Complexity analysis of a unifying algorithm for model checking interval temporal logic ☆

Laura Bozzelli [a], Angelo Montanari [b,*], Adriano Peron [a,*]

[a] University of Napoli "Federico II", Napoli, Italy
[b] University of Udine, Udine, Italy

A B S T R A C T

Model checking (MC) for Halpern and Shoham's interval temporal logic HS has been recently shown to be decidable. An intriguing open question is its exact complexity for full HS: it is at least **EXPSPACE**-hard, and the only known upper bound, which exploits an abstract representation of Kripke structure paths (descriptor), is non-elementary.
In this paper, we provide a uniform framework to MC for full HS and meaningful fragments of it, with a specific type of descriptor for each fragment. Then, we devise a general MC alternating algorithm, parameterized by the descriptor's type, which has a polynomially bounded number of alternations and whose running time is bounded by the length of minimal representatives of descriptors (certificates). We analyze its complexity and give tight bounds on the length of certificates. For two types of descriptor, we obtain exponential upper and lower bounds; for the other ones, we provide non-elementary lower bounds.

Keywords:
Interval temporal logic
Model checking
Computational complexity

1. Introduction

Model checking (MC) is a well-established formal method to automatically check the global correctness of finite-state reactive systems. Finite systems are usually modeled as labeled state-transition graphs (finite Kripke structures), while the properties of interest are specified in standard *Point-based* Temporal Logics (PTLs), such as, for instance, the linear-time temporal logic LTL [1] and the branching-time temporal logics CTL and CTL* [2]. *Interval temporal logics* (ITLs) provide an alternative setting for reasoning about time [3–5]. They assume intervals, instead of points, as their primitive temporal entities allowing one to specify temporal properties that involve, e.g., actions with duration, accomplishments, and temporal aggregations, which are inherently "interval-based", and thus cannot be naturally expressed by PTLs. ITLs find application in a variety of fields (see, e.g., [4,6,7]), including artificial intelligence (reasoning about action and change, qualitative reasoning, planning, and natural language processing), theoretical computer science (specification and verification of programs), and temporal and spatio-temporal databases. Among ITLs, the landmark is *Halpern and Shoham's modal logic of time intervals* (HS) [3] which features one modality for each of the 13 possible ordering relations between pairs of intervals (the so-called Allen's relations [8]), apart from equality. The satisfiability problem for HS turns out to be undecidable over all relevant classes of linear orders, and most of its fragments (with some meaningful exceptions [9–11]) are undecidable as well [12–14].

☆ We acknowledge the support from the GNCS project *Strategic Reasoning and Automatic Synthesis of Multi-Agent Systems* (164850-667).
* Corresponding authors.
E-mail addresses: lr.bozzelli@gmail.com (L. Bozzelli), angelo.montanari@uniud.it (A. Montanari), adrperon@unina.it (A. Peron).

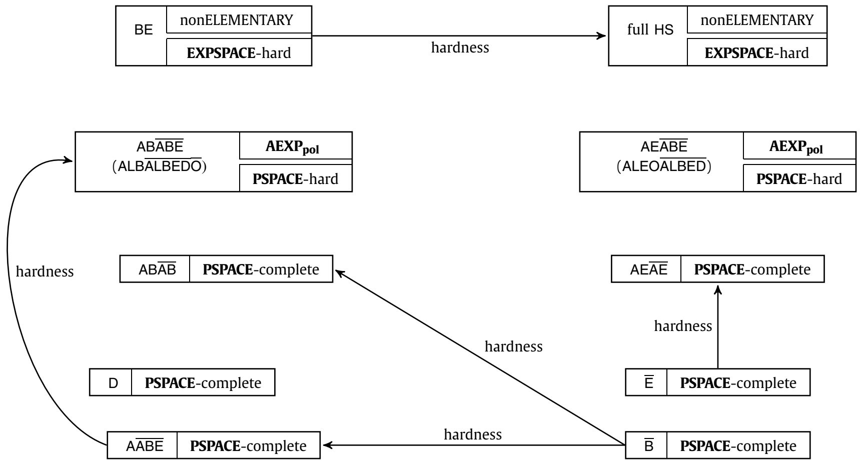

Fig. 1. Complexity of the MC problem for HS fragments under the state-based semantics.

MC of (finite) Kripke structures against HS has been investigated only very recently [6,15–24]. The idea is to interpret each finite path of a Kripke structure as an interval, whose labeling is defined on the basis of the labeling of the component states, i.e., a proposition letter holds over an interval if and only if it holds over each component state (*homogeneity assumption* [25]). In this paper, we focus on the MC problem for HS under the *state-based semantics* (time branches both in the future and in the past), whose decidability has been proved in [16]. In this setting, temporal modalities for Allen's relations *Begin-with* (B), *End-with* (E), and *During* (D) have a "linear-time" character: they allow one to select proper prefixes (B), proper suffixes (E), and inner sub-paths (D) of the current path, respectively. A graphical account of Allen's relations and the corresponding modalities is given in Table 1. Modalities associated with the other Allen's relations are, instead, "branching-time": they allow one *either* to non-deterministically extend a prefix (resp., suffix, sub-path) of the current path in the future or in the past, *or* to non-deterministically select an independent path whose starting point (resp., ending point) is reachable from (resp., can reach) the ending point (resp., starting point) of the current path. Expressiveness of HS under the state-based semantics has been studied in [18], together with two other decidable variants of it: the *computation-tree-based* semantics, that allows time to branch only in the future, and the *trace-based* one, that disallows time branching. The computation-tree-based variant of HS is expressively equivalent to finitary CTL* (the variant of CTL* with quantification over finite paths), while the trace-based one is equivalent to LTL, but at least exponentially more succinct. The considered state-based variant is more expressive than the computation-tree-based one and expressively incomparable with both LTL and CTL*.

As for the complexity of the state-based MC problem, it is at least **EXPSPACE**-hard [17] for full HS, while the only known upper bound is non-elementary [16]. Such an upper bound is obtained by defining a finite abstraction over the (possibly infinite) set of finite paths of a Kripke structure. The abstraction is parameterized by a natural number h and associates with each path a bounded abstract representation, called *h-level BE-descriptor*. The h-level BE-descriptor of a path (see [16]) is a tree-like structure of depth h which conveys meaningful information about the prefixes and suffixes of the path, namely, the first and last state, the internal states, and, in case $h > 0$, recursively, the BE-descriptors of level $(h-1)$ of such sub-paths. The distinctive property of descriptors is that for a given $h \geq 0$, the abstraction partitions the (possible infinite) set of paths of a Kripke structure into a finite number of equivalence classes such that: (i) paths in the same equivalence class, that is, paths with the same h-level BE-descriptor, are indistinguishable with respect to the fulfillment of HS formulas with nesting depth of modalities for prefixes (B) and suffixes (E) at most h, and (ii) each equivalence class admits a bounded minimal representative (*h-level BE-certificate*), whose length is at most a tower of exponentials of height h.

An MC procedure for full HS based on BE-descriptors is outlined in [16], but some important features of it have not been analyzed in detail. In particular, the succinctness of BE-descriptors has not been investigated so far. This is a fundamental issue as the computational cost of the MC procedure based on BE-descriptors depends on the length of h-level BE-certificates, where h is the joint nesting depth of modalities for B and E in the given formula. In subsequent papers [19,20,22–24], research focuses on some syntactic fragments of HS: the fragment featuring only the modality for the *During* relation (D), and fragments featuring modalities for a subset of Allen's relations *meets* (A), *begin-with* (B), *end-with* (E), and their transposed relations $\overline{A}$, $\overline{B}$, and $\overline{E}$, respectively. For these fragments, different MC algorithms have been devised, which make use of ad hoc path contraction techniques to obtain bounded minimal size representatives for paths. The complete picture of known complexity results is given in Fig. 1. The most investigated fragments of HS are AB$\overline{\text{AB}}$E and AE$\overline{\text{AB}}$E (see [22–24]). These fragments have been introduced and studied in [22], where **EXPSPACE** membership of MC has been established by a quite involved technique. Then, they have been studied in a setting which relaxes the homogeneity assumption and exploits regular expressions to define the behavior of proposition letters over intervals in terms of their component states (see [23,24]). In this more expressive setting, it has been proved that the MC problem for AB$\overline{\text{AB}}$E and AE$\overline{\text{AB}}$E belongs to the complexity class **AEXP**$_{\mathbf{pol}}$ of those problems decided by exponential-time bounded alternating Turing

Machines with a polynomially bounded number of alternations (such a class is included in **EXPSPACE** and it captures the exact complexity of some relevant problems, e.g., the first-order theory of real addition with order [26]). This result provides an improved upper bound (membership in $\mathbf{AEXP_{pol}}$) to MC against $\mathsf{AB\overline{A}\overline{B}E}$ and $\mathsf{AE\overline{A}\overline{B}E}$ when the homogeneity assumption is enforced.

It is worth pointing out that known complexity (upper and lower) bounds for full HS coincide with those for the fragment only featuring the two linear-time modalities for E and B, thus suggesting that the complexity of the fragments strictly depends on the featured combination of linear-time modalities. Following this intuition, we develop a uniform framework for the state-based MC problem for syntactic fragments of HS which is parametric in the featured non-empty subset $\mathcal{B}$ of non-interdefinable linear Allen's relations in $\{B, E, D\}$ (apart from $\{D\}$), which is called *linear-time basis*. The linear-time basis is then combined with modalities in a subset (which depends on the specific basis) of the branching-time Allen's relations, namely, the relations in the set $\{A, L, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$. In particular, the fragment for the complete basis $\{B, E\}$ expresses the full logic HS, while, for the bases $\{B\}$ end $\{E\}$, we consider the fragments $\mathsf{ALB\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}\overline{O}}$ and $\mathsf{ALEO\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}}$, respectively (these two fragments are as expressive as $\mathsf{AB\overline{A}\overline{B}E}$ and $\mathsf{AE\overline{A}\overline{B}E}$, respectively).

The proposed framework is based on a uniform technique for model checking the considered fragments which generalizes the descriptor-based approach for full HS proposed in [16] and, at the same time, allows us to fix complexity upper bounds on some already-investigated fragments which are comparable with the ones obtained by ad hoc techniques [23,24]. As in [16], we exploit the idea of h-level descriptors. For each linear-time basis $\mathcal{B}$ and every natural number h, we introduce the notion of *h-level $\mathcal{B}$-descriptor*, which coincides with the notion of h-level BE-descriptor for the complete basis $\mathcal{B} = \{B, E\}$. As already established for the basis $\{B, E\}$ [16], we show that for all the other bases, paths with the same h-level $\mathcal{B}$-descriptor (i) are indistinguishable with respect to the fulfillment of HS formulas in the corresponding fragment with nesting depth of modalities for $\mathcal{B}$ at most h (*h-depth satisfaction property*), and (ii) in case the basis is distinct from $\{D\}$, they admit a bounded minimal representative (*h-level $\mathcal{B}$-certificate*). It is an open question whether the descriptors for the basis $\{D\}$ enjoy *bounded* h-level $\mathcal{B}$-certificates (*bounded path property*). Moreover, as an additional interesting contribution, we show that for each basis $\mathcal{B}$, the set of branching-time modalities in the related fragment is maximal with respect to the fulfillment of the h-depth satisfaction property. More precisely, we prove that for any extension of the fragment for $\mathcal{B}$ with a branching-time modality non-expressible in the fragment (the addition of the modality for O to $\mathsf{ALB\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}\overline{O}}$ and of the modality for $\overline{O}$ to $\mathsf{ALEO\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}}$), there are a Kripke structure and a formula (in the extended fragment) that, for each h, distinguishes some distinct paths with the same h-level $\mathcal{B}$-descriptor.

We exploit the previous results (in particular, the h-depth satisfaction property and the bounded path property) to devise an *alternating-time* algorithm, parameterized in the basis $\mathcal{B}\,(\neq \{D\})$, to model check the associated fragment. The algorithm runs in time bounded by the maximal length of h-level $\mathcal{B}$-certificates of the input Kripke structure, with h being the $\mathcal{B}$-nesting depth of the input formula, and the number of alternations between existential and universal choices is at most equal to the size of the input formula. As the most relevant contribution, for each basis $\mathcal{B}$, we provide tight bounds on the length of h-level $\mathcal{B}$-certificates. For the bases $\{B\}$ and $\{E\}$, we obtain singly-exponential upper and lower bounds. Hence, by the proposed alternating algorithm, we obtain membership in $\mathbf{AEXP_{pol}}$ of MC for $\mathsf{ALB\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}\overline{O}}$ and $\mathsf{ALEO\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}}$. Such a complexity result (membership in $\mathbf{AEXP_{pol}}$) has already been obtained in [23,24] for $\mathsf{AB\overline{A}\overline{B}E}$ and $\mathsf{AE\overline{A}\overline{B}E}$ by exploiting a more involved finite abstraction of paths (recall that $\mathsf{AB\overline{A}\overline{B}E}$ and $\mathsf{AE\overline{A}\overline{B}E}$ are as expressive as $\mathsf{ALB\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}\overline{O}}$ and $\mathsf{ALEO\overline{A}\overline{L}\overline{B}\overline{E}\overline{D}}$, respectively). On the other hand, for all bases $\mathcal{B}$ distinct from $\{B\}$ and $\{E\}$, we state a non-elementary lower bound. In particular, the result obtained for the basis $\{B, E\}$ negatively answers a question left open in [16] about the possibility of fixing an elementary upper bound on the size of BE-descriptors, and, at the same time, it provides new insight on the MC problem for full HS: if elementary procedures exist, they certainly have to exploit structures less powerful than descriptors.

The rest of the paper is organized as follows. In Section 2, we recall the state-based MC framework for HS. In Section 3, for each basis $\mathcal{B}$, we introduce the notion of $\mathcal{B}$-descriptor proving the relevant properties of the abstraction based on $\mathcal{B}$-descriptors. In particular, we prove maximality of the studied fragments with respect to the considered abstraction. Then, in Section 4, for each basis $\mathcal{B}(\neq \{D\})$, we describe the MC algorithm for the associated fragment. In Section 5, we give tight bounds on the length of $\mathcal{B}$-certificates. Finally, Section 6 provides an assessment of the work done, and outlines future research directions.

2. Preliminaries

In this section, we first introduce basic notation, and then recall the logic HS [3] and the state-based MC framework for verifying HS formulas [16].

Let $\mathbb{N}$ be the set of natural numbers. For all $i, j \in \mathbb{N}$, with $i \leq j$, $[i, j]$ denotes the set of natural numbers h such that $i \leq h \leq j$. For all $n, h \in \mathbb{N}$, $Tower(n, h)$ denotes a tower of exponentials of height h and argument n: more precisely, $Tower(n, 0) = n$ and $Tower(n, h+1) = 2^{Tower(n,h)}$.

2.1. Finite words

Let Σ be a finite alphabet. The set of all finite words over Σ is denoted by Σ^*, and $\Sigma^+ := \Sigma^* \setminus \{\varepsilon\}$, where ε is the empty word. Let w be a finite word over Σ. We denote by $|w|$ the length of w. For all $i, j \in \mathbb{N}$, with $i \leq j$, $w(i)$ is the $(i+1)$-th letter of w while $w[i, j]$ denotes the infix of w given by $w(i) \cdots w(j)$. If $w \neq \varepsilon$, then we denote by $\mathsf{fst}(w)$ and $\mathsf{lst}(w)$ the

Table 1
Allen's primitive relations and corresponding HS modalities.

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

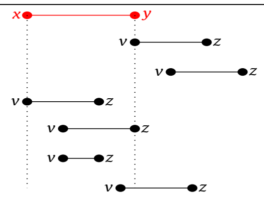

first and last symbol of w, and by internal(w) the set of letters in Σ occurring in $w[1, n-1]$, with $|w| = n+1$. The set Pref(w) of *non-empty proper prefixes* of w is the set of non-empty finite words u such that $w = u \cdot v$ for some finite word $v \neq \varepsilon$. The set Suff(w) of *non-empty proper suffixes* of w is the set of non-empty words u such that $w = v \cdot u$ for some finite word $v \neq \varepsilon$. A *sub-word* (resp., *internal sub-word*) of w is a word w' such that w is of the form $w = u \cdot w' \cdot v$ for some words (resp., for some non-empty words) u and v.

2.2. The interval temporal logic HS

An interval algebra to reason about intervals and their relative orders was proposed by Allen in [8], while a systematic logical study of interval representation and reasoning was done a few years later by Halpern and Shoham, who introduced the interval temporal logic HS featuring one modality for each Allen relation, but equality (it is obviously unnecessary to be considered) [3]. 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 $b\overline{\mathcal{R}}a$ iff $a\mathcal{R}b$) and equality.

Let $\mathcal{AP}$ be a finite set of atomic propositions. HS formulas ψ over $\mathcal{AP}$ are defined as follows:

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

where $p \in \mathcal{AP}$ and $\langle X \rangle$ is the existential temporal modality for the (non-trivial) Allen's relation $X \in \{A, L, B, E, D, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$. The size $|\psi|$ of a formula ψ is the number of distinct subformulas of ψ. We also exploit the standard logical connectives $\vee$ (disjunction) and $\rightarrow$ (implication) as abbreviations, and for any temporal modality $\langle X \rangle$, the dual universal modality $[X]$ defined as: $[X]\,\psi := \neg\langle X \rangle\neg\psi$. An HS formula ψ is in *negative normal form* (*NNF*) if negation is applied only to atomic formulas in $\mathcal{AP}$. By using De Morgan's laws and for any existential modality $\langle X \rangle$, the dual universal modality $[X]$, we can convert in linear-time an HS formula ψ into an equivalent formula in *NNF*, called the *NNF* of ψ. For a formula ψ in *NNF*, the *dual* $\widetilde{\psi}$ of ψ is the *NNF* of $\neg\psi$.

Given a set $U \subseteq \{A, L, B, E, D, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$ of Allen's relations, the *joint nesting depth of U in a formula* ψ denoted by $\mathrm{depth}_U(\psi)$ is defined as: (*i*) $\mathrm{depth}_U(p) = 0$, for any $p \in \mathcal{AP}$; (*ii*) $\mathrm{depth}_U(\neg\psi) = \mathrm{depth}_U(\psi)$; (*iii*) $\mathrm{depth}_U(\psi \wedge \varphi) = \max\{\mathrm{depth}_U(\psi), \mathrm{depth}_U(\varphi)\}$; (*iv*) $\mathrm{depth}_U(\langle X \rangle\,\psi) = 1 + \mathrm{depth}_U(\psi)$ if $X \in U$, and $\mathrm{depth}_U(\langle X \rangle\,\psi) = \mathrm{depth}_U(\psi)$ otherwise.

Given any subset of Allen's relations $\{X_1, .., X_n\}$, we denote by $\mathsf{X_1 \cdots X_n}$ the HS fragment featuring existential (and universal) modalities for $X_1, .., X_n$ only.

The logic HS can be regarded as a multi-modal logic and its semantics can be defined over a multi-modal Kripke structure, called *abstract interval model* (AIM for short), where intervals are treated as atomic objects and Allen's relations as binary relations over intervals. In the following we shall instantiate the abstract interval model taking the set of paths of a Kripke structure as set of intervals.

Definition 2.1 *(Abstract interval models [16]).* An *abstract interval model* (AIM) over $\mathcal{AP}$ is a tuple $\mathcal{A} = (\mathcal{AP}, \mathbb{I}, \{Y_{\mathbb{I}}\}_{Y \in \{A,L,B,E,D,O\}}, Lab_{\mathbb{I}})$, where $\mathbb{I}$ is a possibly infinite set of worlds (abstract intervals), $Y_{\mathbb{I}}$ is a binary relation over $\mathbb{I}$ for each $Y \in \{A, L, B, E, D, O\}$, and $Lab_{\mathbb{I}} : \mathbb{I} \mapsto 2^{\mathcal{AP}}$ is a labeling function which assigns a set of proposition letters from $\mathcal{AP}$ to each abstract interval.

