diamond\varphi \stackrel{\text{def}}{=} \text{true} \mathbf{U} \varphi$ ,  $\varphi_1 \rightarrow \varphi_2 \stackrel{\text{def}}{=} \neg\varphi_1 \vee \varphi_2$ , and  $\square\varphi \stackrel{\text{def}}{=} \neg\diamond\neg\varphi$ , etc.

As before, we use trace assignments for the semantics of A-HLTL. Given  $(\Pi, t)$  where  $\Pi$  is a trace assignment and  $t$  a trajectory, we use  $(\Pi, t) + 1$  for the successor of  $(\Pi, t)$  defined as  $(\Pi', t')$  where  $t' = t^1$ , and  $\Pi'(\pi) = \Pi(\pi) + 1$  if  $\pi \in t(0)$  and  $\Pi'(\pi) = \Pi(\pi)$  otherwise. We use  $(\Pi, t) + k$  as the  $k$ -th successor of  $(\Pi, t)$ .

The satisfaction of an asynchronous HyperLTL formula  $\varphi$  over a trace assignment  $\Pi$  and a set of traces  $T$ , denoted by  $\Pi \models_T \varphi$  is defined as follows:

$\Pi \models_T \exists \pi. \varphi$	iff	for some $\sigma \in T : \Pi[\pi \mapsto (\sigma, 0)] \models_T \varphi$
$\Pi \models_T \forall \pi. \varphi$	iff	for all $\sigma \in T : \Pi[\pi \mapsto (\sigma, 0)] \models_T \varphi$
$\Pi \models_T \mathbf{E}\psi$	iff	for some $t \in \text{TRJ}_{\text{Dom}(\Pi)}. (\Pi, t) \models \psi$
$\Pi \models_T \mathbf{A}\psi$	iff	for all $t \in \text{TRJ}_{\text{Dom}(\Pi)}. (\Pi, t) \models \psi$
$(\Pi, t) \models a_\pi$	iff	$a \in \Pi(\pi)(0)$
$(\Pi, t) \models \neg \psi$	iff	$(\Pi, t) \not\models \psi$
$(\Pi, t) \models \psi_1 \vee \psi_2$	iff	$(\Pi, t) \models \psi_1$ or $(\Pi, t) \models \psi_2$
$(\Pi, t) \models \bigcirc \psi$	iff	$(\Pi, t) + 1 \models \psi$
$(\Pi, t) \models \psi_1 \mathcal{U} \psi_2$	iff	for some $i \geq 0 : (\Pi, t) + i \models \psi_2$ and for all $j < i : (\Pi, t) + j \models \psi_1$

We say that a set  $T$  of traces satisfies a closed sentence  $\varphi$ , denoted by  $T \models \varphi$ , if  $\Pi_\emptyset \models_T \varphi$ . We say that a Kripke structure  $\mathcal{K}$  satisfies an A-HLTL formula  $\varphi$  (and write  $\mathcal{K} \models \varphi$ ) if and only if we have  $\text{Traces}(\mathcal{K}, S_{\text{init}}) \models \varphi$ .

### 3.2 Examples of A-HLTL

We illustrate the expressive power of A-HLTL by introducing the asynchronous version of well-known properties.

*Linearizability* [15] requires that any history of execution of a concurrent data structure (i.e., sequence of invocation and response by different threads) matches some sequential order of invocations and responses:

$$\varphi_{\text{LNZ}} \stackrel{\text{def}}{=} \forall \pi. \exists \pi'. \mathbf{E}. \square (\text{history}_\pi \leftrightarrow \text{history}_{\pi'})$$

where *history* denotes method invocations (and not the actual execution of the internal instructions of the concurrent library) by the different threads and the response observed, trace  $\pi$  ranges over the concurrent data structure and  $\pi'$  ranges over its sequential counterpart.

*Goguen and Meseguer's noninterference (GMNI)* [13] stipulates that, for all traces, the low-observable output must not change when all high inputs are removed:

$$\varphi_{\text{GMNI}} \stackrel{\text{def}}{=} \forall \pi. \exists \pi'. \mathbf{E}. (\square \lambda_{\pi'}) \wedge \square (lo_\pi \leftrightarrow lo_{\pi'})$$

where  $\lambda_{\pi'}$  expresses that all of the high inputs in the current state of  $\pi'$  have dummy value  $\lambda$ , and denotes low-observable output proposition.

*Not never terminates* [17] requires that for every initial state, there is a terminating trace and a non-terminating trace:

$$\varphi_{\text{NNT}} \stackrel{\text{def}}{=} \forall \pi. \exists \pi'. \exists \pi''. \mathbf{E}. (\pi[0] = \pi'[0] = \pi''[0]) \rightarrow (\diamond \text{term}_{\pi'} \wedge \square \neg \text{term}_{\pi''})$$

*Termination-insensitive noninterference* [24] requires that for two executions that start from a low-observable states, information leaks are permitted if they are transmitted purely by the program's termination behavior. That is, the program may diverge on some high inputs and terminate on others:

$$\varphi_{\text{TIN}} \stackrel{\text{def}}{=} \forall \pi. \forall \pi'. \text{E}. (l_\pi \leftrightarrow l_{\pi'}) \rightarrow \left( (\Box \neg \text{term}_\pi \vee \Box \neg \text{term}_{\pi'}) \vee \Diamond (\text{term}_\pi \wedge \text{term}_{\pi'} \wedge l_\pi \leftrightarrow l_{\pi'}) \right)$$

*Termination-sensitive noninterference* [3] Termination-sensitive noninterference is the same as termination insensitive, except that it forbids one trace to diverge and the other to terminate:

$$\varphi_{\text{TSN}} \stackrel{\text{def}}{=} \forall \pi. \forall \pi'. \text{E}. (l_\pi \leftrightarrow l_{\pi'}) \rightarrow \left( (\Box \neg \text{term}_\pi \wedge \Box \neg \text{term}_{\pi'}) \vee \Diamond (\text{term}_\pi \wedge \text{term}_{\pi'} \wedge l_\pi \leftrightarrow l_{\pi'}) \right)$$

## 4 Model-Checking A-HLTL

In this section, we show the decidability of the model-checking problem for two classes of A-HLTL formulas using two different algorithms:

- (1) a *stuttering* construction in which we modify the Kripke structure  $\mathcal{K}$  to accept all stuttering expansions of paths in  $\mathcal{K}$ ; and
- (2) an *acceleration* construction in which the modified Kripke structure accelerates jumping directly to the synchronization points.

In both cases the problem is reduced to model-checking HyperLTL formulas, which is known to be decidable [6, 11]. We describe each construction separately.

### 4.1 The Stuttering Construction

We consider first A-HLTL formulas of the form  $\forall \pi_1 \dots \pi_n. \text{E}. \psi$ . We will then extend our results to the  $\exists^*$  fragment, to handle the A trajectory modality and to a larger collection of predicates. The class of temporal formulas  $\psi$  that we handle are called *admissible* formulas, and are defined as the Boolean combination of:

1. any number of state formulas, which may relate propositions  $p_{\pi_i}$  of different traces arbitrarily;
2. any number temporal formulas (called *monadic temporal formulas*), each of which only uses one trace variable and is invariant under stuttering (guaranteed for example by forbidding the use of  $\bigcirc$ ), and
3. one *phase formula*, which is an invariant that can relate different traces in a restricted way (see below).

Given an admissible formula  $\psi$ , we use  $\psi_{ph}$  for its phase formula, and we use  $\psi[\psi_{ph} \triangleleft \xi]$  for the formula that results from  $\psi$  by replacing  $\psi_{ph}$  with  $\xi$ . Since  $\psi_{ph}$  occurs only once in  $\psi$ , we use the fact that  $\psi_{ph}$  appears with a single polarity.



We present here the construction for positive polarity which is the case in all practical formulas (the case for negative polarity is analogous).

The algorithm has two parts. First, we generate the *stuttering* Kripke structure  $\mathcal{K}^{st}$  whose paths are the stuttering expansions of paths in the original Kripke structure  $\mathcal{K}$ . Then, we modify the admissible formula  $\psi$  into  $\psi_{sync}$  such that  $\mathcal{K} \models \forall \pi_1 \dots \pi_n. E.\psi$  if and only if  $\mathcal{K}^{st} \models \forall \pi_1 \dots \pi_n. \psi_{sync}$ . We describe each of the concepts separately.

*Phase formulas.* We first define *atomic phase formulas* ( $\bigwedge_{p \in P} p_{\pi_i} \leftrightarrow p_{\pi_j}$ ) which are characterized by  $(\pi_i, \pi_j, P)$ , where  $P \subseteq \text{AP}$  and  $\pi_i$  and  $\pi_j$  are two different trace variables. We use *color* to refer to a valuation of the variables in  $P$ . Essentially, an atomic phase formula asserts that all propositions in  $P$  coincide in both traces at all points in time, that is, both traces exhibit the same sequence of colors. Since the passage of time proceeds at different speeds in the different traces—according to the trajectory—atomic phase formulas state the traces for  $\pi_i$  and  $\pi_j$  are sequences of phases of the same color, where corresponding phases may have different lengths. A phase formula is formed from atomic formulas as follows:

$$\Box \left( \bigwedge_{p \in P^1} p_{\pi_i^1} \leftrightarrow p_{\pi_j^1} \wedge \dots \wedge \bigwedge_{p \in P^k} p_{\pi_i^k} \leftrightarrow p_{\pi_j^k} \right)$$

We use  $\mathcal{P} : \{(\pi_i^1, \pi_j^1, P^1), \dots, (\pi_i^k, \pi_j^k, P^k)\}$  for the collection of predicates and trace variables that characterize a phase formula.

*Stuttering Kripke structure.* We start from  $\mathcal{K}$  and create  $\mathcal{K}^{st}$  that accepts the stuttering expansions of traces in  $\mathcal{K}$ . First, the alphabet of atomic propositions is enriched with a fresh proposition  $st$ , that is  $\text{AP}^{st} = \text{AP} \cup \{st\}$ , to encode whether the state represents a real move or a stuttering move. Given  $\mathcal{K} = \langle S, S_{init}, \delta, L \rangle$ , the stuttering Kripke structure is  $\mathcal{K}^{st} = \langle S^{st}, S_{init}^{st}, \delta^{st}, L^{st} \rangle$  where:

- $S^{st} = S \cup \{s^{st} \mid s \in S\}$  contains two copies of each state in  $S$ , where we use  $s^{st}$  to denote the stuttering state that corresponds to  $s$ ;
- $\delta^{st} = \delta \cup \{(s, s^{st})\} \cup \{(s^{st}, s^{st})\} \cup \{(s^{st}, s') \mid \text{for every } (s, s') \in \delta\}$ .
- $L^{st}(s) = L(s)$  for  $s \in S$ , and  $L^{st}(s^{st}) = L(s) \cup \{st\}$ .

The construction generates a Kripke structure which is linear in the size of the original Kripke structure. It is easy to see that every stuttering expansion of a path of  $\mathcal{K}$  has a corresponding path in  $\mathcal{K}^{st}$ , where the repeated version of state  $s$  is captured by state  $s^{st}$ . Conversely every path  $\rho'$  in  $\mathcal{K}^{st}$  whose trace satisfies  $\Box \diamond \neg st$  can be turned into its “stuttering compression” by removing all stuttering states, which is a path of  $\mathcal{K}$ . Note that the constraint  $\Box \diamond \neg st$  guarantees that there are infinitely many non-stuttering positions in  $\rho'$ , so  $\rho$  is well-defined. This construction makes a one-to-one correspondence between a trajectory and a tuple of traces of  $\mathcal{K}$ , with a corresponding tuple of traces of  $\mathcal{K}^{st}$ .

