

Which fragments of the interval temporal logic HS are tractable in model checking? ☆

Laura Bozzelli^a, Alberto Molinari^{b,*}, Angelo Montanari^b, Adriano Peron^a,
Pietro Sala^c

^a Department of Electronic Engineering and Information Technologies, University of Napoli “Federico II”, Italy

^b Department of Mathematics, Computer Science, and Physics, University of Udine, Italy

^c Department of Computer Science, University of Verona, Italy

A B S T R A C T

Since the 80s, model checking (MC) has been applied to the automatic verification of hardware/software systems. Point-based temporal logics, such as LTL, CTL, CTL*, and the like, are commonly used in MC as the specification language; however, there are some inherently interval-based properties of computations, e.g., temporal aggregations and durations, that cannot be properly dealt with by these logics, as they model a state-by-state evolution of systems.

Recently, an MC framework for the verification of interval-based properties of computations, based on Halpern and Shoham’s interval temporal logic (HS, for short) and its fragments, has been proposed and systematically investigated. In this paper, we focus on the boundaries that separate tractable and intractable HS fragments in MC. We first prove that MC for the logic BE of Allen’s relations *started-by* and *finished-by* is provably intractable, being EXPSpace-hard. Such a lower bound immediately propagates to full HS. Then, in contrast, we show that other noteworthy HS fragments, i.e., the logic $A\bar{A}B\bar{B}$ (resp., $A\bar{A}E\bar{E}$) of Allen’s relations *meets*, *met-by*, *starts* (resp., *finishes*), and *started-by* (resp., *finished-by*), are well-behaved, and turn out to have the same complexity as LTL (PSPACE-complete). Halfway are the fragments $A\bar{A}B\bar{B}E$ and $A\bar{A}E\bar{B}E$, whose EXPSpace membership and PSPACE hardness are already known. Here, we give an original proof of EXPSpace membership, that substantially simplifies the complexity of the constructions previously used for such a result. Contraction techniques—suitably tailored to each HS fragment—are at the heart of our results, enabling us to prove a pair of remarkable small-model properties.

Keywords:

Model checking

Interval temporal logics

Complexity

1. Introduction

In the context of automatic system verification, model checking (MC) is one of the most widely applied techniques. It allows a user to automatically check whether or not some desired properties of a system, specified by a temporal logic formula, hold over a model of the system, usually represented by a Kripke structure, through an exhaustive enumeration

☆ This paper gives an organic and systematic account of the results originally presented in the conference papers [6] and [7].

* Corresponding author.

E-mail addresses: lr.bozzelli@gmail.com (L. Bozzelli), molinari.alberto@gmail.com (A. Molinari), angelo.montanari@uniud.it (A. Montanari),

of all the states reachable by the computations of the system. If some of them is not fulfilled, MC algorithms produce a counterexample, extremely useful for debugging purposes. MC has been applied in a variety of application domains, including, e.g., planning [16], communication and security protocols [2,3], embedded reactive systems [13], computer device drivers [40], testing of railway control systems [4,33], and verification of clinical guidelines [15].

Point-based temporal logics, such as LTL, CTL, and CTL* [14,34]—that allow one to state properties of computation states—are usually adopted in MC as the specification language, as they are easy to understand, and suitable for many practical purposes. However, some relevant temporal properties, like those involving actions with duration, accomplishments, and temporal aggregations, are inherently “interval-based” and thus cannot be properly expressed by point-based temporal logics. This is the case, for instance, of statements like “the robot has to get back to the base in 10 minutes” and “the average speed of the moving device cannot exceed the established threshold”. Here, interval temporal logics (ITLs) come into play. In this paper, we focus on MC algorithms for ITLs.

ITLs take intervals, instead of points, as their primitive entities [19,32,38,39]. They have been applied in a variety of fields, including artificial intelligence (reasoning about action and change, qualitative reasoning, planning, configuration and multi-agent systems, and computational linguistics [5,16,25,35]), theoretical computer science (formal verification [32,41]), and databases (temporal and spatio-temporal databases [17]). Halpern and Shoham’s modal logic of time intervals HS [19] is probably the most famous ITL. It features one modality for each of the 13 possible ordering relations between pairs of intervals (the so-called Allen’s relations [1]), apart from equality. The problem of satisfiability checking for HS was proved to be undecidable for all relevant (classes of) linear orders [19], and most HS fragments are undecidable as well [9,21, 26]. However the logic of temporal neighborhood $\overline{A\overline{A}}$ and the logic of sub-intervals D [10–12,31] are noteworthy decidable exceptions.

1.1. Main contributions

In this paper, we address and solve some open issues about the MC problem for HS, which has recently entered the research agenda [22–24,27–30]. In order to check interval properties of computations, one needs to collect information about states into computation stretches: each finite path (trace) of a Kripke structure is interpreted as an interval, whose labeling is defined on the basis of the component states.

In [27], Molinari et al. dealt with MC for full HS over finite Kripke structures (under the homogeneity assumption [36], according to which a proposition letter holds in a trace if and only if it holds at each occurring state). They introduced the problem and proved that it is nonelementarily decidable and PSPACE-hard. Later, several fragments of HS were studied, which, similarly to what happens with satisfiability, exhibit better computational complexities. Here, we focus on the border between computationally good and bad HS fragments.

We first show that the combined use of modalities for interval prefixes B and suffixes E (modalities for Allen’s relations *started-by* and *finished-by*, respectively) is critical. More precisely, we prove that MC for the HS fragment BE, whose modalities can express properties of both interval prefixes and suffixes simultaneously, is EXPSpace-hard, and this lower bound immediately propagates to full HS. The result is obtained by a polynomial-time reduction from a domino-tiling problem for grids with rows of single exponential length [20] to the MC problem for BE.

Then, we show that the complexity of MC for HS fragments where properties of prefixes and suffixes of intervals are dealt with separately is markedly lower.

In [30], the authors proved that the MC problem for the logic $\overline{A\overline{A}}$ of temporal neighborhood, which only considers properties of future and past intervals (using modalities for meet A and met-by \overline{A} , respectively) is in P^{NP}. Moreover, as shown in [28], the addition of modalities for the transposed relations \overline{B} (for B) and \overline{E} (for E) to $\overline{A\overline{A}}$, respectively allowing for interval extensions to the right and to the left, results in the fragment $\overline{A\overline{A}\overline{B}\overline{E}}$ whose MC problem is PSPACE-complete.

Here, we prove that MC for the HS fragment $\overline{A\overline{A}\overline{B}\overline{B}}$ (resp., $\overline{A\overline{A}\overline{E}\overline{E}}$), that allows one to express properties of future and past intervals, interval prefixes (resp., suffixes), and right (resp., left) interval extensions, is in PSPACE. Since MC for the HS fragment featuring only one modality for right (resp., left) interval extensions is PSPACE-hard [30], PSPACE-completeness immediately follows.

Moreover, as a “byproduct”, we show that if we restrict HS to modalities either for interval prefixes B or for interval suffixes E only, MC turns out to be co-NP-complete. These results are achieved by means of a small-model property based on the notion of *induced trace*. Intuitively, given a trace (finite path) ρ in a Kripke structure and a formula φ of $\overline{A\overline{A}\overline{B}\overline{B}}$ / $\overline{A\overline{A}\overline{E}\overline{E}}$, we prove that it is always possible to build, by iteratively contracting the trace ρ , another (induced) trace, whose length is *polynomially bounded* in the size of the formula and the Kripke structure, which preserves the satisfiability of φ with respect to ρ .

The lower bound for BE MC shows that there is no way to provide an MC algorithm for the extension of $\overline{A\overline{A}\overline{B}\overline{B}}$ with E (resp., of $\overline{A\overline{A}\overline{E}\overline{E}}$ with B) with a “good” computational complexity. The picture is not so clear for the extension of $\overline{A\overline{A}\overline{B}\overline{B}}$ with \overline{E} (resp., $\overline{A\overline{A}\overline{E}\overline{E}}$ with \overline{B}).

The membership of $\overline{A\overline{A}\overline{B}\overline{B}\overline{E}}$ (resp., $\overline{A\overline{A}\overline{E}\overline{E}\overline{B}}$) to EXPSpace was already shown in [29] and the PSPACE-hardness of MC for $\overline{A\overline{A}\overline{B}\overline{B}}$ gives the best (unmatching) lower bound. The EXPSpace MC algorithm developed in [29] exploits a small exponential-size model property, that, given any trace of the Kripke structure and a fixed bound k on the nesting depth of B modalities in $\overline{A\overline{A}\overline{B}\overline{B}\overline{E}}$ formulas, allows one to find a *trace representative*, whose length is exponentially bounded by the size of the Kripke structure, which preserves satisfiability of the original trace for any $\overline{A\overline{A}\overline{B}\overline{B}\overline{E}}$ formula fulfilling the nesting depth constraint k .

As a matter of fact, the proof of the existence of the trace representative is rather involved and it exploits very technical arguments.

In this paper, we provide a much more understandable and compact proof of the membership to ExpSpace of the MC problem for $A\bar{A}B\bar{B}E$, which makes use of (a suitable extension of) the notion of induced traces we exploit in the proof of the small model property for $A\bar{A}\bar{B}\bar{B}$. We expect such a proof technique to be helpful in coping with the problem of finding a lower upper bound to the MC problem for BE , which is known to be nonelementarily decidable only.

1.2. Related work

In [22–24], Lomuscio and Michaliszyn addressed the MC problem for some fragments of HS extended with epistemic modalities. Their semantic assumptions are different from those made in [27], making it difficult to compare the two research lines. In both cases, HS formulas are evaluated over finite paths/intervals of a Kripke structure. However, while in [27] homogeneity is assumed, in [22,23] truth of proposition letters over an interval depends only on its endpoints.

In [22], they focused on the HS fragment BED of Allen’s relations *started-by*, *finished-by*, and *contains* (since modality $\langle D \rangle$ is definable in terms of modalities $\langle B \rangle$ and $\langle E \rangle$, BED is actually as expressive as BE), extended with epistemic modalities. They considered a *restricted form of MC (local MC)*, which checks the specification against a single (finite) initial computation interval. Their goal was indeed to reason about a given computation of a multi-agent system, rather than on all its admissible computations. They proved that the considered MC problem is PSPACE -complete. Moreover, they showed that the same problem restricted to the pure temporal fragment BED , that is, the one obtained by removing epistemic modalities, is in P . These results do not come as a surprise: modalities $\langle B \rangle$ and $\langle E \rangle$ allow one to access only sub-intervals of the initial one, whose number is quadratic in the length (number of states) of the initial interval.

In [23], they demonstrated that the picture drastically changes with other fragments of HS that allow one to access infinitely many intervals. In particular, they proved that the MC problem for the HS fragment $\bar{A}BL$ of Allen’s relations *meets*, *starts*, and *before* (since modality $\langle L \rangle$ is definable in terms of modality $\langle A \rangle$, $\bar{A}BL$ is actually as expressive as $\bar{A}\bar{B}$), extended with epistemic modalities, is decidable with a non-elementary upper bound. Note that, thanks to modalities $\langle A \rangle$ and $\langle \bar{B} \rangle$, formulas of $\bar{A}BL$ can possibly refer to infinitely many (future) intervals.

Finally, in [24], they showed how to use regular expressions in order to specify the way intervals of a Kripke structure get labeled. Such an extension leads to a significant increase in expressiveness, as the labeling of an interval is no more determined by that of its endpoints, but it depends on the ordered sequence of states the interval consists of. They proved that there is not a corresponding increase in computational complexity, as the complexity bounds given in [22,23] still hold with the new semantic variant: (local) MC for BED is still in PSPACE , and MC is non-elementarily decidable for $\bar{A}BL$.

1.3. Outline of the paper

The paper is organized as follows. In Section 2, we introduce the basic notions of the MC problem for HS (in particular Kripke structures, *abstract interval models*, and the interpretation of HS formulas over traces), and, then, we briefly summarize known complexity results about the MC problem for HS fragments. In Section 3, we state the ExpSpace -hardness of MC for the HS fragment BE by a reduction from a suitable domino-tiling problem. In Section 4, we first introduce the notion of induced trace and then we prove, via a contraction technique, a small polynomial-size model property for $A\bar{A}\bar{B}\bar{B}$ and $A\bar{A}E\bar{E}$, which allows us to devise a PSPACE MC algorithm for them. In addition, we consider the one-modality fragments B and E , and prove their co-NP-completeness. In Section 5, we focus on the fragment $A\bar{A}B\bar{B}E$ and the symmetric fragment $A\bar{A}E\bar{B}E$. We first define the equivalence relation of *trace bisimilarity*, and then we introduce the notion of *prefix sampling*. With these tools, we prove a small exponential-size model property for $A\bar{A}B\bar{B}E$ (and $A\bar{A}E\bar{B}E$), resulting into an easier and more compact proof of the membership of the two fragments to ExpSpace . We provide an assessment of the work done and outline future research directions in the conclusions.

2. Preliminaries

We start with some notation. Let \mathbb{N} be the set of natural numbers. For all $i, j \in \mathbb{N}$, with $i \leq j$, we denote by $[i, j]$ the set of $h \in \mathbb{N}$ such that $i \leq h \leq j$.

Let Σ be an alphabet and ρ be a finite word over Σ . We denote by $|\rho|$ the length of ρ . For all $1 \leq i \leq j \leq |\rho|$, $\rho(i)$ represents the i -th letter of ρ (we also say that i is a ρ -position), while $\rho(i, j)$ denotes the finite subword of ρ given by $\rho(i) \cdots \rho(j)$; ε is the empty word.

If $|\rho| = n$, we define $\text{fst}(\rho) = \rho(1)$ and $\text{lst}(\rho) = \rho(n)$. The sets of all proper prefixes and suffixes of ρ are $\text{Pref}(\rho) = \{\rho(1, i) \mid 1 \leq i \leq n - 1\}$ and $\text{Suff}(\rho) = \{\rho(i, n) \mid 2 \leq i \leq n\}$, respectively; ρ^i represents the suffix $\rho(i, |\rho|)$. (For the sake of convenience, we will sometimes write ρ^1 for ρ .)

The concatenation of two words ρ and ρ' is denoted by $\rho \cdot \rho'$. Moreover, if $\text{lst}(\rho) = \text{fst}(\rho')$, $\rho \star \rho'$ denotes $\rho(1, n - 1) \cdot \rho'$, with $n = |\rho|$ (\star -concatenation). In the following, when we write $\rho \star \rho'$, we implicitly assume that $\text{lst}(\rho) = \text{fst}(\rho')$.

Table 1
Allen's relations and corresponding HS modalities.