In the interval setting, $\mathbb{I}$ is interpreted as a set of intervals and $Y_{\mathbb{I}}$ as Allen's relation Y for each $Y \in \{A, L, B, E, D, O\}$; $Lab_{\mathbb{I}}$ assigns to each interval in $\mathbb{I}$ the set of atomic propositions that hold over it. Given an interval $I \in \mathbb{I}$, the truth of an HS formula over I is inductively defined as follows:

- $\mathcal{A}, I \models p$ if $p \in Lab_{\mathbb{I}}(I)$, for any $p \in \mathcal{AP}$;
- $\mathcal{A}, I \models \neg\psi$ iff $\mathcal{A}, I \not\models \psi$;
- $\mathcal{A}, I \models \psi \wedge \psi'$ iff $\mathcal{A}, I \models \psi$ and $\mathcal{A}, I \models \psi'$;
- $\mathcal{A}, I \models \langle Y \rangle\,\psi$, for $Y \in \{A, L, B, E, D, O\}$, if $I\, Y_{\mathbb{I}}\, J$ and $\mathcal{A}, J \models \psi$ for some $J \in \mathbb{I}$;
- $\mathcal{A}, I \models \langle \overline{Y} \rangle\,\psi$, for $Y \in \{A, L, B, E, D, O\}$, if $J\, Y_{\mathbb{I}}\, I$ and $\mathcal{A}, J \models \psi$ for some $J \in \mathbb{I}$.

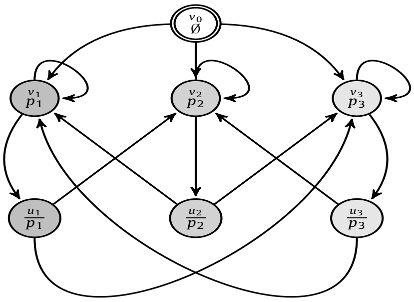

Fig. 2. The Kripke structure $\mathcal{K}_{Sched}$.

State-based model checking against HS. In the context of MC, finite state systems are usually modeled as finite Kripke structures over a finite set $\mathcal{AP}$ of atomic propositions which represent predicates over the states of the system.

Definition 2.2. A *Kripke structure* over $\mathcal{AP}$ is a tuple $\mathcal{K} = (\mathcal{AP}, S, R, Lab, s_0)$, where S is a set of states, $R \subseteq S \times S$ is a transition relation, $Lab : S \mapsto 2^{\mathcal{AP}}$ is a labeling function assigning to each state s the set of propositions that hold over it, and $s_0 \in S$ is the initial state. We say that $\mathcal{K}$ is finite if S is finite.

Let $\mathcal{K} = (\mathcal{AP}, S, R, Lab, s_0)$ be a Kripke structure. A path π of $\mathcal{K}$ is a non-empty finite word over S such that for all $0 \leq i < |\pi|$, $(\pi(i), \pi(i+1)) \in R$. A *sub-path* (resp., *internal sub-path*) of π is a path of $\mathcal{K}$ which is a subword (resp., internal subword) of π. A path is *initial* if it starts from the initial state of $\mathcal{K}$.

Example 2.1. In Fig. 2, we depict 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. In the initial state v_0 no process is served. In the states v_i, with $i \in \{1, 2, 3\}$, the i-th process is served (the proposition p_i holds in those states). The loop on v_i represents the use of the resource. A transition from the state v_i to u_i represents the unlock of the granted resource (the proposition $\overline{p_i}$ holds in that state). The scheduler cannot serve the same process twice in two successive rounds, and then v_i is not directly reachable from $\overline{v}_i$. A transition from u_i to v_j with $j \neq i$, represents the fact that the process j-th has issued a request for the resource and is served.

We now recall the state-based approach [16] for model checking Kripke structures against HS formulas which consists in defining a mapping from a Kripke structure $\mathcal{K}$ to an AIM $\mathcal{A}_{\mathcal{K}}$, where the abstract intervals correspond to the paths of the Kripke structure, the Allen's relations are interpreted over the paths of $\mathcal{K}$ in a natural way, and the following assumption is adopted: a proposition holds over an interval if and only if it holds over all its subintervals (*homogeneity principle*).

Definition 2.3. Let $\mathcal{K} = (\mathcal{AP}, S, R, Lab, s_0)$ be a Kripke structure and R^+ the strict transitive closure of R. The AIM induced by $\mathcal{K}$ is $\mathcal{A}_{\mathcal{K}} = (\mathcal{AP}, \mathbb{I}, \{Y_{\mathbb{I}}\}_{Y \in \{A,L,B,E,D,O\}}, Lab_{\mathbb{I}})$, where $\mathbb{I}$ is the set of paths of $\mathcal{K}$, and:

- *meet:* $A_{\mathbb{I}} = \{(\pi, \pi') \in \mathbb{I} \times \mathbb{I} \mid \mathsf{lst}(\pi) = \mathsf{fst}(\pi')\}$,
- *later:* $L_{\mathbb{I}} = \{(\pi, \pi') \in \mathbb{I} \times \mathbb{I} \mid (\mathsf{lst}(\pi), \mathsf{fst}(\pi')) \in R^+\}$,
- *begins-with:* $B_{\mathbb{I}} = \{(\pi, \pi') \in \mathbb{I} \times \mathbb{I} \mid \pi' \in \mathsf{Pref}(\pi)\}$,
- *ends-with:* $E_{\mathbb{I}} = \{(\pi, \pi') \in \mathbb{I} \times \mathbb{I} \mid \pi' \in \mathsf{Suff}(\pi)\}$,
- *during:* $D_{\mathbb{I}} = \{(\pi, \pi') \in \mathbb{I} \times \mathbb{I} \mid \pi'$ is an internal sub-path of $\pi\}$,
- *overlaps:* $O_{\mathbb{I}} = \{(\pi, \pi') \in \mathbb{I} \times \mathbb{I} \mid \pi' = \nu \cdot \nu'$ for some paths ν and ν' such that $\nu \in \mathsf{Suff}(\pi)$and $|\nu| > 1\}$,
- *homogeneity principle:* for all $p \in \mathcal{AP}$, $Lab_{\mathbb{I}}^{-1}(p) = \{\pi \in \mathbb{I} \mid p \in \bigcap_{i=0}^{i<|\pi|} Lab(\pi(i))\}$.

Note that for a finite Kripke structure $\mathcal{K}$, the number of paths in $\mathcal{K}$ may be infinite (this happens when $\mathcal{K}$ has loops), hence the number of intervals in $\mathcal{A}_{\mathcal{K}}$ may be infinite.

Definition 2.4 *(State-based model checking against HS).* A Kripke structure $\mathcal{K}$ over $\mathcal{AP}$ is a *model* of an HS formula ψ over $\mathcal{AP}$, written $\mathcal{K} \models \psi$, if for all *initial* paths π of $\mathcal{K}$, $\mathcal{A}_{\mathcal{K}}, \pi \models \psi$. In the following, we also write $\mathcal{K}, \pi \models \psi$ to mean $\mathcal{A}_{\mathcal{K}}, \pi \models \psi$. The *(finite) model checking problem (against HS)* consists in checking whether $\mathcal{K} \models \psi$ for a given HS formula ψ and a finite Kripke structure $\mathcal{K}$.

Example 2.2. We give now an example of property to be checked over the Kripke structure $\mathcal{K}_{Sched}$ of Example 2.1. We start defining a formula $Activity_i$ with $i \in \{1, 2, 3\}$ which precisely characterizes a path of $\mathcal{K}_{Sched}$ corresponding with the use and unlock of the shared resource by the i-th process. $Activity_i$ is defined as follows: $Activity_i := \neg p_i \wedge [B] p_i$. The formula ensures that the path underlying the interval has the form $v_i^+ \cdot u_i$ (all the proper prefixes satisfy p_i but the whole interval does not). $\mathcal{K}_{Sched}$ satisfies the property that any two activities of a process are interleaved with at least an activity of a different process, i.e.

$$\mathcal{K}_{Sched} \models \bigwedge_{i\in\{1,2,3\}} [D]\big((\langle \mathrm{B}\rangle\, Activity_i \wedge \langle \mathrm{E}\rangle\, Activity_i) \implies \langle \mathrm{D}\rangle \bigvee_{j\in\{1,2,3\},\, j\neq i} Activity_j\big)$$

Any subpath of an initial path starting and ending with an activity of the i-th process has a internal activity of another process. Notice that the formula use only linear-time modalities.

It is worth noting that we assume the *non-strict semantics of HS*, which admits intervals (paths) consisting of a single point (all the results proved in the paper hold for the strict semantics as well). Under such an assumption, all HS-temporal modalities can be expressed in terms of $\langle \mathrm{B}\rangle$, $\langle \mathrm{E}\rangle$, $\langle \overline{\mathrm{B}}\rangle$, and $\langle \overline{\mathrm{E}}\rangle$ [5]. As an example, $\langle \mathrm{D}\rangle$ can be expressed in terms of $\langle \mathrm{B}\rangle$ and $\langle \mathrm{E}\rangle$ as $\langle \mathrm{D}\rangle\,\psi := \langle \mathrm{B}\rangle\,\langle \mathrm{E}\rangle\,\psi$, while $\langle \mathrm{A}\rangle$ can be expressed in terms of $\langle \mathrm{E}\rangle$ and $\langle \overline{\mathrm{B}}\rangle$ as $\langle \mathrm{A}\rangle\,\psi := ([\mathrm{E}]\perp \wedge (\psi \vee \langle \overline{\mathrm{B}}\rangle\,\psi)) \vee \langle \mathrm{E}\rangle([\mathrm{E}]\perp \wedge (\psi \vee \langle \overline{\mathrm{B}}\rangle\,\psi))$.

Finally, observe that the temporal modalities for the Allens's relations in $\{B, E, D\}$ have a "linear-time" semantics, i.e., they allow to select only slices (subpaths) of the current timeline (path). The semantics of the temporal modalities associated with the other Allen's relations (i.e., the ones in $\{A, L, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$) is instead "branching-time" (i.e., they allow to non-deterministically extend the current timeline in the future or in the past). Hence, a non-empty subset of non-interdefinable Allen's relations in $\{B, E, D\}$ is called a *linear-time basis* $\mathcal{B}$ of HS. The possible bases are $\{B\}$, $\{E\}$, $\{D\}$, $\{B, D\}$, $\{B, E\}$, and $\{E, D\}$.

2.3. The studied HS fragments

In this subsection, we introduce the syntactical fragments which will be investigated in the rest of the paper. For each linear-time basis $\mathcal{B}$, we define a fragment of HS, denoted by $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$, obtained by combining the linear-time modalities of the basis $\mathcal{B}$ of HS with the branching-time modalities for the Allen's relations in the subset $\mathcal{F}_{\mathcal{B}}$ (depending of $\mathcal{B}$) of $\mathcal{F} = \{A, L, \overline{A}, \overline{L}, \overline{B}, \overline{E}\}$ defined as follows:

- $\mathcal{F}_{\mathcal{B}} = \mathcal{F}$ if $\mathcal{B} \in \{\{B\}, \{E\}, \{B, E\}\}$;
- $\mathcal{F}_{\{B,D\}} = \mathcal{F} \setminus \{\overline{B}\}$ and $\mathcal{F}_{\{D,E\}} = \mathcal{F} \setminus \{\overline{E}\}$;
- $\mathcal{F}_{\{D\}} = \mathcal{F} \setminus \{\overline{B}, \overline{E}\}$.

As we shall see in the next section, the basis $\mathcal{B}$ determines the structure of the descriptors (h-level $\mathcal{B}$-descriptors, for each $h \geq 0$) and the unifying MC algorithm (illustrated in Section 4) exploits a crucial property of descriptors: for a given $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula ψ, all the paths of the given Kripke structure which have the same $\mathrm{depth}_{\mathcal{B}}(\psi)$-level $\mathcal{B}$-descriptor are indistinguishable by formula ψ. In Section 3, we will show that the set of branching-time modalities in $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ is maximal with respect to the fulfillment of the previous property for the fixed basis $\mathcal{B}$. Formally, we will prove that for any extension of $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with a branching-time modality non-expressible in $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$, one can construct a formula able to distinguish, for each h, some distinct paths with the same h-level $\mathcal{B}$-descriptor.

It is worth noticing that the MC procedure in Section 4 operates on all the fragments $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ except for the one for the basis $\{D\}$. The reason is that the algorithm exploits another crucial property of descriptors which will be stated in Section 3 for all bases $\mathcal{B} \neq \{D\}$, namely, the existence of a bound on the length of the minimal representatives among the paths of a Kripke structure with the same h-level $\mathcal{B}$-descriptor for a given $h \geq 0$. We do not know whether the descriptors based on the linear-time base $\{D\}$ enjoy such a property.

We conclude this section by observing that for the (existential) branching-time modalities associated with the Allen's relations which are not in $\mathcal{F} = \{A, L, \overline{A}, \overline{L}, \overline{B}, \overline{E}\}$, that is, $\langle \overline{\mathrm{D}}\rangle$, $\langle \mathrm{O}\rangle$, and $\langle \overline{\mathrm{O}}\rangle$, the following statements hold:

- $\langle \overline{\mathrm{D}}\rangle$ can be expressed in terms of $\langle \overline{\mathrm{B}}\rangle$ and $\langle \overline{\mathrm{E}}\rangle$: $\langle \overline{\mathrm{D}}\rangle\,\psi \equiv \langle \overline{\mathrm{B}}\rangle\,\langle \overline{\mathrm{E}}\rangle\,\psi$;
- $\langle \mathrm{O}\rangle$ can be expressed in terms of $\langle \overline{\mathrm{B}}\rangle$ and $\langle \mathrm{E}\rangle$: $\langle \mathrm{O}\rangle\,\psi \equiv \langle \mathrm{E}\rangle(\langle \mathrm{E}\rangle\top \wedge \langle \overline{\mathrm{B}}\rangle\,\psi)$;
- $\langle \overline{\mathrm{O}}\rangle$ can be expressed in terms of $\langle \overline{\mathrm{E}}\rangle$ and $\langle \mathrm{B}\rangle$: $\langle \overline{\mathrm{O}}\rangle\,\psi \equiv \langle \mathrm{B}\rangle(\langle \mathrm{B}\rangle\top \wedge \langle \overline{\mathrm{E}}\rangle\,\psi)$.

Hence, we can draw the following correspondences:

- the complete basis $\{B, E\}$ corresponds to the full logic HS;
- $\mathsf{HS}_{\{B\}}(\mathcal{F}_{\{B\}})$ corresponds to $\mathsf{ALB\overline{ALBEDO}}$;
- $\mathsf{HS}_{\{E\}}(\mathcal{F}_{\{E\}})$ corresponds to $\mathsf{ALEO\overline{ALBED}}$;
- $\mathsf{HS}_{\{B,D\}}(\mathcal{F}_{\{B,D\}})$ corresponds to $\mathsf{ALBD\overline{ALEO}}$;
- $\mathsf{HS}_{\{D,E\}}(\mathcal{F}_{\{D,E\}})$ corresponds to $\mathsf{ALEDO\overline{ALB}}$.

Notice that $\mathsf{HS}_{\{B\}}(\mathcal{F}_{\{B\}})$ and $\mathsf{HS}_{\{E\}}(\mathcal{F}_{\{E\}})$ are as expressive as $\mathsf{AB\overline{ABE}}$ and $\mathsf{AE\overline{ABE}}$, respectively, investigated in [22–24]. However, we consider $\langle \mathrm{L}\rangle$ and $\langle \overline{\mathrm{L}}\rangle$ as independent modalities for the following reasons:

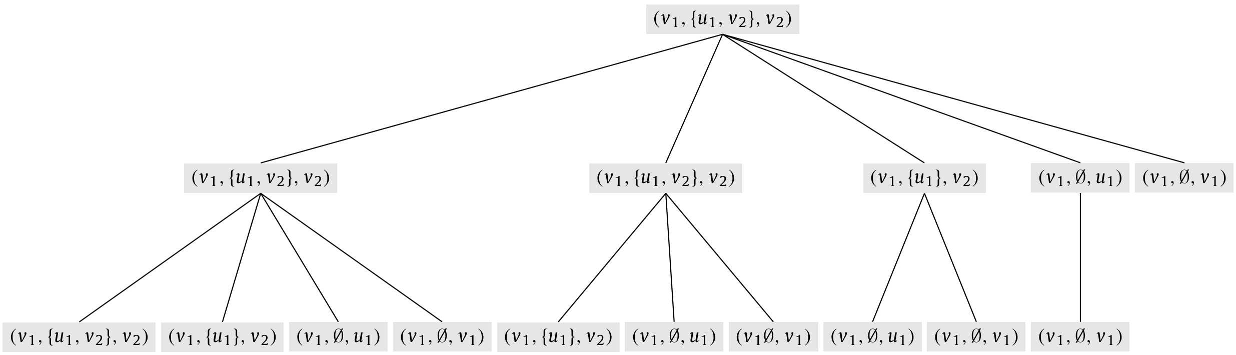

Fig. 3. The 2-level $\{B\}$-descriptor for the path $v_1u_1v_2v_2v_2v_2u_2$ of the Kripke structure of Example 2.1 representing an activity of the first process followed by one of the second.

- in the non-strict semantics, to express $\langle L\rangle$ and $\langle\overline{L}\rangle$ in $\mathsf{AB}\overline{\mathsf{ABE}}$ (resp., $\mathsf{AE}\overline{\mathsf{ABE}}$), the use of the linear-time modality $\langle B\rangle$ (resp., $\langle E\rangle$) is necessary;
- for a given linear-time basis $\mathcal{B}$, an important complexity parameter in the MC algorithm in Section 4 is the joint $\mathcal{B}$-nesting depth of the given formula.

3. Descriptor-based abstraction of paths in a finite Kripke structure

In this section, we introduce a finite abstraction of the paths of a finite Kripke structure, parameterized by a linear-time basis $\mathcal{B}$ of HS and a natural number $h \geq 0$, which will be exploited in Section 4 for devising a uniform MC algorithm for all the fragments $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with $\mathcal{B}$ distinct from $\{D\}$. The considered abstraction generalizes the one based on descriptors originally introduced in [16] for full HS. For each linear-time basis $\mathcal{B}$ and natural number $h \geq 0$, we associate with each path of the given finite Kripke structure $\mathcal{K}$ a tree object (abstraction) ranging over a finite set, called *h-level $\mathcal{B}$ descriptor*. The important property of this abstraction is that it partitions the (possibly infinite) set of $\mathcal{K}$-paths into a finite number of equivalence classes such that paths in the same equivalence class, i.e., paths with the same h-level $\mathcal{B}$-descriptor, are indistinguishable from all the formulas in $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with joint $\mathcal{B}$-nesting depth at most h (*h-depth satisfaction property*). The second fundamental property that holds for each basis $\mathcal{B} \neq \{D\}$ is that each equivalence class contains a bounded witness path whose length is bounded by the number of distinct h-level $\mathcal{B}$ descriptors (*bounded path property*).[1] In particular, to design an MC algorithm for the logic $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with $\mathcal{B} \neq \{D\}$, we shall consider *h-level $\mathcal{B}$-certificates*, namely paths which enjoy the property of having minimal length with respect to the set of paths having the same h-level $\mathcal{B}$-descriptor. In Section 5, we will provide tight bounds on the length of h-level $\mathcal{B}$-certificates.

In this section, we also formally state maximality of the considered fragments $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with respect to the $\mathcal{B}$-description-based abstraction. More precisely, we show that for each basis $\mathcal{B}$ and branching-time modality $\langle X\rangle$ non-expressible in $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$, there are a formula φ in the extension of $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with modality $\langle X\rangle$ and a finite Kripke structure $\mathcal{K}$ such that for each $h \geq 1$, formula φ is able to distinguish distinct paths of $\mathcal{K}$ with the same h-level $\mathcal{B}$-descriptor.