*State and monadic formulas are not affected by trajectories.* State formulas are relational formulas that are evaluated at the beginning of the computation. Temporal monadic formulas only refer to one trace variable and are stuttering invariant by definition. Therefore, none of these formulas are affected by the stuttering

induced by a trajectory, as the relative stuttering among traces does not affect their truth valuation. We first note that given a trace assigned for each of the trace variables in  $\text{Vars}(\varphi)$  the truth value of state formulas and monadic formulas does not depend on the trajectory chosen.

*Phase alignment of asynchronous sequences.* We use the stuttering in  $\mathcal{K}^{st}$  to encode the relative progress of traces as dictated by a trajectory. We will now introduce synchronous HyperLTL formulas to reason in  $\mathcal{K}^{st}$  about the corresponding states during the asynchronous evaluation in  $\mathcal{K}$ . The important concept is that of “phase changes”, which are the points in a trace  $\sigma$  at which the valuation of the predicates  $P$  in an atomic phase formula  $(\pi_i, \pi_j, P)$  change. Let  $\Pi$  be a trace assignment for traces in  $\mathcal{K}$  that maps  $\pi_i$  to a pointed trace  $(\sigma, l)$ . We say that in assignment  $\Pi$ , trace variable  $\pi_i$  is *about to change phase* with respect to  $(\pi_i, \pi_j, P)$  if for some  $p \in P$  either  $p \in \sigma(l)$  but  $p \notin \sigma(l+1)$  or  $p \notin \sigma(l)$  but  $p \in \sigma(l+1)$ . Note that in  $\mathcal{K}^{st}$  the next relevant letter (the one corresponding to  $\sigma(l+1)$  is the first letter that is not a stuttering letter). Formula  $change_P(\pi_i)$  captures that the next non-stuttering step of  $\pi_i$  is a phase change (with respect to predicates in  $P$  and therefore with respect to atomic phase formula  $\alpha$ ):

$$change_P(\pi_i) \stackrel{\text{def}}{=} \bigvee_{p \in P} p_{\pi_i} \not\leftrightarrow \bigcirc(st_{\pi_i} \mathcal{U} p_{\pi_i})$$

A phase change for  $\pi_i$  in atomic phase formula  $(\pi_i, \pi_j, P)$  implies that  $\pi_j$  must also proceed to change phase. The second observation is that when  $\pi_i$  and  $\pi_j$  are not changing phases, any choice that the trajectory makes will preserve the valuation of the atomic phase formula.

We now capture formally this intuition as formulas. Predicate  $move(\pi_i) \stackrel{\text{def}}{=} \bigcirc(\neg st_{\pi_i})$  indicates whether trace variable  $\pi_i$  will move (and not stutter) at a given instant of the computation. The following temporal formula captures the consistency criteria of phase changes as a synchronized decision for moving traces  $\pi_i$  and  $\pi_j$  related by an atomic phase formula  $(\pi_i, \pi_j, P)$ :

$$align_{(\pi_i, \pi_j, P)} \stackrel{\text{def}}{=} \left( \begin{array}{l} (move(\pi_i) \wedge move(\pi_j)) \rightarrow (change_P(\pi_i) \leftrightarrow change_P(\pi_j)) \wedge \\ (move(\pi_i) \wedge \neg move(\pi_j)) \rightarrow \neg change_P(\pi_i) \quad \quad \quad \wedge \\ (\neg move(\pi_i) \wedge move(\pi_j)) \rightarrow \neg change_P(\pi_j) \end{array} \right)$$

We will reduce the model-checking problem in A-HLTL to checking in  $\mathcal{K}^{st}$  that tuples of traces that align phase changes—for all atomic phase formulas—satisfy all sub-formulas of the specification  $\psi$ . The following two formulas express that all atomic phase formulas align, and that all traces are fair (all traces eventually move):

$$phase \stackrel{\text{def}}{=} \bigwedge_{(\pi_i, \pi_j, P) \in \mathcal{P}} align_{(\pi_i, \pi_j, P)} \quad \quad \quad fair \stackrel{\text{def}}{=} \bigwedge_{\pi_i \in \{\pi_1 \dots \pi_n\}} \square \diamond \neg st_i$$

We will then check in  $\mathcal{K}^{st}$  that all stuttering traces that align phases and are fair satisfy the desired formula  $\psi$ , that is  $(\Box phase \wedge fair) \rightarrow \psi$ . Note that all those tuples of traces that do not align phases are ruled out in the antecedent.

A final technical detail in the construction is that we must guarantee that for all tuples of paths of  $\mathcal{K}$  there are stuttering expansions that are fair and align phases, and that they have the same number of phases. Otherwise, there are paths of  $\mathcal{K}$  that cannot be aligned, which inevitably leads to a violation of  $\psi_{ph}$ . It could be the case that some tuple of traces of  $\mathcal{K}$  cannot possibly align the phase changes corresponding to all atomic phase formulas. This can happen in two cases: (1) when two traces have different number of phases, and (2) when there is a circular dependency between the atomic formulas that force the trajectory to synchronize the traces in incompatible orders. The first case is captured by:

$$missalign \stackrel{\text{def}}{=} \bigvee_{(\pi_i, \pi_j, P)} \left( \Box \neg change_P(\pi_i) \not\leftrightarrow \left( \Box \neg change_P(\pi_j) \right) \right)$$

The second case is captured by the following formula, where  $cycles(\psi_{ph})$  are the sequences of atomic formulas that form a simple cycle, that is  $[(\pi^0, \pi^1, P^0), (\pi^1, \pi^2, P^1) \dots (\pi^k, \pi^0, P^k)]$  such that the second trace variable is the first trace variable of the next atomic phase formula, circularly (see Ex. 1 below):

$$block \stackrel{\text{def}}{=} \bigvee_{C \in cycles(\psi_{ph})} \left( \bigwedge_{(\pi_i, \pi_j, P) \in C} change_P(\pi_i) \wedge \neg change_P(\pi_j) \right)$$

Essentially, *block* encodes whether the set of traces involved cannot proceed without violating *phase*, because *align* forbids all traces involved to move. Hence, the formula  $phase \mathcal{U} (missalign \vee block)$  captures to those traces of  $\mathcal{K}^{st}$  that contain an aligned prefix of computation that lead to a miss-alignment or a block. The proof of correctness shows that given a tuple of traces of  $\mathcal{K}$ , if there is a trajectory that aligns the phase changes (which must exist if there is a trajectory that makes  $\psi_{ph}$  true), then all trajectories that respect  $\Box phase$  will also align the phase changes (and also satisfy  $\psi_{ph}$ ).

We are finally ready to describe the synchronous phase formula  $\psi_{sync}$ . First, this formula is only evaluated against tuples of fair traces, which correspond to the stuttering extensions of paths of  $\mathcal{K}$ . Then, the phase formula  $\psi_{ph}$  is translated into a formula that captures (1) that following a phase alignment cannot lead to a block or to two traces changing phases a different number of times, and (2) that if phases are aligned then  $\psi_{ph}$  holds. Formally,

$$\psi_{sync} \stackrel{\text{def}}{=} fair \rightarrow \psi[\psi_{ph} \triangleleft \psi'], \quad \text{where } \psi' = \left( \neg(phase \mathcal{U} (missalign \vee block)) \wedge \left( \Box phase \rightarrow \psi_{ph} \right) \right)$$

*Example 1.* We illustrate the previous definitions with the Kripke structures  $\mathcal{K}_1$ ,

$\mathcal{K}_2$  and  $\mathcal{K}_3$  in Fig. 5 and their stuttering variants  $\mathcal{K}_1^{st}$ ,  $\mathcal{K}_2^{st}$  and  $\mathcal{K}_3^{st}$ . Consider formula  $\forall\pi_1.\forall\pi_2.E.\Box(a_{\pi_1} \leftrightarrow a_{\pi_2})$ . Consider the following trace assignments:

$\Pi^1(\pi_1) \mapsto \{ \} \{st\} \{a\} \dots$	$\Pi^2(\pi_1) \mapsto \{ \} \{a\} \{a\} \dots$	$\Pi^3(\pi_1) \mapsto \{a\} \{ \} \{ \} \dots$
$\Pi^1(\pi_2) \mapsto \{ \} \{ \} \{a\} \dots$	$\Pi^2(\pi_2) \mapsto \{ \} \{ \} \{a\} \dots$	$\Pi^3(\pi_2) \mapsto \{a\} \{ \} \{a\} \dots$

Consider the trace assignment  $\Pi^1$  on the left, where  $\pi_1$  is a trace of  $\mathcal{K}_1^{st}$  corresponding to the path of  $\mathcal{K}_1$  that visits  $s_1$ , and  $\pi_2$  corresponds to the path that visits  $s_2$ . This trace assignment aligns the atomic phase formula  $(\pi_1, \pi_2, \{a\})$  at all positions. In particular, at position 0, we have  $change_{\{a\}}(\pi_1)$ , but  $\neg change_{\{a\}}(\pi_2)$ , and  $\neg move(\pi_1)$  and  $move(\pi_2)$ , as  $align_{\{a\}}$  requires.

Consider now the trace assignment  $\Pi^2$  in the middle, where again  $\pi_1$  corresponds to the path in  $\mathcal{K}_1^{st}$  that visits  $s_1$  and  $\pi_2$  the path that visits  $s_2$ . In this case, we have  $\neg align_{\{a\}}$  at position 0 because  $change_{\{a\}}(\pi_1)$  and  $\neg change_{\{a\}}(\pi_2)$  hold, and both  $move(\pi_1)$  and  $move(\pi_2)$ . Consider now  $\Pi^3$  on the right, where  $\pi_1$  corresponds to the path of  $\mathcal{K}_2^{st}$  that visits  $s_3$  and  $\pi_2$  to the path of  $\mathcal{K}_2^{st}$  that visits  $s_4$ . In this case  $align_{\{a\}}$  holds at 0 and  $missalign$  holds at 1 because at 1,  $\Box\neg change_{\{a\}}(\pi_1)$  holds, but not  $\Box\neg change_{\{a\}}(\pi_2)$ . Therefore,  $phase\mathcal{U}$  ( $missalign \vee block$ ) holds for  $\Pi^3$ . Finally, consider  $\forall\pi_1.\forall\pi_2.\forall\pi_3.E.\Box(a_{\pi_1} \leftrightarrow a_{\pi_2} \wedge b_{\pi_2} \leftrightarrow b_{\pi_3} \wedge c_{\pi_3} \leftrightarrow c_{\pi_2})$  and the trace assignment  $\Pi$  of  $\mathcal{K}_3^{st}$  shown below on the

$\Pi(\pi_1) \mapsto \{ \} \{ \} \{a\} \{c\} \dots$
$\Pi(\pi_2) \mapsto \{ \} \{ \} \{b\} \{a\} \dots$
$\Pi(\pi_3) \mapsto \{ \} \{ \} \{c\} \{b\} \dots$

left. In this case  $phase$  holds at position 0 and  $block$  holds at position 1. This is because  $change_{\{a\}}(\pi_1)$  and  $\neg change_{\{a\}}(\pi_2)$ ,  $change_{\{b\}}(\pi_2)$  and  $\neg change_{\{b\}}(\pi_3)$ , and  $change_{\{c\}}(\pi_3)$  and also  $\neg change_{\{c\}}(\pi_1)$ . This illustrates that it will not

be possible to align all three atomic phase formulas.

