The Logic of Prefixes and Suffixes is Elementary under Homogeneity *

Dario Della Monica
University of Udine
dario.dellamonica@uniud.it

Angelo Montanari
University of Udine
angelo.montanari@uniud.it

Gabriele Puppis
University of Udine
gabriele.puppis@uniud.it

Pietro Sala
University of Verona
pietro.sala@univr.it

***Abstract*—In this paper, we study the finite satisfiability problem for the logic BE under the homogeneity assumption. BE is the cornerstone of Halpern and Shoham's interval temporal logic, and features modal operators corresponding to the prefix (a.k.a. "Begins") and suffix (a.k.a. "Ends") relations on intervals. In terms of complexity, BE lies in between the "Chop" logic C, whose satisfiability problem is known to be non-elementary, and the PSPACE-complete interval logic D of the sub-interval (a.k.a. "During") relation. BE was shown to be EXPSPACE-hard, and the only known satisfiability procedure is primitive recursive, but not elementary. Our contribution consists of tightening the complexity bounds of the satisfiability problem for BE, by proving it to be EXPSPACE-complete. We do so by devising an equi-satisfiable normal form with boundedly many nested modalities. The normalization technique resembles Scott's quantifier elimination, but it turns out to be much more involved due to the limitations enforced by the homogeneity assumption.**

***Index Terms*—interval temporal logic, homogeneity assumption, satisfiability problem, quantifier elimination**

I. INTRODUCTION

In this paper, we study the computational complexity of the satisfiability problem for the logic BE of the prefix and suffix interval relations. The considered interpretation setting is the one with intervals over finite linear orders, under the homogeneity assumption (see below). The logic BE is at the core of the galaxy of interval temporal logics [9] and has interesting connections with standard point-based temporal logics [2]. In general, formulas of interval temporal logics can express properties of *pairs* of time points, rather than properties of single time points, and are evaluated as sets of such pairs, that is, as binary relations on points. They are very expressive in comparison to point-based ones, and it does not come as a surprise that, in general, there is no reduction of their satisfiability problem to satisfiability of classical monadic second-order logic.

The *logic* BE has two (unary) modalities, $\langle B\rangle$ and $\langle E\rangle$, that quantify over prefixes and suffixes of the current interval, respectively. These modalities can be viewed as the logical counterparts of Allen's binary relations *Begins* and *Ends* [1]. In particular, the logic BE can be considered as a fragment of Halpern and Shoham's interval temporal logic [9], denoted HS, which features one modal operator for each of the twelve non-trivial Allen's relations.

The satisfiability problem for BE turns out to be *undecidable* over all relevant classes of interval structures [7], [10]. One can however escape this bleak landscape by constraining the semantics, in particular, the interpretation of the propositional letters. An interesting example is given by the *homogeneity* assumption, according to which a propositional letter holds at an interval if and only if it holds at all of its points. In other words, according to the homogeneous semantics, the labelling of an arbitrary interval in a model is uniquely determined by those of the singleton intervals contained in it.

An advantage of the homogeneity assumption is that it makes it possible to define a natural interpretation of interval logics over Kripke structures. For example, this comes in handy when studying the *model-checking problem*, which is defined as the problem of deciding whether a given formula is valid over all (homogeneous) interval structures generated by a given Kripke structure. As such, the problem can be seen as a variant of the classical validity/satisfiability problem, and many decidability and complexity results can be transferred from one problem to another. In [12] it was shown that the model-checking and satisfiability problems for BE, and in fact for full HS logic, become decidable when one restricts to *homogeneous* interval structures.

Despite its simple syntax and the homogeneity assumption, the logic BE turns out to be quite expressive and succinct. In [2], Bozzelli et al. have shown that, when interpreted over finite words, LTL (Linear Temporal Logic) and BE, under homogeneity, define the same class of star-free regular languages, but with the latter formalism being at least exponentially more succinct than the former. This is also reflected in the complexity of the satisfiability problem for BE, which was shown to be EXPSPACE-hard [3, Theorem 3.1].[1] On the other hand, the only known decision procedure [12] for satisfiability

* We acknowledge the support from the 2022 Italian INdAM-GNCS project *"Elaborazione del Linguaggio Naturale e Logica Temporale per la Formalizzazione di Testi"*, ref. no. CUP_E55F22000270001. We would also like to thank Alberto Molinari, Laura Bozzelli, and Adriano Peron for many useful discussions. Finally, Angelo Montanari would like to acknowledge the work done on the problem, together with Gabriele Puppis and Pietro Sala, when he was on leave at LaBRI in Bordeaux.