The rest of the section is organized as follows. In Subsection 3.1, we formally define h-level $\mathcal{B}$-descriptors and provide some example. In Subsection 3.2, we state the h-depth satisfaction property and the bounded path property. Finally, in Subsection 3.3, we show the maximality of the fragments $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with respect to the $\mathcal{B}$-descriptor abstraction.

3.1. Descriptors

In this subsection, we formally define the notion of *h-level $\mathcal{B}$-descriptor* for a linear-time basis $\mathcal{B}$ and a natural number $h \geq 0$. The definition of h-level $\mathcal{B}$-descriptors exploits *h-level Σ-terms* and *h-level bipartite Σ-terms*, where Σ denotes a given finite alphabet. Intuitively, an h-level Σ-term corresponds to an unordered finite tree of height h such that subtrees rooted at distinct children of the same node are *not* isomorphic. An h-level bipartite Σ-term is similar but, additionally, we require that each edge in the tree has a color from a set of two colors. Formally, the set of h-level Σ-terms t is inductively defined as follows:

- if $h = 0$, then $t = a$ for some $a \in \Sigma$; otherwise, t has the form (a, T) where $a \in \Sigma$ and T is a (possibly empty) subset of $(h-1)$-level Σ-terms.

The set of h-level bipartite Σ-terms t is inductively defined as follows:

[1] As already pointed out, we do not know if a bounded path property holds for the basis $\{D\}$ as well.

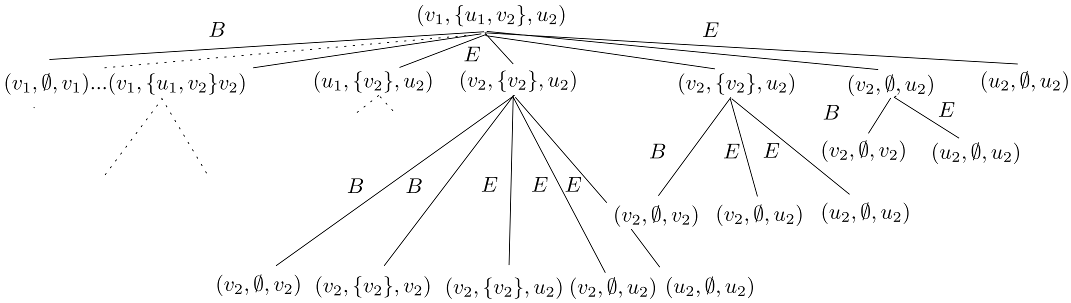

Fig. 4. A fragment of the 2-level $\{BE\}$-descriptor for the path $v_1u_1v_2v_2v_2v_2u_2$ of the Kripke structure of Example 2.1. The rightmost expanded sub-tree rooted in the node $(v_2, \{v_2\}, u_2)$ describes both the suffix $v_2v_2v_2u_2$ and the suffix $v_2v_2v_2u_2$. The suffix $v_2v_2u_2$ is described by the right sub-tree and the suffix v_2u_2 is described by the sub-tree rooted in the node $(v_2, \emptyset, u_2)$.

- if $h = 0$, then $t = a$ for some $a \in \Sigma$; otherwise t is of the form (a, T_1, T_2) where $a \in \Sigma$ and T_1 and T_2 are (possibly empty) subsets of $(h-1)$-level Σ-terms.

We say that a is the root of t. The size of an h-level (bipartite) Σ-term is the number of nodes in the associated tree representation. The following statement holds.

Remark 3.1. The number of distinct h-level Σ-terms (resp., h-level bipartite Σ-terms) over Σ is $Tower(\Theta(|\Sigma|), h)$.

Definition 3.1 *(Descriptors).* Let Σ be a finite alphabet and $\mathcal{B}$ be a linear-time basis of HS. Given a non-empty finite word w over Σ and $h \geq 0$, the h-level $\mathcal{B}$-descriptor $\mathcal{B}_h(w)$ of w is the h-level $(\Sigma \times 2^{\Sigma} \times \Sigma)$-term (resp., h-level bipartite $(\Sigma \times 2^{\Sigma} \times \Sigma)$-term) if $|\mathcal{B}| = 1$ (resp., $|\mathcal{B}| = 2$) inductively satisfying the following conditions. For the base case, i.e. $h = 0$, then $\mathcal{B}_0(w) = (\mathsf{fst}(w), \mathsf{internal}(w), \mathsf{lst}(w))$. For the induction step, i.e. $h > 0$, we have:

- Case $\mathcal{B} = \{B\}$ (resp., $\mathcal{B} = \{E\}$, resp., $\mathcal{B} = \{D\}$): $\mathcal{B}_h(w) = (\mathcal{B}_0(w), T)$ where T is the set of $(h-1)$-level $\mathcal{B}$-descriptors of the non-empty proper prefixes (resp., non-empty proper suffixes, resp., non-empty internal sub-words) of w.
- Case $\mathcal{B} = \{B, E\}$: $\mathcal{B}_h(w) = (\mathcal{B}_0(w), T_B, T_E)$ with T_B (resp., T_E) the set of $(h-1)$-level $\mathcal{B}$-descriptors of the non-empty proper prefixes (resp., non-empty proper suffixes) of w.
- Case $\mathcal{B} = \{B, D\}$ (resp., $\mathcal{B} = \{D, E\}$): as in $\mathcal{B} = \{B, E\}$ by replacing T_E (resp., T_B) with the set of $(h-1)$-level $\mathcal{B}$-descriptors of the non-empty internal sub-words of w.

For a linear-time basis $\mathcal{B} = \{X, Y\}$ (resp., $\mathcal{B} = \{X\}$), an h-level $\mathcal{B}$-descriptor is also called h-level XY-descriptor (resp., h-level X-descriptor) and for a non-empty word, we write $XY_h(w)$ (resp., $X_h(w)$) to mean $\mathcal{B}_h(w)$. For a finite Kripke structure $\mathcal{K}$, a basis $\mathcal{B}$, and $h \geq 0$, we denote by $\mathcal{B}_h(\mathcal{K})$ the *finite* set of h-level $\mathcal{B}$-descriptors associated with the paths of $\mathcal{K}$.

An example of 2-level $\{B\}$-descriptor is depicted in Fig. 3 for a path of the Kripke structure of Example 2.1. An example of $\{B, E\}$-descriptor is given Fig. 4. Intuitively, the h-level $\mathcal{B}$-descriptor $\mathcal{B}_h(\pi)$ of a Kripke structure path π has enough information for checking the fulfillment of $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formulas with joint $\mathcal{B}$-nesting depth at most h:

- to check the fulfillment of proposition letters, $\mathcal{B}_h(\pi)$ keeps tracks at each node of the set of states visited by the current sub-path of π;
- to deal with the branching-time modalities for Allen's relations in $\mathcal{F}_{\mathcal{B}}$, $\mathcal{B}_h(\pi)$ keeps tracks at each node also of the first and last states of the current sub-path;
- finally, to check the fulfillment of the linear-time modalities for the basis $\mathcal{B}$, $\mathcal{B}_h(\pi)$ keeps information about all the sub-paths of the current sub-path π' which can be obtained from π' by applying Allen's relations in the basis $\mathcal{B}$.

3.2. Properties of descriptors

In this subsection, for each basis $\mathcal{B}$ distinct from $\{D\}$, we state two properties of $\mathcal{B}$-descriptors which are fundamental to design an MC algorithm for the logic $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$. The first crucial property (*h-depth satisfaction property*), which holds for the basis $\{D\}$ as well, is that for a finite Kripke structure $\mathcal{K}$ and a natural number $h \geq 0$, paths of $\mathcal{K}$ which have the same h-level $\mathcal{B}$-descriptor satisfy the same formulas in $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with $\mathcal{B}$-nesting depth at most h (see Proposition 3.2). To prove this property, we need a preliminary *stabilization result*, namely, we have to show that the property of two paths π and π' of having the same h-level $\mathcal{B}$-descriptor is preserved by right (resp., left) concatenation with another path of $\mathcal{K}$ whenever $\mathcal{B}$ is distinct from $\{D\}$ and $\{B, D\}$ (resp., from $\{D\}$ and $\{D, E\}$). This right-stabilization (resp., left-stabilization) result is used for handling the branching-time modality $\langle\overline{\mathrm{B}}\rangle$ (resp., $\langle\overline{\mathrm{E}}\rangle$) in the fragments $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$, which include such a modality.

Proposition 3.1. *Let $h \geq 0$, $\mathcal{B} \neq \{D\}$ be a basis, and π and π' be two paths of a finite Kripke structure $\mathcal{K}$ such that $\mathcal{B}_h(\pi) = \mathcal{B}_h(\pi')$. Then, for all paths π_L and π_R of $\mathcal{K}$ such that $\pi_L \cdot \pi$ and $\pi \cdot \pi_R$ are defined, the following properties hold:*

(1) if $\mathcal{B} \neq \{B, D\}$, then $\mathcal{B}_h(\pi \cdot \pi_R) = \mathcal{B}_h(\pi' \cdot \pi_R)$;
(2) if $\mathcal{B} \neq \{D, E\}$, then $\mathcal{B}_h(\pi_L \cdot \pi) = \mathcal{B}_h(\pi_L \cdot \pi')$.

Proof. First note that since $\mathcal{B}_h(\pi) = \mathcal{B}_h(\pi')$, it holds that $\mathsf{fst}(\pi) = \mathsf{fst}(\pi')$ and $\mathsf{lst}(\pi) = \mathsf{lst}(\pi')$. Hence, $\pi_L \cdot \pi$ (resp., $\pi \cdot \pi_R$) is a path of $\mathcal{K}$ if and only if $\pi_L \cdot \pi'$ (resp., $\pi' \cdot \pi_R$) is a path. For the basis $\mathcal{B} = \{B, E\}$, the result has been proved in [16]. Here, we focus on the cases where either $\mathcal{B} = \{E\}$ (E-descriptors) or $\mathcal{B} = \{D, E\}$ (DE-descriptors). The cases where either $\mathcal{B} = \{B\}$ or $\mathcal{B} = \{B, D\}$ are similar. We prove Properties (1) and (2) by induction on h. The base case $h = 0$ is straightforward. Now, assume that $h > 0$.

First, let us consider Property (1) when either $\mathcal{B} = \{E\}$ or $\mathcal{B} = \{D, E\}$. We prove the inclusion $\mathcal{B}_h(\pi \cdot \pi_R) \subseteq \mathcal{B}_h(\pi' \cdot \pi_R)$ (the other inclusion being similar). Let ν_E be a non-empty proper suffix of $\pi \cdot \pi_R$ and ν_D be an internal sub-path of $\pi \cdot \pi_R$. We need to prove that:

(1.1) there is a non-empty proper suffix ν'_E of $\pi' \cdot \pi_R$ such that ν_E and ν'_E have the same $(h-1)$-level $\mathcal{B}$-descriptor;
(1.2) if $\mathcal{B} = \{D, E\}$ then there is an internal sub-path ν'_D of $\pi' \cdot \pi_R$ such that ν_D and ν'_D have the same $(h-1)$-level DE-descriptor.

For Condition (1.1), if ν_E is a suffix of π_R, then we set $\nu'_E = \nu_E$ and the result holds. Otherwise, there is a non-empty proper suffix ξ of π such that $\nu_E = \xi \cdot \pi_R$. Since $\mathcal{B}_h(\pi) = \mathcal{B}_h(\pi')$ and $\mathcal{B} \in \{\{E\}, \{D, E\}\}$, there is a non-empty proper suffix ξ' of π' such that ξ and ξ' have the same $(h-1)$-level $\mathcal{B}$-descriptor. Thus, by setting $\nu' = \xi' \cdot \pi_R$, by the inductive hypothesis, the result follows. Consider now Condition (1.2). We distinguish three cases:

1. ν_D is either a prefix or an internal sub-path of π_R: we set $\nu'_D = \nu_D$ and the result trivially holds.
2. ν_D is an internal sub-path (resp., proper suffix) of π: by hypothesis, there is an internal sub-path (resp., proper suffix) ν'_D of π' such that ν_D and ν'_D have the same $(h-1)$-level DE-descriptor. Being ν'_D an internal sub-path of $\pi' \cdot \pi_R$, the result follows.
3. there is a proper suffix ξ of π and a proper prefix η of π_R such that $\nu_D = \xi \cdot \eta$. By hypothesis, there is a proper suffix ξ' of π' such that ξ and ξ' have the same $(h-1)$-level DE-descriptor. By the induction hypothesis, $\xi \cdot \eta$ and $\xi' \cdot \eta$ have the same $(h-1)$-level DE-descriptor. Hence, by setting $\nu'_D = \xi' \cdot \eta$, the result follows.

Now, let us prove Property (2) for the basis $\{E\}$. Let ν be a non-empty proper suffix of $\pi_L \cdot \pi$ (the case where we choose a non-empty proper suffix of $\pi_L \cdot \pi'$ is similar). We need to show that there is a non-empty proper suffix ν' of $\pi_L \cdot \pi'$ such that ν and ν' have the same $(h-1)$-level E-descriptor. If $\nu = \pi$ or ν is a proper suffix of π, then, by hypothesis, there is a non-empty suffix ν' of π' such that ν and ν' have the same $(h-1)$-level E-descriptor. Thus, since ν' is a non-empty proper suffix of $\pi_L \cdot \pi'$, the result follows. Otherwise, there is a non-empty proper suffix ξ of π_L such that $\nu = \xi \cdot \pi$. Let $\nu' = \xi \cdot \pi'$. Since π and π' have the same $(h-1)$-level E-descriptor, by the inductive hypothesis, we obtain that ν and ν' have the same $(h-1)$-level E-descriptor, concluding the proof. □

We note that when the basis $\mathcal{B}$ is either $\{D\}$ or $\{B, D\}$ (resp., either $\{D\}$ or $\{E, D\}$), the property of two paths π and π' to have the same h-level $\mathcal{B}$-descriptor is not in general preserved by right (resp., left) concatenation with another path of $\mathcal{K}$. As an example, for the bases $\{D\}$ and $\{B, D\}$, let $\mathcal{K}$ be a Kripke structure consisting of three states s_1, s_2, and s_3 such that (s_i, s_j) is an edge of $\mathcal{K}$ for all $1 \leq i, j \leq 3$. Let us consider the two paths $\pi = s_1 s_2^3 (s_2 s_3)^3 s_1$ and $\pi' = s_1 s_2^3 (s_3 s_2)^3 s_1$. One can check that π and π' have the same 1-level BD-descriptor and the same 1-level D-descriptor. On the other hand, $\pi \cdot s_1$ and $\pi' \cdot s_1$ have distinct 1-level BD-descriptors and distinct 1-level D-descriptors: in particular, while $\pi \cdot s_1$ has the internal sub-word $s_3 s_1$, there is no internal sub-word ν' of $\pi' \cdot s_1$ such that $\mathsf{fst}(\nu') = s_3$, $\mathsf{lst}(\nu') = s_1$, and $\mathsf{internal}(\nu') = \emptyset$.

By exploiting Proposition 3.1, we can now prove the following key property.

Proposition 3.2 *(h-Depth satisfaction property). Let $h \geq 0$, $\mathcal{B}$ a basis, and π and π' be two paths of a finite Kripke structure $\mathcal{K}$ such that $\mathcal{B}_h(\pi) = \mathcal{B}_h(\pi')$. Then, for each $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula ψ with $\mathrm{depth}_{\mathcal{B}}(\psi) \leq h$, it holds that $\mathcal{K}, \pi \models \psi$ iff $\mathcal{K}, \pi' \models \psi$.*

Proof. We prove the thesis by a nested induction on the structure of the formula ψ and the joint nesting depth $\mathrm{depth}_{\mathcal{B}}(\psi)$. For the base case, ψ is an atomic proposition, and since π and π' have the same h-level $\mathcal{B}$-descriptor (in particular, π and π' visit the same states), the result follows. Now, let us consider the inductive case. The cases where the root modality of ψ is a Boolean connective directly follow from the inductive hypothesis. As for the cases where the root modality is in $\{\langle \mathrm{A} \rangle, \langle \overline{\mathrm{A}} \rangle, \langle \mathrm{L} \rangle, \langle \overline{\mathrm{L}} \rangle\}$, the result follows from the fact that, being $\mathcal{B}_h(\pi) = \mathcal{B}_h(\pi')$, we have that $\mathsf{fst}(\pi) = \mathsf{fst}(\pi')$ and $\mathsf{lst}(\pi) = \mathsf{lst}(\pi')$. If, instead, the root of ψ is in $\{\langle \overline{\mathrm{B}} \rangle, \langle \overline{\mathrm{E}} \rangle\}$, the result easily follows from Proposition 3.1 and the inductive hypothesis on the structure of the formula. It remains to consider the cases where the root modality is in $\mathcal{B}$. Here, we focus one the case where $\psi = \langle \mathrm{D} \rangle \varphi$ (the other cases, where the root modality is $\langle \mathrm{B} \rangle$ or $\langle \mathrm{E} \rangle$ are similar). Hence, either $\mathcal{B} = \{D\}$, or $\mathcal{B} = \{D, B\}$, or $\mathcal{B} = \{D, E\}$. We prove the implication $\mathcal{K}, \pi \models \psi \Rightarrow \mathcal{K}, \pi' \models \psi$ (the converse implication can be dealt with similarly). Assume that $\mathcal{K}, \pi \models \psi$, where $\psi = \langle \mathrm{D} \rangle\, \varphi$. Since $0 < \mathrm{depth}_{\mathcal{B}}(\psi) \leq h$, it holds that $h > 0$, and since $\mathcal{K}, \pi \models \langle \mathrm{D} \rangle\, \varphi$, there is an internal sub-path ν of π such that $\mathcal{K}, \nu \models \varphi$. Since π and π' have the same h-level $\mathcal{B}$-descriptor, there exists an internal sub-path ν' of ρ' such that ν and ν' have the same $(h-1)$ level $\mathcal{B}$-descriptor. Being $\mathrm{depth}_{\mathcal{B}}(\varphi) \leq h-1$, by the inductive hypothesis $\mathcal{K}, \nu' \models \varphi$ holds. Hence, $\mathcal{K}, \pi' \models \langle \mathrm{D} \rangle\, \varphi$ and the thesis follows. □

By exploiting Proposition 3.1, we can now state *a bounded path property* for each basis $\mathcal{B} \neq \{D\}$ and natural number $h \geq 0$, which intuitively provides a bounded witness path for each h-level $\mathcal{B}$-descriptor associated with an arbitrary path of a finite Kripke structure. The bounded path property is crucial in Section 4 to design the MC algorithm for the logic $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$. We do not know if such a property holds for the basis $\{D\}$ as well.

Proposition 3.3 *(Bounded path property). Let $\mathcal{B} \neq \{D\}$ be a basis, $\mathcal{K}$ be a finite Kripke structure, $h \geq 0$, and π be a $\mathcal{K}$-path. Then, there exists a path π' with the same h-level $\mathcal{B}$-descriptor as π and whose length is bounded by $|\mathcal{B}_h(\mathcal{K})|$ (i.e., the number of distinct h-level $\mathcal{B}$-descriptors of the $\mathcal{K}$-paths).*

Proof. Let $|\pi| = n$ and $\mathcal{B}$ be a basis distinct from $\{D\}$. We first consider the case where $\mathcal{B} \neq \{B, D\}$. Since there are n distinct non-empty prefixes of π, if $n > |\mathcal{B}_h(\mathcal{K})|$, then π can be written in the form $\pi = \nu \cdot \nu' \cdot \nu''$, where $|\nu| > 0$, $|\nu'| > 0$, and ν and $\nu \cdot \nu'$ have the same h-level $\mathcal{B}$-descriptor. By Property (1) of Proposition 3.1, the strictly smaller path $\nu \cdot \nu''$ has the same h-level $\mathcal{B}$-descriptor as π. We can iterate such a contraction process until there are no more pairs of prefixes with the same h-level $\mathcal{B}$-descriptor proving the statement for $\mathcal{B} \neq \{B, D\}$. The case for the basis $\{B, D\}$ is handled similarly by considering the non-empty suffixes of π and by applying Property (2) of Proposition 3.1. □

By exploiting Propositions 3.2 and 3.3, we can introduce bounded minimal representatives ($\mathcal{B}$-certificates) of paths which are used in the MC algorithm defined in Section 4.