We are now ready to state the main result of this section.

**Theorem 1.** *Let  $\mathcal{K}$  be a Kripke structure and  $\psi$  an admissible formula. Then,  $\mathcal{K} \models \forall\pi_1 \dots \pi_n.E.\psi$  if and only if  $\mathcal{K}^{st} \models \forall\pi_1 \dots \pi_n.\psi_{sync}$ .*

Dually, to show that the  $\exists^*$  fragment is decidable, we consider replacing  $\psi_{ph}$  by the formula

$$\psi_{esync} \stackrel{\text{def}}{=} fair \wedge \psi[\psi_{ph} \triangleleft (\Box phase \wedge \psi_{ph})]$$

**Theorem 2.** *Let  $\mathcal{K}$  be a Kripke structure and  $\psi$  an admissible formula. Then  $\mathcal{K} \models \exists\pi_1 \dots \pi_n.E.\psi$  if and only if  $\mathcal{K}^{st} \models \exists\pi_1 \dots \pi_n.\psi_{esync}$ .*




Fig. 5: Kripke structure  $\mathcal{K}_1$  (left),  $\mathcal{K}_2$  (middle) and  $\mathcal{K}_3$  (right).

The proof of Theorem 2 takes a witness tuple and trajectory in  $\mathcal{K}$  and shows that the induced tuple in  $\mathcal{K}^{st}$  is *fair*, satisfies  $\Box phase$  and that the valuation of  $\psi_{ph}$  is preserved. Similarly, as before, tuples of traces of  $\mathcal{K}^{st}$  that are fair and follow phase alignments induce a trajectory on their stuttering compression that also preserve  $\psi_{ph}$ .

**Corollary 1.** *The problems of model-checking  $\forall^*$  admissible A-HLTL formulas and  $\exists^*$  admissible A-HLTL formulas is decidable.*

We finally consider the negation of phase formulas, called *co-phase formulas*, which are formulas of the form  $\Diamond \neg R$  where  $R$  a conjunction of atomic phase formulas. Interestingly, deciding co-admissible formulas (consisting of Boolean combinations of state-formulas, monadic temporal formulas and one co-phase formula in positive polarity) is easier than before, as one can turn the co-phase formula into a monadic formula enumerating all the violations of the atomic phase formulas ( $p \in P$  such that  $p_{\pi_i} \not\leftrightarrow p_{\pi_j}$ ) turns the atomic phase formula into  $(\Diamond p_{\pi_i} \wedge \Diamond \neg p_{\pi_j}) \vee (\Diamond \neg p_{\pi_i} \wedge \Diamond p_{\pi_j})$ . It follows that model-checking co-admissible formulas is also decidable (for both  $\forall^*$  and  $\exists^*$ ). Note that an admissible formula in negative polarity is a co-admissible formula in positive polarity (and vice versa). Finally, since  $\mathcal{K} \models \forall \pi_1 \dots \forall \pi_n. A.\psi$  if and only if  $\mathcal{K} \not\models \exists \pi_1 \dots \exists \pi_n. E.\neg\psi$ , it follows that model-checking is also decidable for the A modality for both admissible and co-admissible formulas (in both polarities), and for both the  $\forall^*$  and  $\exists^*$  fragments.

**Theorem 3.** *Model-checking  $\forall^*$  or  $\exists^*$  admissible and co-admissible formulas is decidable both for formulas with E and formulas with A.*

## 4.2 The Accelerating Construction

The admissible formula in the stuttering construction can express many formulas of interest, but the quantifier structure admits no quantifier alternation. We now consider a second decidable fragment for A-HLTL formulas consisting of formulas with arbitrary quantification  $\mathbb{Q}_1 \pi_1. \mathbb{Q}_2 \pi_2. \dots \mathbb{Q}_n \pi_n E.\psi$  such that  $\mathbb{Q}_i \in \{\forall, \exists\}$ , but where  $\psi$  is an admissible formula where all atomic phase formulas use the same atomic predicates  $P \subseteq AP$ . We call these admissible formulas *simple admissible formulas*. The proof of decidability proceeds this time by creating the *accelerated* Kripke structure  $\mathcal{K}^{acc}$ , where paths jump in one step to the next phase change, and reducing to a HyperLTL model-checking problem on  $\mathcal{K}^{acc}$ .

*Accelerated Kripke structure.* The main idea of the acceleration construction is to convert a finite sequence of transitions in  $\mathcal{K}$  that only change phase in the last transition into a single transition in  $\mathcal{K}^{acc}$ . Also, an infinite sequence of transitions with no phase change is transformed into a self-loop around a sink state. The alphabet remains the same, AP. Given  $\mathcal{K} = \langle S, S_{init}, \delta, L \rangle$ , the accelerated Kripke structure is  $\mathcal{K}^{acc} = \langle S^{acc}, S_{init}, \delta^{acc}, L^{acc} \rangle$  where:

- $S^{acc} = S \cup \{s_{\perp} \mid s \in S\}$  contains two copies of each state in  $S$ , where we use  $s_{\perp}$  to denote the sink state associated with  $s$ . We use  $color(s)$  for the phase of  $s$ , that is, the concrete valuation in  $s$  of the Boolean predicates in  $P$  of the atomic phase formula.
- For every states  $s, s' \in S$  such that  $color(s) \neq color(s')$ , if there is a finite path  $ss_2s_3 \dots s_n s'$  in  $\mathcal{K}$  such that  $color(s) = color(s_2) = \dots = color(s_n)$ , then we add a transition  $(s, s')$  to  $\delta^{acc}$ . These transitions model the jump at the frontier of phase changes. Additionally, if  $s$  can be a sink we add a transition  $(s, s_{\perp})$  and a self-loop from  $s_{\perp}$  to itself.
- $L^{acc}(s) = L(s)$  for  $s \in S$ , and  $L^{acc}(s_{\perp}) = L(s)$ .

This construction can, with standard techniques, be enriched to encode the satisfaction of the temporal monadic formulas along paths of  $\mathcal{K}$ , and then also accelerate the fairness conditions (annotating the accepting states reached along the accelerated paths) into  $\mathcal{K}^{acc}$ .

*Relating paths to accelerated paths.* We now define two auxiliary functions to aid in the proof.

- The first function,  $acc$ , maps paths in  $\mathcal{K}$  into paths in  $\mathcal{K}^{acc}$ . Let  $s$  be an arbitrary state of  $\mathcal{K}$  and  $\rho : ss_1s_2s_3 \dots$  an outgoing path from  $s$ . Either there are infinitely many phase changes in  $\rho$  or only finitely many changes. We create the path  $\rho' = acc(\rho)$  as follows. The initial state of  $\rho$ , that is,  $s$ , is preserved. The states  $s_{i_j}$  in  $\sigma$  that are color changes (that is  $color(s_{i_{j-1}}) \neq color(s_{i_j})$ ) are also preserved, while the states  $s_k$  with  $color(s_{k-1}) = color(s_k)$  are removed from  $\rho$ . If there are only finitely many color changes in  $\rho$ , with  $r$  being the last state preserved, then we pad the path with  $r_{\perp}^{\omega}$ , so  $\rho'$  is also an infinite path. It is easy to see that  $\rho'$  is a path of  $\mathcal{K}^{acc}$  outgoing  $s$ . It is also easy to see that the phase changes in  $\rho$  and  $\rho'$  are the same.
- The second map,  $dec$ , takes a path  $\rho' : ss'_1s'_2 \dots$  of  $\mathcal{K}^{acc}$  and maps it to a path of  $\mathcal{K}$  as follows. For every transition  $(s'_i, s'_{i+1})$  in  $\rho'$  such that  $s'_{i+1}$  is not of the form  $r_{\perp}$ , there is a finite path  $r_1r_2 \dots r_m$  in  $\mathcal{K}$  from  $s'_i$  into  $s'_{i+1}$  that visits only states with the same color as  $s'_i$ , except  $s_{i+1}$  that is a color change. In  $\rho$ , we insert  $r_1r_2 \dots r_m$  between  $s'_i$  and  $s'_{i+1}$ . Now, if for some  $j$ ,  $s'_j$  is of the form  $r_{\perp}$  then  $s'_k = r_{\perp}$  for all  $k > j$ . In  $\mathcal{K}$  there must an infinite path from  $s'_j$  that only visits the same color as  $s'_j$ . We remove all successor states after the first such  $r_{\perp}$  state and replace it with one such infinite path.

Given a trace assignment  $\Pi$  for formula  $\mathbb{Q}_1\pi_1 \dots \mathbb{Q}_n\pi_n.E.\psi$  that assigns  $\Pi(\pi_i) = (\sigma_i, 0)$  for every  $i$  and a path assignment  $\Pi'$  for formula  $\mathbb{Q}_1\pi_1 \dots \mathbb{Q}_n\pi_n.\psi$  that assigns  $\Pi'(\pi_i) = (\sigma'_i, 0)$ , we write  $acc(\Pi) = \Pi'$  if the paths that generate the corresponding traces are related by  $acc$ . Similarly we defined  $dec(\Pi') = \Pi$ . It is easy to show from the construction above that if  $\Pi \models E\psi$  then  $acc(\Pi) \models \psi$ , and if  $\Pi' \models \psi$  then  $dec(\Pi') \models E\psi$ .

The main result for the accelerating construction follows immediately from this observation and allows to reduce the model-checking problem to HyperLTL.

**Theorem 4.** *Let  $\mathcal{K}$  be an arbitrary Kripke structure,  $\mathbb{Q}_1\pi_1\dots\mathbb{Q}_n\pi_n.E.\psi$  such that  $\psi$  is a simple admissible formula. Then  $\mathcal{K} \models \mathbb{Q}_1\pi_1\dots\mathbb{Q}_n\pi_n.E.\psi$  if and only if  $\mathcal{K}^{acc} \models \mathbb{Q}_1\pi_1\dots\mathbb{Q}_n\pi_n.\psi$ .*

### 4.3 Decidable Practical A-HLTL formulas

We revisit the properties expressed in Section 3.2.