979-8-3503-3587-3/23/$31.00 ©2023 IEEE

[1] In fact, the cited result focuses on the model-checking problem for BE, which takes as input, not only a formula, but also a Kripke structure. It happens that the Kripke structure used in the proof of the EXPSPACE lowerbound generates every possible homogeneous interval structure, and hence the result can be immediately transferred to the validity/satisfiability problem for BE.

of BE formulas is basically the one for full HS, and is not elementary.

It is also worth contrasting the expressiveness and complexity of BE, under homogeneity, with those of two close relatives of it: the Chop logic C [16] and the logic D of the sub-interval relation [4]. The *logic* C has a binary modality $\langle C \rangle$ that allows one to split the current interval in two parts and predicate separately on them. The *logic* D has a unary modality $\langle D \rangle$ that allows one to predicate about sub-intervals of the current interval. It is easy to see that, in terms of expressiveness, BE lies in between D and C, in the sense that modality $\langle D \rangle$ can be defined in BE, i.e., $\langle D \rangle \varphi$ is equivalent to $\langle B \rangle \langle E \rangle \varphi$, and modalities $\langle B \rangle$ and $\langle E \rangle$ can in turn be defined in C, e.g., $\langle B \rangle \varphi$ is equivalent to $\varphi \langle C \rangle \mathrm{true}$. Under the homogeneity assumption, the satisfiability problem for C is non-elementarily decidable, precisely, tower-complete, in view of the existence of straightforward reductions to and from language-emptiness of generalized star-free regular expressions [13], [15], while the satisfiability problem for D was shown to be PSpace-complete by a suitable contraction method [4]. It is also worth pointing out that if the homogeneity assumption is removed, the satisfiability problem for D becomes undecidable [11].

Based on the observations above, it is crucial to close the gap between the complexity lowerbound and upperbound of the satisfiability problem for BE. Significant effort has been invested in recent years towards both raising the ExpSpace lowerbound, e.g., using variants of Stockmeyer's counters [15], and developing elementary satisfiability procedures. Despite these efforts, the complexity gap remained unchanged and proved to be an intriguing challenge. The special status of BE is witnessed by the fact that the many results about the complexity of the satisfiability and/or model-checking problems for proper fragments of HS, under the homogeneity assumption, concern logics that include neither modality $\langle B \rangle$ nor modality $\langle E \rangle$ or feature only one of them (an up-to-date picture can be found in [5], [6]).

In this paper, we manage to prove that the satisfiability problem for the logic BE, under homogeneity, is elementarily decidable, and precisely ExpSpace-complete. This result is established using a rather unexpected normalization technique, which consists of transforming an arbitrary BE formula into an equi-satisfiable one with boundedly many nested modalities. Specifically, we will show that one can compute, in polynomial time, normalized formulas with nesting depth of modalities at most 4, and with at most 2 alternations between universal and existential modalities. The transformation of BE formulas into normalized ones can be also viewed as a quantifier elimination technique à-la Scott [14]. In this perspective, however, the transformation has to deal with an increased difficulty: due to the homogeneity assumption, the elements over which we predicate cannot be labelled in an arbitrary way. In view of this difficulty, it is quite surprising that an equi-satisfiable normalized formula can be computed in polynomial time from any given arbitrary BE formula.

The rest of the paper is organized as follows. In Section II, we introduce the logic BE and we point out the relevant implications of the homogeneity assumption. In Section III, we define the transformation of BE formulas into normalized ones. In Section IV, we derive an optimal satisfiability procedure and analyse its complexity. Conclusions provide an assessment of the work done and outline future research directions. For reader convenience, technical terms and notation in the electronic version of the paper are linked to their definitions, which can then be accessed with a mouse click.