Allen relation	HS	Definition w.r.t. interval structures	Example
MEETS	$\langle A \rangle$	$[x, y] \mathcal{R}_A [v, z] \iff y = v$	
BEFORE	$\langle L \rangle$	$[x, y] \mathcal{R}_L [v, z] \iff y < v$	
STARTED-BY	$\langle B \rangle$	$[x, y] \mathcal{R}_B [v, z] \iff x = v \wedge z < y$	
FINISHED-BY	$\langle E \rangle$	$[x, y] \mathcal{R}_E [v, z] \iff y = z \wedge x < v$	
CONTAINS	$\langle D \rangle$	$[x, y] \mathcal{R}_D [v, z] \iff x < v \wedge z < y$	
OVERLAPS	$\langle O \rangle$	$[x, y] \mathcal{R}_O [v, z] \iff x < v < y < z$	

2.1. The interval temporal logic HS

In 1983, Allen proposed an interval algebra to reason about intervals and their relative order [1]; then, a few years later, Halpern and Shoham started a systematic logical study of interval representation and reasoning: they introduced the interval temporal logic HS, which features one modality for each Allen interval relation, but equality [19]. Table 1 depicts 6 of the 13 Allen's relations, together with the corresponding HS (existential) modalities. The other 7 relations are the 6 inverse relations (given a binary relation \mathcal{R} , the inverse relation $\overline{\mathcal{R}}$ is such that $b\overline{\mathcal{R}}a$ if and only if $a\mathcal{R}b$) and equality.

The language of HS consists of a set of proposition letters \mathcal{AP} , the constant \top (*true*), the Boolean connectives \neg and \wedge , and a temporal modality for each (non-trivial) Allen's relation, i.e., $\langle A \rangle$, $\langle L \rangle$, $\langle B \rangle$, $\langle E \rangle$, $\langle D \rangle$, $\langle O \rangle$, $\langle \overline{A} \rangle$, $\langle \overline{L} \rangle$, $\langle \overline{B} \rangle$, $\langle \overline{E} \rangle$, $\langle \overline{D} \rangle$, and $\langle \overline{O} \rangle$. HS formulas are formally defined by the grammar

$$\psi ::= \top \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \langle X \rangle \psi,$$

where $p \in \mathcal{AP}$ and $X \in \{A, L, B, E, D, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$. In the following, we will also exploit the other usual logical connectives (disjunction \vee , implication \rightarrow , and double implication \leftrightarrow) as abbreviations, and the constant *false* \perp for $\neg\top$. Furthermore, for any existential modality $\langle X \rangle$, the dual universal modality $[X]\psi$ is defined as $\neg\langle X \rangle\neg\psi$.

The joint nesting depth of B and E in a formula ψ , denoted by $d_{BE}(\psi)$, is defined as:

- $d_{BE}(p) = 0$, for any $p \in \mathcal{AP}$;
- $d_{BE}(\neg\psi) = d_{BE}(\psi)$;
- $d_{BE}(\psi \wedge \phi) = \max\{d_{BE}(\psi), d_{BE}(\phi)\}$;
- $d_{BE}(\langle X \rangle \psi) = 1 + d_{BE}(\psi)$, when $X = B$ or $X = E$;
- $d_{BE}(\langle X \rangle \psi) = d_{BE}(\psi)$, when $X \neq B$ and $X \neq E$.

Given any subset of Allen's relations $\{X_1, \dots, X_n\}$, we denote by $X_1 \dots X_n$ the HS fragment closed under Boolean connectives that features (existential and universal) modalities for X_1, \dots, X_n only. If we consider formulas ψ of HS fragments devoid of E (resp., B), the B -nesting depth (resp., E -nesting depth) of ψ , denoted as $d_B(\psi)$ (resp., $d_E(\psi)$), accounts for modality B (resp., E) only, and $d_B(\psi) = d_{BE}(\psi)$ (resp., $d_E(\psi) = d_{BE}(\psi)$).

W.l.o.g., we assume the *non-strict semantics* of HS, which admits intervals consisting of a single point.¹ Under such an assumption, all HS modalities can be expressed in terms of modalities $\langle B \rangle$, $\langle E \rangle$, $\langle \overline{B} \rangle$, and $\langle \overline{E} \rangle$ [38]. HS can thus be regarded as a multi-modal logic with these 4 primitive modalities and its semantics can be defined over a multi-modal Kripke structure, called *abstract interval model*, where intervals are treated as atomic objects and Allen's relations as binary relations between pairs of intervals. Since later we will focus on the HS fragments $A\overline{A}\overline{E}\overline{B}\overline{E}$, $A\overline{A}\overline{E}\overline{E}$ which do not feature $\langle B \rangle$, and $A\overline{A}\overline{B}\overline{B}\overline{E}$, $A\overline{A}\overline{B}\overline{B}$ which do not feature $\langle E \rangle$, we explicitly add both $\langle A \rangle$ and $\langle \overline{A} \rangle$ to the considered set of HS modalities.

Definition 2.1 (*Abstract interval model* [27]). An *abstract interval model* is a tuple $\mathcal{A} = (\mathcal{AP}, \mathbb{I}, A_{\mathbb{I}}, B_{\mathbb{I}}, E_{\mathbb{I}}, \sigma)$, where \mathcal{AP} is a set of proposition letters, \mathbb{I} is a possibly infinite set of atomic objects (worlds), $A_{\mathbb{I}}$, $B_{\mathbb{I}}$, and $E_{\mathbb{I}}$ are three binary relations over \mathbb{I} , and $\sigma : \mathbb{I} \mapsto 2^{\mathcal{AP}}$ is a (total) labeling function, assigning a set of proposition letters to each world.

In the interval setting, \mathbb{I} is interpreted as a set of intervals and $A_{\mathbb{I}}$, $B_{\mathbb{I}}$, and $E_{\mathbb{I}}$ as Allen's relations A (*meets*), B (*started-by*), and E (*finished-by*), respectively; σ assigns to each interval in \mathbb{I} the set of proposition letters that hold over it.

Given an abstract interval model $\mathcal{A} = (\mathcal{AP}, \mathbb{I}, A_{\mathbb{I}}, B_{\mathbb{I}}, E_{\mathbb{I}}, \sigma)$ and an interval $I \in \mathbb{I}$, the truth of an HS formula over I is inductively defined as follows:

- $\mathcal{A}, I \models \top$;
- $\mathcal{A}, I \models p$ iff $p \in \sigma(I)$, for any $p \in \mathcal{AP}$;

¹ All the results we prove in the paper hold for the strict semantics as well.



Fig. 1. The Kripke structure \mathcal{X}_Q .

- $\mathcal{A}, I \models \neg\psi$ iff it is not true that $\mathcal{A}, I \models \psi$ (also denoted as $\mathcal{A}, I \not\models \psi$);
- $\mathcal{A}, I \models \psi \wedge \phi$ iff $\mathcal{A}, I \models \psi$ and $\mathcal{A}, I \models \phi$;
- $\mathcal{A}, I \models \langle X \rangle \psi$, for $X \in \{A, B, E\}$, iff there is $J \in \mathbb{I}$ s.t. $I X_{\mathbb{I}} J$ and $\mathcal{A}, J \models \psi$;
- $\mathcal{A}, I \models \langle \bar{X} \rangle \psi$, for $\bar{X} \in \{\bar{A}, \bar{B}, \bar{E}\}$, iff there is $J \in \mathbb{I}$ s.t. $J X_{\mathbb{I}} I$ and $\mathcal{A}, J \models \psi$.

2.2. Kripke structures and abstract interval models

In the context of MC, finite state systems are usually modeled as finite Kripke structures. In [27], the authors define a mapping from Kripke structures to abstract interval models, that allows one to specify interval properties of computations by means of HS formulas.

Definition 2.2 (*Kripke structure and trace*). A Kripke structure is a tuple $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$, where \mathcal{AP} is a set of proposition letters, W is a set of states, $\delta \subseteq W \times W$ is a left-total relation between pairs of states, $\mu : W \mapsto 2^{\mathcal{AP}}$ is a total labeling function, and $w_0 \in W$ is the initial state. \mathcal{X} is finite if its set of states is finite.

A trace of \mathcal{X} is a non-empty finite word ρ over W such that $(\rho(i), \rho(i+1)) \in \delta$ for all $i \in [1, |\rho| - 1]$; $\text{states}(\rho)$ is the set of states occurring in ρ . A trace ρ is called *initial* if it starts from the initial state w_0 of \mathcal{X} , that is, $\text{fst}(\rho) = w_0$. We denote by $\text{Trc}_{\mathcal{X}}$ the set of all traces of \mathcal{X} .

Intuitively, for all $w \in W$, $\mu(w)$ is the set of proposition letters that hold on the state w , while δ is the transition relation that describes the evolution of the system over time.

Fig. 1 depicts the finite Kripke structure $\mathcal{X}_Q = (\{p, q\}, \{v_0, v_1\}, \delta, \mu, v_0)$, where $\delta = \{(v_0, v_0), (v_0, v_1), (v_1, v_0), (v_1, v_1)\}$, $\mu(v_0) = \{p\}$, and $\mu(v_1) = \{q\}$. A double circle identifies the initial state v_0 .

An abstract interval model (over $\text{Trc}_{\mathcal{X}}$) can be naturally associated with a Kripke structure \mathcal{X} by considering the set of intervals as the set of traces of \mathcal{X} . Since \mathcal{X} has loops (δ is left-total), the number of traces in $\text{Trc}_{\mathcal{X}}$, and thus the number of intervals, is infinite.

Definition 2.3 (*Induced abstract interval model* [27]). The abstract interval model induced by a finite Kripke structure $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$ is the tuple

$$\mathcal{A}_{\mathcal{X}} = (\mathcal{AP}, \mathbb{I}, A_{\mathbb{I}}, B_{\mathbb{I}}, E_{\mathbb{I}}, \sigma), \text{ where}$$

- $\mathbb{I} = \text{Trc}_{\mathcal{X}}$;
- $A_{\mathbb{I}} = \{(\rho, \rho') \in \mathbb{I} \times \mathbb{I} \mid \text{fst}(\rho) = \text{fst}(\rho')\}$;
- $B_{\mathbb{I}} = \{(\rho, \rho') \in \mathbb{I} \times \mathbb{I} \mid \rho' \in \text{Pref}(\rho)\}$;
- $E_{\mathbb{I}} = \{(\rho, \rho') \in \mathbb{I} \times \mathbb{I} \mid \rho' \in \text{Suff}(\rho)\}$;
- $\sigma : \mathbb{I} \mapsto 2^{\mathcal{AP}}$ is such that, for all $\rho \in \mathbb{I}$, $\sigma(\rho) = \bigcap_{w \in \text{states}(\rho)} \mu(w)$.

Relations $A_{\mathbb{I}}$, $B_{\mathbb{I}}$, and $E_{\mathbb{I}}$ are interpreted as the Allen's relations A , B , and E , respectively. Moreover, according to the definition of σ , $p \in \mathcal{AP}$ holds over $\rho = v_1 \cdots v_n$ if and only if it holds over all the states v_1, \dots, v_n of ρ . This conforms to the *homogeneity principle* [36], according to which a proposition letter holds over an interval if and only if it holds over all its subintervals.

Definition 2.4 (*Satisfiability over traces and model checking*). Let \mathcal{X} be a Kripke structure and ψ be an HS formula. We say that a trace $\rho \in \text{Trc}_{\mathcal{X}}$ satisfies ψ , denoted as $\mathcal{X}, \rho \models \psi$, if it holds that $\mathcal{A}_{\mathcal{X}}, \rho \models \psi$.

The Kripke structure \mathcal{X} is a model of ψ , denoted as $\mathcal{X} \models \psi$, if for all *initial* traces $\rho' \in \text{Trc}_{\mathcal{X}}$, it holds that $\mathcal{X}, \rho' \models \psi$. The *model checking* (MC) problem for HS consists of checking whether $\mathcal{X} \models \psi$ for a finite Kripke structure \mathcal{X} and an HS formula ψ .

As we already pointed out, MC is not trivially decidable since $\text{Trc}_{\mathcal{X}}$ is infinite.

Notice that the assumed semantics of HS is a state-based semantics which allows branching both in the future (by modalities A and \bar{B}) and in the past (by modalities \bar{A} and \bar{E}). For a comparison of different possible semantics for HS—i.e., state-based semantics, linear semantics, and computation-tree-based semantics—and an expressiveness comparison with standard point-based temporal logics LTL, CTL, and CTL*, we refer the reader to [8].




Fig. 2. The Kripke structure \mathcal{K}_{Sched} .

Example 2.5. In Fig. 2, we give an example of a finite Kripke structure \mathcal{K}_{Sched} that models the behavior of a scheduler serving three processes which are continuously requesting the use of a common resource. The initial state (denoted by a double circle) is v_0 : no process is served in that state. In any other state v_i and \bar{v}_i , with $i \in \{1, 2, 3\}$, the i -th process is served (this is denoted by the fact that p_i holds in those states). For the sake of readability, edges are marked either by r_i , for *request*(i), or by u_i , for *unlock*(i). Edge labels do not have a semantic value, that is, they are neither part of the structure definition, nor proposition letters, and they are simply used to ease reference to edges. Process i is served in state v_i , then, after job completion, a transition u_i from v_i to \bar{v}_i is taken. The scheduler cannot serve the same process twice in two successive rounds and, for that reason, v_i is not directly reachable from \bar{v}_i . A transition r_j , with $j \neq i$, from \bar{v}_i to v_j is then taken and process j is served.