- *Linearizability.* The property  $\varphi_{\text{LNZ}}$  is of the form  $\forall\pi.\exists\pi'.E.\Box(\text{history}_\pi \leftrightarrow \text{history}_{\pi'})$  where the temporal formula is a simple admissible formula. Therefore  $\varphi_{\text{LNZ}}$  is decidable by the accelerating construction.
- *Goguen and Meseguer's non-interference.* The property  $\varphi_{\text{GMNI}}$  is expressed by  $\forall\pi.\exists\pi'.E.(\Box\lambda_{\pi'} \wedge \Box(lo_\pi \leftrightarrow lo_{\pi'}))$ , that is, a Boolean combination of a monadic temporal formula and a simple admissible formula. Therefore,  $\varphi_{\text{GMNI}}$  is decidable by the acceleration algorithm.
- *Not never terminates.* Formula  $\varphi_{\text{NNT}}$  is simply a Boolean combination of state formulas and monadic temporal formulas:  $\forall\pi.\exists\pi'.\exists\pi''.E.(\pi[0] = \pi'[0] = \pi''[0]) \rightarrow (\Diamond \text{term}_{\pi'} \wedge \Box\neg\text{term}_{\pi''})$ , so it is again decidable by the acceleration construction.
- *Termination-insensitive noninterference.* To handle  $\varphi_{\text{TIN}}$  we rewrite the formula as follows

$$\varphi_{\text{TIN}} \stackrel{\text{def}}{=} \forall\pi.\forall\pi'.E.(l_\pi \leftrightarrow l_{\pi'}) \rightarrow \left( \begin{array}{l} (\Box\neg\text{term}_\pi \vee \Box\neg\text{term}_{\pi'}) \vee \\ \Box((l_\pi \wedge \text{term}_\pi) \leftrightarrow (l_{\pi'} \wedge \text{term}_{\pi'})) \end{array} \right)$$

Note that  $(l_\pi \wedge \text{term}_\pi)$  can be turned into a state predicate of  $\pi$ . This formula is equivalent because the last case is evaluated precisely to  $l_\pi \leftrightarrow l_{\pi'}$  when both traces terminate. This formula can be handled by the stuttering construction.

- *Termination-sensitive noninterference.* Similarly, to handle  $\varphi_{\text{TSN}}$  we rewrite the formula as

$$\varphi_{\text{TSN}} \stackrel{\text{def}}{=} \forall\pi.\forall\pi'.E.(l_\pi \leftrightarrow l_{\pi'}) \rightarrow \left( \begin{array}{l} (\Box\neg\text{term}_\pi \wedge \Box\neg\text{term}_{\pi'}) \vee \\ \Box((l_\pi \wedge \text{term}_\pi) \leftrightarrow (l_{\pi'} \wedge \text{term}_{\pi'})) \end{array} \right)$$

This is again equivalent because the last case again is the only relevant case when both paths terminate. Again, this case is covered by the stuttering construction.

## 5 Undecidability and Lower-bound Complexity

In this section, we show that the general problem of model-checking A-HLTL is undecidable. Then, we show a polynomial reduction from the synchronous HyperLTL model-checking into A-HLTL model-checking, which shows that even for those A-HLTL formulas for which the model-checking is decidable, this problem is no easier than the corresponding problem for HyperLTL, which is known to be PSPACE-hard in the size of the Kripke structure.

**Theorem 5.** *Let  $\mathcal{K}$  be a Kripke structure and  $\varphi$  be an asynchronous HyperLTL formula. The problem of determining whether or not  $\mathcal{K} \models \varphi$  is undecidable.*

*Proof (sketch).* We reduce the complement of the *post correspondence problem* (PCP) [22, 25] to the A-HLTL model checking problem. PCP consists of a set of dominos, for example, of the form  $[\frac{w}{v}] = \{[\frac{b}{ca}], [\frac{a}{ab}], [\frac{ca}{a}], [\frac{abc}{c}]\}$  and the problem is to decide whether there is a sequence of dominos (with possible repetitions), such that the upper and lower finite strings of the dominos are equal. A solution to the above set of dominos is the sequence  $[\frac{a}{ab}][\frac{b}{ca}][\frac{ca}{a}][\frac{a}{ab}][\frac{abc}{c}]$ . We map a given set of dominos to a Kripke structure that allows arranging the dominos in a sequence (see Fig. 6 for an example), where  $v$  and  $w$  indicate lower and upper words, respectively,  $dom^i$  is for each domino  $[\frac{w_i}{v_i}]$ , and proposition  $lc$  marks whether or not a new letter is processed. The A-HLTL formula in our reduction is the following such that  $dom_{\pi_w} \stackrel{\text{def}}{=} \bigvee_{i \in [1..k]} dom_{\pi_w}^i$ :

$$\begin{aligned} \varphi_{\overline{PCP}} &\stackrel{\text{def}}{=} \forall \pi_w \forall \pi_v. \mathbf{E}. \left( \varphi_{type} \rightarrow (\varphi_{domino} \vee \varphi_{word}) \right) \\ \text{where } \varphi_{type} &\stackrel{\text{def}}{=} \left( (w_{\pi_w} \wedge \neg v_{\pi_w}) \mathcal{U} \text{end}_{\pi_w} \right) \wedge \left( (\neg w_{\pi_v} \wedge v_{\pi_v}) \mathcal{U} \text{end}_{\pi_v} \right) \\ \varphi_{domino} &\stackrel{\text{def}}{=} \square (dom_{\pi_w} \leftrightarrow dom_{\pi_v}) \wedge \diamond \bigvee_{i=1}^k dom_{\pi_w}^i \not\leftrightarrow dom_{\pi_v}^i \\ \varphi_{word} &\stackrel{\text{def}}{=} \square (lc_{\pi_w} \leftrightarrow lc_{\pi_v}) \wedge \diamond \bigvee_{l \in \Sigma_{PCP}} (l_{\pi_w} \not\leftrightarrow l_{\pi_v}) \end{aligned}$$

The intention of formula  $\varphi_{\overline{PCP}}$  is that the Kripke structure is a model of the formula if and only if the original PCP problem has *no* solution. Intuitively, formula  $\varphi_{type}$  forces trace  $\pi_w$  (respectively,  $\pi_v$ ) to traverse only the traces labeled by  $w$  (respectively,  $v$ ) to build a  $w$ -word (respectively,  $v$ -word). Formula  $\varphi_{domino}$  establishes that the trajectory aligns the positions at which the domino indices

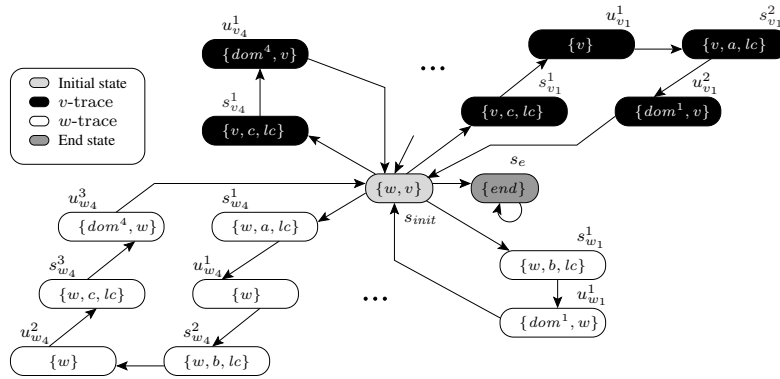


Fig. 6: Mapping from PCP to model checking A-HLTL (only construction for dominos  $[\frac{w_1}{v_1}] = [\frac{b}{ca}]$  and  $[\frac{w_4}{v_4}] = [\frac{abc}{c}]$  are shown).



are checked and at last once the index is different. Finally, formula  $\varphi_{word}$  captures if  $\pi_w$  and  $\pi_v$  are aligned to compare the letters, at least one pair of the letters prescribed by the existential trajectory are different. In the detailed proof in the appendix, we show that the constructed Kripke structure satisfies formula  $\varphi_{pcp}$  if and only if the answer to deciding PCP is negative.  $\square$

Theorem 5 above implies that there is no algorithm to decide the model-checking problem correctly for every formula and every system. However, as we saw in Section 4 for some formulas the model-checking problem is decidable. We now show that in these cases the problem is at least as hard as model-checking HyperLTL, which is known to be PSPACE-hard [6, 23].

**Theorem 6.** *Given a HyperLTL formula  $\varphi$  and a Kripke structure  $\mathcal{K}$  there is a A-HLTL formula  $\varphi'$  and a Kripke structure  $\mathcal{K}'$  such that  $\mathcal{K}'$  is linear in the size of  $\mathcal{K}$ ,  $\varphi'$  is polynomial on the size of  $\varphi$  and  $\mathcal{K} \models \varphi$  if and only if  $\mathcal{K}' \models \varphi'$ .*

The proof proceeds as follow. Given  $\mathcal{K}$  we build a Kripke structure  $\mathcal{K}'$  that alternates between real states in  $\mathcal{K}$  and synchronization states. Then the formula is transformed to force alternations at every other step, therefore forcing the trajectory to synchronize (see Appendix A.2 for details). Since the model-checking problem for HyperLTL is PSPACE-hard on the size of the Kripke structure, the same follows for A-HLTL.

**Corollary 2.** *For asynchronous HyperLTL formulas, the model checking problem is PSPACE-hard in the size of the system.*

## 6 Case Studies and Evaluation

We evaluated our algorithm for the  $\forall_{\pi}^*E$  A-HLTL fragment on several examples. After reducing the asynchronous model checking problem to a synchronous one, we use MCHYPER [9, 10] to check our property. MCHYPER is a model checker for synchronous HyperLTL that can handle formulas with up to one quantifier alternation. It computes the self composition of the system and composes it with the formula automaton. ABC [5] is then used as the backend tool checking the reachability of a violation.

Our reduction from the asynchronous to the synchronous semantics follows the stuttering construction described in Section 4.1. To model check a system against an A-HLTL formula, we first add a stuttering input to the system that forces the system to stutter in the current state. The transformed formula ensures that the stuttering guarantees synchronous phase changes. Currently, the transformation stage is only partially automated while the synchronous verification is fully automated. Future work includes the full automation of all steps. We now describe the various case studies<sup>4</sup>. All our experiments were performed on a MacBook Pro with a 3.3GHz processor and 16GB of RAM running MacOS 11.1.

<sup>4</sup> All experimental data will be publicly available.

## 6.1 Compiler Optimizations

We modeled the source and target programs of different compiler optimization techniques (from [19]) as finite state machines encoded as circuits, and used asynchronous hyperproperties to prove the correspondence between both programs. We analyzed the following optimizations:

- Common Branch Factorization (CBF), where expressions occurring in both branches of a conditional are factored out;
- Loop Peeling (LP), which consists in unrolling of a loop that is executed at least once;
- Dead Branch Elimination (DBE), that is, removing conditional checks and their branches that are unreachable; and
- Expression Flattening (EF), which splits complex computations into several explicit steps.

Besides evaluating each optimization individually, we also examined several combinations of these optimizations. Each optimization affects the alignment between source and target program, so synchronous hyperproperties fail to recognize the correspondence between both programs. Using asynchronous hyperproperties instead allows us to compensate for this misalignment by stuttering the programs accordingly. Essentially, each optimization is checked against the following A-HLTL formula in which  $\pi$  represents traces from the source program and  $\pi'$  traces from the target program:

$$\forall\pi.\forall\pi'.\mathbf{E}.\left(\bigwedge_{i\in I} i_\pi \leftrightarrow i_{\pi'}\right) \rightarrow \left(\Box \bigwedge_{o\in O} o_\pi \leftrightarrow o_{\pi'}\right)$$

This formula states that for all pairs of traces that initially agree on the inputs from the set  $I$  there exists a trajectory that aligns the phase changes of the outputs in set  $O$ . We use the stuttering construction and MCHYPER to verify that in all cases the source and target programs go through the same phases of possibly different length. The results of this case study are summarized in Table 1(a). We note that A-HLTL model-checking subsumes the approach in [19] based on construction of a *buffer automaton* to reason about the alignment of executions.