II. Preliminaries

Let the *time domain* be a finite prefix of the natural numbers $(N,<)$. Intervals over N are denoted by $[x,y]$, for $x,y \in N$ and $x \leq y$, and the set of all intervals over N is denoted $\mathbb{I}(N)$. We let $<_B$ (resp., $<_E$) be the proper *prefix* (resp., *suffix*) relation on intervals, defined by $J <_B I$ if and only if $\min(I) = \min(J) \leq \max(J) < \max(I)$ (resp., $J <_E I$ if and only if $\min(I) < \min(J) \leq \max(J) = \max(I)$).

Formulas of the *logic* BE are constructed starting from propositional letters belonging to a finite non-empty set Σ, called *signature*, using classical Boolean connectives and modal operators. The latter operators are used to quantify over prefixes and suffixes of the current interval. Formally, BE formulas satisfy the following grammar:

$$\varphi \ ::= \ p \ \ (\text{for } p \in \Sigma) \ \mid \ \neg\varphi \ \mid \ \varphi \vee \varphi \ \mid \ \langle B \rangle \varphi \ \mid \ \langle E \rangle \varphi.$$

Semantics is given in terms of an interval structure $\mathcal{S}$ and one of its intervals I. Formally, an *interval structure* over a signature Σ is a pair $\mathcal{S} = (\mathbb{I}(N), \sigma)$, where $\sigma : \mathbb{I}(N) \to \wp(\Sigma)$ is a labelling of intervals by subsets of Σ. Whether a BE formula φ *holds at* an interval I of $\mathcal{S}$, denoted $\mathcal{S}, I \vDash \varphi$, is determined by the following rules:

- $\mathcal{S}, I \vDash p$ if $p \in \sigma(I)$;
- $\mathcal{S}, I \vDash \neg\varphi$ if $\mathcal{S}, I \nvDash \varphi$;
- $\mathcal{S}, I \vDash \varphi_1 \vee \varphi_2$ if $\mathcal{S}, I \vDash \varphi_1$ or $\mathcal{S}, I \vDash \varphi_2$;
- $\mathcal{S}, I \vDash \langle B \rangle \varphi$ if $\mathcal{S}, J \vDash \varphi$ for some $J <_B I$;
- $\mathcal{S}, I \vDash \langle E \rangle \varphi$ if $\mathcal{S}, J \vDash \varphi$ for some $J <_E I$.

A formula is *valid* if it holds at every interval of every interval structure; similarly, it is *satisfiable* if it holds at some interval of some interval structure. Two formulas φ and φ' are *equivalent* if for every interval structure $\mathcal{S}$ and every interval I in it, $\mathcal{S}, I \vDash \varphi$ iff $\mathcal{S}, I \vDash \varphi'$. They are *equi-satisfiable* if either they are both satisfiable or none of them is. The notions of validity, satisfiability, and equivalence can be relativized to a specific class of interval structures (possibly even to a single interval structure). As an example, we say that a formula φ is *valid over* a class $\mathscr{C}$ of interval structures if $\mathcal{S}, I \vDash \varphi$ for all $\mathcal{S} \in \mathscr{C}$ and all $I \in \mathcal{S}$. In the particular case where $\mathscr{C}$ contains a single interval structure $\mathcal{S}$, we will say that a formula φ is *valid over* $\mathcal{S}$ if $\mathcal{S}, I \vDash \varphi$ for all $I \in \mathcal{S}$.

It is possible to add syntactic sugar to the logic BE. As an example, we will often use shorthands like $\varphi_1 \wedge \varphi_2 = \neg(\neg\varphi_1 \vee \neg\varphi_2)$, false $= p \wedge \neg p$ (for any $p \in \Sigma$), true $= \neg$false, and $[X]\varphi = \neg\langle X \rangle \neg\varphi$, for $X \in \{B, E\}$. Some other useful shorthands are $\pi = [B]\,\mathrm{false}$, which constrains the interval where it is evaluated to be a singleton, and $[G]\varphi = \varphi \wedge [B]\,\varphi \wedge [E]\,\varphi \wedge [B]\,[E]\,\varphi$, which constrains all sub-intervals (including the current interval, its proper prefixes, and its proper suffixes) to satisfy φ. The shorthands π and $[G]$ can be viewed as derived nullary and unary modal operators, respectively, and can be added as syntactic sugar to BE.