We now show how some meaningful properties to be checked against \mathcal{K}_{Sched} can be expressed in HS, and, in particular, by formulas of $A\bar{A}\bar{E}\bar{E}$ (a fragment that will be studied in detail in Section 4). In all formulas, we force the validity of the considered property over all legal computation sub-intervals by using modality $[E]$ (all computation sub-intervals are suffixes of at least one initial trace). The truth of the next statements can easily be checked (in the following $\langle E \rangle^k$ stands for k occurrences of modality $\langle E \rangle$):

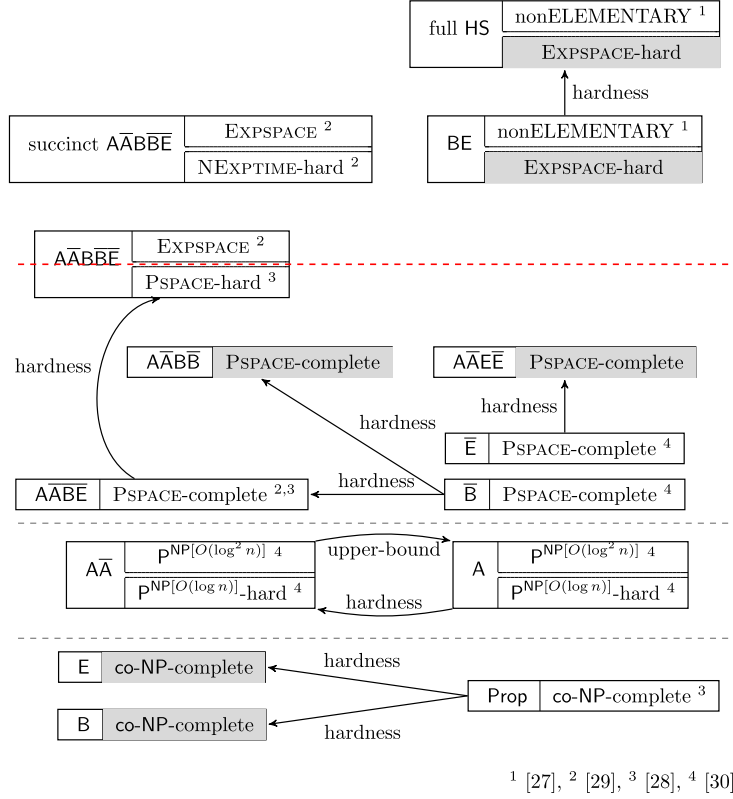
- $\mathcal{K}_{Sched} \models [E](\langle E \rangle^3 \top \rightarrow (\chi(p_1, p_2) \vee \chi(p_1, p_3) \vee \chi(p_2, p_3)))$, where $\chi(p, q)$ stands for $\langle E \rangle \langle \bar{A} \rangle p \wedge \langle E \rangle \langle \bar{A} \rangle q$;
- $\mathcal{K}_{Sched} \not\models [E](\langle E \rangle^{10} \top \rightarrow \langle E \rangle \langle \bar{A} \rangle p_3)$;
- $\mathcal{K}_{Sched} \not\models [E](\langle E \rangle^5 \rightarrow (\langle E \rangle \langle \bar{A} \rangle p_1 \wedge \langle E \rangle \langle \bar{A} \rangle p_2 \wedge \langle E \rangle \langle \bar{A} \rangle p_3))$.

The first formula states that in any suffix, having length at least 4, of an initial trace at least 2 proposition letters are witnessed. \mathcal{K}_{Sched} satisfies the formula since a process cannot be executed twice in a row. The second formula states that in any suffix of an initial trace having length at least 11 process 3 is executed at least once in some internal states (*non-starvation*). \mathcal{K}_{Sched} does not satisfy the formula since the scheduler can avoid executing a process ad libitum. The third formula requires that in any suffix of an initial trace having length at least 6 p_1, p_2, p_3 are all witnessed. The only way to satisfy this property is to constrain the scheduler to execute the three processes in a strictly periodic manner (*strict alternation*), that is, $p_i p_j p_k p_i p_j p_k p_i p_j p_k \dots$, $i, j, k \in \{1, 2, 3\}, i \neq j \neq k \neq i$, but this is not the case. \square

2.3. The general picture

We now describe known and new complexity results about the MC problem for HS fragments (see Fig. 3 for a graphical account).

In [27], the authors show that, given a finite Kripke structure \mathcal{K} and a bound k on the structural complexity of HS formulas, that is, on the (joint) nesting depth of $\langle E \rangle$ and $\langle B \rangle$ modalities, it is possible to obtain a *finite* representation for $\mathcal{A}_{\mathcal{K}}$, which is equivalent to $\mathcal{A}_{\mathcal{K}}$ with respect to the fulfillment of HS formulas with structural complexity less than or equal to k . Then, by exploiting such a representation, they prove that the MC problem for (full) HS is decidable, providing an algorithm with non-elementary complexity. Moreover, they show that the problem for the fragment $A\bar{A}\bar{B}\bar{E}$, and thus for full HS, is PSPACE-hard (EXPSpace-hard if a suitable succinct encoding of formulas is exploited). In [29], the authors study the HS fragments $A\bar{A}\bar{B}\bar{B}\bar{E}$ and $A\bar{A}\bar{E}\bar{B}\bar{E}$, devising for both the fragments an EXPSpace MC algorithm which finds, for each trace of the input Kripke structure, a satisfaction-preserving trace of bounded exponential length, i.e., a *trace representative*. In this way, the algorithm needs to check only trace representatives instead of traces of unbounded length. Then, in [28], they prove that the problem for these two fragments is PSPACE-hard (if a suitable succinct encoding of formulas is exploited, the



¹ [27], ² [29], ³ [28], ⁴ [30]

Fig. 3. Complexity of the MC problem for HS fragments: known results are depicted in white boxes, new ones in gray boxes. In red: the line between tractable and intractable fragments. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

algorithm remains in EXPSpace, but a NEXPTIME lower bound can be given [29]). Finally, they show that formulas satisfying a constant bound on their B -nesting (resp., E -nesting) depth can be checked in polynomial working space [29].

In [28,30] the authors identify some well-behaved HS fragments, namely, $AAB\bar{E}$, B , \bar{E} , $A\bar{A}$, A , and \bar{A} , which are still expressive enough to capture meaningful interval properties of state-transition systems and whose MC problem has a computational complexity markedly lower than that of full HS. In particular, they prove that the problem is PSPACE-complete for the first three fragments, and in between $P^{NP[O(\log n)]}$ and $P^{NP[O(\log^2 n)]}$ [18,37] for the last three ($P^{NP[O(\log n)]}$ and $P^{NP[O(\log^2 n)]}$ are the complexity classes of the problems decided by deterministic polynomial-time Turing machines, making $O(\log n)$ and $O(\log^2 n)$ queries to an NP oracle, respectively).

The first contribution of this paper to the overall picture is the strengthening of the lower bound to the complexity of MC for full HS obtained by proving the EXPSpace-hardness of the fragment BE (Section 3).

Then, we prove PSPACE-completeness of the MC problem for $A\bar{A}\bar{B}\bar{E}$ and $A\bar{A}\bar{E}\bar{E}$, thus improving the previously-known EXPSpace upper bound given in [29] (Section 4). The surprising aspect of this result is that these two fragments are (from a computational point of view) as complex as the almost minimal fragments \bar{B} and \bar{E} .

Next, we show that the one-modality fragments B and E are in co-NP, and thus co-NP-complete, namely, they have the same computational complexity of Prop (the pure propositional fragment of HS) [28] (Section 4).

Finally, we focus on the MC problem for the fragments $A\bar{A}\bar{B}\bar{E}$ and $A\bar{A}\bar{E}\bar{E}$, and we prove in a much simpler way (compared to that of [29]) its membership to EXPSpace (Section 5).

3. EXPSpace-hardness of BE

In this section, we prove that the MC problem for the HS fragment BE, whose modalities can express properties of both interval prefixes and suffixes, is EXPSpace-hard. The result is obtained by a polynomial-time reduction from a domino-tiling problem for grids with rows of single exponential length [20] to the MC problem for BE. Since, MC for full HS is clearly at least as hard as MC for BE, such a lower bound immediately propagates to full HS, improving the known lower bound.

We start with the definition of the domino-tiling problem. An instance \mathcal{I} of a domino-tiling problem for grids with rows of single exponential length is a tuple $\mathcal{I} = (C, \Delta, n, d_{init}, d_{final})$, where C is a finite set of colors, $\Delta \subseteq C^4$ is a set of tuples $(c_{down}, c_{left}, c_{up}, c_{right})$ of four colors, called *domino-types*, $n > 0$ is a natural number encoded in *unary*, and $d_{init}, d_{final} \in \Delta$ are two distinguished domino-types (respectively, the initial and final domino-types). The *size* of \mathcal{I} is defined as $|C| + |\Delta| + n$.




Fig. 4. A (generic) instance of the domino-tiling problem, where d_i^j denotes $f(i, j)$.




Fig. 5. Encoding of a tiling as a word, where d_i^j denotes $f(i, j)$.

Intuitively, a tiling of a grid is a color labeling of the edges of each grid square element (see Fig. 4). Formally, a *tiling* of \mathcal{I} is a mapping $f : [0, k] \times [0, 2^n - 1] \rightarrow \Delta$, for some $k \geq 0$, that satisfies the following constraints:

- two adjacent cells in a row have the same color on the shared edge, namely, for all $(i, j) \in [0, k] \times [0, 2^n - 2]$, $[f(i, j)]_{right} = [f(i, j + 1)]_{left}$;
- two adjacent cells in a column have the same color on the shared edge, namely, for all $(i, j) \in [0, k - 1] \times [0, 2^n - 1]$, $[f(i, j)]_{up} = [f(i + 1, j)]_{down}$;
- $f(0, 0) = d_{init}$ (initialization) and $f(k, 2^n - 1) = d_{final}$ (acceptance).

It is well-known that checking the existence (resp., the non-existence) of a tiling of \mathcal{I} is an EXPSpace-complete problem [20].

We now show how the domino-tiling problem can be reduced in polynomial time to the MC problem for BE. In particular, we show how to build in polynomial time a finite Kripke structure $\mathcal{K}_{\mathcal{I}}$ and a BE formula $\varphi_{\mathcal{I}}$ such that there exists an initial trace of $\mathcal{K}_{\mathcal{I}}$ satisfying $\varphi_{\mathcal{I}}$ if and only if there exists a tiling of \mathcal{I} . Hence, $\mathcal{K}_{\mathcal{I}} \models \neg \varphi_{\mathcal{I}}$ if and only if there is no tiling of \mathcal{I} .

The encoding of tilings exploits the set of proposition letters $\mathcal{AP} = \Delta \cup \{\$, 0, 1\}$. Proposition letters in $\{0, 1\}$ are used for the binary encoding of the value of an n -bit counter numbering the cells of one row of a tiling, while the proposition letter $\$$ is used as a separator.

In particular, a cell with content $d \in \Delta$ and column number $j \in [0, 2^n - 1]$ is encoded by the word of length $n + 1$ over \mathcal{AP} given by $db_1 \dots b_n$, where $b_1 \dots b_n$ is the binary encoding of the column number j (b_n is the most significant bit). A row is then represented by the word listing the encodings of cells from left to right, and a tiling f consisting of $k + 1$ rows is encoded by the finite word $r_0 \$ r_1 \dots \$ r_k$, where r_i is the encoding of the i -th row of f , for all $i \in [0, k]$. See Fig. 5 for a graphical account of a word encoding of a tiling.

The Kripke structure $\mathcal{K}_{\mathcal{I}}$ is trivially defined as

$$\mathcal{K}_{\mathcal{I}} = (\mathcal{AP}, \mathcal{AP}, \mathcal{AP} \times \mathcal{AP}, \mu, d_{init}),$$

where $\mu(p) = \{p\}$, for any $p \in \mathcal{AP}$. Thus, the initial traces of $\mathcal{K}_{\mathcal{I}}$ correspond to the finite words over \mathcal{AP} which start with the initial domino type d_{init} .

In order to build the BE formula $\varphi_{\mathcal{I}}$, we use some auxiliary formulas, namely, $length_i$, $beg(p)$, $end(p)$, ϕ_{cell} , and $\theta_j(b, b')$, where $i \in [1, 2n + 2]$, $j \in [2, n + 1]$, $p \in \mathcal{AP}$, and $b, b' \in \{0, 1\}$.

The formula $length_i$ has size linear in i and it characterizes the traces of length i . It can be expressed as follows:

$$length_i := (\underbrace{\langle B \rangle \dots \langle B \rangle}_{i-1} \top) \wedge (\underbrace{[B] \dots [B]}_i \perp).$$

The formula $beg(p)$ (resp., $end(p)$) captures the traces of \mathcal{K} which start (resp., end) in the state p :

$$beg(p) := (p \wedge length_1) \vee \langle B \rangle (p \wedge length_1), \quad end(p) := (p \wedge length_1) \vee \langle E \rangle (p \wedge length_1).$$




Fig. 6. Encoding of a trace ρ starting with a cell $c = (d1001)$ and ending with a cell $c' = (d'1101)$ (here, $n = 4$). The formula $\theta_2(1, 1)$ is satisfied by ρ , while $\theta_3(1, 0)$ is not.

The formula ϕ_{cell} captures the traces of $\mathcal{X}_{\mathcal{T}}$ which encode cells:

$$\phi_{\text{cell}} := \text{length}_{n+1} \wedge \left(\bigvee_{d \in \Delta} \text{beg}(d) \right) \wedge [E](\text{beg}(0) \vee \text{beg}(1)).$$

Finally, for all $j \in [2, n+1]$ and $b, b' \in \{0, 1\}$, the formula $\theta_j(b, b')$ is defined as:

$$\theta_j(b, b') := \langle B \rangle (\text{length}_j \wedge \text{end}(b)) \wedge \langle E \rangle (\text{length}_{n-j+2} \wedge \text{beg}(b')).$$

The formula $\theta_j(b, b')$ is satisfied by a trace ρ if $|\rho| \geq j+1$, $|\rho| \geq n-j+3$, $\rho(j) = b$, and $\rho(|\rho| - n + j - 1) = b'$. In particular, for a trace ρ starting with a cell c and ending with a cell c' , $\theta_j(b, b')$ is satisfied by ρ if the $(j-1)$ -th bit of c is b and the $(j-1)$ -th bit of c' is b' . See Fig. 6 for an example.

Additionally, we use the derived operator $\langle G \rangle$ and its dual $[G]$, which allow us to select arbitrary substraces of the given trace, including the trace itself:

$$\langle G \rangle \psi := \psi \vee \langle B \rangle \psi \vee \langle E \rangle \psi \vee \langle B \rangle \langle E \rangle \psi.$$

The formula $\varphi_{\mathcal{I}}$ is defined as follows:

$$\varphi_{\mathcal{I}} := \varphi_b \wedge \varphi_{\text{req}} \wedge \varphi_{\text{inc}} \wedge \varphi_{rr} \wedge \varphi_{rc}.$$

The conjunct φ_b checks that the given trace starts with a cell with content d_{init} and column number 0, and ends with a cell with content d_{final} and column number $2^n - 1$:

$$\varphi_b := \langle B \rangle \phi_{\text{cell}} \wedge \text{beg}(d_{\text{init}}) \wedge \langle E \rangle (\phi_{\text{cell}} \wedge \text{beg}(d_{\text{final}})) \wedge \bigwedge_{j=2}^{n+1} \theta_j(0, 1).$$

The conjunct φ_{req} ensures the following two requirements: (i) each occurrence of $\$$ in the given trace is followed by a cell with column number 0 and (ii) each cell c in the given trace is followed either by another cell, or by the separator $\$$, and in the latter case c has column number $2^n - 1$. The first requirement is encoded by the formula:

$$[G]((\text{length}_{n+2} \wedge \text{beg}(\$)) \longrightarrow \langle E \rangle (\phi_{\text{cell}} \wedge [E]\text{beg}(0)));$$

the second one by the formula:

$$[G] \left\{ (\text{length}_{n+2} \wedge \bigvee_{d \in \Delta} \text{beg}(d)) \longrightarrow \left(\langle B \rangle \phi_{\text{cell}} \wedge (\text{end}(\$) \vee \bigvee_{d \in \Delta} \text{end}(d)) \wedge (\text{end}(\$) \longrightarrow [E](\text{beg}(\$) \vee \text{beg}(1))) \right) \right\}.$$

The conjunct φ_{inc} checks that adjacent cells along the given trace have consecutive column numbers:

$$\varphi_{\text{inc}} := [G] \left(\phi_{\text{two_cells}} \longrightarrow \bigvee_{j=2}^{n+1} [\theta_j(0, 1) \wedge \bigwedge_{h=2}^{j-1} \theta_h(1, 0) \wedge \bigwedge_{h=j+1}^{n+1} \bigvee_{b \in \{0, 1\}} \theta_h(b, b)] \right),$$

where $\phi_{\text{two_cells}}$ is given by $\text{length}_{2n+2} \wedge \langle B \rangle \phi_{\text{cell}} \wedge \langle E \rangle \phi_{\text{cell}}$.