## 6.2 SPI Bus Protocol

The Serial Peripheral Interface (SPI) is a bus protocol that supports a single main component’s communication with multiple secondary components. Each secondary can be selected individually by the main via the secondary’s own *ss* (“secondary select”) input signal. If a secondary is enabled (that is, if  $\neg ss$  holds as the secondary select is “active low”), it reads the *mosi* (main out, secondary in) signal and writes to the *miso* (main in, secondary out) wire.

We verify the behavior of a single SPI secondary component that receives an input which it sends to the main component upon request. This behavior should always be the same, independent of when the secondary is enabled or how fast the bus protocol’s “serial clock” (*sclk*) set by the main component ticks

Optimizations	System Size		Time (s)	Property	System Size		Time (s)
	#LS	#ANDS			#LS	#ANDS	
EF	12	64	0.6	SPI-correct	30	175	65.7
DBE	16	128	0.8		33	296	155.8
CBF	16	145	2.7	SPI-term			
LP	28	514	365.9				
CBF+DBE	16	137	11.4				
CBF+DBE+EF	20	175	10.0				
CBF+EF	20	180	1.7				
EF+LP	41	8642	1315.2				

(a) Compiler Optimizations

(b) SPI

Table 1: Verification times and system sizes in number of latches (#LS) and AND-gates (#ANDS) for the case studies.

compared to the secondary’s internal clock. The A-HLTL formula we check is the following (see observational determinism in Section 1):

$$\forall \pi. \forall \pi'. E. \left( \begin{array}{c} \bigwedge_{i \in \{in, init\}} i_{\pi} \leftrightarrow i_{\pi'} \\ \wedge \\ SPI \text{ input assumptions} \end{array} \right) \rightarrow \square \left( \begin{array}{c} (miso_{\pi} \wedge \neg sclk_{\pi} \wedge \neg ss_{\pi}) \\ \leftrightarrow \\ (miso_{\pi'} \wedge \neg sclk_{\pi'} \wedge \neg ss_{\pi'}) \end{array} \right)$$

This formula (called **SPI-correct** in Table 1(b)) ensures that for all pairs of traces  $\pi$  and  $\pi'$  that agree on the initial configuration, on the input, and additional *SPI input assumptions*, there is a trajectory that aligns their relevant behavior. We consider it relevant that both secondaries agree on their *miso* output whenever they are enabled and the *sclk* is low. Checking *miso* only when the *sclk* is low is sufficient as changes on *miso* only occur at falling edges of the *sclk*. The *SPI input assumptions* are required to guarantee the implicit assumptions of the protocol, for example, that the *sclk* behaves as an infinitely ticking clock. By introducing additional variables and applying logical transformations, we obtain an equivalent formula that syntactically lies in the fragment of the stuttering construction. Again, we reduce this model checking problem to the synchronous semantics and use MCHYPER to perform the verification.

In a second experiment, we modified the system to send the value only once and checked it for termination insensitive noninterference **SPI-term** (see Sections 3.2 and 4.3). In our setup, we use the variable *term* to flag that the secondary has sent the full value. In the premise of the formula, we require that the input value is equal on both traces and again assume that the inputs conform to the SPI protocol. The conclusion checks if both secondaries have sent the same values by using additional variables that are set together with *term*. The results of this case study are summarized in Table 1(b).

## 7 Related Work

The study of specific hyperproperties, such as noninterference, dates back to the seminal work by Goguen and Meseguer [13] in the 1980s. The first systematic study of hyperproperties is due to Clarkson and Schneider [7].

It is well-known that classic specification languages like LTL cannot express hyperproperties. There are two principal methods with which the standard logics have been extended to express hyperproperties:

- The first method is the quantification over variables that identify specific paths or traces. The temporal logics LTL, CTL\* have been extended with quantification over traces and paths, resulting in the temporal logics HyperLTL and HyperCTL\* [6]. There are also extensions of the  $\mu$ -calculus, most recently, the temporal fixpoint calculus  $H_\mu$  [14], which adds path quantifiers to the polyadic  $\mu$ -calculus [2].
- The second method is the addition of the equal-level predicate  $E$  to first-order and second-order logics, like MPL, MSO, FOL, and S1S, which results in the logics FOL[E], S1S[E], MPL[E], MSO[E] [8, 12].

HyperCTL\*, MPL[E], and MSO[E] are branching-time logics, we therefore focus in the following on the linear-time logics HyperLTL,  $H_\mu$ , FOL[E], and S1S[E]. Among these logics, HyperLTL is the only logic for which practical model-checking algorithms are known [9, 10, 16]. For HyperLTL, the algorithms have been implemented in the model checkers MCHyper and bounded model checker HyperQube. As discussed in this paper, HyperLTL is limited to synchronous hyperproperties.

FOL[E] can express a limited form of asynchronous hyperproperties. As shown in [8], FOL[E] is subsumed by HyperLTL with additional quantification over predicates. Using such predicates as “markers,” one can relate different positions in different traces. However, only a finite number of such predicates is available in each formula. S1S[E] is known to be strictly more expressive than FOL[E] [8], and conjectured to subsume  $H_\mu$  [14]. For S1S[E] and  $H_\mu$ , the model checking problem is in general undecidable; for  $H_\mu$ , two fragments, the  $k$ -synchronous,  $k$ -context bounded fragments, have been identified for which model checking remains decidable [14]. It is not known, however, if any of the commonly used hyperproperties, like observational determinism, noninterference, or linearizability, can be encoded in these fragments.

Like S1S[E] and  $H_\mu$ , asynchronous HyperLTL has an (in general) undecidable model checking problem. However, in this paper we have identified decidable fragments of asynchronous HyperLTL that can express observational determinism, noninterference, and linearizability. A-HLTL is thus the first logic for hyperproperties that can express the major asynchronous hyperproperties of interest within decidable fragments. Furthermore, asynchronous HyperLTL is the first logic for asynchronous hyperproperties with a practical model checking algorithm.

## 8 Conclusion

We have introduced A-HLTL, a temporal logic to describe asynchronous hyperproperties. This logic extends HyperLTL with *trajectory* modalities, which control when a trace proceeds and when it stutters. Synchronous HyperLTL corresponds to a trajectory that always moves all paths in a lock-step manner. This notion of trajectory allows to define formulas that are invariant under stuttering, paving the way for relevant model-checking optimizations such a partial order reduction and abstraction-refinement techniques in the context of hyperproperties. We show that model-checking A-HLTL formulas is in general undecidable, and identify two fragments of A-HLTL formulas, which cover a rich set of security requirements and can be decided by a reduction to HyperLTL model-checking. This in turn has allowed us to reuse the existing model-checker McHyper.

Future work includes the study of larger decidable fragments (that encompass both fragments studied here), extending the logic allowing several trajectory modalities, as well as their implementation in practical tools. Extending bounded model-checking [16] to A-HLTL is another interesting research problem. Asynchronous hyperproperties are important for applying a logic-based verification approach to verify hyperproperties for *software* programs, because the relative speed of the execution of programs depends on many factors like the compiler, hardware, execution platform and concurrent running programs, that the analysis must tolerate. Therefore, future work includes adapting techniques for infinite-state software model-checking, like deductive methods, abstraction, etc to verify A-HLTL properties of software systems.

## References

1. B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21:181–185, 1985.
2. H. R. Andersen. A polyadic modal  $\mu$ -calculus, 1994. Technical Report.
3. A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-insensitive noninterference leaks more than just a bit. In *Proceedings of 13th European Symposium on Research in Computer Security Computer Security (ESORICS)*, pages 333–348, 2008.
4. B. Bonakdarpour, C. Sánchez, and G. Schneider. Monitoring hyperproperties by combining static analysis and runtime verification. In *Proceedings of the 8th Leveraging Applications of Formal Methods, Verification and Validation (ISoLA)*, pages 8–27, 2018.
5. R. K. Brayton and A. Mishchenko. ABC: an academic industrial-strength verification tool. In *Proceedings of CAV*, volume 6174 of *LNCS*, pages 24–40. Springer, 2010.
6. M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez. Temporal logics for hyperproperties. In *Proc. of the 3rd Int’l Conf. on Principles of Security and Trust (POST’14)*, volume 8414 of *LNCS*, pages 265–284. Springer, 2014.
7. M. R. Clarkson and F. B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
8. N. Coenen, B. Finkbeiner, C. Hahn, and J. Hofmann. The hierarchy of hyperlogics. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2019.
9. N. Coenen, B. Finkbeiner, C. Sánchez, and L. Tentrup. Verifying hyperliveness. In I. Dillig and S. Tasiran, editors, *Computer Aided Verification*, pages 121–139, Cham, 2019. Springer International Publishing.
10. B. Finkbeiner, M. Rabe, and C. Sánchez. Algorithms for model checking HyperLTL and HyperCTL\*. In *In Proc. of the 27th Int’l Conf. on Computer Aided Verification (CAV’15)*, volume 9206 of *LNCS*, pages 30–48. Springer, 2015.
11. B. Finkbeiner, M. N. Rabe, and C. Sánchez. A temporal logic for hyperproperties. *CoRR*, abs/1306.6657, 2013.
12. B. Finkbeiner and M. Zimmermann. The first-order logic of hyperproperties. In *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, pages 30:1–30:14, 2017.
13. J. A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 11–20, 1982.
14. J. O. Gutsfeld, M. Müller-Olm, and C. Ohrem. Automata and fixpoints for asynchronous hyperproperties. *Proc. ACM Program. Lang.*, 5(POPL):1–29, 2021.
15. M. Herlihy and J. M. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, 1990.
16. T.-H. Hsu, C. Sánchez, and B. Bonakdarpour. Bounded model checking for hyperproperties. In *Proceedings of the 27th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS)*, 2021. To appear.
17. L. Lamport. “sometime” is sometimes “not never” - on the temporal logic of programs. In *Proceedings of the Seventh Annual ACM Symposium on Principles of Programming Languages (PoPL)*, pages 174–185. ACM Press, 1980.

18. Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems*. Springer-Verlag, 1995.
19. K. S. Namjoshi and L. M. Tabajara. Witnessing secure compilation. In D. Beyer and D. Zufferey, editors, *Verification, Model Checking, and Abstract Interpretation*, pages 1–22, Cham, 2020. Springer International Publishing.
20. A. Pnueli. The temporal logic of programs. In *Symposium on Foundations of Computer Science (FOCS)*, pages 46–57, 1977.
21. A. Pnueli, M. Siegel, and E. Singerman. Translation validation. In *Proceedings of the 4th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS)*, pages 151–166, 1998.
22. E. L. Post. A variant of a recursively unsolvable problem. *Bulleting of the American Mathematical Society*, 52:264–268, 1946.
23. M. N. Rabe. *A Temporal Logic Approach to Information-flow Control*. PhD thesis, Saarland University, 2016.
24. A. Sabelfeld and D. Sands. A per model of secure information flow in sequential programs. *Higher Order Symbolic Computation*, 14(1):59–91, 2001.
25. M. Sipser. *Introduction to the Theory of Computation*. MIT Press, Boston, MA, USA, 2012.
26. Y. Wang, M. Zarei, B. Bonakdarpour, and M. Pajic. Statistical verification of hyperproperties for cyber-physical systems. *ACM Transactions on Embedded Computing Systems (TECS)*, 18(5s):92:1–92:23, 2019.
27. S. Zdancewic and A. C. Myers. Observational determinism for concurrent program security. In *Proceedings of the 16th IEEE Computer Security Foundations Workshop (CSFW)*, page 29, 2003.