Homogeneity assumption. We recall from [10], [12] that the satisfiability problems for the logic BE is undecidable, unless one restricts to homogeneous interval structures. An interval structure $\mathcal{S} = (\mathbb{I}(N), \sigma)$ is *homogeneous* if its labelling satisfies the condition $\sigma(I) = \bigcap_{x \in I} \sigma([x,x])$ for all $I \in \mathbb{I}(N)$. Intuitively, this means that the labelling σ is uniquely determined by its restriction to singleton intervals. Let us take a closer look at the implications of homogeneity.

First of all, we have that every formula $\langle B\rangle(p_1 \wedge p_2)$ is equivalent to $\langle B\rangle p_1 \wedge \langle B\rangle p_2$, and similarly for $\langle E\rangle$. Note, however, that homogeneity does not imply similar properties for arbitrary formulas φ_1, φ_2 replacing the propositional letters p_1, p_2. As an example, the formulas $\langle B\rangle(p \wedge \neg p)$ and $(\langle B\rangle p) \wedge (\langle B\rangle \neg p)$ are not equivalent.

Homogeneity can also be exploited to efficiently rewrite any BE formula into an equivalent one where every occurrence of a propositional letter is conjoined with π. Based on this observation, we introduce the following mild normal form:

Definition 1. *A* BE *formula ψ is in* homogeneous normal form *if every occurrence of a propositional letter p in ψ appears inside the subformula $\pi \wedge p$.*

Basically, the homogeneous normal form restricts propositional letters to be only evaluated at singleton intervals. As an example, the formula $(\pi \wedge q) \vee \langle B\rangle(\pi \wedge \neg(\pi \wedge p))$ is in homogeneous normal form and holds at an interval I iff I consists of a single point labelled by q or the left endpoint of I is not labelled by p.

Proposition 2. *One can transform in linear time any* BE *formula ψ into one in homogeneous normal form that is equivalent to ψ when interpreted over* homogeneous *interval structures.*

Proof. It suffices to replace every occurrence of a propositional letter p in ψ by the formula $\texttt{everywhere}(p) = (\pi \wedge p) \vee \big(\langle B\rangle(\pi \wedge p) \wedge \langle E\rangle(\pi \wedge p) \wedge [B](\pi \vee \langle E\rangle(\pi \wedge p))\big)$. The resulting formula is equivalent to ψ since, over homogeneous interval structures, p is equivalent to $\texttt{everywhere}(p)$. □

We denote by BE_π the fragment of logic BE that contains only formulas in homogeneous normal form. *From this point forward, we will exclusively work with BE_π formulas, with the understanding that this assumption may occasionally go unstated. Accordingly, we will treat (sub)formulas of the form $\pi \wedge p$ as atomic.*

III. A bounded-nesting normal form for BE_π

In this section, we describe a transformation of arbitrary BE_π formulas into equi-satisfiable ones with boundedly many nested modalities. The transformation is somehow reminiscent of the so-called Scott normal form for the two-variable fragment of first-order logic [14], since it results in a formula, over an extended set of propositional letters, that is satisfiable if and only if the original formula was. The increased difficulty here is that the valuation of the new propositional letters emerging from the transformation must satisfy the homogeneity assumption. This is to say that we cannot identify intervals satisfying a certain (sub)formula φ by labelling them with a fresh propositional letter q_φ. Rather, we will identify these intervals by appropriately correlating fresh labels assigned to their endpoints. Our transformation will exploit in a crucial way the fact that, under homogeneity, valuations of formulas at two overlapping intervals have "less degrees of freedom" than valuations of the same formulas at disjoint intervals.

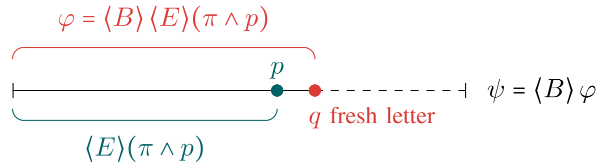

Fig. 1. Example of normalization of a formula $\psi = \langle B\rangle\varphi$.

Definition 3. *The* modal depth *(or simply* depth*) of a BE_π formula is the maximum number of nested modal operators $\langle B\rangle$ and $\langle E\rangle$ in it, not counting those defining the operator π. A BE_π formula is in* shallow normal form *if it is of the form $\psi \wedge [G]\xi$, where both ψ and ξ have depth at most 2.*

Concerning the above definition, we recall that $[G]\xi$ is a shorthand for $\xi \wedge [B]\xi \wedge [E]\xi \wedge [B][E]\xi$, so a formula in shallow normal form has depth at most 4. However, not all depth-4 formulas are in shallow normal form.