Note that φ_{req} and φ_{inc} ensure that column numbers are correctly encoded.

The conjunct φ_{rr} checks that adjacent cells in a row have the same color on the shared edge:

$$\varphi_{rr} := [G] \left(\phi_{\text{two_cells}} \longrightarrow \bigvee_{(d, d') \in \Delta \times \Delta \mid d_{\text{right}} = d'_{\text{left}}} (\text{beg}(d) \wedge \langle E \rangle (\text{length}_{n+1} \wedge \text{beg}(d'))) \right).$$

Finally, the conjunct φ_{rc} checks that adjacent cells in a column have the same color on the shared edge. For this, it suffices to require that the following condition holds:

- for each subtrace of the given one containing exactly one occurrence of $\$$, starting with a cell c , and ending with a cell c' , if c and c' have the same column number, then $d_{\text{up}} = d'_{\text{down}}$, where d (resp., d') is the content of c (resp., c').




Fig. 7. A trace π induced by ρ .

Accordingly, the formula φ_{rc} is defined as follows, where we use the formulas $\theta_j(b, b)$, with $j \in [2, n + 1]$ and $b \in \{0, 1\}$, to express that c and c' have the same column number:

$$\varphi_{rc} := [G] \left\{ \left(\phi_{one}(\$) \wedge \langle B \rangle \phi_{cell} \wedge \langle E \rangle \phi_{cell} \wedge \bigwedge_{j=2}^{n+1} \bigvee_{b \in \{0,1\}} \theta_j(b, b) \right) \right. \\ \left. \longrightarrow \bigvee_{(d,d') \in \Delta \times \Delta \mid d_{up} = d'_{down}} (\text{beg}(d) \wedge \langle E \rangle (\text{length}_{n+1} \wedge \text{beg}(d'))) \right\},$$

where $\phi_{one}(\$)$ is defined as $(\langle B \rangle \text{end}(\$)) \wedge \neg(\langle B \rangle \text{end}(\$) \wedge \langle B \rangle \text{end}(\$))$.

It is worth pointing out that $\varphi_{\mathcal{I}}$ has size polynomial in the size of \mathcal{I} .

By construction, a trace ρ of $\mathcal{X}_{\mathcal{I}}$ satisfies $\varphi_{\mathcal{I}}$ if and only if ρ encodes a tiling. Since the initial traces of $\mathcal{X}_{\mathcal{I}}$ are the finite words over \mathcal{AP} starting with d_{init} , it follows that there exists a tiling of \mathcal{I} if and only if there exists an initial trace of $\mathcal{X}_{\mathcal{I}}$ which satisfies $\varphi_{\mathcal{I}}$. The above-given reduction proves the following theorem.

Theorem 3.1. *The MC problem for BE formulas over finite Kripke structures is EXPSpace-hard (under polynomial-time reductions).*

4. A polynomial-size model-trace property for $A\bar{A}\bar{B}\bar{B}$ and $A\bar{A}\bar{E}\bar{E}$

In this section, we show that the MC problem for the fragments $A\bar{A}\bar{B}\bar{B}$ and $A\bar{A}\bar{E}\bar{E}$ is PSPACE-complete. Moreover, we prove that MC for the smaller fragments B and E is co-NP-complete.

We first prove the membership to PSPACE of the MC problem for the fragments $A\bar{A}\bar{B}\bar{B}$ and $A\bar{A}\bar{E}\bar{E}$ by showing that they enjoy a *polynomial-size model-trace property*, that is, we show that if a trace ρ of a finite Kripke structure \mathcal{X} satisfies a given formula φ of $A\bar{A}\bar{E}\bar{E}$ or $A\bar{A}\bar{B}\bar{B}$, then there exists a trace π , whose length is polynomial in the sizes of φ and \mathcal{X} , starting from and leading to the same states as ρ , that satisfies φ . In the following, we focus on the fragment $A\bar{A}\bar{E}\bar{E}$, being the case of the fragment $A\bar{A}\bar{B}\bar{B}$ completely symmetric.

Let $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$ be a finite Kripke structure. We start by introducing the basic notions of *induced trace* and *well-formed trace*, that will be extensively used to prove the polynomial-size model-trace property. Intuitively, we say that a trace $\pi \in \text{Trc}_{\mathcal{X}}$ is induced by a trace $\rho \in \text{Trc}_{\mathcal{X}}$ if it can be obtained from ρ by suitably contracting it, that is, by concatenating some substraces of ρ . Well-formedness adds a condition on the suffixes of an induced trace.

Definition 4.1 (*Induced and well-formed trace*). Let $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$ be a finite Kripke structure and let $\rho \in \text{Trc}_{\mathcal{X}}$, with $|\rho| = n$. We say that a trace $\pi \in \text{Trc}_{\mathcal{X}}$ is *induced* by ρ if there exists an increasing sequence of ρ -positions $i_1 < \dots < i_k$, with $i_1 = 1$ and $i_k = n$, such that $\pi = \rho(i_1) \dots \rho(i_k)$. For $j = 1, \dots, k$, the π -position j and the ρ -position i_j are called corresponding positions. Moreover, we say that an induced trace π is *well-formed* (with respect to ρ) if for all π -positions j , with corresponding ρ -positions i_j , and all proposition letters $p \in \mathcal{AP}$, it holds that $\mathcal{X}, \pi^j \models p$ if and only if $\mathcal{X}, \rho^{i_j} \models p$.

As an example, let us consider Fig. 7. The trace $\pi = \rho(1)\rho(4)\rho(5)\rho(7)\rho(10)$ is induced by ρ , provided that both ρ and π are traces of a Kripke structure \mathcal{X} , and the positions 1, 2, 3, 4, and 5 of π correspond to the positions $i_1 = 1$, $i_2 = 4$, $i_3 = 5$, $i_4 = 7$ and $i_5 = 10$ of ρ . Note that if π is induced by ρ , then $\text{fst}(\pi) = \text{fst}(\rho)$, $\text{lst}(\pi) = \text{lst}(\rho)$, and $|\pi| \leq |\rho|$ ($|\pi| = |\rho|$ if and only if $\pi = \rho$).

Well-formedness implies that the suffix of π starting from position j and that of ρ starting from the corresponding position i_j agree over all the proposition letters in \mathcal{AP} , that is, they have the same “satisfaction pattern” of proposition letters. In particular, for all $p \in \mathcal{AP}$, $\mathcal{X}, \pi \models p$ if and only if $\mathcal{X}, \rho \models p$. It can be easily checked that the well-formedness relation is *transitive*.

The following proposition shows how it is possible to contract a trace while preserving the same satisfaction pattern of proposition letters with respect to suffixes. Such a criterion represents a “basic step” in a contraction process that will allow us to prove the polynomial-size model-trace property.

Proposition 4.2. *Let $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$ be a finite Kripke structure. For any trace $\rho \in \text{Trc}_{\mathcal{X}}$, there exists a well-formed (with respect to ρ) trace $\pi \in \text{Trc}_{\mathcal{X}}$ such that $|\pi| \leq |W| \cdot (|\mathcal{AP}| + 1)$.*




Fig. 8. The contraction step of Proposition 4.2 (with $\mathcal{AP}(\rho, k) = \{p \in \mathcal{AP} \mid \mathcal{X}, \rho^k \models p\}$).

Proof. Let $\rho \in \text{Trc}_{\mathcal{X}}$, with $|\rho| = n$. If $n \leq |W| \cdot (|\mathcal{AP}| + 1)$, the thesis trivially holds. Let us assume $n > |W| \cdot (|\mathcal{AP}| + 1)$. We show that there exists $\pi \in \text{Trc}_{\mathcal{X}}$, with $|\pi| < n$, which is well-formed with respect to ρ .

Since $n > |W| \cdot (|\mathcal{AP}| + 1)$, there is some state $w \in W$ occurring in ρ at least $|\mathcal{AP}| + 2$ times. Assume that for all ρ -positions i and j , with $j > i$, if $\rho(i) = \rho(j) = w$, then there exists some $p \in \mathcal{AP}$ such that $\mathcal{X}, \rho^j \models p$ and $\mathcal{X}, \rho^i \not\models p$. This assumption leads to a contradiction, as the suffixes of ρ may feature at most $|\mathcal{AP}| + 1$ distinct satisfaction patterns of proposition letters (due to the homogeneity principle in Definition 2.3), while there are at least $|\mathcal{AP}| + 2$ occurrences of w . As a consequence, there are two ρ -positions i and j , with $j > i$, such that $\rho(i) = \rho(j) = w$ and, for all $p \in \mathcal{AP}$, $\mathcal{X}, \rho^j \models p$ if and only if $\mathcal{X}, \rho^i \models p$. (See Fig. 8 for a graphical account.) It is easy to see that $\pi = \rho(1, i) \star \rho(j, n) \in \text{Trc}_{\mathcal{X}}$ is well-formed with respect to ρ and $|\pi| < n$. If $|\pi| \leq |W| \cdot (|\mathcal{AP}| + 1)$, the thesis is proved; otherwise, the same basic step can be iterated a finite number of times, and the thesis follows by transitivity of the well-formedness relation. \square

The next definition identifies some distinguished positions in a trace, called *witness positions*. As we will see in the proof of Theorem 4.4, if we perform a contraction (see the proof of Proposition 4.2, and its graphical account in Fig. 8) between a pair of such positions, we get a trace which is equivalent to the original one with respect to the satisfiability of the considered $\text{A}\bar{\text{A}}\bar{\text{E}}\bar{\text{E}}$ formula. In the following, we restrict ourselves to formulas in *negation normal form* (NNF), namely, formulas where negation is applied only to proposition letters. By using De Morgan's laws and the dual modalities $[E]$, $[\bar{E}]$, $[A]$, and $[\bar{A}]$ of $\langle E \rangle$, $\langle \bar{E} \rangle$, $\langle A \rangle$, and $\langle \bar{A} \rangle$, respectively, we can trivially convert in linear time a formula into an equivalent one in NNF, having at most double length.

Definition 4.3 (*Witness position*). Let $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$ be a finite Kripke structure, $\rho \in \text{Trc}_{\mathcal{X}}$, and φ be a formula of $\text{A}\bar{\text{A}}\bar{\text{E}}\bar{\text{E}}$. Let us denote by $E(\varphi, \rho)$ the set of subformulas of the form $\langle E \rangle \psi$ of φ such that $\mathcal{X}, \rho \models \langle E \rangle \psi$. The set $Wt(\varphi, \rho)$ of *witness positions* of ρ for φ is the *minimal* set of ρ -positions satisfying the following constraint: for each $\langle E \rangle \psi \in E(\varphi, \rho)$, the greatest ρ -position $i > 1$ such that $\mathcal{X}, \rho^i \models \psi$ belongs to $Wt(\varphi, \rho)$.²

It is easy to see that the cardinalities of $E(\varphi, \rho)$ and of $Wt(\varphi, \rho)$ are at most $|\varphi| - 1$. We are now ready to prove the polynomial-size model-trace property.

Theorem 4.4 (*Polynomial-size model-trace property for $\text{A}\bar{\text{A}}\bar{\text{E}}\bar{\text{E}}$*). Let $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$ be a finite Kripke structure, $\rho, \sigma \in \text{Trc}_{\mathcal{X}}$, and φ be an $\text{A}\bar{\text{A}}\bar{\text{E}}\bar{\text{E}}$ formula in NNF such that $\mathcal{X}, \rho \star \sigma \models \varphi$. Then, there exists π , induced by ρ , such that $\mathcal{X}, \pi \star \sigma \models \varphi$, and $|\pi| \leq |W| \cdot (|\varphi| + 1)^2$.

As a preliminary remark, we note that the theorem holds in particular if $|\sigma| = 1$, and thus $\rho \star \sigma = \rho$ and $\pi \star \sigma = \pi$. In such a case, if $\mathcal{X}, \rho \models \varphi$, then $\mathcal{X}, \pi \models \varphi$, where π is induced by ρ and $|\pi| \leq |W| \cdot (|\varphi| + 1)^2$. The more general statement of Theorem 4.4 is needed for technical reasons in the soundness/completeness proofs of the algorithms for MC given in the following.

Proof. W.l.o.g., we restrict ourselves to the proposition letters occurring in φ , thus having $|\mathcal{AP}| \leq |\varphi|$. Let $Wt(\varphi, \rho \star \sigma)$ be the set of witness positions of $\rho \star \sigma$ for φ , let $\{i_1, \dots, i_k\}$ be the ordering of $Wt(\varphi, \rho \star \sigma)$ such that $i_1 < \dots < i_k$, and let $i_0 = 1$ and $i_{k+1} = |\rho \star \sigma|$. Hence, $1 = i_0 < i_1 < \dots < i_k \leq i_{k+1} = |\rho \star \sigma|$.

If $|\rho| \leq |W| \cdot (|\varphi| + 1)^2$, the thesis trivially holds. Let us assume that $|\rho| > |W| \cdot (|\varphi| + 1)^2$. We show that there exists a trace π induced by ρ , with $|\pi| < |\rho|$, such that $\mathcal{X}, \pi \star \sigma \models \varphi$.