## A Full Proofs of Undecidability and Complexity

### A.1 Undecidability of Asynchronous HyperLTL

We show now the full proof of undecidability of full A-HLTL via reduction from PCP (Theorem 5)

**Theorem 5.** *Let  $\mathcal{K}$  be a Kripke structure and  $\varphi$  be an asynchronous HyperLTL formula. The problem of determining whether or not  $\mathcal{K} \models \varphi$  is undecidable.*

*Proof.* In order to show the undecidability of model-checking A-HLTL, we reduce the complement of the *post correspondence problem (PCP)*, which is a well-known undecidable problem [22, 25], to our model checking problem. The PCP is as follows. Let  $\Sigma_{pcp}$  be an alphabet of size at least two. The input to the decision problem is a collection of dominos of the form:

$$\left[ \frac{w}{v} \right] = \left\{ \left[ \frac{w_1}{v_1} \right], \left[ \frac{w_2}{v_2} \right], \dots, \left[ \frac{w_k}{v_k} \right] \right\}$$

where for all  $i \in [1, k]$ , we have  $v_i, w_i \in \Sigma_{pcp}^*$ . The decision problem consists of checking whether there exists a finite sequence of dominos of the form

$$\left[ \frac{w_{i_1}}{v_{i_1}} \right] \left[ \frac{w_{i_2}}{v_{i_2}} \right] \dots \left[ \frac{w_{i_m}}{v_{i_m}} \right]$$

where each index  $i_j \in [1, k]$ , such that the upper and lower finite strings of the dominos are equal, i.e.,

$$v_{i_1} v_{i_2} \dots v_{i_m} = w_{i_1} w_{i_2} \dots w_{i_m}.$$

For some collection of dominos, it is impossible to find a match. An example is the following:

$$\left[ \frac{w}{v} \right] = \left\{ \left[ \frac{abc}{ab} \right], \left[ \frac{ca}{a} \right], \left[ \frac{acc}{ba} \right] \right\}$$

We now start from an arbitrary instance  $\left[ \frac{w}{v} \right] = \left\{ \left[ \frac{w_1}{v_1} \right], \left[ \frac{w_2}{v_2} \right], \dots, \left[ \frac{w_k}{v_k} \right] \right\}$  of the PCP and create a Kripke structure  $\mathcal{K}_{pcp} = \langle S, S_{init}, \delta, L \rangle$  and an asynchronous HyperLTL formula  $\varphi_{\overline{pcp}}$  (i.e., no solution to PCP), and later show that the PCP instance has *no* solution if and only if  $\mathcal{K}_{pcp} \models \varphi_{\overline{pcp}}$ . To illustrate the construction, we use the following instance of PCP as a running example. The alphabet is  $\Sigma_{pcp} = \{a, b, c\}$  and the set of dominos [25] is:

$$\left[ \frac{w}{v} \right] = \left\{ \left[ \frac{b}{ca} \right], \left[ \frac{a}{ab} \right], \left[ \frac{ca}{a} \right], \left[ \frac{abc}{c} \right] \right\}$$

*Kripke structure.* We map the PCP instance to  $\mathcal{K}_{pcp}$  as follows (see Fig. 6 for example):

- *Atomic propositions:* The set AP of atomic propositions include the following:



- Two propositions  $v$  and  $w$  to indicate that a state belongs to a path that builds a  $v$ -word in the instance of PCP, or to a path that builds a  $w$ -word.
- A proposition  $dom^i$  for each domino from  $[\frac{w_i}{v_i}]$ , where  $i \in [1, k]$ , that indicates the domino that contains a word from  $v_i$  or  $w_i$  that has been processed.
- A *letter clock* proposition  $lc$ , which serves the purpose of marking whether or not a new letter is processed.
- A proposition  $end$ , which serves the purpose of marking the end of sequencing dominos.
- A proposition for each letter in the PCP alphabet  $\Sigma_{pcp}$ .

Thus, we have:

$$AP = \Sigma_{pcp} \cup \{v, w, lc, end\} \cup \{dom^i \mid i \in [1, k]\}$$

- *States and transitions:* We include the following states  $S$  and transitions  $\delta$ :
  - We include a single initial state  $s_{init}$  with no labels and an end state  $s_e$  labeled by  $\{end\}$  with a self-loop  $(s_e, s_e)$ .
  - For each word  $w_i = w_i^1 \cdots w_i^l$  in a domino  $[\frac{w_i}{v_i}]$  of the PCP instance, where  $i \in [1, k]$ , we include a path of length  $2l$  that starts and ends in  $s_{init}$ , and whose states are only reachable in the given path. More specifically,
    - \* For each  $j \in [1, l]$ , we introduce two states  $s_{w_i}^j$  and  $u_{w_i}^j$ . Each state  $s_{w_i}^j$ , where  $i \in [1, k]$ , is labeled by  $\{w, w_i^j, lc\}$ . This state indicates a new letter in the word  $w_i$ . Each state  $u_{w_i}^j$ , where  $j \in [1, l]$ , is labeled by  $\{w\}$ . This state does not correspond to a letter and its purpose is to server as separation between two letters within the same word. In this manner, a word is encoded by a sequence of letter/no-letter states. Finally, state  $u_{w_i}^l$  is labeled by  $\{dom^i, w\}$ , which is used for synchronizing the end of words  $w_i$  and  $v_i$  in domino  $i$ .
    - \* We add a transition  $(s_{w_i}^j, u_{w_i}^j)$  for each  $j \in [1, l]$  and a transition  $(u_{w_i}^j, s_{w_i}^{j+1})$  for each  $j \in [1, l]$ .
    - \* We also add transitions  $(s_{init}, s_{w_i}^1)$ ,  $(u_{w_i}^l, s_{init})$ .
  - For each word  $v_i = v_i^1 \cdots v_i^l$  in a domino  $[\frac{w_i}{v_i}]$  of the PCP instance, where  $i \in [1, k]$ , we include a path in the same way as we did for  $w_i$ , except that all instances of “ $w$ ” are replaced by “ $v$ ”.

*Asynchronous HyperLTL formula* The formula that will be checked against the Kripke structure is the following:

$$\varphi_{pcp} \stackrel{\text{def}}{=} \forall \pi_w \forall \pi_v. \mathbf{E}. \left( \varphi_{type} \rightarrow (\varphi_{domino} \vee \varphi_{word}) \right) \quad (1)$$

where

$$\begin{aligned}\varphi_{type} &\stackrel{\text{def}}{=} \left( (w_{\pi_w} \wedge \neg v_{\pi_w}) \mathcal{U} \text{end}_{\pi_w} \right) \wedge \left( (\neg w_{\pi_v} \wedge v_{\pi_v}) \mathcal{U} \text{end}_{\pi_v} \right) \\ \varphi_{domino} &\stackrel{\text{def}}{=} \square (dom_{\pi_w} \leftrightarrow dom_{\pi_v}) \wedge \diamond \bigvee_{i=1}^k dom_{\pi_w}^i \not\leftrightarrow dom_{\pi_v}^i \\ \varphi_{word} &\stackrel{\text{def}}{=} \square (lc_{\pi_w} \leftrightarrow lc_{\pi_v}) \wedge \diamond \bigvee_{l \in \Sigma_{pcp}} (l_{\pi_w} \not\leftrightarrow l_{\pi_v})\end{aligned}$$

and

$$dom_{\pi_w} \stackrel{\text{def}}{=} \bigvee_{i \in [1..k]} dom_{\pi_w}^i$$

The intention of formula  $\varphi_{\overline{pcp}}$  is that the Kripke structure is a model of the formula if and only if the original PCP problem has *no* solution. Intuitively, formula  $\varphi_{type}$  forces trace  $\pi_w$  (respectively,  $\pi_v$ ) to traverse only the traces labeled by  $w$  (respectively,  $v$ ) to build a  $w$ -word (respectively,  $v$ -word). When such traces are built (i.e., the left side of the implication), PCP has no solution when at least one of the following conditions hold: either (1) the order of dominos in  $\pi_w$  and  $\pi_v$  is not legitimate, or (2) the letters do not match. More formally, formula  $\varphi_{domino}$  establishes that the trajectory aligns the positions at which the domino indices are checked and at last once the index is different. Finally, formula  $\varphi_{word}$  captures if  $\pi_w$  and  $\pi_v$  are aligned to compare the letters, at least one pair of the letters prescribed by the existential trajectory are different.

We now show that the given PCP instance has no matching if and only if  $\mathcal{K}_{pcp} \not\models \varphi_{\overline{pcp}}$ :

- ( $\Rightarrow$ ) Supposed that the PCP instance has no solution. That is, no sequence of the dominos of the form:

$$\begin{bmatrix} w_{i_1} \\ v_{i_1} \end{bmatrix} \begin{bmatrix} w_{i_2} \\ v_{i_2} \end{bmatrix} \dots \begin{bmatrix} w_{i_m} \\ v_{i_m} \end{bmatrix}$$

constructs a matching where  $w_{i_1} w_{i_2} \dots w_{i_m} = v_{i_1} v_{i_2} \dots v_{i_m}$ . In our running example, the following is a sequence of dominos that does not render a good matching:

$$\begin{bmatrix} a \\ ab \end{bmatrix} \begin{bmatrix} b \\ ca \end{bmatrix} \begin{bmatrix} ca \\ a \end{bmatrix} \begin{bmatrix} abc \\ c \end{bmatrix} \quad (2)$$

But the following is a solution to the problem:

$$\begin{bmatrix} a \\ ab \end{bmatrix} \begin{bmatrix} b \\ ca \end{bmatrix} \begin{bmatrix} ca \\ a \end{bmatrix} \begin{bmatrix} a \\ ab \end{bmatrix} \begin{bmatrix} abc \\ c \end{bmatrix} \quad (3)$$

We now show that if the PCP problem has no solution, then  $\mathcal{K}_{pcp} \models \varphi_{\overline{pcp}}$ . Let  $\pi_w$  and  $\pi_v$  be any two arbitrary traces of the  $\mathcal{K}_{pcp}$ . If these traces do




Fig. 7: Obtaining PCP matching from an instantiation of trajectory  $\tau_2$  in Formula 1.

not satisfy formula  $\varphi_{type}$ , then they are representing two words that are actually not viable  $w$  and  $v$  words. This makes  $\varphi_{\overline{pcp}}$  vacuously true. So, let us consider the case where  $\pi_w$  and  $\pi_v$  satisfy  $\varphi_{type}$ . For example, Fig. 7 shows an instantiation of traces  $\pi_w$  and  $\pi_v$  with sequence (2) that satisfies  $\varphi_{type}$ .

Now, we consider two cases: either the words represented by  $\pi_w$  and  $\pi_v$  do not establish a legitimate sequence of dominos or they do. If they do not, then traces  $\pi_w$  and  $\pi_v$  satisfy formula  $\varphi_{domino}$  and, hence, we have  $\mathcal{K}_{pcp} \models \varphi_{\overline{pcp}}$ . Now, consider the case where  $\pi_w$  and  $\pi_v$  do establish a legitimate sequence of dominos. Since there is no solution to PCP, the words represented by  $\pi_w$  and  $\pi_v$  cannot match. That means there is at least one pair of letters that do not match, which in turn implies the existence of a trajectory that finds this mismatch and indeed satisfies formula  $\varphi_{word}$ . Put it another way, since there is no solution to PCP, the sequence that reaches a letter mismatch construct the witness to the existential trajectory quantifier. Fig. 7 shows the trajectory that identifies the mismatched letters in our running example. Therefore, we have  $\mathcal{K}_{pcp} \models \varphi_{\overline{pcp}}$ .