Definition 3.2 *($\mathcal{B}$-certificate).* Given a basis $\mathcal{B} \neq \{D\}$, a finite Kripke structure $\mathcal{K}$, and $h \geq 0$, an *h-level $\mathcal{B}$-certificate of $\mathcal{K}$* is a path π of $\mathcal{K}$ such that there is no path π' so that $|\pi'| < |\pi|$ and π and π' have the same h-level $\mathcal{B}$-descriptor. Given an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula φ, a *$\mathcal{B}$-certificate for $(\mathcal{K}, \varphi)$* is an h-level $\mathcal{B}$-certificate of $\mathcal{K}$, where $h = \mathrm{depth}_{\mathcal{B}}(\varphi)$.

Notice that, by Proposition 3.3, an upper bound on the length of $\mathcal{B}$-certificates for $(\mathcal{K}, \varphi)$ is $|\mathcal{B}_h(\mathcal{K})|$, with $h = \mathrm{depth}_{\mathcal{B}}(\varphi)$.

3.3. Maximality of the fragments $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ with respect to the descriptor-based abstraction

In this subsection, we show that for each basis $\mathcal{B}$, the fragment $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ is maximal with respect to the finite abstraction based on $\mathcal{B}$-descriptors. Formally, we prove that there is a finite Kripke structure $\mathcal{K}$ over a set of four propositions such that for each basis $\mathcal{B}$ and extension $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{X\})$ of the fragment $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ (with a branching-time modality $\langle \mathrm{X} \rangle$ that cannot be expressed in $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$), one can construct an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{X\})$ formula that, for all $n \geq 1$, distinguishes some distinct paths of $\mathcal{K}$ with the same n-level $\mathcal{B}$-descriptor.

Let $\mathcal{AP} = \{p_0, p_1, p_2, p_3\}$ be a set of 4 proposition letters and $\mathcal{K}$ be a finite Kripke structure consisting of 4 states s_0, s_1, s_2, and s_3, where s_0 is the initial state, (s_i, s_j) is an edge of $\mathcal{K}$, for all $0 \leq i, j \leq 3$, and the labeling Lab is such that $Lab(s_i) = \{p_i\}$ for all $0 \leq i \leq 3$. We first exhibit two families $(\pi_n)_{n \geq 1}$ and $(\pi'_n)_{n \geq 1}$ of $\mathcal{K}$-paths such that π_n and π'_n are distinct and have the same n-level $\mathcal{B}$-descriptor, for all $n \geq 1$ and basis $\mathcal{B} \in \{\{B\}, \{D\}, \{B, D\}\}$.

Proposition 3.4. *For all $n \geq 1$, let π_n and π'_n be the paths of $\mathcal{K}$ given by*

$$\pi_n := s_0 s_2^{n+2} (s_2 s_3)^{n+2} s_1 \qquad \pi'_n = s_0 s_2^{n+2} (s_3 s_2)^{n+2} s_1$$

Then, the following statements hold:

1. *for each basis $\mathcal{B} \in \{\{B\}, \{D\}, \{B, D\}\}$, π_n and π'_n have the same n-level $\mathcal{B}$-descriptor;*
2. *for each basis $\mathcal{B} \in \{\{E\}, \{D\}, \{D, E\}\}$, $(\pi_n)^R$ and $(\pi'_n)^R$ have the same n-level $\mathcal{B}$-descriptor,*

where w^R is the reverse of w, for any finite word w.

Proof. We focus on Property 1 (the proof of Property 2 is similar: it follows from Property 1 and duality in the role of prefixes and suffixes). We first prove the following claim.

Claim. Let $n \geq 1$, $\mathcal{B} \in \{\{B\}, \{D\}, \{B, D\}\}$, $\rho_0 \in \{\varepsilon, s_0\}$, and $\rho_3 \in \{\varepsilon, s_3\}$. Then, for all $k \in [0, n-1]$, the following paths have the same k-level $\mathcal{B}$-descriptor.

- $\rho_0 \cdot s_2^{n+3} \cdot (s_3 s_2)^i \cdot \rho_3$ and $\rho_0 \cdot s_2^{n+2} \cdot (s_3 s_2)^i \cdot \rho_3$, for all $i \in [0, n+1]$;

Table 2
Extension $\mathcal{F}_{\mathcal{B}}^{ext}$ of the set $\mathcal{F}_{\mathcal{B}}$, for each basis $\mathcal{B}$.

	$\mathcal{F}_{\mathcal{B}}$	$\mathcal{F}_{\mathcal{B}}^{ext}$
$\mathcal{B} = \{B, E\}$	$\{A, L, \overline{A}, \overline{L}, \overline{B}, \overline{E}\}$	$\{A, L, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$
$\mathcal{B} = \{B\}$	$\{A, L, \overline{A}, \overline{L}, \overline{B}, \overline{E}\}$	$\{A, L, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$
$\mathcal{B} = \{E\}$	$\{A, L, \overline{A}, \overline{L}, \overline{B}, \overline{E}\}$	$\{A, L, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}\}$
$\mathcal{B} = \{B, D\}$	$\{A, L, \overline{A}, \overline{L}, \overline{E}\}$	$\{A, L, \overline{A}, \overline{L}, \overline{E}, \overline{O}\}$
$\mathcal{B} = \{D, E\}$	$\{A, L, \overline{A}, \overline{L}, \overline{B}\}$	$\{A, L, O, \overline{A}, \overline{L}, \overline{B}\}$
$\mathcal{B} = \{D\}$	$\{A, L, \overline{A}, \overline{L}\}$	$\{A, L, \overline{A}, \overline{L}\}$

- $\rho_0 \cdot s_2^{n+2} \cdot (s_2 s_3)^{n+2}$ and $\rho_0 \cdot s_2^{n+2} \cdot (s_3 s_2)^{n+1} \cdot s_3$;
- $\rho_0 \cdot s_2^{n+2} \cdot (s_2 s_3)^{n+1} \cdot s_2$ and $\rho_0 \cdot s_2^{n+2} \cdot (s_3 s_2)^{n+2}$;
- $(s_2 s_3)^{n+2}$ and $(s_2 s_3)^{n+1}$;
- $(s_3 s_2)^{n+1}$ and $(s_3 s_2)^{n+2}$.

Proof of the claim. The proof is by a straightforward induction on $k \in [0, n-1]$. □

Let $n \geq 1$ and $\mathcal{B} \in \{\{B\}, \{D\}, \{B, D\}\}$. By the previous claim and definition of the paths π_n and π'_n, it easily follows that $\mathcal{B}_0(\pi_n) = \mathcal{B}_0(\pi'_n)$, and (i) for each prefix (resp., internal sub-path) ν of π_n, there is a prefix (resp., internal sub-path) ν' of π'_n such that $\mathcal{B}_{n-1}(\pi_n) = \mathcal{B}_{n-1}(\pi'_n)$, and (i) for each prefix (resp., internal sub-path) ν' of π'_n, there is a prefix (resp., internal sub-path) ν of π_n such that $\mathcal{B}_{n-1}(\pi_n) = \mathcal{B}_{n-1}(\pi'_n)$. Hence, π_n and π'_n have the same n-level $\mathcal{B}$-descriptor, and the result follows. □

Let $\mathcal{F}^{ext}$ be the set of Allen's relations associated with all the branching-time modalities, that is, $\mathcal{F}^{ext} = \{A, L, O, \overline{A}, \overline{L}, \overline{B}, \overline{E}, \overline{D}, \overline{O}\}$. For each basis $\mathcal{B}$, we consider an extension $\mathcal{F}_{\mathcal{B}}^{ext}$ of $\mathcal{F}_{\mathcal{B}}$ with some relations in $\mathcal{F}^{ext}$ as shown in Table 2. In Subsection 2.3, we have proved that for each basis $\mathcal{B}$, the fragment $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}}^{ext})$ corresponds to $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$.

By exploiting Proposition 3.4, we show that for each basis $\mathcal{B}$ and each Allen relation X for branching-time modalities such that $X \notin \mathcal{F}_{\mathcal{B}}^{ext}$, there is an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{X\})$ formula $\varphi_{\mathcal{B},X}$ so that for all $n \geq 1$, $\varphi_{\mathcal{B},X}$ distinguishes distinct paths of $\mathcal{K}$ with the same n-level $\mathcal{B}$-descriptor.

Proposition 3.5. *For all $n \geq 1$, let π_n and π'_n be the paths of the Kripke structure $\mathcal{K}$ in Proposition 3.4. Then, for all bases $\mathcal{B}$ and Allen's relations X associated with branching-time modalities such that $X \notin \mathcal{F}_{\mathcal{B}}^{ext}$, there is an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{X\})$ formula $\varphi_{\mathcal{B},X}$ satisfying the following properties:*

- *if $\mathcal{B} \in \{\{B\}, \{B, D\}\}$, then $\mathcal{K}, \pi_n \not\models \varphi_{\mathcal{B},X}$ and $\mathcal{K}, \pi'_n \models \varphi_{\mathcal{B},X}$, for all $n \geq 1$;*
- *if $\mathcal{B} \in \{\{E\}, \{D, E\}\}$, then $\mathcal{K}, (\pi_n)^R \not\models \varphi_{\mathcal{B},X}$ and $\mathcal{K}, (\pi'_n)^R \models \varphi_{\mathcal{B},X}$, for all $n \geq 1$;*
- *if $\mathcal{B} = \{D\}$, then* either *$\mathcal{K}, \pi_n \not\models \varphi_{\mathcal{B},X}$ and $\mathcal{K}, \pi'_n \models \varphi_{\mathcal{B},X}$, for all $n \geq 1$,* or *$\mathcal{K}, (\pi_n)^R \not\models \varphi_{\mathcal{B},X}$ and $\mathcal{K}, (\pi'_n)^R \models \varphi_{\mathcal{B},X}$, for all $n \geq 1$.*

Proof. Since Allen's relations $A, \overline{A}, L$, and $\overline{L}$ are in the set $\mathcal{F}_{\mathcal{B}}$ (and thus in $\mathcal{F}_{\mathcal{B}}^{ext}$ as well) for each basis $\mathcal{B}$, we can assume that $X \in \{D, O, \overline{B}, \overline{E}, \overline{O}\}$. In the proof, we crucially exploit the fact that for all $n \geq 1$, the suffix of length 2 of π_n is $s_3 s_1$ while the suffix of length 2 of π'_n is $s_2 s_1$. We consider each Allen relation $X \in \{D, O, \overline{B}, \overline{E}, \overline{O}\}$ in the order $\overline{B}, \overline{E}, D, O, \overline{O}$.

Case $X = \overline{B}$. The bases $\mathcal{B}$ such that $\overline{B} \notin \mathcal{F}_{\mathcal{B}}^{ext}$ are $\{D\}$ and $\{B, D\}$. We define an $\mathsf{AD\overline{A}\overline{B}}$ formula $\varphi_{\overline{B}}$ (hence, $\varphi_{\overline{B}}$ is in the fragment $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{\overline{B}\})$, for each $\mathcal{B} \in \{\{D\}, \{B, D\}\}$), which distinguishes the paths π_n and π'_n, for all $n \geq 1$:

$$\varphi_{\overline{B}} := \langle \overline{\mathrm{B}} \rangle \Big(\langle \mathrm{A} \rangle\, p_1 \wedge [\mathrm{D}](\langle \overline{\mathrm{A}} \rangle\, p_1 \rightarrow \langle \mathrm{A} \rangle\, p_1) \wedge \langle \mathrm{D} \rangle \big(\langle \mathrm{A} \rangle\, p_1 \wedge \langle \mathrm{D} \rangle \top \wedge [\mathrm{D}]\, p_2 \wedge [\mathrm{D}]\,[\mathrm{D}] \bot \big) \Big)$$

Let π be a path $\mathcal{K}$ such that $|\pi| > 2$, $\mathsf{lst}(\pi) = s_1$, and $s_1 \notin \mathsf{internal}(\pi)$. Note that the paths π_n and π'_n satisfy this condition for all $n \geq 1$. By construction, it holds that π satisfies $\varphi_{\overline{B}}$ if and only if there is a right extension $\pi \cdot \pi_R$ (with π_R being non-empty) of π such that

(i) $\pi \cdot \pi_R$ ends at state s_1 and each internal sub-path of $\pi \cdot \pi_R$ which starts at state s_1 ends at state s_1 as well;
(ii) there is an internal sub-path ν of $\pi \cdot \pi_R$ of length 3 or 4 ending at state s_1 such that each internal sub-path of ν visits only state s_2.

Condition (i) implies that $\pi_R = s_1^k$ for some $k \geq 1$ (i.e., π_R visits only state s_1). Since $s_1 \notin \mathsf{internal}(\pi)$, it follows that the internal sub-path ν in Condition (ii) corresponds to a suffix of π. This suffix ν of π has length 3 or 4 and $\mathsf{internal}(\nu) = \{s_2\}$. By construction of paths π_n and π'_n, we obtain that $\mathcal{K}, \pi_n \not\models \varphi_{\overline{B}}$ and $\mathcal{K}, \pi'_n \models \varphi_{\overline{B}}$ for all $n \geq 1$, and the result for Allen relation $\overline{B}$ follows.

Case $X = \overline{E}$**.** The bases $\mathcal{B}$ such that $\overline{E} \notin \mathcal{F}_{\mathcal{B}}^{ext}$ are $\{D\}$ and $\{D, E\}$. Let $\varphi_{\overline{E}}$ be the following $\mathsf{AD\overline{A}E}$ formula:

$$\varphi_{\overline{E}} := \langle\overline{\mathrm{E}}\rangle\Big(\langle\overline{\mathrm{A}}\rangle\, p_1 \;\wedge\; [\mathrm{D}](\langle\mathrm{A}\rangle\, p_1 \rightarrow \langle\overline{\mathrm{A}}\rangle\, p_1) \wedge \langle\mathrm{D}\rangle\big(\langle\overline{\mathrm{A}}\rangle\, p_1 \wedge \langle\mathrm{D}\rangle\top \wedge [\mathrm{D}]\, p_2 \wedge [\mathrm{D}]\,[\mathrm{D}]\,\bot\big)\Big)$$

Note that $\varphi_{\overline{E}}$ is obtained from $\varphi_{\overline{B}}$ by replacing each occurrence of modality $\langle\overline{\mathrm{B}}\rangle$ (resp., $\langle\mathrm{A}\rangle$, $\langle\overline{\mathrm{A}}\rangle$) by $\langle\overline{\mathrm{E}}\rangle$ (resp., $\langle\overline{\mathrm{A}}\rangle$, $\langle\mathrm{A}\rangle$). Since for all $n \geq 1$, $\mathcal{K}, \pi_n \not\models \varphi_{\overline{B}}$ and $\mathcal{K}, \pi_n' \models \varphi_{\overline{B}}$, by duality, we obtain that $\mathcal{K}, (\pi_n)^R \not\models \varphi_{\overline{E}}$ and $\mathcal{K}, (\pi_n')^R \models \varphi_{\overline{E}}$, for all $n \geq 1$. Thus, being $\mathsf{AD\overline{A}E}$ a fragment of $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{\overline{E}\})$, for each $\mathcal{B} \in \{\{D\}, \{D, E\}\}$, the result for Allen relation $\overline{E}$ follows.

Case $X = \overline{D}$**.** The bases $\mathcal{B}$ such that $\overline{D} \notin \mathcal{F}_{\mathcal{B}}^{ext}$ are $\{D\}$, $\{B, D\}$, and $\{D, E\}$. Here, we focus on the bases $\{D\}$ and $\{B, D\}$ (the proof for the basis $\{D, E\}$ is similar and we omit the details). We define an $\mathsf{AD\overline{A}\overline{D}}$ formula $\varphi_{\overline{D}}$ (hence, $\varphi_{\overline{D}}$ is a formula in the fragment $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{\overline{D}\})$, for each $\mathcal{B} \in \{\{D\}, \{B, D\}\}$), which distinguishes the paths π_n and π_n', for all $n \geq 1$:

$$\begin{aligned}\varphi_{\overline{D}} := \;&\langle\overline{\mathrm{D}}\rangle\Big(\langle\mathrm{A}\rangle\, p_1 \;\wedge\; [\mathrm{D}](\langle\overline{\mathrm{A}}\rangle\, p_1 \rightarrow \langle\mathrm{A}\rangle\, p_1) \;\wedge\; \langle\overline{\mathrm{A}}\rangle\, p_0 \;\wedge\; [\mathrm{D}](\langle\mathrm{A}\rangle\, p_0 \rightarrow \langle\overline{\mathrm{A}}\rangle\, p_0)\\ &\langle\mathrm{D}\rangle\big(\langle\mathrm{A}\rangle\, p_1 \wedge \langle\mathrm{D}\rangle\top \wedge [\mathrm{D}]\, p_2 \wedge [\mathrm{D}]\,[\mathrm{D}]\,\bot\big)\Big)\end{aligned}$$

Let π be a path $\mathcal{K}$ such that $|\pi| > 2$, $\mathsf{fst}(\pi) = s_0$, $\mathsf{lst}(\pi) = s_1$, and $s_0, s_1 \notin \mathsf{internal}(\pi)$. Note that the paths π_n and π_n' satisfy this condition for all $n \geq 1$. By construction, it holds that π satisfies $\varphi_{\overline{D}}$ if and only if there is a left-right extension $\pi_L \cdot \pi \cdot \pi_R$ of π, where both π_L and π_R are non-empty, such that

(i) $\pi_L \cdot \pi \cdot \pi_R$ ends at state s_1 and each internal sub-path of $\pi_L \cdot \pi \cdot \pi_R$ which starts at state s_1 ends at state s_1 as well;
(ii) $\pi_L \cdot \pi \cdot \pi_R$ starts at state s_0 and each internal sub-path of $\pi_L \cdot \pi \cdot \pi_R$ which ends at state s_0 starts at state s_0 as well;
(iii) there is an internal sub-path ν of $\pi_L \cdot \pi \cdot \pi_R$ of length 3 or 4 ending at state s_1 such that each internal sub-path of ν visits only state s_2.

Conditions (i) and (ii) imply that $\pi_L = s_0^h$ and $\pi_R = s_1^k$, for some $h, k \geq 1$. Since $s_1 \notin \mathsf{internal}(\pi)$, it follows that the internal sub-path ν in Condition (iii) corresponds to a suffix of π. This suffix ν of π has length 3 or 4 and $\mathsf{internal}(\nu) = \{s_2\}$. By construction of paths π_n and π_n' (recall that the suffix of length 2 of π_n is $s_3 s_1$ while the suffix of length 2 of π_n' is $s_2 s_1$), we obtain that $\mathcal{K}, \pi_n \not\models \varphi_{\overline{D}}$ and $\mathcal{K}, \pi_n' \models \varphi_{\overline{D}}$ for all $n \geq 1$, and the result follows.

Case $X = O$**.** Here, we have to consider the bases $\{B\}$, $\{B, D\}$, and $\{D\}$. Let φ_{BO} and φ_{DO} be the BO and $\mathsf{ADO\overline{A}}$ formulas, respectively, defined as follows:

$$\begin{aligned}\varphi_{BO} &:= \langle\mathrm{O}\rangle(\langle\mathrm{B}\rangle\, p_2 \;\wedge\; [\mathrm{B}]\,[\mathrm{B}]\,[\mathrm{B}]\,\bot)\\ \varphi_{DO} &:= \langle\mathrm{O}\rangle\big(\langle\mathrm{D}\rangle(\langle\overline{\mathrm{A}}\rangle\, p_2 \wedge \langle\mathrm{A}\rangle\, s_1) \;\wedge\; [\mathrm{D}]\,[\mathrm{D}]\,\bot\big)\end{aligned}$$

Recall that modality $\langle\mathrm{O}\rangle$ allows one to select a right extension of a proper suffix of length at least 2 of the given path. Thus, formula φ_{BO} is fulfilled by a path π of $\mathcal{K}$ if and only if there is a suffix ν of π of length 2 such that $\mathsf{fst}(\nu) = s_2$. On the other hand, formula φ_{DO} is fulfilled by a path π of $\mathcal{K}$ if and only if there is a suffix ν of π of length 2 or 3 which can be extended to the right into a path $\nu \cdot \nu_R$ of length 4 such that the unique internal sub-path of $\nu \cdot \nu_R$ of length 2 starts at state s_2 and ends at state s_1. By construction, we have that for all $\psi \in \{\varphi_{BO}, \varphi_{DO}\}$, $\mathcal{K}, \pi_n \not\models \psi$ and $\mathcal{K}, \pi_n' \models \psi$, and the result for Allen relation O follows.