W.l.o.g., we can assume that, for some $j \geq 0$, $i_0 < i_1 < \dots < i_j$ are ρ -positions, while $i_{j+1} < \dots < i_{k+1}$ are $(\rho \star \sigma)$ -positions not in ρ . Then, either (i) there exists $t \in [0, j - 1]$ such that $i_{t+1} - i_t > |W| \cdot (|\varphi| + 1)$ or (ii) $|\rho(i_j, |\rho|)| > |W| \cdot (|\varphi| + 1)$. By way of contradiction, suppose that neither (i) nor (ii) holds. We need to distinguish two cases. If $\rho \star \sigma = \rho$, then $|\rho| = (i_{k+1} - i_0) + 1 \leq (k + 1) \cdot |W| \cdot (|\varphi| + 1) + 1$; otherwise $(|\rho| < |\rho \star \sigma|)$, $|\rho| = (i_j - i_0) + |\rho(i_j, |\rho|)| \leq j \cdot |W| \cdot (|\varphi| + 1) + |W| \cdot (|\varphi| + 1) \leq (k + 1) \cdot |W| \cdot (|\varphi| + 1)$. The contradiction follows since $(k + 1) \cdot |W| \cdot (|\varphi| + 1) + 1 \leq |\varphi| \cdot |W| \cdot (|\varphi| + 1) + 1 \leq |W| \cdot (|\varphi| + 1)^2$.

Let $(\alpha, \beta) = (i_t, i_{t+1})$ in case (i), and $(\alpha, \beta) = (i_j, |\rho|)$ in case (ii) and let $\rho' = \rho(\alpha, \beta)$. In both cases, $|\rho'| > |W| \cdot (|\varphi| + 1) \geq |W| \cdot (|\mathcal{AP}| + 1)$, being $|\mathcal{AP}| \leq |\varphi|$. By Proposition 4.2, there exists a trace π' of \mathcal{X} , well-formed with respect

² Note that such a ρ -position exists by definition of $E(\varphi, \rho)$.




Fig. 9. Representation of the contraction step of Theorem 4.4—case (i).

to ρ' , such that $|\pi'| \leq |W| \cdot (|\mathcal{AP}| + 1) < |\rho'|$. Let π be the trace induced by ρ obtained by replacing the subtrace ρ' of ρ by π' (see Fig. 9 for a graphical account). Since $|\pi| < |\rho|$, it remains to prove that $\mathcal{X}, \pi \star \sigma \models \varphi$.

Let us denote $\pi \star \sigma$ by $\bar{\pi}$ and $\rho \star \sigma$ by $\bar{\rho}$. Let $H : [1, |\bar{\pi}|] \rightarrow [1, |\bar{\rho}|]$ be the function mapping positions of $\bar{\pi}$ into positions of $\bar{\rho}$ in such a way that positions “outside” π' , i.e., outside the interval $[\alpha, \alpha + |\pi'| - 1]$, are mapped into their original positions in $\bar{\rho}$ while those “inside” π' , i.e., in $[\alpha, \alpha + |\pi'| - 1]$, are mapped into the corresponding positions in $\bar{\rho}'$ (exploiting well-formedness of π' with respect to ρ'):

$$H(m) = \begin{cases} m & \text{if } m < \alpha; \\ \alpha + \ell_{m-\alpha+1} - 1 & \text{if } \alpha \leq m < \alpha + |\pi'|; \\ m + (|\rho'| - |\pi'|) & \text{if } m \geq \alpha + |\pi'|, \end{cases} \quad (1)$$

where ℓ_s is the ρ' -position corresponding to the π' -position s , with $\ell_s \in [1, |\rho'|]$ and $s \in [1, |\pi'|]$.

It is easy to check that the map H satisfies the following properties:

1. H is strictly monotonic, i.e., for all $j, j' \in [1, |\bar{\pi}|]$, $j < j'$ iff $H(j) < H(j')$;
2. for all $j \in [1, |\bar{\pi}|]$, $\bar{\pi}(j) = \bar{\rho}(H(j))$;
3. $H(1) = 1$ and $H(|\bar{\pi}|) = |\bar{\rho}|$;
4. $Wt(\varphi, \bar{\rho}) \subseteq \{H(j) \mid j \in [1, |\bar{\pi}|]\}$, i.e., all witness positions are preserved;
5. for each $j \in [1, |\bar{\pi}|]$ and $p \in \mathcal{AP}$, $\mathcal{X}, \bar{\pi}^j \models p$ iff $\mathcal{X}, \bar{\rho}^{H(j)} \models p$.

The fact that $\mathcal{X}, \bar{\pi} \models \varphi$ is an immediate consequence of the following claim, considering that $H(1) = 1$, $\mathcal{X}, \bar{\rho} \models \varphi$, $\bar{\rho}^1 = \bar{\rho}$, and $\bar{\pi}^1 = \bar{\pi}$.

Claim. For all $j \in [1, |\bar{\pi}|]$, all subformulas ψ of φ , and all $u \in \text{Trc}_{\mathcal{X}}$, it holds that if $\mathcal{X}, u \star \bar{\rho}^{H(j)} \models \psi$, then $\mathcal{X}, u \star \bar{\pi}^j \models \psi$.

Proof. Assume that $\mathcal{X}, u \star \bar{\rho}^{H(j)} \models \psi$. Note that $u \star \bar{\rho}^{H(j)}$ is defined if and only if $u \star \bar{\pi}^j$ is defined. We prove by induction on the structure of ψ that $\mathcal{X}, u \star \bar{\pi}^j \models \psi$. Since φ is in NNF, only the following cases occur:

- $\psi = p$ or $\psi = \neg p$ for some $p \in \mathcal{AP}$. By Property 5 of H , $\mathcal{X}, \bar{\pi}^j \models p$ iff $\mathcal{X}, \bar{\rho}^{H(j)} \models p$. Hence, $\mathcal{X}, u \star \bar{\pi}^j \models p$ if and only if $\mathcal{X}, u \star \bar{\rho}^{H(j)} \models p$, and the result holds.
- $\psi = \theta_1 \wedge \theta_2$ or $\psi = \theta_1 \vee \theta_2$, for some $\text{A}\bar{\text{A}}\bar{\text{E}}\bar{\text{E}}$ formulas θ_1 and θ_2 : the result directly follows from the inductive hypothesis.
- $\psi = [E]\theta$. We need to show that for each proper suffix η of $u \star \bar{\pi}^j$, $\mathcal{X}, \eta \models \theta$. We distinguish two cases:
 - η is not a proper suffix of $\bar{\pi}^j$. Hence, η is of the form $u^h \star \bar{\pi}^j$ for some $h \in [2, |u|]$. Since $\mathcal{X}, u \star \bar{\rho}^{H(j)} \models [E]\theta$, then $\mathcal{X}, u^h \star \bar{\rho}^{H(j)} \models \theta$. By the inductive hypothesis, $\mathcal{X}, u^h \star \bar{\pi}^j \models \theta$.
 - η is a proper suffix of $\bar{\pi}^j$. Hence, $\eta = \bar{\pi}^h$ for some $h \in [j+1, |\bar{\pi}|]$. By Property 1 of H , $H(h) > H(j)$, and since $\mathcal{X}, u \star \bar{\rho}^{H(j)} \models [E]\theta$, we have that $\mathcal{X}, \bar{\rho}^{H(h)} \models \theta$. By the inductive hypothesis, $\mathcal{X}, \bar{\pi}^h \models \theta$.
Therefore, $\mathcal{X}, u \star \bar{\pi}^j \models [E]\theta$.
- $\psi = \langle E \rangle \theta$. We need to show that there exists a proper suffix of $u \star \bar{\pi}^j$ satisfying θ . Since $\mathcal{X}, u \star \bar{\rho}^{H(j)} \models \psi$, there exists a proper suffix η' of $u \star \bar{\rho}^{H(j)}$ such that $\mathcal{X}, \eta' \models \theta$. We distinguish two cases:
 - η' is not a proper suffix of $\bar{\rho}^{H(j)}$. Hence, η' is of the form $u^h \star \bar{\rho}^{H(j)}$ for some $h \in [2, |u|]$. By the inductive hypothesis, $\mathcal{X}, u^h \star \bar{\pi}^j \models \theta$. Hence, $\mathcal{X}, u \star \bar{\pi}^j \models \langle E \rangle \theta$.
 - η' is a proper suffix of $\bar{\rho}^{H(j)}$. Hence, $\eta' = \bar{\rho}^i$ for some $i \in [H(j)+1, |\bar{\rho}|]$, and $\mathcal{X}, \bar{\rho}^i \models \theta$. Let i' be the greatest position of $\bar{\rho}$ such that $\mathcal{X}, \bar{\rho}^{i'} \models \theta$. Hence $i' \geq i$ and, by Definition 4.3, $i' \in Wt(\varphi, \bar{\rho})$. By Property 4 of H , $i' = H(h)$ for some $\bar{\pi}$ -position h . Since $H(h) > H(j)$, it holds that $h > j$ (Property 1). By the inductive hypothesis, $\mathcal{X}, \bar{\pi}^h \models \theta$, and we obtain that $\mathcal{X}, u \star \bar{\pi}^j \models \langle E \rangle \theta$.
Therefore, in both the cases, $\mathcal{X}, u \star \bar{\pi}^j \models \langle E \rangle \theta$.

- $\psi = [\bar{E}]\theta$ or $\psi = \langle \bar{E} \rangle \theta$: the thesis holds as a direct consequence of the inductive hypothesis.
- $\psi = [A]\theta$, $\psi = \langle A \rangle \theta$, $\psi = [\bar{A}]\theta$, or $\psi = \langle \bar{A} \rangle \theta$. Since $u \star \bar{\pi}^j$ and $u \star \bar{\rho}^{H(j)}$ start at the same state and lead to the same state (by Properties 2 and 3 of H), the thesis trivially follows. This concludes the proof of the claim. \square

Algorithm 1 $\text{ModCheck}(\mathcal{X}, \psi)$.

```
1: for all initial traces  $\sigma \in \text{Trc}_{\mathcal{X}}$  such that  $|\sigma| \leq |W| \cdot (2|\psi| + 3)^2$  do
2:   if  $\text{Check}(\mathcal{X}, \psi, \sigma) = 0$  then
3:     return 0: " $\mathcal{X}, \sigma \not\models \psi$ "
4: return 1: " $\mathcal{X} \models \psi$ "
```

◁ Counterexample found

Algorithm 2 $\text{Check}(\mathcal{X}, \psi, \sigma)$.

```
1: if  $\psi = p$ , for  $p \in \mathcal{AP}$  then
2:   if  $p \in \bigcap_{s \in \text{states}(\sigma)} \mu(s)$  then
3:     return 1 else return 0
4: else if  $\psi = \neg\varphi$  then
5:   return  $1 - \text{Check}(\mathcal{X}, \varphi, \sigma)$ 
6: else if  $\psi = \varphi_1 \wedge \varphi_2$  then
7:   if  $\text{Check}(\mathcal{X}, \varphi_1, \sigma) = 0$  then
8:     return 0
9:   else
10:    return  $\text{Check}(\mathcal{X}, \varphi_2, \sigma)$ 
11: else if  $\psi = \langle A \rangle \varphi$  then
12:   for all  $\pi \in \text{Trc}_{\mathcal{X}}$  such that  $\text{fst}(\pi) = \text{lst}(\sigma)$ , and  $|\pi| \leq |W| \cdot (2|\varphi| + 1)^2$  do
13:     if  $\text{Check}(\mathcal{X}, \varphi, \pi) = 1$  then
14:       return 1
15:   return 0
16: else if  $\psi = \langle E \rangle \varphi$  then
17:   for each proper suffix  $\pi$  of  $\sigma$  do
18:     if  $\text{Check}(\mathcal{X}, \varphi, \pi) = 1$  then
19:       return 1
20:   return 0
21: else if  $\psi = \langle \bar{E} \rangle \varphi$  then
22:   for all  $\pi \in \text{Trc}_{\mathcal{X}}$  s.t.  $\text{lst}(\pi) = \text{fst}(\sigma)$ , and  $2 \leq |\pi| \leq |W| \cdot (2|\varphi| + 1)^2$  do
23:     if  $\text{Check}(\mathcal{X}, \varphi, \pi \star \sigma) = 1$  then
24:       return 1
25:   return 0
26: ...
```

◁ $\psi = \langle \bar{A} \rangle \varphi$ is analogous to $\psi = \langle A \rangle \varphi$

We have shown that $\mathcal{X}, \bar{\pi} \models \varphi$, with $|\bar{\pi}| < |\rho|$. Now, if $|\bar{\pi}| \leq |W| \cdot (|\varphi| + 1)^2$, the thesis is proved; otherwise, the above contraction step can be iterated a finite number of times, until the bound is reached, proving the thesis of Theorem 4.4. \square

By exploiting the polynomial-size model-trace property stated by Theorem 4.4, it is easy to define a PSPACE MC algorithm for A \bar{A} E \bar{E} . The main MC procedure for A \bar{A} E \bar{E} formulas is $\text{ModCheck}(\mathcal{X}, \psi)$ (Algorithm 1). All the initial traces σ , obtained by visiting the unraveling of \mathcal{X} from w_0 up to depth $|W| \cdot (2|\psi| + 3)^2$, are checked with respect to ψ by the function $\text{Check}(\mathcal{X}, \psi, \sigma)$ (Algorithm 2) which decides whether $\mathcal{X}, \sigma \models \psi$. The Check function is iteratively called until either some initial trace is found that does not satisfy ψ or all bounded initial traces satisfy ψ (and thus $\mathcal{X} \models \psi$). The call of $\text{Check}(\mathcal{X}, \psi, \sigma)$ (Algorithm 2) decides whether $\mathcal{X}, \sigma \models \psi$ by recursively calling itself on the sub-formulas of ψ either over σ or over (bounded) traces obtained by unraveling \mathcal{X} forward (starting from $\text{lst}(\sigma)$) for occurrences of the modality $\langle A \rangle$ and backward (starting from $\text{fst}(\sigma)$) for occurrences of the modalities $\langle \bar{A} \rangle$ and $\langle \bar{E} \rangle$.

Note that the considered bound on the length of initial traces is actually $|W| \cdot (2|\psi| + 3)^2 \geq |W| \cdot (|\text{NNF}(\neg\psi)| + 1)^2$ (first line of the ModCheck procedure). The reason is that the correctness proof of the algorithm exploits the polynomial bound of Theorem 4.4 for the formula $\neg\psi$ that has to be reduced in NNF.

We now prove soundness and completeness of the proposed procedures starting from the function Check of Algorithm 2.

Lemma 4.5. *Let ψ be an A \bar{A} E \bar{E} formula, \mathcal{X} be a finite Kripke structure, and $\sigma \in \text{Trc}_{\mathcal{X}}$. Then, $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$ if and only if $\mathcal{X}, \sigma \models \psi$.*

Proof. The proof is by induction on the structure of ψ . The base case where $\psi = p$, for some $p \in \mathcal{AP}$, directly follows from the definition (line 2 of Algorithm 2). The cases in which $\psi = \neg\varphi$ and $\psi = \varphi_1 \wedge \varphi_2$ are also trivial and thus omitted. We focus on the remaining cases.