- ( $\Rightarrow$ ) Now, suppose we have  $\mathcal{K}_{pcp} \models \varphi_{\overline{pcp}}$ . This implies that for any pair of traces  $\pi_v$  and  $\pi_w$  that represent legitimate  $w$  and  $v$  words (i.e., where  $\varphi_{type}$  holds), there exists a trajectory that satisfies  $(\varphi_{domino} \vee \varphi_{word})$ . This means that for the chosen pair of traces  $\pi_v$  and  $\pi_w$ , the existential trajectory has found either illegitimate alignment of dominos, or letters that do not match. In our example, the existential trajectory can be instantiated by the trajectory shown in Fig. 7. Such a witness trajectory shows that for a pair of traces that represent legitimate  $w$  and  $v$  words, either they do not establish a legitimate sequence of dominos or not all letters along the words match. Since this holds for any such pair of traces that explore any ordering of  $w$  and  $v$  words, it means that there is no solution to PCP.

This shows that the  $\mathcal{K}_{pcp} \models \varphi_{\overline{pcp}}$  if and only if the PCP instance  $\frac{w}{v}$  has no solution. Consequently, the model-checking problem for asynchronous HyperLTL is undecidable.  $\square$

## A.2 Lower Bound Complexity

We show now construction of the lower bound complexity using a reduction to HyperLTL model-checking. The main results is Corollary 2 that follows from the following theorem.

**Theorem 6.** *Given a HyperLTL formula  $\varphi$  and a Kripke structure  $\mathcal{K}$  there is a A-HLTL formula  $\varphi'$  and a Kripke structure  $\mathcal{K}'$  such that  $\mathcal{K}'$  is linear in the size of  $\mathcal{K}$ ,  $\varphi'$  is polynomial on the size of  $\varphi$  and*

$$\mathcal{K} \models \varphi \text{ if and only if } \mathcal{K}' \models \varphi'$$

*Proof.* We reduce the HyperLTL model-checking problem, which is known to be PSPACE-hard to asynchronous HyperLTL model-checking.

Consider an instance of the HyperLTL model-checking problem consist of a Kripke structure  $\mathcal{K} = \langle S, S_{init}, \delta, L \rangle$  and a HyperLTL formula of the form:

$$\varphi = Q_1\pi_1.Q_2\pi_2 \cdots Q_n\pi_n.\psi$$

where each  $Q_i \in \{\forall, \exists\}$  is a trace quantifier ( $i \in [1, n]$ ). We map this instance to asynchronous HyperLTL model-checking problem as follows:


– **Kripke structure.** We now transform  $\mathcal{K}$  into another Kripke structure  $\mathcal{K}' = \langle S', s_{init}, \delta', L' \rangle$  as follows:

- First, we set the set of atomic propositions as  $AP \cup \{sync\}$ .
- For each transition  $(s_0, s_1) \in \delta$ , we introduce a new state  $u_{(s_0, s_1)}$  labeled only by proposition *sync*. That is:

$$S' = S \cup \{u_{(s_0, s_1)} \mid (s_0, s_1) \in \delta\}$$

The set of labels of states in  $S$  remains unchanged.

- We replace transition  $(s_0, s_1)$  by two transitions  $(s_0, u_{(s_0, s_1)})$  and  $(u_{(s_0, s_1)}, s_1)$ :



That is,  $\delta' = \{(s_0, u_{(s_0, s_1)}), (u_{(s_0, s_1)}, s_1) \mid (s_0, s_1) \in \delta\}$ .

– **Formula.** We first map subformulas of  $\psi$  using the following  $g$  function:

$$\begin{aligned} g(a_\pi) &= a_\pi \\ g(\neg\psi) &= \neg g(\psi) \\ g(\psi_1 \vee \psi_2) &= g(\psi_1) \vee g(\psi_2) \\ g(\bigcirc\psi) &= \bigcirc\bigcirc g(\psi) \\ g(\psi_1 \mathcal{U} \psi_2) &= \left( \bigwedge_{i=1}^n \neg sync_{\pi_i} \rightarrow g(\psi_1) \right) \mathcal{U} \left( g(\psi_2) \wedge \bigwedge_{i=1}^n \neg sync_{\pi_i} \right) \end{aligned}$$

Intuitively, function  $g$  ensures that (1) state formulas are evaluated in the same fashion as in HyperLTL semantics (2) the next operator is evaluated

within two steps (instead of one) due to the existence of intermediate *sync* states, and (3) both sides of the until operator are only evaluated over *sync* states. The asynchronous HyperLTL formula will be the following:

$$\varphi' \stackrel{\text{def}}{=} Q_1\pi_1.Q_2\pi_2 \cdots Q_n\pi_n.E.(g(\varphi) \wedge \psi_{sync})$$

where

$$\psi_{sync} \stackrel{\text{def}}{=} \square_{\tau} \left( \bigwedge_{i,j=1}^n (sync_{\pi_i} \leftrightarrow sync_{\pi_j}) \wedge \left( \bigwedge_{i=1}^n sync_{\pi_i} \leftrightarrow \bigcirc \bigwedge_{i=1}^n \neg sync_{\pi_i} \right) \right)$$

Intuitively, formula  $\psi_{sync}$  enforces the synchronous semantics. That is, each step of trajectory  $\tau$  alternates between all *sync* and all  $\neg sync$  states across all traces. Formula  $\psi_{map}$  ensures that all subformulas of the given HyperLTL formula are evaluated only in *sync* states across all traces.

We now show that  $\mathcal{K} \models \varphi$  if and only if  $\mathcal{K}' \models \varphi'$ :

- ( $\Rightarrow$ ) Suppose that  $\mathcal{K} \models \varphi$  holds. To show that  $\mathcal{K}' \models \varphi'$ , we instantiate all existential trace quantifiers in  $\varphi'$  as we do for  $\mathcal{K}$  to satisfy  $\varphi$  (except that traces in  $\text{Traces}(\mathcal{K}', s_{init})$  have the additional intermediate *sync* state in between each two consecutive states of traces of  $\mathcal{K}$ ). Moreover, we instantiate  $\tau$  in  $\varphi'$  with a trajectory that advances all traces in all its steps (i.e., identical to the semantics of synchronous HyperLTL). It is straightforward to see that this trajectory satisfies  $\psi_{sync}$ , as in each step of the trajectory, *sync* is alternatively true and false either across all traces or none. We now show that  $\psi_{map}$  is also satisfied by structural induction:
  - If  $\psi_{map}$  is a state formula (i.e., an atomic proposition, negation or disjunction), then it is trivially satisfied, as  $g(\psi) = \psi$ .
  - If  $\bigcirc\psi$  holds in  $\mathcal{K}$ , then  $\psi$  holds in two steps in  $\tau$ , as an intermediate *sync* state is inserted in each transition of  $\delta$  of  $\mathcal{K}'$ . This implies that  $\bigcirc_{\tau}\bigcirc_{\tau}\psi$  holds in  $\mathcal{K}'$ .
  - Suppose  $\psi_1 \mathcal{U} \psi_2$  holds in  $\mathcal{K}$ . Then, since there exists some step  $i \geq 0$  in instantiated traces from  $\text{Traces}(\mathcal{K}, s_{init})$  such that  $\psi_2$  holds. This means that in the instantiated traces from  $\text{Traces}(\mathcal{K}', s_{init})$ ,  $g(\psi_2)$  holds in  $2i$  steps, where *sync* is false across all traces. This the right side of the until holds in  $\mathcal{K}'$ . Also, for all  $j \in [0, i)$ , in instantiated traces from  $\text{Traces}(\mathcal{K}, s_{init})$ ,  $\psi_1$  holds. This means that in instantiated traces from  $\text{Traces}(\mathcal{K}', s_{init})$ ,  $g(\psi_1)$  also holds if *sync* is false in states across all traces. Hence, we have  $\mathcal{K}' \models \varphi'$ .
- ( $\Leftarrow$ ) Suppose that we have  $\mathcal{K}' \models \varphi'$ . To show that  $\mathcal{K} \models \varphi$ , we instantiate all existential trace quantifiers in  $\varphi$  as we do for  $\mathcal{K}'$  to satisfy  $\varphi'$  (except that in traces of  $\text{Traces}(\mathcal{K}, s_{init})$ , intermediate *u* states in between each two consecutive states of traces of  $\mathcal{K}'$  are eliminated). Since,  $\psi_{sync}$  holds in  $\varphi'$ , there exists a trajectory  $\tau$  which advances all traces in all its steps such that in each step of the trajectory, *sync* is alternatively true and false either across all traces or none. We now show that  $\varphi$  holds in  $\mathcal{K}$  by structural induction:

- If  $\psi$  is a state formula (i.e., an atomic proposition, negation or disjunction), then it is trivially satisfied, as  $\psi = g(\psi)$ .
- First, notice that the number of nested next operators in  $\varphi'$  is always even. If  $\bigcirc\bigcirc\psi$  holds in  $\mathcal{K}'$ , then  $\psi$  holds in two steps in  $\tau$ . When the  $u$  state is eliminated,  $\psi$  will hold in the next state. This implies that  $\bigcirc\psi$  holds in  $\mathcal{K}$ .
- Suppose  $\psi_1 \mathcal{U} \psi_2$  holds in  $\mathcal{K}'$ . Then, there exists some step  $i \geq 0$  in instantiated traces from  $\text{Traces}(\mathcal{K}', s_{init})$  such that  $\psi_2$  holds. In this step, *sync* has to be false states across all traces. This means that in the instantiated traces from  $\text{Traces}(\mathcal{K}, s_{init})$ ,  $g^{-1}(\psi_2)$  (obtained by removing the right side of the conjunction when computing  $g$  of an until) holds in  $i/2$  steps. Also,  $\psi_1$  holds in all states  $j \in [0, i)$ , in instantiated traces from  $\text{Traces}(\mathcal{K}', s_{init})$ . This means that in all  $j \in [0, i/2)$  in the instantiated traces from  $\text{Traces}(\mathcal{K}, s_{init})$ ,  $g^{-1}(\psi_1)$  (obtained by removing the left side of the implication when computing  $g$ ) holds.

It is easy to see that the size of  $\mathcal{K}'$  is linear in the size of  $\mathcal{K}$  and that  $\varphi'$  is polynomial on the size of  $\varphi$ .  $\square$

**Corollary 2.** *For asynchronous HyperLTL formulas, the model-checking problem is PSPACE-hard in the size of the system.*

## B Full Decidability Proofs

### B.1 The Stuttering Construction

Since there is only one occurrence of  $\psi_{ph}$  in  $\psi$  we present here the case where the occurrence is with positive polarity as in Section 4. The case for negative polarity is analogous. We will use the fact that  $\psi_{ph}$  occurs with positive polarity to spuriously evaluate  $\psi'$  (the formula that  $\psi_{ph}$  gets translated into) into *true* for those tuples of paths of  $\mathcal{K}^{st}$  for which the evaluation is not relevant (non-aligned traces). The final verdict will depend on those paths that are aligned.