Theorem 4. *Given any BE_π formula ψ, one can compute in in polynomial time an equi-satisfiable formula $\psi^\star$ that is in shallow normal form.*

To highlight one of the key ideas underlying the proof of the theorem, which we postpone to the next subsections, we give an example of normalization of a formula.

Example 5. *Consider the formula $\psi = \langle B\rangle\varphi$ over the signature $\Sigma = \{p\}$, where $\varphi = \langle B\rangle\langle E\rangle(\pi \wedge p)$. Figure 1 shows an example of an interval structure satisfying ψ; in particular, it highlights intervals witnessing φ (in red) and $\langle E\rangle(\pi \wedge p)$ (in blue). Note that ψ has depth 3 and is not in shallow normal form. To rewrite ψ into an equi-satisfiable formula in shallow normal form, we introduce a new propositional letter q with the purpose of marking the right endpoints of the intervals that satisfy φ and that are minimal w.r.t. the prefix relation (we call these intervals prefix-minimal, for short). Note that the right endpoints of these intervals are immediately to the right of the p-labelled points. We thus consider interval structures over the expanded signature $\Sigma' = \{p, q\}$ that make the following formula valid:*

$$\xi \;=\; \underbrace{\big(\neg\pi \;\wedge\; \neg\langle B\rangle\neg\pi\big)}_{\text{interval has exactly two points}} \;\rightarrow\; \underbrace{\big(\langle B\rangle(\pi \wedge p) \;\leftrightarrow\; \langle E\rangle(\pi \wedge q)\big)}_{q \text{ is to the right whenever } p \text{ is to the left}}$$

We can verify that, over interval structures that make ξ valid, every prefix-minimal interval that satisfies φ also satisfies $\varphi' = \langle E\rangle(\neg\pi) \wedge \langle E\rangle(\pi \wedge q)$, and, conversely, every interval that satisfies φ' also satisfies φ. This implies that, again over interval structures that make ξ valid, the depth-3 formula $\psi = \langle B\rangle\varphi$ is equivalent to the depth-2 formula $\psi' = \langle B\rangle\varphi'$. Moreover, since the labelling of any interval structure over $\Sigma = \{p\}$ can always be expanded with the fresh letter q so as to satisfy $[G]\xi$, we conclude that ψ is equi-satisfiable as the formula $\psi^\star = \psi' \wedge [G]\xi$. Since ξ has depth 1, $\psi^\star$ is also in shallow normal form.

The normalization procedure for an arbitrary formula ψ iterates a rewriting similar to the one presented in Example 5. More precisely, we start by replacing every outermost subformula of ψ of depth $d > 2$ and of the form $\langle B\rangle\varphi$ (resp., $\langle E\rangle\varphi$) with an equi-satisfiable formula $\langle B\rangle\varphi'$ (resp., $\langle E\rangle\varphi'$) of depth 2. This rewriting step extends the signature with new propositional letters, which are constrained while preserving equi-satisfiability using formulas similar to the $[G]\xi$ of Example 5. Constraints will contain occurrences of the original subformula φ, and thus need to be normalized in their turn in order to eventually obtain formulas of depth at most 2. More details and formal arguments about the normalization procedure of Theorem 4 will be provided in the next subsections.

We conclude this part by observing an immediate consequence of Theorem 4. We recall from [12] the existence of a rather simple, but non-elementary procedure for deciding satisfiability of a BE formula ψ under homogeneity. A close inspection to the description of this procedure shows that it has non-deterministic time complexity $\mathcal{O}(tow(h, |\psi|))$, where $tow(h,n) = 2^{2^{\cdot^{\cdot^{\cdot^{n}}}}}$ is the tower of h exponents ending with n and h is the maximum number of nested modal operators in the input formula ψ. As the shorthand π can be directy handled in constant time, the parameter h of the said complexity bound can be identified with our notion of modal depth for BE_π formulas. In particular, when we consider a formula ψ in shallow normal form, the parameter h is at most 4. Together with Proposition 2 and Theorem 4, this gives a first rough complexity bound to the satisfiability problem for BE logic under the homogeneity assumption:

Corollary 6. *The satisfiability problem for* BE *logic restricted to homogeneous interval structures is elementarily decidable, i.e., at least in* 4NEXPTIME*.*

We shall provide later, in Section IV, a more careful complexity analysis, showing that the satisfiability problem for BE logic under homogeneity is actually EXPSPACE-complete.