- $\psi = \langle A \rangle \varphi$. If $\mathcal{X}, \sigma \models \psi$, then there exists a trace $\rho \in \text{Trc}_{\mathcal{X}}$ such that $\text{lst}(\sigma) = \text{fst}(\rho)$ and $\mathcal{X}, \rho \models \varphi$. By Theorem 4.4, there exists a trace $\pi \in \text{Trc}_{\mathcal{X}}$, with $|\pi| \leq |W| \cdot (|\varphi| + 1)^2$ and $\text{fst}(\pi) = \text{fst}(\rho) (= \text{lst}(\sigma))$, such that $\mathcal{X}, \pi \models \varphi'$, where φ' is the NNF of φ . Thus, being $|\pi| \leq |W| \cdot (2|\varphi| + 1)^2$, such trace π is considered in the for-loop at line 12. By the inductive hypothesis, $\text{Check}(\mathcal{X}, \varphi, \pi) = 1$ and thus $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$. Vice versa, if $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$, then there exists a trace $\pi \in \text{Trc}_{\mathcal{X}}$, with $\text{lst}(\sigma) = \text{fst}(\pi)$, such that $\text{Check}(\mathcal{X}, \varphi, \pi) = 1$. By the inductive hypothesis, $\mathcal{X}, \pi \models \varphi$, hence $\mathcal{X}, \sigma \models \psi$.

- $\psi = \langle \bar{A} \rangle \varphi$ is analogous to the previous case.
- $\psi = \langle E \rangle \varphi$. If $\mathcal{X}, \sigma \models \psi$, there exists a trace $\pi \in \text{Suff}(\sigma)$ such that $\mathcal{X}, \pi \models \varphi$. By the inductive hypothesis, $\text{Check}(\mathcal{X}, \varphi, \pi) = 1$. Since all the proper suffixes of σ are checked (line 17), $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$.
Vice versa, if $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$, then for some $\pi \in \text{Suff}(\sigma)$, it holds that $\text{Check}(\mathcal{X}, \varphi, \pi) = 1$. By the inductive hypothesis $\mathcal{X}, \pi \models \varphi$ implying that $\mathcal{X}, \sigma \models \psi$.
- $\psi = \langle \bar{E} \rangle \varphi$. If $\mathcal{X}, \sigma \models \psi$, then there exists a trace $\rho \in \text{Trc}_{\mathcal{X}}$, with $|\rho| \geq 2$, such that $\mathcal{X}, \rho \star \sigma \models \varphi$. By Theorem 4.4, there exists a trace $\pi \in \text{Trc}_{\mathcal{X}}$ induced by ρ , with $|\pi| \leq |W| \cdot (|\varphi'| + 1)^2$, such that $\mathcal{X}, \pi \star \sigma \models \varphi'$, where φ' is the NNF of φ . Such trace π is considered in the for-loop at line 22, since $|\pi| \leq |W| \cdot (2|\varphi| + 1)^2$ and $|\pi| \geq 2$ as it is induced by ρ . By the inductive hypothesis, $\text{Check}(\mathcal{X}, \varphi, \pi \star \sigma) = 1$ implying that $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$.
Vice versa, if $\text{Check}(\mathcal{X}, \psi, \sigma) = 1$, then there exists a trace $\pi \in \text{Trc}_{\mathcal{X}}$, with $|\pi| \geq 2$, such that $\text{Check}(\mathcal{X}, \varphi, \pi \star \sigma) = 1$. By the inductive hypothesis, $\mathcal{X}, \pi \star \sigma \models \varphi$, hence $\mathcal{X}, \sigma \models \psi$. \square

We prove now soundness and completeness of Algorithm 1.

Theorem 4.6. *Let ψ be an $A\bar{A}\bar{E}\bar{E}$ formula and \mathcal{X} be a finite Kripke structure. Then, $\text{ModCheck}(\mathcal{X}, \psi) = 1$ if and only if $\mathcal{X} \models \psi$.*

Proof. (\Leftarrow) If $\mathcal{X} \models \psi$, then, for all initial traces $\rho \in \text{Trc}_{\mathcal{X}}$, we have that $\mathcal{X}, \rho \models \psi$. By Lemma 4.5, it follows that $\text{Check}(\mathcal{X}, \psi, \rho) = 1$. Now, since the for-loop at line 1 considers a subset of the initial traces, it holds that $\text{ModCheck}(\mathcal{X}, \psi) = 1$.

(\Rightarrow) If $\text{ModCheck}(\mathcal{X}, \psi) = 1$, then, for any initial trace ρ considered by the for-loop at line 1, that is, with $|\rho| \leq |W| \cdot (2|\psi| + 3)^2$, it holds that $\text{Check}(\mathcal{X}, \psi, \rho) = 1$. Let us assume by contradiction that $\mathcal{X} \not\models \psi$, that is, there exists an initial trace $\rho' \in \text{Trc}_{\mathcal{X}}$ such that $\mathcal{X}, \rho' \models \neg\psi$, or, equivalently, $\mathcal{X}, \rho' \models \bar{\psi}$, where $\bar{\psi}$ is the NNF of $\neg\psi$. Thus, by Theorem 4.4, there exists an initial trace σ , with $|\sigma| \leq |W| \cdot (|\bar{\psi}| + 1)^2 \leq |W| \cdot (2|\psi| + 3)^2$, such that $\mathcal{X}, \sigma \models \bar{\psi}$, namely, $\mathcal{X}, \sigma \not\models \psi$. By Lemma 4.5, it holds that $\text{Check}(\mathcal{X}, \psi, \sigma) = 0$, leading to a contradiction and proving that $\mathcal{X} \models \psi$. \square

The model checking procedures require *polynomial working space*, since:

- ModCheck needs to store only a trace no longer than $|W| \cdot (2|\psi| + 3)^2$ (obviously, many traces are generated while visiting the unraveling of \mathcal{X} , but only one at a time needs to be stored);
- every recursive call to Check (possibly) needs space for a trace no longer than $|W| \cdot (2|\varphi| + 1)^2$, where φ is a subformula of ψ such that $|\varphi| \leq |\psi| - 1$;
- at most one call to ModCheck and $|\psi|$ calls to Check can be simultaneously active.

Therefore, the maximum space needed by the given algorithms is $(|\psi| + 1) \cdot O(\log |W|) \cdot (|W| \cdot (2|\psi| + 3)^2)$ bits, where $O(\log |W|)$ bits are needed to represent a state of \mathcal{X} .

Theorem 4.6, along with the above space analysis and the fact that MC for the fragment \bar{E} is known to be PSPACE-hard [30], entail the following corollary.

Corollary 4.7. *The MC problem for $A\bar{A}\bar{E}\bar{E}$ formulas over finite Kripke structures is PSPACE-complete.*

The same result, that is, PSPACE-completeness, clearly holds also for any sub-fragment of $A\bar{A}\bar{E}\bar{E}$ which features the modality \bar{E} .

We now conclude the section by showing that the MC problem for the fragments B and E is in co-NP, that is, they have the same complexity as the purely propositional fragment Prop. We focus on the fragment E, as the case of B is completely symmetric. As we shall see, the MC algorithm heavily rests on the polynomial-size model-trace property.

The algorithm is based on the *non-deterministic* procedure $\text{CounterExE}(\mathcal{X}, \psi)$ (Algorithm 3) which searches for counterexamples to the input E formula ψ (initial traces satisfying $\neg\psi$). If such a counterexample is found, clearly $\mathcal{X} \not\models \psi$. First, the procedure generates in a *non-deterministic* way an initial trace ρ , whose length is at most $|W| \cdot (2|\psi| + 3)^2$, by means of $\underline{A_trace}(\mathcal{X}, w_0, |\psi|)$. Then, the *deterministic* function $\text{CheckE}(\mathcal{X}, \psi, \rho)$, reported in Algorithm 4, evaluates ψ over ρ . If CheckE returns \perp , a counterexample has been found and CounterExE returns **Yes** (thus the non-deterministic computation of the algorithm is successful). Otherwise, it returns **No** (the computation fails).

As for the function CheckE , the following statement holds.

Proposition 4.8. *Let ψ be an E formula, \mathcal{X} be a finite Kripke structure, and ρ be a trace of \mathcal{X} . Then, $\text{CheckE}(\mathcal{X}, \psi, \rho) = \top$ if and only if $\mathcal{X}, \rho \models \psi$.*

CheckE exploits a Boolean table T with an entry for each pair consisting of a subformula of ψ and the starting position of a suffix of ρ (the size of T is then $|\psi| \times |\rho|$). The function scans all the subformulas φ of the input ψ by increasing length, and it stores in the Boolean entry $T[\varphi, i]$, for $1 \leq i \leq |\rho|$, whether $\mathcal{X}, \rho^i \models \varphi$ or not. Note that the result of the evaluation of ψ over ρ is stored in $T[\psi, 1]$, as $\rho^1 = \rho$. Since subformulas of ψ are considered by increasing length order,

Algorithm 3 CounterExE(\mathcal{X}, ψ).

```
1:  $\rho \leftarrow \text{A\_trace}(\mathcal{X}, w_0, |\psi|)$   $\triangleleft$  a trace of  $\mathcal{X}$  from  $w_0$  of length  $\leq |W| \cdot (2|\psi| + 3)^2$ 
2: if CheckE( $\mathcal{X}, \psi, \rho$ ) =  $\perp$  then
3:   return Yes: " $\mathcal{X}, \rho \not\models \psi$ "  $\triangleleft$  Counterexample found
4: else
5:   return No: " $\mathcal{X}, \rho \models \psi$ "  $\triangleleft$  Counterexample not found
```

Algorithm 4 CheckE(\mathcal{X}, ψ, ρ).

```
1:  $T \leftarrow \text{New\_Bool\_Table}(|\psi|, |\rho|)$   $\triangleleft$  creates new table of  $|\psi| \times |\rho|$  Boolean entries
2: for all subformulas  $\varphi$  of  $\psi$  by increasing length do
3:   if  $\varphi = p$ , for  $p \in \mathcal{AP}$  then
4:      $T[p, |\rho|] \leftarrow p \in \mu(\text{lst}(\rho))$ 
5:     for  $i = |\rho| - 1, \dots, 1$  do
6:        $T[p, i] \leftarrow T[p, i + 1]$  and  $p \in \mu(\rho(i))$ 
7:   else if  $\varphi = \neg\varphi_1$  then
8:     for  $i = |\rho|, \dots, 1$  do
9:        $T[\varphi, i] \leftarrow \text{not } T[\varphi_1, i]$ 
10:  else if  $\varphi = \varphi_1 \wedge \varphi_2$  then
11:    for  $i = |\rho|, \dots, 1$  do
12:       $T[\varphi, i] \leftarrow T[\varphi_1, i]$  and  $T[\varphi_2, i]$ 
13:  else if  $\varphi = (E)\varphi_1$  then
14:     $T[\varphi, |\rho|] \leftarrow \perp$ 
15:    for  $i = |\rho| - 1, \dots, 1$  do
16:       $T[\varphi, i] \leftarrow T[\varphi, i + 1]$  or  $T[\varphi_1, i + 1]$ 
17: return  $T[\psi, 1]$ 
```

during an iteration starting at line 2, when a subformula φ of ψ is being processed, it holds that $T[\xi, i]$ is defined, for all other subformulas ξ processed in some previous iteration, and $T[\xi, i] = \top$ if and only if $\mathcal{X}, \rho^i \models \xi$. This implies that, at the end, $T[\psi, 1] = \top$ if and only if $\mathcal{X}, \rho \models \psi$.

We can now prove that the procedure CounterExE is sound and complete. If CounterExE(\mathcal{X}, ψ) has a successful computation, then there exists an initial trace ρ such that CheckE(\mathcal{X}, ψ, ρ) = \perp . This means that $\mathcal{X}, \rho \not\models \psi$, and thus $\mathcal{X} \not\models \psi$. Conversely, if $\mathcal{X} \not\models \psi$ then there exists an initial trace ρ such that $\mathcal{X}, \rho \not\models \psi$. By Theorem 4.4, there exists an initial trace π , whose length is bounded by $|W| \cdot (|\psi'| + 1)^2 \leq |W| \cdot (2|\psi| + 3)^2$, such that $\mathcal{X}, \pi \models \psi'$, where ψ' is the NNF of $\neg\psi$. Now, some non-deterministic instance of A_trace($\mathcal{X}, w_0, |\psi|$) generates exactly such π , being $|\pi| \leq |W| \cdot (2|\psi| + 3)^2$. Moreover, CheckE(\mathcal{X}, ψ, π) = \perp , and thus CounterExE(\mathcal{X}, ψ) has a successful computation.

CounterExE(\mathcal{X}, ψ) is in NP, as the generated trace(s) ρ has (have) a length polynomial in $|W|$ and $|\psi|$, and can thus be computed in polynomial time. Subsequently, CheckE performs a polynomial number of steps, since all it has to do is filling in the table T , which features $|\psi| \cdot |\rho|$ entries.

Corollary 4.9. *The MC problem for E formulas over finite Kripke structures is co-NP-complete.*

Proof. Since CounterExE(\mathcal{X}, ψ) has a successful computation if and only if $\mathcal{X} \not\models \psi$, and such a procedure runs in (non-deterministic) polynomial time, the MC problem belongs to co-NP. The co-NP-hardness derives immediately from that of the purely propositional HS fragment Prop, as proved in [28]. \square

5. An exponential-size model-trace property for $\overline{\text{AABB}}\overline{\text{E}}$ and $\overline{\text{AAEB}}\overline{\text{E}}$

In this section, we prove that the fragments $\overline{\text{AABB}}\overline{\text{E}}$ and $\overline{\text{AAEB}}\overline{\text{E}}$ enjoy an *exponential-size model-trace property*. Such a property ensures that, for each $h \geq 0$ and trace ρ of a finite Kripke structure $\mathcal{X} = (\mathcal{AP}, W, \delta, \mu, w_0)$, it is always possible to find another trace of \mathcal{X} induced by ρ , whose length is at most $(|W| + 2)^{h+2}$, which is indistinguishable from ρ with respect to the satisfiability of any $\overline{\text{AABB}}\overline{\text{E}}$ (resp., $\overline{\text{AAEB}}\overline{\text{E}}$) formula φ with B -nesting depth $d_B(\varphi)$ (resp., E -nesting depth $d_E(\varphi)$) at most h .

To prove such a property, we first introduce the notion of *h-prefix bisimilarity* (resp., *h-suffix bisimilarity*), which defines an equivalence relation over $\text{Trc}_{\mathcal{X}}$ ensuring that equivalent traces satisfy the same $\overline{\text{AABB}}\overline{\text{E}}$ (resp., $\overline{\text{AAEB}}\overline{\text{E}}$) formulas with B -nesting (resp., E -nesting) depth at most h .