For the following discussion we fix  $\mathcal{K}$  as the original Kripke structure,  $\mathcal{K}^{st}$  the corresponding stuttering structure and  $\psi$  as the admissible formula. Given a tuple of traces  $(\sigma_1, \dots, \sigma_n)$  of  $\mathcal{K}$  and a trajectory  $t$ , we define the following sequence of tuples of pointed traces

$$((\sigma_1, p_1^0), \dots, (\sigma_n, p_n^0)) \quad , \quad ((\sigma_1, p_1^1), \dots, (\sigma_n, p_n^1)) \quad \dots$$

such that (1)  $p_i^0 = 0$  and (2)  $p_i^{k+1} = p_i^k + 1$  if  $\pi_i \in t(k)$  and  $p_i^{k+1} = p_i^k$  if  $\pi_i \notin t(k)$ . In turn, this sequence induces a tuple of traces  $(\sigma_1^*, \dots, \sigma_n^*)$  of  $\mathcal{K}^{st}$  as follows:

$$\begin{aligned} \sigma_i^*(0) &= \sigma_i(0) \\ \sigma_i^*(k+1) &= \begin{cases} \sigma_i(p_i^{k+1}) & \text{if } p_i^{k+1} = p_i^k \\ \sigma_i(p_i^{k+1}) \cup \{st\} & \text{if } p_i^{k+1} = p_i^k + 1 \end{cases} \end{aligned}$$

Note that  $\sigma_i^*$  is encoding the progress of  $\sigma_i$  according to trajectory  $t$ . It is easy to see that  $(\sigma_1^*, \dots, \sigma_n^*)$  are fair traces of  $\mathcal{K}^{st}$  (due to the fairness requirement of  $t$ ) and that  $\sigma_i^*$  is a stuttering expansion of  $\sigma_i$ . We call  $(\sigma_1^*, \dots, \sigma_n^*)$  the traces induced by  $(\sigma_1, \dots, \sigma_n)$  and  $t$ , which we denote by

$$(\sigma_1^*, \dots, \sigma_n^*) = \text{expand}((\sigma_1, \dots, \sigma_n), t).$$

Similarly, given a tuple of fair traces  $(\sigma_1^*, \dots, \sigma_n^*)$  of  $\mathcal{K}^{st}$  there is a unique tuple  $(\sigma_1, \dots, \sigma_n)$  of traces of  $\mathcal{K}$  and trajectory  $t$  defined as follows. Given  $\sigma_i^*$ , let  $i_1 < i_2 < \dots$  be the ordered set of positions for which  $st \notin \sigma_i^*$  (this set is guaranteed to be infinite because  $\sigma_i^*$  is fair). Then  $\sigma_i$  is defined as  $\sigma_i(k) = \sigma_i^*(i_k)$ . Also,  $t(k)$  contains  $\pi_i$  if and only if  $st \notin \sigma_i^*(k)$ . We call  $(\sigma_1, \dots, \sigma_n)$  and  $t$  the trace and trajectory induced by  $(\sigma_1^*, \dots, \sigma_n^*)$ , which we denote by

$$((\sigma_1, \dots, \sigma_n), t) = \text{compress}((\sigma_1^*, \dots, \sigma_n^*)).$$

The following result follows directly from the definitions, and implies that  $(\text{compress} \circ \text{expand})$  and  $(\text{expand} \circ \text{compress})$  are identities.

**Lemma 1.** *Let  $(\sigma_1, \dots, \sigma_n)$  be a tuple of traces of  $\mathcal{K}$ ,  $t$  a trajectory and  $(\sigma_1^*, \dots, \sigma_n^*)$  a tuple of fair traces of  $\mathcal{K}^{st}$ . Then  $((\sigma_1, \dots, \sigma_n), t) = \text{compress}((\sigma_1^*, \dots, \sigma_n^*))$  if and only if  $(\sigma_1^*, \dots, \sigma_n^*) = \text{expand}((\sigma_1, \dots, \sigma_n), t)$ .*

Abusing notation, we use  $\Pi' = \text{expand}(\Pi, t)$  for trace assignments  $\Pi$  in  $\mathcal{K}$ , trajectory  $t$  an  $\Pi'$  in  $\mathcal{K}^{st}$  if  $(\Pi'(\pi_1), \dots, \Pi'(\pi_n)) = \text{expand}((\Pi(\pi_1), \dots, \Pi(\pi_n)), t)$ . We define  $(\Pi, t) = \text{compress}(\Pi')$  analogously.

We first prove some auxiliary lemmas.

**Lemma 2.**  $(\Pi, t) \models \psi_{ph}$  if and only if  $\text{expand}(\Pi, t) \models (\Box \text{phase} \wedge \psi_{ph})$ .

*Proof.* The proof (in both directions) proceeds by a simple induction on the steps  $k$  of the traces.  $\square$

**Lemma 3.** Let  $\Pi'$  and  $\Pi''$  be stuttering equivalent with  $\Pi' \models \Box \text{phase}$  and  $\Pi' \models \Box \text{phase}$ . Then  $\Pi' \models \psi_{ph}$  if and only if  $\Pi'' \models \psi_{ph}$ .

*Proof.* The proof proceeds by proving that the valuations of  $\psi_{ph}$  at color-changing positions and non color-changing positions coincide for all atomic phase formulas.  $\square$

**Lemma 4.** Let  $\Pi'$  be such that  $\Pi' \models \text{phase } \mathcal{U} (\text{missalign} \vee \text{block})$  and let  $\Pi''$  be such that  $\Pi'$  and  $\Pi''$  are stuttering equivalent. Then  $\Pi'' \not\models \Box \text{phase}$ .

*Proof.* Since  $\Pi' \models \text{phase } \mathcal{U} (\text{missalign} \vee \text{block})$  let  $k$  be the position at which  $(\Pi' + k) \models \text{missalign} \vee \text{block}$ . We consider the two cases separately.

- $(\Pi' + k) \models \text{missalign}$ . Let  $\pi_i$  and  $\pi_j$  be the paths that are miss-aligned. This means that in from  $\Pi'$  and  $\Pi' + k$ ,  $\Pi'(\pi_i)$  and  $\Pi'(\pi_j)$  change  $P$ -phases the same number of times, and one of them (say  $\pi_i$ ) will not change color and the other will. If  $\Pi''$  is stuttering equivalent to  $\Pi'$ , then  $\Pi''(\pi_i)$  and  $\Pi''(\pi_j)$  will also change a different colors a different number of times, which contradicts that  $\Pi'' \models \Box \text{phase}$ .
- $(\Pi' + k) \models \text{block}$ . In this case, the cycle implies each trace involved performs at  $k$ , a fix-number of color-changes at its  $k$ -th position for each of its each two atomic phase formulas, and that the circular constraint is mutually unsatisfiable. If  $\Pi''$  is stuttering equivalent to  $\Pi'$ ,  $\Box \text{phase}$  would imply that each trace involved in the cycle will be with the same number of color changes after reaching the position equivalent to  $k$  in  $\Pi'$ , therefore violating the constraint as well.  $\square$

A consequence of Lemma 2 and Lemma 4 is the following corollary.

**Corollary 3.**  $\Pi' \models \neg(\text{phase } \mathcal{U} (\text{missalign} \vee \text{block}))$  and  $(\Pi, t) = \text{compress}(\Pi')$  then there is no trajectory  $t_2$  for which  $(\Pi, t_2) \models \psi_{ph}$ .

*Proof.* Otherwise, Lemma 2 would imply that  $\text{expand}(\Pi, t_2) \models \Box \text{phase} \wedge \psi_{ph}$  and  $\text{expand}(\Pi, t_2)$  is a stuttering equivalent to  $\Pi'$ .  $\square$

We are ready to prove the main result.

**Theorem 1.** Let  $\mathcal{K}$  be a Kripke structure and  $\psi$  an admissible formula. Then  $\mathcal{K} \models \forall \pi_1 \dots \pi_n. E.\psi$  if and only if  $\mathcal{K}^{st} \models \forall \pi_1 \dots \pi_n. \psi_{sync}$ .



*Proof.* We prove the two directions separately.

- “ $\Rightarrow$ ”. Assume  $\mathcal{K} \models \forall \pi_1 \dots \pi_n. \mathbf{E}.\psi$ . Let  $\Pi'$  be an arbitrary trace assignment for  $\mathcal{K}^{st}$ . By contradiction, assume  $\Pi' \not\models \psi_{sync}$ . Let  $(\Pi, t) = \text{compress}(\Pi')$ . By our assumption, there must exist a  $t_2$  such that  $(\Pi, t_2) \models \psi$ . Since the valuation for the state and monadic sub-formulas of  $\psi$  is the same for  $(\Pi, t)$ ,  $(\Pi, t_2)$  and  $\Pi'$ , it must be the case that  $\Pi' \not\models \psi'$  and  $(\Pi, t_2) \models \psi_{ph}$ . Now, if  $\Pi' \not\models \psi'$ , then either
  - $\Pi' \models \text{phase } \mathcal{U} \text{ (missalign } \vee \text{ block)}$ . In this case Corollary 3 contradicts that  $(\Pi, t_2) \models \psi_{ph}$ .
  - $\Pi' \not\models \Box \text{phase} \rightarrow \psi_{ph}$ . In other words  $\Pi' \models \Box \text{phase}$  but  $\Pi' \not\models \psi_{ph}$ . However, Lemma 3 implies that  $\Pi' \models \psi_{ph}$  because  $\text{expand}(\Pi, t_2)$  and  $\Pi'$  are stuttering equivalent and  $\text{expand}(\Pi, t_2) \models \Box \text{phase} \wedge \psi_{ph}$ .
- “ $\Leftarrow$ ”. We assume now that  $\mathcal{K}^{st} \models \forall \pi_1 \dots \pi_n. \psi_{sync}$ . Let  $\Pi$  be an arbitrary trace assignment of  $\mathcal{K}$  and let  $\Pi'$  be an arbitrary fair stuttering expansion. Then, by assumption  $\Pi' \models \psi_{sync}$ .
  - If  $\Pi' \not\models \psi'$  then (all state and monadic sub-formulas of  $\psi$  coincide for  $\Pi$  and  $\Pi'$ ), then  $(\Pi, t) \models \psi$  for any  $t$  and we are done.
  - If  $\Pi' \models \psi'$ , in particular  $\Pi' \models \Box \text{phase} \wedge \psi_{ph}$ . Then, by Lemma 2,  $\text{compress}(\Pi') \models \psi_{ph}$  and  $\text{compress}(\Pi') = (\Pi, t)$  which witnesses the appropriate trajectory  $t$ .  $\square$

**Theorem 2.** *Let  $\mathcal{K}$  be a Kripke structure and  $\psi$  an admissible formula. Then  $\mathcal{K} \models \exists \pi_1 \dots \pi_n. \mathbf{E}.\psi$  if and only if  $\mathcal{K}^{st} \models \exists \pi_1 \dots \pi_n. \psi_{sync}$ .*

*Proof.* This follows directly from Lemma 2.  $\square$

**Theorem 3.** *Model-checking  $\forall^*$  or  $\exists^*$  admissible and co-admissible formulas is decidable both for formulas with  $\mathbf{E}$  and formulas with  $\mathbf{A}$ .*