Case $X = \overline{O}$**.** This case is the dual of the previous case and the bases to examine are $\{E\}$, $\{D, E\}$, and $\{D\}$. Let $\varphi_{E\overline{O}}$ and $\varphi_{D\overline{O}}$ be the $\mathsf{E\overline{O}}$ and $\mathsf{AD\overline{A}\overline{O}}$ formulas, respectively, defined as follows:

$$\begin{aligned}\varphi_{E\overline{O}} &:= \langle\overline{\mathrm{O}}\rangle(\langle\mathrm{E}\rangle\, p_2 \;\wedge\; [\mathrm{E}]\,[\mathrm{E}]\,[\mathrm{E}]\,\bot)\\ \varphi_{D\overline{O}} &:= \langle\overline{\mathrm{O}}\rangle\big(\langle\mathrm{D}\rangle(\langle\mathrm{A}\rangle\, p_2 \wedge \langle\overline{\mathrm{A}}\rangle\, s_1) \;\wedge\; [\mathrm{D}]\,[\mathrm{D}]\,\bot\big)\end{aligned}$$

By construction, for all $\psi \in \{\varphi_{E\overline{O}}, \varphi_{D\overline{O}}\}$, $\mathcal{K}, (\pi_n)^R \not\models \psi$ and $\mathcal{K}, (\pi_n')^R \models \psi$, and the result for Allen relation $\overline{O}$ follows. This concludes the proof. □

By Propositions 3.4–3.5, we obtain the main result of the subsection.

Theorem 3.1. *There is a finite Kripke structure $\mathcal{K}$ over a set of 4 proposition letters such that, for all bases $\mathcal{B}$ and Allen relations X associated with branching-time modalities with $X \notin \mathcal{F}_{\mathcal{B}}^{ext}$, there is a $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}} \cup \{X\})$ formula $\varphi_{\mathcal{B},X}$ so that for all $n \geq 1$, $\mathcal{K}, \varphi_{\mathcal{B},X} \models \nu_n$ and $\mathcal{K}, \varphi_{\mathcal{B},X} \not\models \nu_n'$ for some paths ν_n and ν_n' of $\mathcal{K}$ with the same n-level $\mathcal{B}$-descriptor.*

4. Decision procedures based on descriptors

In this section, by exploiting Propositions 3.1–3.3, we design an alternating-time MC algorithm for the logic $\mathsf{HS}_{\mathcal{B}}(\mathcal{F})$, for each basis $\mathcal{B} \neq \{D\}$. Recall that $\mathcal{F}_{\mathcal{B}} = \mathcal{F}$ if $\mathcal{B} \in \{\{B\}, \{E\}, \{B, E\}\}$, $\mathcal{F}_{\{B,D\}} = \mathcal{F} \setminus \{\overline{B}\}$, and $\mathcal{F}_{\{D,E\}} = \mathcal{F} \setminus \{\overline{B}\}$, where $\mathcal{F} = \{A, \overline{A}, \overline{B}, \overline{E}, L, \overline{L}\}$. The structure of the algorithm is similar to the exponential-time bounded alternating procedure given in [24] for MC against the fragments $\mathsf{AB\overline{A}\overline{B}\overline{E}}$ and $\mathsf{AE\overline{A}\overline{B}\overline{E}}$ (which have the same expressiveness as $\mathsf{HS}_{\{B\}}(\mathcal{F}_{\{B\}})$ and $\mathsf{HS}_{\{E\}}(\mathcal{F}_{\{E\}})$, respectively), where the notion of certificate is different from the one exploited here.

The MC algorithm we propose is described by the procedure described in Fig. 5, which is parametric in the basis $\mathcal{B}$. For the sake of complexity analysis, such a procedure can be easily translated into a time-bounded Alternating Turing Machine (ATM) deciding the MC problem for the logic $\mathsf{HS}_{\mathcal{B}}(\mathcal{F})$.

We assume that $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formulas are in *NNF*. As complexity measures of a formula φ, we consider the size $|\varphi|$ and the standard *alternation depth*, denoted by $\Upsilon(\varphi)$, between the existential $\langle \mathrm{X} \rangle$ and universal modalities $[\mathrm{X}]$ occurring in the *NNF* of φ for $X \in \{\overline{B}, \overline{E}\}$. Formally, we establish the following result, where $\mathsf{MC}_{\mathcal{B}}$ is the set of pairs $(\mathcal{K}, \varphi)$ consisting of a finite Kripke structure $\mathcal{K}$ and an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula φ such that $\mathcal{K} \models \varphi$.

Proposition 4.1. *For each basis $\mathcal{B} \neq \{D\}$, the parametric procedure in Fig. 5 describes a time-bounded ATM deciding $\mathsf{MC}_{\mathcal{B}}$ which, given an input $(\mathcal{K}, \varphi)$, has a number of alternations (between existentially and universal choices) at most $\Upsilon(\varphi) + 2$ and runs in time $M_{\mathcal{B}}(\mathcal{K}, \varphi)^{O(|\varphi|^d)}$, where $M_{\mathcal{B}}(\mathcal{K}, \varphi)$ is the maximal length of a $\mathcal{B}$-certificate for the input, and $d = 2$ if $D \in \mathcal{B}$ and $d = 1$ otherwise.*

In the following, we describe the procedure and then, in Subsection 4.1, we prove the statement of Proposition 4.1. To this end, we introduce some auxiliary notation. Let us fix a finite Kripke structure $\mathcal{K}$ with transition relation R and an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula φ in *NNF*. We denote by $\mathsf{SD}(\varphi)$ the set consisting of the subformulas ψ of φ and the *duals* $\widetilde{\psi}$. Let π be a $\mathcal{B}$-certificate for $(\mathcal{K}, \varphi)$ and $h = \mathrm{depth}_{\mathcal{B}}(\varphi)$. For each $X \in \mathcal{B} \cup \mathcal{F}_{\mathcal{B}}$, *an X-witness of π for $(\mathcal{K}, \varphi)$* is a path π' of $\mathcal{K}$ satisfying the following conditions:

- case $X \in \mathcal{B}$: π' is a non-empty proper prefix (resp., non-empty proper suffix, resp., non-empty internal subpath) of π if $X = B$ (resp., $X = E$, resp., $X = D$);
- case $X \in \mathcal{F}_{\mathcal{B}}$: π' is a $\mathcal{B}$-certificate for $(\mathcal{K}, \varphi)$ such that:
  - case $X = A$: $\mathsf{fst}(\pi') = \mathsf{lst}(\pi)$;
  - case $X = \overline{A}$: $\mathsf{lst}(\pi') = \mathsf{fst}(\pi)$;
  - case $X = L$: $(\mathsf{lst}(\pi), \mathsf{fst}(\pi')) \in R^+$;
  - case $X = \overline{L}$: $(\mathsf{lst}(\pi'), \mathsf{fst}(\pi)) \in R^+$;
  - case $X = \overline{B}$: π' has the same h-level $\mathcal{B}$-descriptor of a path of the form $\pi \cdot \pi''$ for some *$\mathcal{B}$-certificate* π'' of $(\mathcal{K}, \varphi)$;
  - case $X = \overline{E}$: π' has the same h-level $\mathcal{B}$-descriptor of a path of the form $\pi'' \cdot \pi$ for some *$\mathcal{B}$-certificate* π'' of $(\mathcal{K}, \varphi)$.

Note that when $X \in \{A, L\}$ (resp., $X \in \{\overline{A}, \overline{L}\}$), the set of X-witnesses of π for $(\mathcal{K}, \varphi)$ coincides with the set of X-witnesses of the one-length path $\mathsf{lst}(\pi)$ (resp., $\mathsf{fst}(\pi)$) for $(\mathcal{K}, \varphi)$. By Propositions 3.1–3.3, we can prove the following key property.

Proposition 4.2. *Let $\mathcal{B} \neq \{D\}$ be a basis, φ an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula in NNF, $\mathcal{K}$ a finite Kripke structure, and π a $\mathcal{B}$-certificate for $(\mathcal{K}, \varphi)$. Then, for each $\langle \mathrm{X} \rangle\, \psi \in \mathsf{SD}(\varphi)$, $\mathcal{K}, \pi \models \langle \mathrm{X} \rangle\, \psi$ if and only if there is an X-witness π' of π for $(\mathcal{K}, \varphi)$ such that $\mathcal{K}, \pi' \models \psi$.*

Proof. The result for $X \in \mathcal{B} \cup \{A, \overline{A}, L, \overline{L}\}$ directly follows from the semantics of HS and Proposition 3.2. Now, assume that $X \in \{\overline{B}, \overline{E}\}$ and let $h = \mathrm{depth}_{\mathcal{B}}(\varphi)$. We focus on the case where $X = \overline{E}$ (the case $X = \overline{B}$ can be dealt with similarly). Hence, $\mathcal{B} \neq \{D, E\}$. For the left implication, let π' be an $\overline{E}$-witness of π for $(\mathcal{K}, \varphi)$ such that $\mathcal{K}, \pi' \models \psi$. We have that there is a path of the form $\pi'' \cdot \pi$ with $|\pi''| \geq 1$ having the same h-level $\mathcal{B}$-descriptor as π'. Since $\langle \overline{\mathrm{E}} \rangle\, \psi \in \mathsf{SD}(\varphi)$, it holds that $\mathrm{depth}_{\mathcal{B}}(\psi) \leq h$. By Proposition 3.2, it follows that $\mathcal{K}, \pi'' \cdot \pi \models \psi$. Hence, $\mathcal{K}, \pi \models \langle \overline{\mathrm{E}} \rangle\, \psi$, and the result follows.

For the converse implication, assume that $\mathcal{K}, \pi \models \langle \overline{\mathrm{E}} \rangle\, \psi$. Hence, there is a path of the form $\pi'' \cdot \pi$ with $|\pi''| \geq 1$ such that $\mathcal{K}, \pi'' \cdot \pi \models \psi$. Let ν be a $\mathcal{B}$-certificate for $(\mathcal{K}, \varphi)$ with the same h-level $\mathcal{B}$-descriptor as π''. As $\mathcal{B} \neq \{D, E\}$, by applying Property (2) of Proposition 3.1 and Proposition 3.2, we deduce that $\mathcal{K}, \nu \cdot \pi \models \psi$. By applying again Proposition 3.2, there is a $\mathcal{B}$-certificate π' for $(\mathcal{K}, \varphi)$ which has the same h-level $\mathcal{B}$-descriptor as $\nu \cdot \pi$ and such that $\mathcal{K}, \pi' \models \psi$. Thus, since π' is an $\overline{E}$-witness of π for $(\mathcal{K}, \varphi)$, the result follows. $\square$

For the given $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula φ in *NNF*, we denote by $\mathsf{A\overline{A}L\overline{L}}(\varphi)$ the set of formulas in $\mathsf{SD}(\varphi)$ of the form $\langle \mathrm{X} \rangle\, \psi'$ or $[\mathrm{X}]\, \psi'$, with $X \in \{A, \overline{A}, L, \overline{L}\}$. An $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ for $(\mathcal{K}, \varphi)$ is a mapping associating with each state s of $\mathcal{K}$ a maximally consistent set of subformulas of $\mathsf{A\overline{A}L\overline{L}}(\varphi)$. More precisely, for all $s \in S$, $\mathcal{L}(s)$ is such that for all $\psi, \widetilde{\psi} \in \mathsf{A\overline{A}L\overline{L}}(\varphi)$, $\mathcal{L}(s) \cap \{\psi, \widetilde{\psi}\}$ is a singleton. $\mathcal{L}$ is *valid* if for all states $s \in S$ and $\psi \in \mathcal{L}(s)$, $\mathcal{K}, s \models \psi$ (here we consider s as a one-length path). Finally, a *well-formed set for* $(\mathcal{K}, \varphi)$ is a finite set $\mathcal{W}$ consisting of pairs (ψ, π) such that $\psi \in \mathsf{SD}(\varphi)$ and π is a $\mathcal{B}$-certificate of $(\mathcal{K}, \varphi)$. $\mathcal{W}$ is said *universal* if each formula occurring in $\mathcal{W}$ is of the form $[\mathrm{X}]\, \psi$, with $X \in \{\overline{B}, \overline{E}\}$. The *dual* $\widetilde{\mathcal{W}}$ of $\mathcal{W}$ is the well-formed set obtained by replacing each pair $(\psi, \pi) \in \mathcal{W}$ by $(\widetilde{\psi}, \pi)$. A well-formed set $\mathcal{W}$ is *valid* if for each $(\psi, \pi) \in \mathcal{W}$, $\mathcal{K}, \pi \models \psi$.

$check_{\mathcal{B}}(\mathcal{K},\varphi)$ [*$\mathcal{K}$ is a finite Kripke structure and φ is an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula in NNF*]

existentially choose an $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ for $(\mathcal{K},\varphi)$;
for each state s and $\psi \in \mathcal{L}(s)$ do
   case $\psi = \langle X\rangle\,\psi'$: existentially choose an X-witness π of s for $(\mathcal{K},\varphi)$ and call
      $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi',\pi)\})$;
   case $\psi = [X]\,\psi'$: universally choose an X-witness π of s for $(\mathcal{K},\varphi)$ and call
      $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi',\pi)\})$; end for
universally choose a $\mathcal{B}$-certificate π for $(\mathcal{K},\varphi)$ with $\mathsf{fst}(\pi) = s_0$ (s_0 is the initial state of $\mathcal{K}$)
   and call $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\varphi,\pi)\})$;

$checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\mathcal{W})$ [*$\mathcal{W}$ is a well-formed set of $(\mathcal{K},\varphi)$ and $\mathcal{L}$ is an $\mathsf{A\overline{A}L\overline{L}}$-labeling for $(\mathcal{K},\varphi)$*]

while $\mathcal{W}$ is *not* universal do
   deterministically select $(\psi,\pi) \in \mathcal{W}$ such that ψ is not of the form $[\overline{E}]\psi'$ and $[\overline{B}]\psi'$
   update $\mathcal{W} \leftarrow \mathcal{W} \setminus \{(\psi,\pi)\}$;
   case $\psi = p$ (resp., $\psi = \neg p$) with $p \in \mathcal{AP}$: if $\mathcal{K},\pi \not\models p$ (resp., $\mathcal{K},\pi \not\models \neg p$) then *reject*;
   case $\psi = \langle X\rangle\,\psi'$ or $\psi = [X]\,\psi'$ with $X \in \{A, L\}$: if $\psi \notin \mathcal{L}(\mathsf{lst}(\pi))$ then *reject*;
   case $\psi = \langle X\rangle\,\psi'$ or $\psi = [X]\,\psi'$ with $X \in \{\overline{A}, \overline{L}\}$: if $\psi \notin \mathcal{L}(\mathsf{fst}(\pi))$ then *reject*;
   case $\psi = \psi_1 \vee \psi_2$: existentially choose $i = 1, 2$, update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi_i,\pi)\}$;
   case $\psi = \psi_1 \wedge \psi_2$: update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi_1,\pi),(\psi_2,\pi)\}$;
   case $\psi = [X]\,\psi'$ with $X \in \mathcal{B}$: update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi',\pi') \mid \pi'$ is an X-*witness* of π for $(\mathcal{K},\varphi)\}$;
   case $\psi = \langle X\rangle\,\psi'$ with $X \in \mathcal{F}_{\mathcal{B}} \cap \{\overline{E}, \overline{B}\}$: existentially choose an X-*witness* π' of π
      for $(\mathcal{K},\varphi)$, update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi',\pi')\}$; end while
if $\mathcal{W} = \emptyset$ then *accept*
else universally choose $(\psi,\pi) \in \widetilde{\mathcal{W}}$ and call $checkFalse_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi,\pi)\})$

$checkFalse_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\mathcal{W})$ [*$\mathcal{W}$ is a well-formed set of $(\mathcal{K},\varphi)$ and $\mathcal{L}$ is an $\mathsf{A\overline{A}L\overline{L}}$-labeling for $(\mathcal{K},\varphi)$*]

while $\mathcal{W}$ is *not* universal do
   deterministically select $(\psi,\pi) \in \mathcal{W}$ such that ψ is not of the form $[\overline{E}]\psi'$ and $[\overline{B}]\psi'$
   update $\mathcal{W} \leftarrow \mathcal{W} \setminus \{(\psi,\pi)\}$;
   case $\psi = p$ (resp., $\psi = \neg p$) with $p \in \mathcal{AP}$: if $\mathcal{K}, p \not\models \pi$ (resp., $\mathcal{K}, \neg p \not\models \pi$) then *accept*;
   case $\psi = \langle X\rangle\,\psi'$ or $\psi = [X]\,\psi'$ with $X \in \{A, L\}$: if $\psi \notin \mathcal{L}(\mathsf{lst}(\pi))$ then *accept*;
   case $\psi = \langle X\rangle\,\psi'$ or $\psi = [X]\,\psi'$ with $X \in \{\overline{A}, \overline{L}\}$: if $\psi \notin \mathcal{L}(\mathsf{fst}(\pi))$ then *accept*;
   case $\psi = \psi_1 \vee \psi_2$: universally choose $i = 1, 2$, update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi_i,\pi)\}$;
   case $\psi = \psi_1 \wedge \psi_2$: update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi_1,\pi),(\psi_2,\pi)\}$;
   case $\psi = [X]\,\psi'$ with $X \in \mathcal{B}$: update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi',\pi') \mid \pi'$ is an X-*witness* of π for $(\mathcal{K},\varphi)\}$;
   case $\psi = \langle X\rangle\,\psi'$ with $X \in \mathcal{F}_{\mathcal{B}} \cap \{\overline{E}, \overline{B}\}$: universally choose an X-*witness* π' of π
      for $(\mathcal{K},\varphi)$, update $\mathcal{W} \leftarrow \mathcal{W} \cup \{(\psi',\pi')\}$; end while
if $\mathcal{W} = \emptyset$ then *reject*
else existentially choose $(\psi,\pi) \in \widetilde{\mathcal{W}}$ and call $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi,\pi)\})$

Fig. 5. Procedure $check_{\mathcal{B}}$ for a linear-time basis $\mathcal{B} \neq \{D\}$.

The procedure $check_{\mathcal{B}}$ in Fig. 5 defines the ATM required to prove the statement of Proposition 4.1 for a parametric basis $\mathcal{B} \neq \{D\}$. The procedure takes a pair $(\mathcal{K},\varphi)$ as input, where φ is an $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ formula, and performs the following steps:

1. it guesses an $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ for $(\mathcal{K},\varphi)$;
2. it checks that the guessed labeling $\mathcal{L}$ is valid;
3. for every $\mathcal{B}$-certificate π of $(\mathcal{K},\varphi)$ starting from the initial state, it checks that $\mathcal{K},\pi \models \varphi$.