Then, we show how to determine, for a given trace ρ , a subset of positions of ρ that allow us to build another trace ρ' , with length at most $(|W| + 2)^{h+2}$, such that ρ and ρ' are *h-prefix bisimilar* (resp., *h-suffix bisimilar*). We call such a set of ρ -positions *prefix* (resp., *suffix*) *sampling* of ρ . Intuitively, they play a role which is analogous to that of witness positions (Definition 4.3) exploited in the previous section.

Let $h \geq 0$. The notions of *h-prefix bisimilarity* and *h-suffix bisimilarity* between a pair of traces ρ and ρ' of a Kripke structure are defined as follows.

Definition 5.1 (*Prefix-bisimilarity and suffix-bisimilarity*). Let $h \geq 0$. Two traces ρ and ρ' of a finite Kripke structure \mathcal{K} are *h -prefix bisimilar* if the following conditions inductively hold:

- for $h = 0$: $\text{fst}(\rho) = \text{fst}(\rho')$, $\text{lst}(\rho) = \text{lst}(\rho')$, and $\text{states}(\rho) = \text{states}(\rho')$;
- for $h > 0$: ρ and ρ' are 0-prefix bisimilar and for each proper prefix ν of ρ (resp., proper prefix ν' of ρ'), there exists a proper prefix ν'' of ρ' (resp., proper prefix ν'' of ρ) such that ν and ν'' are $(h - 1)$ -prefix bisimilar.

The notion of *h -suffix bisimilarity* is defined in a symmetric way by considering suffixes, instead of prefixes, of traces.

As it will be proved in Proposition 5.5 below, h -prefix (resp., h -suffix) bisimilarity is a sufficient condition for two traces ρ and ρ' to be indistinguishable with respect to the satisfiability of $A\bar{A}B\bar{B}E$ (resp., $A\bar{A}E\bar{B}E$) formulas with B -nesting (resp., E -nesting) depth at most h .

The following property can be easily shown.

Property 5.2. Given a finite Kripke structure \mathcal{K} , for all $h \geq 0$, h -prefix (resp., h -suffix) bisimilarity is an equivalence relation over $\text{Trc}_{\mathcal{K}}$.

The following property states that h -suffix bisimilarity and h -prefix bisimilarity “propagate downwards”.

Property 5.3. Let $h > 0$, and ρ and ρ' be two h -prefix (resp., h -suffix) bisimilar traces of a finite Kripke structure \mathcal{K} . Then, ρ and ρ' are also $(h - 1)$ -prefix (resp., $(h - 1)$ -suffix) bisimilar.

As stated by the following Proposition, the relations of h -prefix and h -suffix bisimilarity are preserved by left and right \star -concatenation with a constant string. The property can be easily proved by induction on $h \geq 0$.

Proposition 5.4. Let $h \geq 0$, and let ρ and ρ' be two h -prefix (resp., h -suffix) bisimilar traces of a finite Kripke structure \mathcal{K} . Then, for each trace ρ'' of \mathcal{K} ,

1. $\rho'' \star \rho$ and $\rho'' \star \rho'$ are h -prefix (resp., h -suffix) bisimilar;
2. $\rho \star \rho''$ and $\rho' \star \rho''$ are h -prefix (resp., h -suffix) bisimilar.

By Proposition 5.4 and a straightforward induction on the structural complexity of formulas, we obtain that h -prefix (resp., h -suffix) bisimilarity preserves the satisfiability of $A\bar{A}B\bar{B}E$ (resp., $A\bar{A}E\bar{B}E$) formulas with B -nesting (resp., E -nesting) depth at most h . In other words, h -prefix (resp., h -suffix) bisimilarity is a sufficient condition for two traces to be indistinguishable with respect to any $A\bar{A}B\bar{B}E$ (resp., $A\bar{A}E\bar{B}E$) formula ψ with $d_B(\psi) \leq h$ (resp., $d_E(\psi) \leq h$).

Proposition 5.5. Let $h \geq 0$, and let ρ and ρ' be two h -prefix (resp., h -suffix) bisimilar traces of a finite Kripke structure \mathcal{K} . Then, for each $A\bar{A}B\bar{B}E$ (resp., $A\bar{A}E\bar{B}E$) formula ψ with $d_B(\psi) \leq h$ (resp., $d_E(\psi) \leq h$), it holds that $\mathcal{K}, \rho \models \psi$ if and only if $\mathcal{K}, \rho' \models \psi$.

In the remaining part of the section, we shall focus on the fragment $A\bar{A}B\bar{B}E$ (the case of $A\bar{A}E\bar{B}E$ is completely symmetric). We shall show how to determine a subset of positions of a trace ρ (a *prefix sampling* of ρ), starting from which it is possible to build another trace ρ' , of bounded exponential length, which is indistinguishable from ρ with respect to the satisfiability of $A\bar{A}B\bar{B}E$ formulas up to a given B -nesting depth (*exponential-size model-trace property*). We start by introducing the notions of *prefix-skeleton sampling* and *h -prefix sampling*, and prove some related properties. In the following, we fix a finite Kripke structure $\mathcal{K} = (\mathcal{AP}, W, \delta, \mu, w_0)$, and, given a set I of natural numbers, by “two consecutive elements of I ” we mean a pair of elements $i, j \in I$ such that $i < j$ and $I \cap [i, j] = \{i, j\}$.

Definition 5.6 (*Prefix-skeleton sampling*). Let ρ be a trace of \mathcal{K} . Given two ρ -positions i and j , with $i \leq j$, the *prefix-skeleton sampling* of $\rho(i, j)$ is the minimal set P of ρ -positions in the interval $[i, j]$ satisfying the conditions:

- $i, j \in P$;
- for each state $w \in W$ occurring in $\rho(i + 1, j - 1)$, the least position $k \in [i + 1, j - 1]$ such that $\rho(k) = w$ belongs to P .

Fig. 10 gives a graphical account of the prefix-skeleton sampling of a trace.

From Definition 5.6, it immediately follows that the prefix-skeleton sampling P of (any) trace $\rho(i, j)$ is such that $|P| \leq |W| + 2$, and if $i < j$, then $i + 1 \in P$.

Definition 5.7 (*h -prefix sampling*). Let ρ be a trace of \mathcal{K} . For each $h \geq 1$, the *h -prefix sampling* of ρ is the minimal set P_h of ρ -positions inductively satisfying the following conditions:




Fig. 10. The set P is the prefix-skeleton sampling of $\rho(i, j) = w_1^4 w_2 w_1 w_3 w_1 w_2 w_3 w_1 w_3$.

- for $h = 1$: P_1 is the prefix-skeleton sampling of ρ ;
- for $h > 1$: (i) $P_h \supseteq P_{h-1}$ and (ii) for all pairs of consecutive positions i, j in P_{h-1} , the prefix-skeleton sampling of $\rho(i, j)$ belongs to P_h .

The following upper bound to the cardinality of prefix samplings easily follows from Definition 5.7.

Property 5.8. Let $h \geq 1$ and ρ be a trace of \mathcal{X} . The h -prefix sampling P_h of ρ is such that $|P_h| \leq (|W| + 2)^h$.

Lemma 5.10 and Theorem 5.11 below show how to derive, from any trace ρ of the given finite Kripke structure \mathcal{X} , another trace ρ' , induced by ρ and h -prefix bisimilar to ρ , such that $|\rho'| \leq (|W| + 2)^{h+2}$. By Proposition 5.5, ρ' is indistinguishable from ρ with respect to the satisfiability of any $A\bar{A}B\bar{B}\bar{E}$ formula ψ with $d_B(\psi) \leq h$.

In order to build ρ' , we first compute the $(h + 1)$ -prefix sampling P_{h+1} of ρ . Then, for all the pairs of consecutive ρ -positions $i, j \in P_{h+1}$, we consider a trace induced by $\rho(i, j)$, with no repeated occurrences of any state, except at most the first and last ones (hence, it is no longer than $|W| + 2$). The trace ρ' is just the ordered concatenation (by means of the \star -concatenation operator) of all these traces. The aforementioned bound on $|\rho'|$ holds as, by Property 5.8, $|P_{h+1}| \leq (|W| + 2)^{h+1}$. Lemma 5.10 states that ρ and ρ' are indeed h -prefix bisimilar. The proof of such lemma exploits the following technical result.

Lemma 5.9. Let $h \geq 1$, ρ be a trace of \mathcal{X} , and let i, j be two consecutive ρ -positions in the h -prefix sampling of ρ . Then, for all ρ -positions $n, n' \in [i + 1, j]$ such that $\rho(n) = \rho(n')$, it holds that $\rho(1, n)$ and $\rho(1, n')$ are $(h - 1)$ -prefix bisimilar.

Proof. The proof is by induction on $h \geq 1$.

Base case: $h = 1$. The 1-prefix sampling of ρ is the prefix-skeleton sampling of ρ . Hence, being i and j consecutive positions in this sampling, for each position $k \in [i, j - 1]$, there is $\ell \leq i$ such that $\rho(\ell) = \rho(k)$. Since $\rho(n) = \rho(n')$, it holds that $\text{states}(\rho(1, n)) = \text{states}(\rho(1, n'))$, and thus $\rho(1, n)$ and $\rho(1, n')$ are 0-prefix bisimilar.

Inductive step: $h > 1$. By definition of h -prefix sampling, there are two consecutive positions i', j' in the $(h - 1)$ -prefix sampling of ρ such that i, j are consecutive positions of the prefix-skeleton sampling of $\rho(i', j')$.

If $i = i'$, then $j = i + 1$, and thus, being $n, n' \in [i + 1, j]$, we get that $n = n'$, and the result trivially holds.

Let $i \neq i'$, thus $i > i'$. As in the base case, we easily deduce that $\rho(1, n)$ and $\rho(1, n')$ are 0-prefix bisimilar. It remains to show that for each proper prefix v of $\rho(1, n)$ (resp., v' of $\rho(1, n')$), there is a proper prefix v'' of $\rho(1, n')$ (resp., v'' of $\rho(1, n)$) such that v and v'' are $(h - 2)$ -prefix bisimilar. Let us consider a proper prefix v of $\rho(1, n)$ (the proof for the other direction is symmetric). It holds that $v = \rho(1, m)$, for some $m < n$. We distinguish two cases:

- $m \leq i$. Hence, $\rho(1, m)$ is a proper prefix of $\rho(1, n')$ and the result follows.
- $m > i$. Since i and j are consecutive positions of the prefix-skeleton sampling of $\rho(i', j')$, $i > i'$, and $m \in [i + 1, j - 1]$ (hence $m < j'$), there is $m' \in [i' + 1, i]$ such that $\rho(m') = \rho(m)$ and m' is in the prefix-skeleton sampling of $\rho(i', j')$. Let $v' = \rho(1, m')$. Clearly, v' is a proper prefix of $\rho(1, n')$ (as $n' \geq i + 1$). Moreover, since $m, m' \in [i' + 1, j']$ and i', j' are consecutive positions in the $(h - 1)$ -prefix sampling of ρ , by the inductive hypothesis, $v = \rho(1, m)$ and $v' = \rho(1, m')$ are $(h - 2)$ -prefix bisimilar.

This concludes the proof of Lemma 5.9. \square

Lemma 5.10. Let $h \geq 1$, let ρ be a trace of \mathcal{X} , and let $\rho' = \rho(i_1)\rho(i_2) \cdots \rho(i_k)$ be a trace induced by ρ , where $1 = i_1 < i_2 < \dots < i_k = |\rho|$ and $P_{h+1} \subseteq \{i_1, \dots, i_k\}$, with P_{h+1} the $(h + 1)$ -prefix sampling of ρ . Then, for all $j \in [1, k]$, $\rho'(1, j)$ and $\rho(1, i_j)$ are h -prefix bisimilar.

Note that, as a straightforward consequence, ρ and ρ' are h -prefix bisimilar.

Proof. Let $Q = \{i_1, \dots, i_k\}$ (hence $P_{h+1} \subseteq Q$) and let $j \in [1, k]$. We prove by induction on j that $\rho'(1, j)$ and $\rho(1, i_j)$ are h -prefix bisimilar. As for the base case ($j = 1$), the result holds, since $i_1 = 1$.

Let us assume that $j > 1$. We first show that $\rho(1, i_j)$ and $\rho'(1, j)$ are 0-prefix bisimilar. Clearly, $\rho(1) = \rho(i_1) = \rho'(1)$, $\rho(i_j) = \rho'(j)$, and $\text{states}(\rho'(1, j)) \subseteq \text{states}(\rho(1, i_j))$. Now, if, by contradiction, there was a state w such that

$w \in \text{states}(\rho(1, i_j)) \setminus \text{states}(\rho'(1, j))$, then for all $l \in Q$, with $l \leq i_j$, $\rho(l) \neq w$. However, the prefix-skeleton sampling P_1 of ρ is contained in Q , and the minimal ρ -position l' such that $\rho(l') = w$ belongs to P_1 . Since $w \in \text{states}(\rho(1, i_j))$, we have $l' \leq i_j$ and we get a contradiction, implying that $\text{states}(\rho'(1, j)) = \text{states}(\rho(1, i_j))$.

It remains to prove that:

- (1) for each proper prefix v' of $\rho'(1, j)$, there exists a proper prefix v of $\rho(1, i_j)$ such that v and v' are $(h-1)$ -prefix bisimilar, and
- (2) for each proper prefix v of $\rho(1, i_j)$, there exists a proper prefix v' of $\rho'(1, j)$ such that v and v' are $(h-1)$ -prefix bisimilar.

As for (1), let v' be a proper prefix of $\rho'(1, j)$. Hence, there exists $m \in [1, j-1]$ such that $v' = \rho'(1, m)$. By the inductive hypothesis, $\rho'(1, m)$ and $\rho(1, i_m)$ are h -prefix bisimilar, and thus $(h-1)$ -prefix bisimilar as well by Property 5.3. Since $\rho(1, i_m)$ is a proper prefix of $\rho(1, i_j)$, by choosing $v = \rho(1, i_m)$, (1) follows.