To perform steps (2)–(3), it exploits the auxiliary ATM procedure $checkTrue_{\mathcal{B}}$ reported in Fig. 5. The procedure $checkTrue_{\mathcal{B}}$ takes as input a well-formed set $\mathcal{W}$ for $(\mathcal{K},\varphi)$ and, assuming that the current $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ is valid, checks whether $\mathcal{W}$ is valid. For each pair $(\psi,\pi) \in \mathcal{W}$ such that ψ is not of the form $[X]\,\psi'$, with $X \in \{\overline{B}, \overline{E}\}$, $checkTrue_{\mathcal{B}}$ directly checks whether $\mathcal{K},\pi \models \psi$. In order to allow a deterministic choice of the current element of the iteration, we assume the set $\mathcal{W}$ to be implemented as an ordered data structure. At each iteration of the while loop in $checkTrue_{\mathcal{B}}$, the current pair $(\psi,\pi) \in \mathcal{W}$ is processed according to the semantics of HS, exploiting the guessed $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ and Proposition 4.2. The processing is either deterministic or based on an existential choice, and the currently processed pair (ψ,π) is either removed from $\mathcal{W}$, or replaced by pairs (ψ',π') such that ψ' is a strict subformula of ψ.

At the end of the while loop, the resulting well formed set $\mathcal{W}$ is either empty or universal. In the former case, the procedure accepts. In the latter case, there is a switch in the current operation mode. For each element (ψ,π) in the dual of $\mathcal{W}$ (note that the root modality of ψ is either $\langle\overline{\mathrm{E}}\rangle$ or $\langle\overline{\mathrm{B}}\rangle$), the auxiliary ATM procedure *checkFalse*$_{\mathcal{B}}$ is invoked, which accepts the input $\{(\psi, \pi)\}$ if and only if $\mathcal{K}, \pi \not\models \psi$. The procedure *checkFalse*$_{\mathcal{B}}$, reported in Fig. 5, is the *dual* of *checkTrue*$_{\mathcal{B}}$: it is simply obtained from *checkTrue*$_{\mathcal{B}}$ by switching *accept* and *reject*, by switching existential choices and universal choices, and by converting the last call to *checkFalse*$_{\mathcal{B}}$ into *checkTrue*$_{\mathcal{B}}$. Thus *checkFalse*$_{\mathcal{B}}$ accepts an input $\mathcal{W}$ iff $\mathcal{W}$ is *not* valid. Formally, the following holds, where for an $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ for $(\mathcal{K}, \varphi)$, $\mathcal{L}_{\mathcal{W}}$ denotes the restriction of $\mathcal{L}$ to the set of formulas in $\mathsf{A\overline{A}L\overline{L}}(\varphi)$ which are subformulas of formulas occurring in $\mathcal{W}$. In other terms, for each state s, $\mathcal{L}_{\mathcal{W}}(s)$ contains all and only the formulas $\psi \in \mathcal{L}(s)$ such that either ψ or its dual $\widetilde{\psi}$ is a subformula of some formula occurring in $\mathcal{W}$.

Lemma 4.1. *Let $\mathcal{W}$ be a well-formed set for $(\mathcal{K}, \varphi)$ and $\mathcal{L}$ be an $\mathsf{A\overline{A}L\overline{L}}$-labeling for $(\mathcal{K}, \varphi)$. If $\mathcal{L}_{\mathcal{W}}$ is valid, then the following statements hold:*

1. *checkTrue*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \mathcal{W})$ *accepts* if and only if $\mathcal{W}$ *is valid;*
2. *checkFalse*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \mathcal{W})$ *accepts* if and only if $\mathcal{W}$ *is* not *valid.*

Proof. The proof is by induction on the maximum over the joint nesting depths of $\{\overline{B}, \overline{E}\}$ in the formulas ψ occurring in $\mathcal{W}$, denoted by $\mathrm{depth}_{\overline{\mathrm{BE}}}(\mathcal{W})$. In the proof, we exploit the crucial fact that, by construction, for the given input $(\mathcal{K}, \varphi, \mathcal{L}, \mathcal{W})$, procedures *checkTrue*$_{\mathcal{B}}$ and *checkFalse*$_{\mathcal{B}}$ only exploit the restriction $\mathcal{L}_{\mathcal{W}}$ of the $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$.

First, assume that there is no subformula of some formula occurring in $\mathcal{W}$ of the form $[\mathrm{X}]\,\psi$ for some $X \in \{\overline{B}, \overline{E}\}$. This case includes the base of the induction, where $\mathrm{depth}_{\overline{\mathrm{BE}}}(\mathcal{W}) = 0$. In this case, the two procedures have no nested call and the result easily follows by construction, the semantics of HS, and Proposition 4.2.

Now assume that there is some subformula of a formula occurring in $\mathcal{W}$ of the form $[\mathrm{X}]\,\psi$, for some $X \in \{\overline{B}, \overline{E}\}$. Hence, $\mathrm{depth}_{\overline{\mathrm{BE}}}(\mathcal{W}) > 0$. We focus on procedure *checkTrue*$_{\mathcal{B}}$ (the proof for procedure *checkFalse*$_{\mathcal{B}}$ is similar), and we show that if *checkTrue*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \mathcal{W})$ accepts, then $\mathcal{W}$ is valid (the proof of the converse implication is similar). Thus, assume that *checkTrue*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \mathcal{W})$ accepts. By construction and Proposition 4.2, there is a *non-empty universal* well-formed set $\mathcal{W}_U$ for $(\mathcal{K}, \varphi)$ such that (i) $\mathrm{depth}_{\overline{\mathrm{BE}}}(\mathcal{W}_U) \leq \mathrm{depth}_{\overline{\mathrm{BE}}}(\mathcal{W})$, (ii) $\mathcal{W}$ is valid if $\mathcal{W}_U$ is valid, and (iii) for each $([\mathrm{X}]\,\psi, \pi) \in \mathcal{W}_U$ (hence, $X \in \{\overline{B}, \overline{E}\}$), *checkFalse*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \{(\langle \mathrm{X}\rangle\, \widetilde{\psi}, \pi)\})$ accepts. Let W_t be the set of X-witnesses π' of π for $(\mathcal{K}, \varphi)$. By Proposition 4.2, $\mathcal{K}, [\mathrm{X}]\,\psi \models \pi$ if and only if for each $\pi' \in W_t$, $\mathcal{K}, \langle \mathrm{X}\rangle\, \widetilde{\psi} \not\models \pi$. On the other hand, by construction, *checkFalse*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \{(\langle \mathrm{X}\rangle\, \widetilde{\psi}, \pi)\})$ accepts if and only if for each $\pi' \in W_t$, *checkFalse*$_{\mathcal{B}}(\mathcal{K}, \varphi, \mathcal{L}, \{(\widetilde{\psi}, \pi')\})$ accepts. Since $\mathrm{depth}_{\overline{\mathrm{BE}}}(\{(\widetilde{\psi}, \pi')\}) < \mathrm{depth}_{\overline{\mathrm{BE}}}(\mathcal{W}_U)$, by the inductive hypothesis, we obtain that $\mathcal{W}_U$ is valid. Hence $\mathcal{W}$ is valid, and the result follows. □

The correctness of the algorithm follows from Propositions 3.2 and 4.2 and Lemma 4.1, and is formally proved in the next subsection.

4.1. Correctness of the alternating procedure check$_{\mathcal{B}}$: proof of Proposition 4.1

In this subsection, we show that the ATM *check*$_{\mathcal{B}}$ satisfies Proposition 4.1. We first prove the part of Proposition 4.1 concerning the number of alternations and the running time.

Proposition 4.3 *(Running Time). Given an input $(\mathcal{K}, \varphi)$, the number of alternations (between existentially and universal choices) of the ATM check$_{\mathcal{B}}$ is at most $\Upsilon(\varphi) + 2$ and it runs in time $M_{\mathcal{B}}(\mathcal{K}, \varphi)^{O(|\varphi|^d)}$, where $M_{\mathcal{B}}(\mathcal{K}, \varphi)$ is the maximal length of a $\mathcal{B}$-certificate for the input, and $d = 2$ if $D \in \mathcal{B}$ and $d = 1$ otherwise.*

Proof. First, we observe that in each iteration of the while loops of procedures *checkTrue*$_{\mathcal{B}}$ and *checkFalse*$_{\mathcal{B}}$, the processed pair (ψ, π) in the current well-formed set $\mathcal{W}$ either is removed from $\mathcal{W}$ or is replaced with pairs (ψ', π') such that ψ' is a strict subformula of ψ and π' is a $\mathcal{B}$-certificate for the input $(\mathcal{K}, \varphi)$. This ensures that the algorithm always terminates. Furthermore, we observe that the number of alternations of the ATM *check*$_{\mathcal{B}}$ between existential choices and universal choices is evidently the number of switches between the calls to procedures *checkTrue*$_{\mathcal{B}}$ and *checkFalse*$_{\mathcal{B}}$ plus two, and the top calls to *checkTrue*$_{\mathcal{B}}$ take as input well-formed sets for $(\mathcal{K}, \varphi)$ of the form $\{(\psi, \pi)\}$, where $\psi \in \mathsf{SD}(\varphi)$. Hence, by construction, it easily follows that the number of alternations of the ATM *check*$_{\mathcal{B}}$ on an input $(\mathcal{K}, \varphi)$ is at most $\Upsilon(\varphi) + 2$. For the running time, let $T(\varphi)$ be the standard tree encoding of φ, where each node is labeled by some subformula of φ. Let $\psi \in \mathsf{SD}(\varphi)$. If ψ is a subformula of φ, we define d_ψ as the maximum over the distances from the root in $T(\varphi)$ of ψ-labeled nodes. If instead ψ is the dual of a subformula of φ, we let $d_\psi := d_{\widetilde{\psi}}$. By construction, each step in an iteration of the while loops in procedures *checkTrue*$_{\mathcal{B}}$ and *checkFalse*$_{\mathcal{B}}$ can be performed in time $O(M_{\mathcal{B}}(\mathcal{K}, \varphi))$. Then, it suffices to show that for all computations ρ of the ATM *check*$_{\mathcal{B}}$ from input $(\mathcal{K}, \varphi)$, the overall number N_ψ of iterations of the while loops (of procedures *checkTrue*$_{\mathcal{B}}$ and *checkFalse*$_{\mathcal{B}}$) along the computation ρ for processing the formula ψ is at most $(2^{|\varphi|} \cdot M_{\mathcal{B}}(\mathcal{K}, \varphi))^{d_\psi}$ if $D \notin \mathcal{B}$, and at most $(2^{|\varphi|} \cdot M_{\mathcal{B}}(\mathcal{K}, \varphi))^{(d_\psi+1)^2}$ otherwise. Assume that $D \in \mathcal{B}$ (the other case is similar). The proof is done by induction on d_ψ. For the base case, assume that $d_\psi = 0$. Therefore, $\psi = \varphi$ or $\psi = \widetilde{\varphi}$, and by construction of the algorithm, N_φ and $N_{\widetilde{\varphi}}$ are at most equal to 1. Hence, the result holds. For the inductive step, assume that $d_\psi > 0$. We consider the case where ψ is a subformula of φ (the case where $\widetilde{\psi}$ is a subformula of φ is similar). Then, the result follows from the following chain of inequalities, where $P(\psi)$ denotes the set of nodes of $T(\varphi)$ which are parents of the nodes labeled by ψ, and for each node x, $fo(x)$ denotes the formula labeling x.

$$N_\psi \leq \sum_{x \in P(\psi)} N_{fo(x)} \cdot (M_{\mathcal{B}}(\mathcal{K},\varphi))^2 \leq \sum_{x \in P(\psi)} (2^{|\varphi|} \cdot M_{\mathcal{B}}(\mathcal{K},\varphi))^{(d_{fo(x)}+1)^2} \cdot (M_{\mathcal{B}}(\mathcal{K},\varphi))^2$$

$$\leq \left(2^{|\varphi|} \cdot M_{\mathcal{B}}(\mathcal{K},\varphi)\right)^{(d_\psi+1)^2}$$

The first inequality directly follows from the construction of the algorithm (note that if $fo(x) = [\mathrm{D}]\,\psi$, then the processing of subformula $fo(x)$ in an iteration of the two while loops generates at most $(M_{\mathcal{B}}(\mathcal{K},\varphi))^2$ new "copies" of ψ). The second inequality follows from the inductive hypothesis and the last inequality follows from the fact that $|P(\psi)| \leq 2^{|\varphi|}$ and $d_{fo(x)} \leq d_\psi - 1$ for all $x \in P(\psi)$. This concludes the proof of Proposition 4.3. □

By exploiting Lemma 4.1, we now prove the next result, which concludes the proof of Proposition 4.1.

Proposition 4.4 *(Correctness). The ATM $check_{\mathcal{B}}$ accepts an input $(\mathcal{K},\varphi)$ iff $\mathcal{K} \models \varphi$.*

Proof. Fix an input $(\mathcal{K},\varphi)$ and an $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ for $(\mathcal{K},\varphi)$. An *$\mathcal{L}$-guessing* for $(\mathcal{K},\varphi)$ is a well-formed set $\mathcal{W}$ for $(\mathcal{K},\varphi)$ which minimally satisfies the following conditions for all states s of $\mathcal{K}$:

- for all $\mathcal{B}$-certificates π for $(\mathcal{K},\varphi)$ with $\mathsf{fst}(\pi) = s_0$ (s_0 is the initial state), $(\varphi,\pi) \in \mathcal{W}$;
- for all $\langle \mathrm{X}\rangle\,\psi \in \mathcal{L}(s)$, $(\psi,\pi) \in \mathcal{W}$ for some X-witness π of s for $(\mathcal{K},\varphi)$;
- for all $[\mathrm{X}]\,\psi \in \mathcal{L}(s)$, $(\psi,\pi) \in \mathcal{W}$ for all X-witnesses π of s for $(\mathcal{K},\varphi)$.

Evidently, by construction of the procedure $check_{\mathcal{B}}$, for each input $(\mathcal{K},\varphi)$, it holds that:

(*) $check_{\mathcal{B}}$ accepts $(\mathcal{K},\varphi) \iff$ there is an $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ and a $\mathcal{L}$-guessing $\mathcal{W}$ for $(\mathcal{K},\varphi)$ such that for all $(\psi,\pi) \in \mathcal{W}$, $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi,\pi)\})$ accepts.

Fix an input $(\mathcal{K},\varphi)$. First, assume that $\mathcal{K} \models \varphi$. Let $\mathcal{L}$ be the *valid* $\mathsf{A\overline{A}L\overline{L}}$-labeling defined as follows for all states s: for all $\psi \in \mathsf{A\overline{A}L\overline{L}}(\varphi)$, $\psi \in \mathcal{L}(s)$ iff $\mathcal{K}, s \models \psi$. By Propositions 3.2 and 4.2, there exists an $\mathcal{L}$-guessing $\mathcal{W}$ for $(\mathcal{K},\varphi)$ such that for all $(\psi,\pi) \in \mathcal{W}$, $\mathcal{K},\pi \models \psi$. By Lemma 4.1, for all $(\psi,\pi) \in \mathcal{W}$, $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi,\pi)\})$ accepts. Hence, by Condition (*), procedure $check_{\mathcal{B}}$ accepts $(\mathcal{K},\varphi)$.

For the converse direction, assume that procedure $check_{\mathcal{B}}$ accepts $(\mathcal{K},\varphi)$. By Condition (*), there exists an $\mathsf{A\overline{A}L\overline{L}}$-labeling $\mathcal{L}$ and an $\mathcal{L}$-guessing $\mathcal{W}$ for $(\mathcal{K},\varphi)$ such that for all $(\psi,\pi) \in \mathcal{W}$, $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi,\pi)\})$ accepts. First, we show that $\mathcal{L}$ is valid. Fix a state s and a formula $\psi \in \mathcal{L}(s)$. We need to prove that $\mathcal{K}, s \models \psi$. The proof is by induction on the joint nesting depth $d_U(\psi)$ in ψ where $U = \{A, \overline{A}, L, \overline{L}\}$. Assume that $\psi = [\mathrm{X}]\,\psi'$ for some formula ψ' and $X \in \{A, \overline{A}, L, \overline{L}\}$ (the other cases where $\psi = \langle \mathrm{X}\rangle\,\psi'$ for $X \in \{A, \overline{A}, L, \overline{L}\}$ being similar). By definition of $\mathcal{L}$-guessing, it holds that $(\psi',\pi) \in \mathcal{W}$ for all X-witnesses π of s for $(\mathcal{K},\varphi)$. Moreover, by the induction hypothesis, one can assume that $\mathcal{L}_{\{(\psi',\pi)\}}$ is valid (note that for the base case, i.e. when ψ' does not contain occurrences of modalities $\langle \mathrm{X}\rangle$ or $[\mathrm{X}]$ with $X \in \{A, \overline{A}, L, \overline{L}\}$, $\mathcal{L}_{\{(\psi',\pi)\}}$ is trivially valid). By hypothesis, $checkTrue_{\mathcal{B}}(\mathcal{K},\varphi,\mathcal{L},\{(\psi',\pi)\})$ accepts. By Lemma 4.1, it follows that $\mathcal{K},\pi \models \psi'$ for all X-witnesses π of s for $(\mathcal{K},\varphi)$. Thus, by Proposition 4.2, we obtain that $\mathcal{K}, s \models \psi$. Hence, $\mathcal{L}$ is valid. By definition of $\mathcal{L}$-guessing, for each certificate π for $(\mathcal{K},\varphi)$ with $\mathsf{fst}(\pi) = s_0$, $(\varphi,\pi) \in \mathcal{W}$. Thus, by hypothesis, Lemma 4.1, and Proposition 3.2, we obtain that $\mathcal{K} \models \varphi$. This concludes the proof of Proposition 4.4 and Proposition 4.1 as well. □

5. Tight bounds on the length of certificates

In this section, for each basis $\mathcal{B}$ (except $\{D\}$), we provide tight bounds on the length of h-level $\mathcal{B}$-certificates.

For the bases $\{B\}$ and $\{E\}$, we prove singly exponential upper bounds in Subsection 5.1 and matching lower bounds in Subsection 5.2. By the exponential upper bounds and Proposition 4.1, we deduce that model checking the logics $\mathsf{HS}_{\{B\}}(\mathcal{F}_{\{B\}})$ and $\mathsf{HS}_{\{E\}}(\mathcal{F}_{\{E\}})$ is in the complexity class $\mathbf{AEXP_{pol}}$ of problems decided by exponential-time bounded alternating Turing Machines with a polynomially bounded number of alternations. This complexity result (membership in $\mathbf{AEXP_{pol}}$) has already been obtained in [23,24] for the fragments $\mathsf{AB\overline{A}\overline{B}\overline{E}}$ and $\mathsf{AE\overline{A}\overline{B}\overline{E}}$ by exploiting a more involved finite abstraction of paths (recall that $\mathsf{AB\overline{A}\overline{B}\overline{E}}$ and $\mathsf{AE\overline{A}\overline{B}\overline{E}}$ have the same expressiveness as $\mathsf{HS}_{\{B\}}(\mathcal{F}_{\{B\}})$ and $\mathsf{HS}_{\{E\}}(\mathcal{F}_{\{E\}})$, respectively). Since MC for $\overline{\mathsf{B}}$ and $\overline{\mathsf{E}}$ is already **PSPACE**-hard, we obtain the following result.

Corollary 5.1. *For the basis $\mathcal{B} = \{B\}$ (resp., $\mathcal{B} = \{E\}$), model-checking the logic $HS_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ is in* $\mathbf{AEXP_{pol}}$ *and at least* **PSPACE**-*hard.*

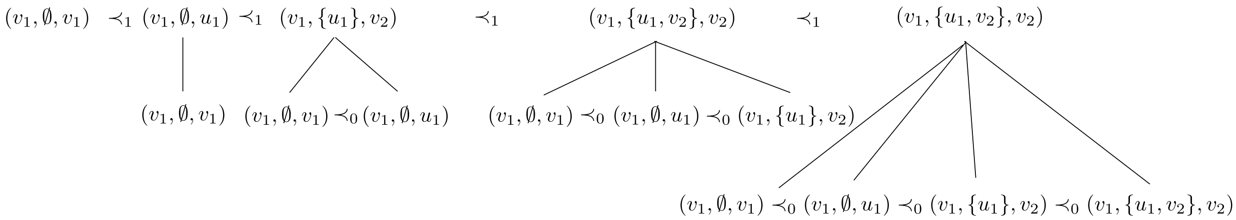

Fig. 6. Ordered 1-prefix descriptors over $\Sigma = \{u_1, v_1, v_2\}$. The $(\Sigma \times 2^\Sigma \times \Sigma)$-terms correspond with the subtrees of the root of the 2-level $\{B\}$-descriptor of Fig. 3.

For all bases $\mathcal{B}$ distinct from $\{B\}$ and $\{E\}$, we state a non-elementary lower bound (see Subsection 5.3) on the length of h-level $\mathcal{B}$-certificates. In particular, the result obtained for the basis $\{B, E\}$ negatively answers a question left open in [16] regarding the possibility of fixing an elementary upper bound on the size of BE-descriptors.

5.1. Exponential upper bound on the length of B-certificates and E-certificates

In this Subsection, we provide an exponential upper bound on the length of h-level B-certificates and E-certificates of a finite Kripke structure.

Theorem 5.1. *Let $\mathcal{K}$ be a finite Kripke structure with set of states S and $h \geq 0$. Then, each h-level B-certificate (resp., h-level E-certificate) has length at most $|S|^{2h+2}$.*

In the remaining part of the Subsection, we prove Theorem 5.1 focusing on B-certificates (the proof for E-certificates is similar and omitted). For a given finite alphabet Σ and $h \geq 0$, we first define a variant of the notion of h-level B-descriptor, called *ordered h-prefix descriptor over* Σ, which is not related to a specific word over Σ. The set OPD_h of ordered h-prefix descriptors over Σ is partitioned into $|\Sigma|$ subsets OPD_h^b (for each $b \in \Sigma$), where each of them is equipped with a strict partial order. We show that:

(i) each strict ascendent chain of elements in OPD_h^b has length at most $O(|\Sigma|^{2h+1})$;
(ii) the h-level B-descriptor of a word $w \in \Sigma^+$ is an element in OPD_h;
(iii) for each $w \in \Sigma^+$, the h-level B-descriptors associated to the prefixes of w can be grouped into at most $|\Sigma|$ non-strict ascendent chains.

Thus, by Proposition 3.1 and reasoning as in Proposition 3.3, we fix the upper bound on the length of h-level B-certificates for a given finite Kripke structure.

Let $h \geq 0$. For an h-level $(\Sigma \times 2^\Sigma \times \Sigma)$-term t with root (a, I, b), we say that a (resp., b) is the *first symbol* (resp., *last symbol*) of t.

Definition 5.1 *(Ordered prefix descriptors).* Let Σ be a finite alphabet and $h \geq 0$. We define by induction on h a pair $(OPD_h, \prec_h)$ consisting of a set OPD_h of h-level $(\Sigma \times 2^\Sigma \times \Sigma)$-terms, called *ordered h-prefix descriptors over* Σ and a binary non-reflexive relation $\prec_h$ over OPD_h.

- $h = 0$: OPD_0 is the set of 0-level $(\Sigma \times 2^\Sigma \times \Sigma)$-terms. Given $(a, I, b), (a', I', b') \in OPD_0$, $(a, I, b) \prec_0 (a', I', b')$ *if* (i) $a = a'$ (equality between the first symbols) and (ii) $I \subseteq I'$, and either $b \neq b'$ or $I \subsetneq I'$.
- $h > 0$: OPD_h is the set of h-level $(\Sigma \times 2^\Sigma \times \Sigma)$-terms $t = ((a, I, b), T)$ such that T is a (possibly empty) set of the form $T = \{t_1, \ldots, t_n\}$ where $t_i \in OPD_{h-1}$, t_i has first symbol a, and $t_1 \prec_{h-1} t_2 \prec_{h-1} \ldots \prec_{h-1} t_n$. The binary non-reflexive relation $\prec_h$ is defined as follows: $((a, I, b), T) \prec_h ((a', I', b'), T')$ *if*
  - $a = a'$, $I \subseteq I'$, $T \subseteq T'$;
  - and *either* $b \neq b'$ *or* $I \subsetneq I'$ *or* $T \subsetneq T'$.

An example of ordered 1-prefix descriptors is given in Fig. 6. By construction for each $b \in \Sigma$, the binary relation $\prec_h$ is a strict partial order over the set OPD_h^b of ordered h-prefix descriptors over Σ having the same last symbol b. Additionally, we show that a strict ascendent chain of elements in OPD_h^b has length at most $|\Sigma|^{2h+1}$.

Proposition 5.1. *Let $h \geq 0$, Σ be a finite alphabet, $b \in \Sigma$, and $t_1, \ldots, t_n$ be ordered h-prefix descriptors having last symbol b such that $t_1 \prec_h t_2 \prec_h \ldots \prec_h t_n$. Then, $n \leq |\Sigma|^{2h+1}$.*

Proof. The proof is by induction on $h \geq 0$. For the base case ($h = 0$), there is $a \in \Sigma$ such that for all $i \in [1, n]$, $t_i = (a, I_i, b)$ for some $I_i \subseteq \Sigma$, and $I_1 \subset I_2 \subset \ldots \subset I_n$. Hence, $n \leq |\Sigma|$ and the result follows.

Table 3
Miniwords for $\Sigma_2 = \{a_1, a_2\}$.

	h=0	h=1	h=2
$\mathbf{w_{1,1,h}}$	a_1	$a_1 \cdot w_{1,1,0} = a_1^2$	$a_1 \cdot w_{1,1,1} = a_1^3$
$\mathbf{w_{1,2,h}}$	a_1a_2	$a_1 \cdot w_{1,1,0} \cdot a_2 \cdot w_{2,1,0} = a_1^2 a_2^2 a_1$	$a_1 \cdot w_{1,1,1} \cdot a_2 \cdot w_{2,1,1} = a_1^3 a_2^3 a_1^2 a_2$
$\mathbf{w_{2,1,h}}$	a_2a_1	$a_2 \cdot w_{2,2,0} \cdot a_1 \cdot w_{1,2,0} = a_2^2 a_1^2 a_2$	$a_2 \cdot w_{2,2,1} \cdot a_1 \cdot w_{1,2,1} = a_2^3 a_1^3 a_2^2 a_1$
$\mathbf{w_{2,2,h}}$	a_2	$a_2 \cdot w_{2,2,0} = a_2^2$	$a_2 \cdot w_{2,2,1} = a_2^3$

Now, let $h > 0$. Hence, there is $a \in \Sigma$ such that for all $i \in [1, n]$, t_i is of the form $t_i = ((a, I_i, b), T_i)$. By hypothesis, $I_1 \subseteq I_2 \subseteq \ldots \subseteq I_n$. Moreover, for each $i \in [1, n]$, T_i can be partitioned into at most $|\Sigma|$ strict ascendent chains of ordered $h-1$-prefix descriptors having the same last symbol. Thus, by the induction hypothesis, we have that $|T_i| \leq |\Sigma| \cdot |\Sigma|^{2(h-1)+1} = |\Sigma|^{2h}$ for all $i \in [1, n]$. Fix an arbitrary $i \in [1, n]$. We claim that for each $j \in [i, n]$ such that $I_j = I_i$, it holds that $|j - i| \leq |\Sigma|^{2h}$. Hence, evidently, the result follows. Fix $i, j \in [1, n]$ such that $i < j$ and $I_j = I_i$. Since $t_i \prec_h t_\ell$ for all $\ell \in [i+1, j]$, we have that $|T_i| < |T_{i+1}| < \ldots < |T_j|$. Hence, $j - i \leq |T_j|$ and since $|T_j| \leq |\Sigma|^{2h}$, the result follows. □

By exploiting Proposition 5.1, we deduce the following proposition, from which the upper bound for the h-level B-certificates in Theorem 5.1 directly follows.

Proposition 5.2. *Let $\mathcal{K}$ be a finite Kripke structure with set of states S, $h \geq 0$, and π a path of $\mathcal{K}$. Then, the following holds:*

1. *for all $i, j \in [0, n]$ where $n = |\pi| - 1$, (i) $B_h(\pi[0, i])$ is an ordered h-prefix descriptor, and (ii) if $j > i$ and $B_h(\pi[0, i]) \neq B_h(\pi[0, j])$, then $B_h(\pi[0, i]) \prec_h B_h(\pi[0, j])$;*
2. *there is a path π' having the same h-level B-descriptor as π such that $|\pi'|$ is at most $|S|^{2h+2}$.*

Proof. Property 1 can be proved by a straightforward induction on $h \geq 0$. Now, let us consider Property 2. By reasoning as in the proof of Proposition 3.3, there is a path π' of $\mathcal{K}$ having the same h-level B-descriptor as π and such that distinct non-empty prefixes of π' have distinct h-level B-descriptors as well. Let s be a state visited by π', then by Property 1, the set of h-level B-descriptors associated with the non-empty prefixes of π' ending at state s form a strict ascending chain (with respect $\prec_h$) whose length n_s coincides with the set of positions i of π' such that $\pi'(i) = s$. By Proposition 5.1, $n_s \leq |S|^{2h+1}$. Since $|\pi'| = \sum_{s \in S(\pi')} n_s$ where $S(\pi')$ is the set of states visited by π', we obtain that $|\pi'| \leq |S|^{2h+2}$. □

5.2. Exponential lower bound on the length of B-certificates and E-certificates

In this Subsection, we provide an exponential lower bound on the length of h-level B-certificates and E-certificates (matching the upper bound stated in the previous Subsection 5.1).

Theorem 5.2. *There is a family $\{\mathcal{K}_n\}_{n \geq 1}$ of finite Kripke structures such that for all $n \geq 1$, $\mathcal{K}_n$ has $O(n)$ states and for all $h \geq 1$, there are h-level B-certificates (resp., h-level E-certificates) of $\mathcal{K}_n$ whose length is at least $\frac{1}{h} \cdot (\frac{n}{h})^h \cdot e^{h-1}$.*

We prove Theorem 5.2 focusing on B-certificates (the proof for E-certificates is similar and omitted). For each $n \geq 1$, let $\Sigma_n = \{a_1, \ldots, a_n\}$ be an alphabet consisting of n distinct symbols $a_1, \ldots, a_n$. We exhibit a family $(w_n^h)_{h \geq 0}$ of non-empty words over Σ_n such that for each $h \geq 0$, the length of w_n^h is at least $\frac{1}{h+1} \cdot (\frac{n}{h+1})^{h+1} \cdot e^h$ and w_n^h is a minimal representative of the $h+1$-level B-descriptor $B_{h+1}(w_n^h)$.

Fix $n \geq 1$. Formally, for all $i, j \in [1, n]$ and $h \geq 0$, we define by induction on $h \geq 0$, a non-empty word $w_{i,j,h}$ over Σ_n called (i, j, h)-*miniword*:

- Case $h = 0$ and $i \leq j$: $w_{i,j,h} := a_i a_{i+1} \ldots a_j$.
- Case $h = 0$ and $i > j$: $w_{i,j,h} := a_i a_{i-1} \ldots a_j$.
- Case $h > 0$ and $i \leq j$: $w_{i,j,h} := a_i \cdot u_i \cdot a_{i+1} \cdot u_{i+1} \cdot \ldots \cdot a_j \cdot u_j$ where for each $\ell \in [i, j]$, u_ℓ is the $(\ell, i, h-1)$-miniword.
- Case $h > 0$ and $i > j$: $w_{i,j,h} := a_i \cdot u_i \cdot a_{i-1} \cdot u_{i-1} \cdot \ldots \cdot a_j \cdot u_j$ where for each $\ell \in [i, j]$, u_ℓ is the $(\ell, i, h-1)$-miniword.

We say that $w_{i,j,h}$ *has level* h. Moreover, for $h > 0$, the miniwords u_ℓ of level $h-1$ in the factorization of $w_{i,j,h}$ are called *secondary subwords* of $w_{i,j,h}$, and a *main position* of $w_{i,j,h}$ is a position which is not associated to a secondary-subword position. When $h = 0$, the set of *main positions* of $w_{i,j,0}$ is the set of all its positions. Note that by construction, for each symbol $a \in \Sigma_n$ occurring in $w_{i,j,h}$, the smallest position ℓ such that $w_{i,j,h}(\ell) = a$ is a main position. As an example, in Table 3 we report the miniwords for the alphabet $\Sigma_2 = \{a_1, a_2\}$ of level h for $0 \leq h \leq 2$.

In the following we show that distinct prefixes of h-level miniwords have distinct h-level B-descriptors as well. This result together with a lower bound on the length of miniwords stated in Proposition 5.4 allows us to state the lower bound for B-certificates of Theorem 5.2.

Proposition 5.3. *Let $n \geq 1$ and $h \geq 0$. Then, for each miniword w over Σ_n of level h, distinct prefixes of w have distinct h-level B-descriptors.*

Proof. Fix $n \geq 1$ and $h \geq 0$. The proof is by induction on $h \geq 0$. The base case ($h = 0$) is trivial since by construction distinct prefixes of a miniword of level 0 have distinct last symbols as well. Now, let $h > 0$. Then, for some $i, j \in [1, n]$, w is the (i, j, h)-miniword. Assume that $i \leq j$ (the case where $i > j$ can be dealt similarly). By construction w is of the form $w = a_i \cdot u_i \cdot a_{i+1} \cdot u_{i+1} \cdot \ldots \cdot a_j \cdot u_j$, where u_ℓ is the $(\ell, i, h-1)$-miniword for all $\ell \in [i, j]$. Let ν and ν' be two distinct non-empty prefixes of w. We need to show that ν and ν' have distinct h-level B-descriptors.

We first consider the case where positions $|\nu|$ and $|\nu'|$ are *not* associated to the same secondary subword of w. There are four possible cases:

1. $|\nu|$ and $|\nu'|$ are main positions of w: hence, $w(|\nu|) \neq w(|\nu'|)$ and the result trivially follows.
2. $|\nu|$ and $|\nu'|$ are associated to distinct secondary subwords u_ℓ and $u_{\ell'}$, respectively: assume that $|\nu| < |\nu'|$, hence $\ell < \ell'$ (the other case being symmetric). By construction, each position of w associated to u_ℓ or preceding u_ℓ does not contain occurrences of $a_{\ell'}$. Therefore, ν and ν' have distinct 0-level B-descriptors (hence distinct h-level B-descriptors as well), and the result holds.
3. $|\nu|$ is associated to a main position with symbol a_ℓ and $|\nu'|$ is associated to a secondary subword $u_{\ell'}$. The case where $\ell \neq \ell'$ is similar to the previous one. Thus, assume that $\ell = \ell'$. By construction $a_\ell \notin \mathsf{internal}(\nu)$ while $a_\ell \in \mathsf{internal}(\nu')$, and the result holds.
4. $|\nu'|$ is associated to a main position with symbol $a_{\ell'}$ and $|\nu|$ is associated to a secondary subword u_ℓ: this case is similar to the previous one.

It remains to consider the more intriguing case where $|\nu|$ and $|\nu'|$ are associated to the same secondary subword u_ℓ for some $\ell \in [i, j]$. Hence, there are two distinct non-empty prefixes u and u' of u_ℓ and a non-empty prefix ρ of w such that $\nu = \rho \cdot u$, $\nu' = \rho \cdot u'$ and $\rho = a_i \cdot u_i \cdot a_{i+1} \cdot u_{i+1} \ldots a_\ell$.

Since u and u' are prefixes of u_ℓ, which is the $(\ell, i, h-1)$-miniword, by the induction hypothesis, u and u' have distinct $h-1$-level B-descriptors. Hence, the result directly follows from the following claim.

Claim Let $k \geq 0$ and ν and ν' be two distinct non-empty prefixes of u_ℓ. If ν and ν' have distinct k-level B-descriptors, then $\rho \cdot \nu$ and $\rho \cdot \nu'$ have distinct $k+1$-level B-descriptors.

Proof of the Claim First, assume that $k = 0$ and ν and ν' have distinct 0-level B-descriptors. If $\mathsf{lst}(\nu) \neq \mathsf{lst}(\nu')$ then the result trivially follows. Now, assume that $\mathsf{lst}(\nu) = \mathsf{lst}(\nu')$. Since $B_0(\nu) \neq B_0(\nu')$ and ν and ν' are prefixes of u_ℓ which is the $(\ell, i, h-1)$-miniword (where $\ell \geq i$), it follows that $h > 1$. Hence, u_ℓ is of the form

$$u_\ell = a_\ell \cdot v_\ell \cdot a_{\ell-1} \cdot v_{\ell-1} \cdot \ldots \cdot a_i \cdot v_i$$

where v_s is a word over $\{a_s, \ldots, a_\ell\}$ for all $s \in [i, \ell]$. Since $B_0(\nu) \neq B_0(\nu')$ and $\mathsf{lst}(\nu) = \mathsf{lst}(\nu')$, the last positions of ν and ν' cannot be associated to the same secondary subword of u_ℓ. Hence, the positions $|\nu|$ and $|\nu'|$ are either (i) associated to distinct secondary subwords of u_ℓ, or (ii) $|\nu|$ (resp., $|\nu'|$) is a main position and $|\nu'|$ (resp., $|\nu|$) is associated to a secondary subword of u_ℓ. We consider the case where $|\nu| < |\nu'|$, $|\nu|$ is a main position of u_ℓ, and $|\nu'|$ is associated with a secondary subword of u_ℓ (the other cases being similar). Then, $u_\ell(|\nu|) = a_s$ for some $s \in [i, \ell]$ and there is no proper prefix of ν visiting a_s. We show that for each non-empty proper prefix z of $\rho \cdot \nu$, $B_0(z) \neq B_0(\rho \cdot \nu)$. Hence, since $\rho \cdot \nu$ is a non-empty proper prefix of $\rho \cdot \nu'$, we obtain that $B_1(\rho \cdot \nu) \neq B_1(\rho \cdot \nu')$ and the result follows. Recall that by construction $a_\ell \notin \mathsf{internal}(\rho)$, while $a_\ell \in \mathsf{internal}(\rho \cdot \nu)$. Thus, if z is a prefix of ρ, the result follows. Otherwise, $z = \rho \cdot y$ where y is a non-empty proper prefix of ν. Since no proper prefix of ν visits a_s, $\mathsf{lst}(z) \neq a_s$ and the result holds in this case as well.

It remains to consider the case where $k > 0$ and ν and ν' have distinct k-level B-descriptors. We need to show that $\rho \cdot \nu$ and $\rho \cdot \nu'$ have distinct $k+1$-level B-descriptors. Assume that ν is a proper prefix of ν' (the case where ν' is a proper prefix of ν is symmetric). If $B_0(\nu) \neq B_0(\nu')$, then by the proof for the case $k = 0$, it holds that $B_1(\rho \cdot \nu) \neq B_1(\rho \cdot \nu')$, hence, $B_{k+1}(\rho \cdot \nu) \neq B_{k+1}(\rho \cdot \nu')$ as well. Now, assume that $B_0(\nu) = B_0(\nu')$. Since $B_k(\nu) \neq B_k(\nu')$ and ν is a proper prefix of ν', there is non-empty proper prefix z' of ν' of the form $z' = \nu \cdot x$ (for some word x) such that for each non-empty proper prefix z of ν it holds that $B_{k-1}(z) \neq B_{k-1}(z')$. By the induction hypothesis on k, we have that for each non-empty proper prefix z of ν, $B_k(\rho \cdot z) \neq B_k(\rho \cdot z')$. Thus, in order to show that $B_{k+1}(\rho \cdot \nu) \neq B_{k+1}(\rho \cdot \nu')$, it suffices to show that for each non-empty prefix y of ρ, we have that $B_k(y) \neq B_k(\rho \cdot z')$. By construction $a_\ell \notin \mathsf{internal}(\rho)$ while $a_\ell \in \mathsf{internal}(\rho \cdot z')$. Hence, the result follows, which concludes the proof of the claim and Proposition 5.3 as well. $\square$

For $\Sigma_n = \{a_1, \ldots, a_n\}$, let $\mathcal{K}(\Sigma_n)$ be the Kripke structure $(\Sigma_n, \Sigma_n, R, Lab, a_1)$, where Lab is the identity and $(a_i, a_j) \in R$ for all $i, j \in [1, n]$. The set of paths in $\mathcal{K}(\Sigma_n)$ is the set of non-empty finite words over Σ_n. Hence, the lower bound in Theorem 5.2 for B-certificates directly follows from the following result which is obtained by exploiting Proposition 5.3.