As for (2), assume that v is a proper prefix of $\rho(1, i_j)$. Therefore, there exists $n \in [1, i_j-1]$ such that $v = \rho(1, n)$. We distinguish two cases:

- $n \in P_{h+1}$. Since $n < i_j$, there exists $m \in [1, j-1]$ such that $n = i_m$. By the inductive hypothesis, $\rho(1, n)$ and $\rho'(1, m)$ are h -prefix bisimilar, and thus $(h-1)$ -prefix bisimilar as well by Property 5.3. Since $\rho'(1, m)$ is a proper prefix of $\rho'(1, j)$, by choosing $v' = \rho'(1, m)$, (2) follows.
- $n \notin P_{h+1}$. It follows that there exist two consecutive positions i' and j' in P_{h+1} , with $i' < j'$, such that $n \in [i'+1, j'-1]$. By definition of $(h+1)$ -prefix sampling, there exist two consecutive positions i'' and j'' in the h -prefix sampling of ρ , with $i'' < j''$, such that i' and j' are two consecutive positions in the prefix-skeleton sampling of $\rho(i'', j'')$. First, we observe that $i' \neq i''$ (otherwise, $j' = i' + 1$, which contradicts the fact that $[i'+1, j'-1] \neq \emptyset$, as $n \in [i'+1, j'-1]$). Thus, by definition of prefix-skeleton sampling applied to $\rho(i'', j'')$, and since $n \in [i'+1, j'-1]$, there must be $\ell \in [i''+1, i']$ such that $\rho(\ell) = \rho(n)$ and ℓ is in the prefix-skeleton sampling of $\rho(i'', j'')$. Hence, $\ell \in P_{h+1}$ by definition of $(h+1)$ -prefix sampling. As a consequence, since $\ell < n < i_j$, there exists $m \in [1, j-1]$ such that $\ell = i_m$. By applying Lemma 5.9, we deduce that $\rho(1, n)$ and $\rho(1, i_m)$ are $(h-1)$ -prefix bisimilar. Moreover, by the inductive hypothesis, $\rho(1, i_m)$ and $\rho'(1, m)$ are $(h-1)$ -prefix bisimilar. Thus, by choosing $v' = \rho'(1, m)$, v' is a proper prefix of $\rho'(1, j)$ which is $(h-1)$ -prefix bisimilar to $v = \rho(1, n)$.

This concludes the proof of Lemma 5.10. \square

We are now ready to prove the exponential-size model-trace property.

Theorem 5.11 (Exponential-size model-trace property for $A\bar{A}B\bar{B}E$). *Let ρ be a trace of a finite Kripke structure \mathcal{X} and let $h \geq 0$. Then, there exists a trace ρ' induced by ρ , whose length is at most $(|W| + 2)^{h+2}$, such that for every $A\bar{A}B\bar{B}E$ formula ψ with $d_B(\psi) \leq h$, it holds that $\mathcal{X}, \rho \models \psi$ if and only if $\mathcal{X}, \rho' \models \psi$.*

Proof. Let P_{h+1} be the $(h+1)$ -prefix sampling of ρ . For all pairs of consecutive ρ -positions i and j in P_{h+1} , there exists a trace induced by $\rho(i, j)$, whose length is at most $|W| + 2$, featuring no repeated occurrences of any internal state. We now define ρ' as the trace of \mathcal{X} obtained by an ordered concatenation of all these induced traces by means of the \star -concatenation operator. It is immediate to see that $\rho' = \rho(i_1)\rho(i_2)\cdots\rho(i_k)$, for some indexes $1 = i_1 < i_2 < \cdots < i_k = |\rho|$, where $\{i_1, \dots, i_k\}$ contains the $(h+1)$ -prefix sampling P_{h+1} of ρ . It holds that $|\rho'| \leq |P_{h+1}| \cdot (|W| + 2)$ and since, by Property 5.8, $|P_{h+1}| \leq (|W| + 2)^{h+1}$, we obtain that $|\rho'| \leq (|W| + 2)^{h+2}$. Moreover, by Lemma 5.10, ρ and ρ' are h -prefix bisimilar. By Proposition 5.5, the thesis follows. \square

Theorem 5.11 allows us to easily devise an EXPSPACE MC algorithm for $A\bar{A}B\bar{B}E$ formulas (and symmetrically for $A\bar{A}E\bar{B}E$ formulas), which can be obtained from Algorithms 1 and 2 by adapting the bounds on the length of considered traces.

In [29], the authors proved—in a much more involved way—the existence of a bound on the length of satisfiability-preserving traces (called there *trace representatives*), which is greater than the present one, i.e., $O(|W|^{2h+4})$.

It is worth observing that the polynomial-size model-trace property for $A\bar{A}B\bar{B}/A\bar{A}E\bar{B}$ of the previous section depends on the specific formula φ we are considering (as input of the MC problem), whereas the exponential-size model-trace property for $A\bar{A}B\bar{B}E/A\bar{A}E\bar{B}E$ states the existence of a shorter trace ρ' equivalent to (a generic) ρ with respect to all formulas up to a given B/E -nesting depth h . As a matter of fact, the former relies on the witness positions, which are defined on φ ; the latter relies on h -prefix/suffix bisimilarity and h -prefix/suffix samplings, which are independent of any formula (they are only based on h , that is, the maximum B/E -nesting depth of formulas we want to consider). Therefore, we can say that the latter small-model states a stronger property; however, this may lead to a bound on the length of equivalent traces higher than necessary: we proved the MC problem for $A\bar{A}B\bar{B}E/A\bar{A}E\bar{B}E$ to be in EXPSPACE, but it is only known to be PSPACE-hard (since \bar{E} or \bar{B} are enough for the PSPACE-hardness, as shown in [30]). We do not know whether this complexity gap is due to the small-model proving a loose bound (that might be strengthened by finding another characterization depending on

the input formula as well), or to a weak complexity lower-bound (here, exploiting the other modalities A , \bar{A} and B/E , along with \bar{B} and \bar{E} jointly, may enable us to prove a stronger one)—or to both at the same time.

6. Conclusions

In this paper, we have studied the complexity scenario astride the line dividing tractable and intractable fragments of Halpern and Shoham’s modal logic of time intervals HS with respect to model checking. On one hand, we have shown that the simultaneous presence of modalities (B) and (E) is sufficient for any HS fragment to be EXPSpace-hard—therefore provably intractable—and this lower bound immediately propagates to full HS. On the other hand, we have studied two well-behaved, PSPACE-complete fragments, $A\bar{A}B\bar{B}$ and $A\bar{A}E\bar{E}$, which are quite promising from the point of view of applications. In between, MC for $A\bar{A}B\bar{B}E$ and $A\bar{A}E\bar{B}E$ turns out to be in EXPSpace and PSPACE-hard.

Membership to PSPACE for the former two fragments and membership to EXPSpace for the latter have been proved by means of small-models, which, in turn, rely on suitable (depending on the specific fragment) contraction techniques applied to the traces of a finite Kripke structure. While the first result is novel, the second slightly reduces the bounds for trace representatives given for the same problem in [29], and substantially simplifies the constructions and the complexity of the proofs.

As for future work, we would like to precisely characterize the complexity of MC for $A\bar{A}B\bar{B}E$ and $A\bar{A}E\bar{B}E$. However, an even larger complexity gap is the one for full HS: we have shown it to be EXPSpace-hard, but the only known upper bound is non-elementary [27].

Finally, we want to relax the homogeneity assumption, which limits the expressiveness of HS and its fragments. One possible direction has been outlined by Lomuscio and Michaliszyn in [24], where the proposition letters that hold on an interval have been defined by means of regular expressions over the states of a Kripke structure. We expect that, under this semantic variant, model checking for full HS will remain decidable and that the complexity of some fragments could increase as a combined effect of the expressive power of regular expressions and of HS modalities.

Acknowledgements

The work by Alberto Molinari, Angelo Montanari, and Pietro Sala has been supported by the GNCS project *Logic and Automata for Interval Model Checking*.

References

- [1] J.F. Allen, Maintaining knowledge about temporal intervals, *Commun. ACM* 26 (11) (1983) 832–843.
- [2] A. Armando, R. Carbone, L. Compagna, LTL model checking for security protocols, in: *CSF*, IEEE Computer Society, 2007, pp. 385–396.
- [3] D. Basin, C. Cremers, C. Meadows, Model checking security protocols, <http://www-oldurls.inf.ethz.ch/personal/basin/pubs/security-modelchecking.pdf>, 2015.
- [4] M. Benerecetti, R.D. Guglielmo, U. Gentile, S. Marrone, N. Mazzocca, R. Nardone, A. Peron, L. Velardi, V. Vittorini, Dynamic state machines for modelling railway control systems, *Sci. Comput. Program.* 133 (2017) 116–153.
- [5] H. Bowman, S.J. Thompson, A decision procedure and complete axiomatization of finite interval temporal logic with projection, *J. Logic Comput.* 13 (2) (2003) 195–239.
- [6] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Interval temporal logic model checking based on track bisimilarity and prefix sampling, in: *ICTCS*, in: *CEUR Workshop Proceedings*, vol. 1720, 2016, pp. 49–61, CEUR-WS.org.
- [7] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Interval temporal logic model checking: the border between good and bad HS fragments, in: *IJCAR*, in: *LNCS*, vol. 9706, Springer, 2016, pp. 389–405.
- [8] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Interval vs. point temporal logic model checking: an expressiveness comparison, in: *FSTTCS*, in: *LIPIcs*, Schloss Dagstuhl—Leibniz-Zentrum fuer Informatik, 2016, pp. 26:1–26:14.
- [9] D. Bresolin, D. Della Monica, V. Goranko, A. Montanari, G. Sciavicco, The dark side of interval temporal logic: marking the undecidability border, *Ann. Math. Artif. Intell.* 71 (1–3) (2014) 41–83.
- [10] D. Bresolin, V. Goranko, A. Montanari, P. Sala, Tableau-based decision procedures for the logics of subinterval structures over dense orderings, *J. Logic Comput.* 20 (1) (2010) 133–166.
- [11] D. Bresolin, V. Goranko, A. Montanari, G. Sciavicco, Propositional interval neighborhood logics: expressiveness, decidability, and undecidable extensions, *Ann. Pure Appl. Logic* 161 (3) (2009) 289–304.
- [12] D. Bresolin, A. Montanari, P. Sala, G. Sciavicco, What’s decidable about Halpern and Shoham’s interval logic? The maximal fragment $AB\bar{B}L$, in: *LICS*, IEEE Computer Society, 2011, pp. 387–396.
- [13] A. Cimatti, *Industrial Applications of Model Checking*, Springer, 2001, pp. 153–168, Ch. 6.
- [14] E.A. Emerson, J.Y. Halpern, “Sometimes” and “not never” revisited: on branching versus linear time temporal logic, *J. ACM* 33 (1) (1986) 151–178.
- [15] L. Giordano, P. Terenziani, A. Bottrighi, S. Montani, L. Donzella, Model checking for clinical guidelines: an agent-based approach, in: *AMIA*, 2006, pp. 289–293.
- [16] F. Giunchiglia, P. Traverso, Planning as model checking, in: *ECP*, in: *LNCS*, vol. 1809, Springer, 1999, pp. 1–20.
- [17] V. Goranko, A. Montanari, G. Sciavicco, A road map of interval temporal logics and duration calculi, *J. Appl. Non-Classical Logics* 14 (1–2) (2004) 9–54.
- [18] G. Gottlob, NP trees and Carnap’s modal logic, *J. ACM* 42 (2) (1995) 421–457.
- [19] J.Y. Halpern, Y. Shoham, A propositional modal logic of time intervals, *J. ACM* 38 (4) (1991) 935–962.
- [20] D. Harel, *Algorithmics: The Spirit of Computing*, 2nd edition, Wesley, 1992.
- [21] K. Lodaya, Sharpening the undecidability of interval temporal logic, in: *ASIAN*, in: *LNCS*, vol. 1961, Springer, 2000, pp. 290–298.
- [22] A. Lomuscio, J. Michaliszyn, An epistemic Halpern–Shoham logic, in: *IJCAI*, in: *IJCAI/AAAI*, 2013, pp. 1010–1016.
- [23] A. Lomuscio, J. Michaliszyn, Decidability of model checking multi-agent systems against a class of EHS specifications, in: *Frontiers in Artificial Intelligence and Applications*, in: *ECAI*, vol. 263, IOS Press, 2014, pp. 543–548.

- [24] A. Lomuscio, J. Michaliszyn, Model checking multi-agent systems against epistemic HS specifications with regular expressions, in: KR, AAAI Press, 2016, pp. 298–308.
- [25] A. Lomuscio, F. Raimondi, MCMAS: a model checker for multi-agent systems, in: TACAS, in: LNCS, vol. 3920, Springer, 2006, pp. 450–454.
- [26] J. Marcinkowski, J. Michaliszyn, The undecidability of the logic of subintervals, *Fund. Inform.* 131 (2) (2014) 217–240.
- [27] A. Molinari, A. Montanari, A. Murano, G. Perelli, A. Peron, Checking interval properties of computations, *Acta Inform.* 53 (6) (2016) 587–619.
- [28] A. Molinari, A. Montanari, A. Peron, Complexity of ITL model checking: some well-behaved fragments of the interval logic HS, in: TIME, IEEE Computer Society, 2015, pp. 90–100.
- [29] A. Molinari, A. Montanari, A. Peron, A model checking procedure for interval temporal logics based on track representatives, in: CSL, in: LIPIcs, Schloss Dagstuhl—Leibniz-Zentrum fuer Informatik, 2015, pp. 193–210.
- [30] A. Molinari, A. Montanari, A. Peron, P. Sala, Model checking well-behaved fragments of HS: the (almost) final picture, in: KR, AAAI Press, 2016, pp. 473–483.
- [31] A. Montanari, G. Puppis, P. Sala, Maximal decidable fragments of Halpern and Shoham’s modal logic of intervals, in: ICALP, in: LNCS, vol. 6199, Springer, 2010, pp. 345–356.
- [32] B. Moszkowski, Reasoning About Digital Circuits, Ph.D. thesis, Dept. of Computer Science, Stanford University, Stanford, CA, 1983.
- [33] R. Nardone, U. Gentile, M. Benerecetti, A. Peron, V. Vittorini, S. Marrone, N. Mazzocca, Modeling Railway Control Systems in Promela, Springer, 2016, pp. 121–136, Ch. 1.
- [34] A. Pnueli, The temporal logic of programs, in: FOCS, IEEE Computer Society, 1977, pp. 46–57.
- [35] I. Pratt-Hartmann, Temporal prepositions and their logic, *Artificial Intelligence* 166 (1–2) (2005) 1–36.
- [36] P. Roeper, Intervals and tenses, *J. Philos. Logic* 9 (1980) 451–469.
- [37] P. Schnoebelen, Oracle circuits for branching-time model checking, in: ICALP, in: LNCS, vol. 2719, Springer, 2003, pp. 790–801.
- [38] Y. Venema, Expressiveness and completeness of an interval tense logic, *Notre Dame J. Form. Log.* 31 (4) (1990) 529–547.
- [39] Y. Venema, A modal logic for chopping intervals, *J. Logic Comput.* 1 (4) (1991) 453–476.
- [40] T. Witkowski, N. Blanc, D. Kroening, G. Weissenbacher, Model checking concurrent Linux device drivers, in: ASE, ACM, 2007, pp. 501–504.
- [41] C. Zhou, M.R. Hansen, Duration Calculus: A Formal Approach to Real-Time Systems, Monographs in Theoretical Computer Science. An EATCS Series, Springer, 2004.