Proposition 5.4. *Let $n \geq 1$, $i, j \in [1, n]$, and $h \geq 0$. For the (i, j, h) miniword $w_{i,j,h}$ over Σ_n, the length of $w_{i,j,h}$ is at least $\frac{1}{h+1} \cdot (\frac{|i-j|+1}{h+1})^{h+1} \cdot e^h$ and there is no smaller word u over Σ_n (i.e., such that $|u| < |w_{i,j,h}|$) having the same $h+1$-level B-descriptor as $w_{i,j,h}$.*

Proof. For the (i, j, h)-miniword $w_{i,j,h}$, let $p = |i - j| + 1$. By construction, the length of $w_{i,j,h}$, denoted by $L(p, h)$, depends only on h and p, and satisfies the recurrence: $L(p, h) = p$ if $h = 0$, and $L(p, h) = p + \sum_{\ell=1}^{\ell=p} L(\ell, h-1)$ otherwise. We first show by induction on $h \geq 0$ that $L(p, h) \geq \frac{p^{h+1}}{(h+1)!}$. The base case ($h = 0$) is obvious. Now, let $h > 0$. By the induction hypothesis and the fact that $\sum_{\ell=1}^{\ell=p} \ell^h \geq \frac{p^{h+1}}{h+1}$ (Faulhaber's formula), we have that $L(p, h) = p + \sum_{\ell=1}^{\ell=p} L(\ell, h-1) \geq \sum_{\ell=1}^{\ell=p} \frac{\ell^h}{h!} \geq \frac{p^{h+1}}{(h+1)!}$. Thus, since $(h+1)! \leq \frac{(h+1)^{h+2}}{e^h}$, the claimed lower bound follows. Now, let T be the set of h-level B-descriptors of the non-empty proper prefixes of $w_{i,j,h}$, and u a non-empty word having the same $h+1$-level B-descriptor as $w_{i,j,h}$. Since the number of non-empty proper prefixes of a non-empty word w is $|w| - 1$, by hypothesis, we have that $|u| - 1 \geq |T|$. On the other hand, by Proposition 5.3, $|w_{i,j,h}| - 1 = |T|$. Hence, $|u| \geq |w_{i,j,h}|$, which concludes the proof of Proposition 5.4. □

5.3. Non-elementary lower bounds on the length of BD-certificates, BE-certificates, and DE-certificates

In this Subsection we establish a non-elementary lower bound on the length of h-level $\mathcal{B}$-certificates for each linear-time basis $\mathcal{B} \in \{\{B, D\}, \{B, E\}, \{D, E\}\}$. As an immediate consequence, we obtain a non-elementary lower bound on the running time of the algorithm for model checking the logic $\mathsf{HS}_{\mathcal{B}}(\mathcal{F}_{\mathcal{B}})$ presented in Section 4.

Theorem 5.3. *There is a family $\{\mathcal{K}_n\}_{n\geq 1}$ of finite Kripke structures such that for all $n \geq 1$, $\mathcal{K}_n$ has $O(n)$ states and for all $k \in [0, n-1]$ and basis $\mathcal{B}$ with $\mathcal{B} \in \{\{B, D\}, \{D, E\}\}$ (resp., $\mathcal{B} = \{B, E\}$), there are k-level (resp., $2k$-level) $\mathcal{B}$-certificates of $\mathcal{K}_n$ having length at least $\Omega(Tower(n, k+1))$.*

In the rest of this Subsection we provide a proof of Theorem 5.3. We first show as an intermediate and crucial step that there is a family $\{\Sigma_n\}_{n\geq 1}$ of finite alphabets such that for all $n \geq 1$, Σ_n has cardinality $O(n)$ and for all $k \in [0, n-1]$, there are $\Omega(Tower(n, k+1))$ words over Σ_n having pairwise distinct k-level D-descriptors (resp., $2k$-level BE-descriptors).

Let us fix $n \geq 1$ and let Σ_n be the finite alphabet having cardinality $O(n)$ given by

$$\Sigma_n = \bigcup_{i\in[2,n]} \{\$_i\} \cup \bigcup_{bit\in\{0,1\}} \bigcup_{i\in[1,n]} \{(\$_i, bit)\} \cup \bigcup_{bit\in\{0,1\}} \bigcup_{i\in[1,n]} \{(i, bit)\}$$

Moreover, for each $h \in [1, n]$, let Σ_n^h be the subset of Σ_n given by

$$\Sigma_n^h = \Sigma_n \setminus \Big(\bigcup_{i\in[h+1,n]} \{\$_i\} \cup \bigcup_{bit\in\{0,1\}} \bigcup_{i\in[h+1,n]} \{(\$_i, bit)\}\Big)$$

For each $h \in [1, n]$, we define a suitable encoding of the natural numbers in $[0, Tower(n, h) - 1]$ by finite words over Σ_n^h, called (n, h)-*blocks*. In particular, for $h > 1$, an (n, h)-block encoding a natural number $m \in [0, Tower(n, h) - 1]$ is a sequence of $Tower(n, h-1)$ $(n, h-1)$-blocks, where the ith $(n, h-1)$-block encodes both the value and (recursively) the position of the ith-bit in the binary representation of m. Formally, the set of (n, h)-blocks is defined by induction on h as follows:

Base Step: $h = 1$. An $(n, 1)$-block is a finite word bl over Σ_n^1 of length $n + 2$ having the form $bl = (\$_1, bit)(1, bit_1) \ldots (n, bit_n)(\$_1, bit)$ such that $bit, bit_1, \ldots, bit_n \in \{0, 1\}$. The *content* of bl is bit, and the *index* of bl is the natural number in $[0, Tower(n, 1) - 1]$ (recall that $Tower(n, 1) = 2^n$) whose binary code is $bit_1 \ldots bit_n$.

Induction Step: $1 < h \leq n$. A (n, h)-*block* is a finite word bl over Σ_n^h having the form $(\$_h, bit) \cdot bl_0 \cdot \$_h \cdot \ldots \cdot bl_{\ell-1} \cdot \$_h \cdot bl_\ell \cdot (\$_h, bit)$ such that $\ell = Tower(n, h-1) - 1$, $bit \in \{0, 1\}$ and for all $i \in [0, \ell]$, bl_i is a $(n, h-1)$-block having index i. The *content* of bl is bit and the *index* of bl is the natural number in $[0, Tower(n, h) - 1]$ whose binary code is given by $bit_0, \ldots, bit_\ell$, where bit_i is the content of the sub-block bl_i for all $0 \leq i \leq \ell$.

By construction, the following holds.

Remark 5.1. For all $n \geq 1$ and $h \in [1, n]$, there are $2 \cdot Tower(n, h)$ distinct (n, h)-blocks.

Example 5.1. Let $n = 2$ and $h = 2$. In this case $Tower(n, h) = 16$ and $Tower(n, h-1) = 4$. We can encode by $(2, 2)$-blocks all the integers in $[0, 15]$. Let us consider the number 14 whose binary code (using $Tower(n, h-1) = 4$ bits) is given by 0111 (the first bit is the least significant). The $(2, 2)$-block with content 0 encoding number 14 is given by $(\$_2, 0) \cdot bl_0 \cdot \$_2 \cdot bl_1 \cdot \$_2 \cdot bl_2 \cdot \$_2 \cdot bl_3 \cdot (\$_2, 0)$, where bl_i is the $(2, 1)$-block encoding the value and the position of the ith bit in 0111. For example, $bl_2 = (\$_1, 1)(1, 0)(2, 1)(\$_1, 1)$ while $bl_3 = (\$_1, 1)(1, 1)(2, 1)(\$_1, 1)$.

We now show that the $(h-1)$-level D-descriptors (resp., $(2h-2)$-level BE-descriptors) associated with distinct (n,h)-blocks are distinct as well.

Lemma 5.1. *Let $n \geq 1$. Then, for each $h \in [1,n]$, distinct (n,h)-blocks have distinct $(h-1)$-level D-descriptors and distinct $(2h-2)$-level BE-descriptors.*

Proof. For the fixed $n \geq 1$, the proof of Lemma 5.1 is by induction on $h \in [1,n]$. For the base case, let $h = 1$. Let bl be an $(n,1)$-block. By construction bl is a word of length $n+2$ of the form $bl = (\$_1, bit)(1, bit_1)\ldots(n, bit_n)(\$_1, bit)$ where $bit, bit_1, \ldots, bit_n \in \{0,1\}$. Hence, the 0-level D-descriptor $D_0(bl)$ (resp., 0-level BE-descriptor $BE_0(bl)$) of bl is the triple $((\$_1, bit), \{(1, bit_1), \ldots, (n, bit_n)\}, (\$_1, bit))$, and the result for $h = 1$ easily follows.

Now, for the induction step, assume that $h \in [2,n]$. Let bl and bl' be two (n,h)-blocks such that $bl \neq bl'$. We need to show that the $(h-1)$-level D-descriptors (resp., $(2h-2)$-level BE-descriptors) of bl and bl' are distinct. First, assume that bl and bl' have distinct content: let $(\$_h, bit)$ (resp., $(\$_h, bit')$) be the first letter of bl (resp., bl'). By hypothesis, $bit \neq bit'$. It follows that $D_0(bl) \neq D_0(bl')$ and $BE_0(bl) \neq BE_0(bl')$. Hence, $D_{h-1}(bl) \neq D_{h-1}(bl')$ and $BE_{2h-2}(bl) \neq BE_{2h-2}(bl')$ and the result follows.

Now, assume that bl and bl' have the same content. Since bl and bl' are distinct (n,h)-blocks, by construction there is $i \in [0, Tower(n, h-1) - 1]$ such that the $(n, h-1)$ sub-block sb_i of bl with index i and the $(n, h-1)$ sub-block sb'_i of bl' with index i have distinct content.

We first consider the D-descriptors. Let $(D_0(bl), T)$ (resp., $(D_0(bl), T')$) be the $(h-1)$-level D-descriptor of bl (resp., bl'). We show that for each non-empty internal subword w of bl, the $(h-2)$-level D-descriptor $D_{h-2}(w)$ of w is distinct from the $(h-2)$-level descriptor $D_{h-2}(sb'_i)$ of sb'_i. Hence, $D_{h-2}(sb'_i) \notin T$. Since $D_{h-2}(sb'_i) \in T'$, we obtain that $T \neq T'$ and the result follows. Fix a non-empty internal subword w of bl. By hypothesis and construction, there is no subword of bl which coincides with sb'_i. We distinguish the following cases:

- w is an $(n, h-1)$-block. Since w is an internal subword of bl and no subword of bl coincides with sb'_i, it hold that $w \neq sb'_i$. By the induction hypothesis, $D_{h-2}(w) \neq D_{h-2}(sb'_i)$.
- w is a proper subword of some $(n, h-1)$-block. By construction $D_0(w)$ is of the form (p, P, p') such that either $p \notin \{(\$_{h-1}, 0), (\$_{h-1}, 1)\}$ or $p' \notin \{(\$_{h-1}, 0), (\$_{h-1}, 1)\}$. Since the 0-level descriptor of an $(n, h-1)$-block is of the form $((\$_{h-1}, bit), P', (\$_{h-1}, bit))$ for some $bit \in \{0,1\}$, we obtain that $D_0(w) \neq D_0(sb'_i)$. Hence, $D_{h-2}(w) \neq D_{h-2}(sb'_i)$.
- There is some $(n, h-1)$ sub-block w' of bl such that w' is a proper subword of w. By construction, w contains some occurrence of a symbol in $\{\$_h, (\$_h, 0), (\$_h, 1)\}$. Since such symbols do not occur in an $(n, h-1)$-block, the result holds in this case as well.

It remains to consider the BE-descriptors. Let $(BE_0(bl), T_P, T_S)$ (resp., $(BE_0(bl'), T'_P, T'_S)$) be the $(2h-2)$-level BE-descriptor of bl (resp., bl'), and $w_{sb'_i}$ be the unique proper prefix of bl' having sb'_i as a proper suffix. We show that for each non-empty proper prefix w_P of bl, $BE_{2h-3}(w_{sb'_i}) \neq BE_{2h-3}(w_P)$. Hence, $BE_{2h-3}(w_{sb'_i}) \notin T_P$. Since $BE_{2h-3}(w_{sb'_i}) \in T'_P$, we obtain that $T_P \neq T'_P$ and the result follows. Fix a non-empty proper prefix w_P of bl. Note that since $h \geq 2$, $BE_{2h-3}(w_P)$ is of the form $(BE_0(w_P), R_P, R_S)$ and $BE_{2h-3}(w_{sb'_i})$ is of the form $(BE_0(w_{sb'_i}), R'_P, R'_S)$. Thus, it suffices to prove that $R_S \neq R'_S$. Since a proper suffix of a proper prefix of a word u is an internal word of u and $BE_{2h-4}(sb'_i) \in R'_S$, we just need to show that for each non-empty internal subword u of bl, $BE_{2h-4}(sb'_i) \neq BE_{2h-4}(u)$. For this we proceed as for the case of the D-descriptors but this time we apply the induction hypothesis on the BE_{2h-4}-descriptors. This concludes the proof of Lemma 5.1. □

Proof of Theorem 5.3. Let $n \geq 1$, a_n be a designated letter in the alphabet Σ_n and $\mathcal{K}_n$ the finite Kripke structure over Σ_n given by $\mathcal{K}_n = (\Sigma_n, \Sigma_n, R_n, Lab_n, a_n)$, where $(a, a') \in R_n$ and $Lab_n(a) = \{a\}$ for all $a, a' \in \Sigma_n$. Hence, the paths of $\mathcal{K}_n$ correspond to the non-empty finite words over Σ_n. We show that for all $k \in [0, n-1]$ and basis $\mathcal{B}$ with $\mathcal{B} \in \{\{B,D\}, \{D,E\}\}$ (resp., $\mathcal{B} = \{B,E\}$), there are $\Omega(Tower(n, k+1))$ distinct k-level (resp., $2k$-level) $\mathcal{B}$-certificates of $\mathcal{K}_n$. Hence, Theorem 5.3 directly follows. By Remark 5.1, there are $2 \cdot Tower(n, k+1)$ distinct $(n, k+1)$-blocks. Thus, for the basis $\{B,E\}$, the result directly follows from Lemma 5.1. For the bases $\{B,D\}$ and $\{D,E\}$, the result follows from Lemma 5.1 and the fact that words having distinct k-level D-descriptors have distinct k-level BD-descriptors (resp., distinct k-level DE-descriptors) as well.

6. Conclusions

We have addressed open complexity issues about the approach to model checking for the logic HS based on abstract representations of paths in Kripke structures, called BE-descriptors. We have developed a unifying framework to model check full HS and large HS fragments obtained by (i) introducing, for each basis $\mathcal{B}$, a specialized type of descriptor ($\mathcal{B}$-descriptor) and (ii) designing an alternating-time MC algorithm with a polynomially bounded number of alternations which is parametric with respect to the chosen basis $\mathcal{B}$ and runs in time bounded by the length of $\mathcal{B}$-descriptor certificates. As a main result, for each basis $\mathcal{B}$, we have provided tight bounds on the length of $\mathcal{B}$-certificates: exponential for the bases $\{B\}$ and $\{E\}$ (which lead to $\mathbf{AEXP_{pol}}$ procedures for the related fragments), and non-elementary for the other bases. Future work will be devoted to solve the difficult open question about the existence of an elementary procedure for the MC problem for full HS, and to settle the exact complexity of MC for the HS fragments for the bases $\{B\}$ and $\{E\}$.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] A. Pnueli, The temporal logic of programs, in: Proc. 18th FOCS, IEEE, 1977, pp. 46–57.
[2] 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.
[3] J.Y. Halpern, Y. Shoham, A propositional modal logic of time intervals, J. ACM 38 (4) (1991) 935–962.
[4] B. Moszkowski, Reasoning about digital circuits, Ph.D. thesis, Dept. of Computer Science, Stanford University, Stanford, CA, 1983.
[5] Y. Venema, Expressiveness and completeness of an interval tense logic, Notre Dame J. Form. Log. 31 (4) (1990) 529–547.
[6] A. Lomuscio, J. Michaliszyn, An epistemic Halpern-Shoham logic, in: Proc. 23rd IJCAI, 2013, pp. 1010–1016.
[7] I. Pratt-Hartmann, Temporal propositions and their logic, Artif. Intell. 166 (1–2) (2005) 1–36.
[8] J.F. Allen, Maintaining knowledge about temporal intervals, Commun. ACM 26 (11) (1983) 832–843.
[9] D. Bresolin, V. Goranko, A. Montanari, P. Sala, Tableau-based decision procedures for the logics of subinterval structures over dense orderings, J. Log. Comput. 20 (1) (2010) 133–166.
[10] D. Bresolin, V. Goranko, A. Montanari, G. Sciavicco, Propositional interval neighborhood logics: expressiveness, decidability, and undecidable extensions, Ann. Pure Appl. Log. 161 (3) (2009) 289–304.
[11] A. Montanari, G. Puppis, P. Sala, A decidable weakening of compass logic based on cone-shaped cardinal directions, Log. Methods Comput. Sci. 11 (4) (2015).
[12] 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.
[13] K. Lodaya, Sharpening the undecidability of interval temporal logic, in: Proc. 6th ASIAN, in: LNCS, vol. 1961, Springer, 2000, pp. 290–298.
[14] J. Marcinkowski, J. Michaliszyn, The undecidability of the logic of subintervals, Fundam. Inform. 131 (2) (2014) 217–240.
[15] A. Lomuscio, J. Michaliszyn, Decidability of model checking multi-agent systems against a class of EHS specifications, in: Proc. 21st ECAI, 2014, pp. 543–548.
[16] A. Molinari, A. Montanari, A. Murano, G. Perelli, A. Peron, Checking interval properties of computations, Acta Inform. 53 (6–8) (2016) 587–619.
[17] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Interval temporal logic model checking: the border between good and bad HS fragments, in: Proc. 8th IJCAR, in: LNAI, vol. 9706, Springer, 2016, pp. 389–405.
[18] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Interval vs. point temporal logic model checking: an expressiveness comparison, ACM Trans. Comput. Log. 20 (1) (2019) 4:1–4:31.
[19] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Model checking for fragments of the interval temporal logic HS at the low levels of the polynomial time hierarchy, Inf. Comput. 262 (Part) (2018) 241–264.
[20] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, P. Sala, Satisfiability and model checking for the logic of sub-intervals under the homogeneity assumption, in: Proc. 44th ICALP, in: LIPIcs, vol. 80, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, pp. 120:1–120:14.
[21] A. Molinari, A. Montanari, A. Peron, Complexity of ITL model checking: some well-behaved fragments of the interval logic HS, in: Proc. 22nd TIME, 2015, pp. 90–100.
[22] A. Molinari, A. Montanari, A. Peron, Model checking for fragments of Halpern and Shoham's interval temporal logic based on track representatives, Inf. Comput. 259 (3) (2018) 412–443.
[23] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, On the complexity of model checking for syntactically maximal fragments of the interval temporal logic HS with regular expressions, in: Proceedings Eighth International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2017, Roma, Italy, 20–22 September 2017, 2017, pp. 31–45.
[24] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, Model checking interval temporal logics with regular expressions, Inf. Comput. 272 (2020) 104498, https://doi.org/10.1016/j.ic.2019.104498.
[25] P. Roeper, Intervals and tenses, J. Philos. Log. 9 (1980) 451–469.
[26] J. Ferrante, C. Rackoff, A decision procedure for the first order theory of real addition with order, SIAM J. Comput. 4 (1) (1975) 69–76.