A. Expanders

A first ingredient of the normalization procedure of BE_π formulas is that of an expander. Intuitively, this is a formula that constrains new propositional letters on the basis of the old ones in an arbitrary (homogeneous) interval structure.

Definition 7. *Let $\Sigma \subseteq \Sigma'$ be two signatures, and let $\mathcal{S} = (\mathbb{I}(N), \sigma)$ and $\mathcal{S}' = (\mathbb{I}(N'), \sigma')$ be interval structures over Σ and Σ', respectively. We say that $\mathcal{S}'$ is an expansion of $\mathcal{S}$ if $N' = N$ and $\sigma'(I) \cap \Sigma = \sigma(I)$ for all intervals $I \in \mathbb{I}(N)$.*

An expander from Σ to Σ' is a BE_π formula ξ over Σ' such that, for every interval structure $\mathcal{S}$ over Σ, there is an expansion $\mathcal{S}'$ of $\mathcal{S}$ over Σ' that makes ξ valid.

We report below a simple lemma about expanders.

Lemma 8. *If ξ is an expander from Σ to Σ', ψ and ψ' are formulas over the signatures Σ and Σ', respectively, and ψ, ψ' are equivalent over all interval structures where ξ is valid, then ψ and $\psi' \wedge [G]\xi$ are equi-satisfiable.*

Proof. Suppose that $\psi' \wedge [G]\xi$ is satisfied by an interval structure $\mathcal{S}'$ over Σ'. Because, ξ is valid over $\mathcal{S}'$, ψ is equivalent to ψ' over $\mathcal{S}'$, and hence $\mathcal{S}'$ satisfies ψ. Conversely, if ψ is satisfied by an interval structure $\mathcal{S}$ over Σ, then there is an expansion $\mathcal{S}'$ of $\mathcal{S}$ that makes ξ valid. This implies that ψ and ψ' are equivalent over $\mathcal{S}'$. Hence $\mathcal{S}'$ satisfies ψ', and $\psi' \wedge [G]\xi$ as well. □

B. Minimal witnessing intervals

Recall that the normalization of a BE_π formula replaces subformulas $\langle B\rangle\varphi$ (resp., $\langle E\rangle\varphi$) of depth $d > 2$ with equivalent formulas $\langle B\rangle\varphi'$ (resp., $\langle E\rangle\varphi'$) of depth 2. In this respect, a simple observation is that, in order to determine which intervals satisfy $\langle B\rangle\varphi$ (resp., $\langle E\rangle\varphi$), one could look at intervals that satisfy φ and that are *minimal* for the prefix (resp., suffix) relation.

Definition 9. *Given a BE_π formula φ, an interval structure $\mathcal{S}$, and an interval I in it, we say that I is prefix-minimal (resp., suffix-minimal) for φ if $\mathcal{S}, I \vDash \varphi$ and $\mathcal{S}, J \nvDash \varphi$ for every $J <_B I$ (resp., $J <_E I$).*

We will see later that prefix/suffix-minimal intervals for φ can be unambiguously identified, once their endpoints are annotated with fresh propositional letters, using a formula φ' of size proportional to that of φ, but with depth just 1. A simplified account of this technique was already given in Example 5. Below, we discuss the approach under a more general perspective and highlight a potential issue with overlapping minimal witnesses.

Example 10. *Suppose that φ is a formula of depth 2. We aim at replacing it with a formula φ' of depth 1, so that $\langle B\rangle\varphi'$ turns out to be equivalent to $\langle B\rangle\varphi$ in an appropriate expansion of the interval structure. As discussed earlier, a natural approach is to focus only on intervals that are prefix-minimal for φ, and mark their endpoints with suitable fresh propositional letters. For example, two prefix-minimal intervals for φ are represented in Figure 2 by the red brackets. We mark their left and right endpoints with fresh propositional letters ℓ and r, respectively, and we assume that the interval structure is expanded so as to satisfy the intended use of ℓ and r. We then define $\varphi' = (\langle B\rangle(\pi \wedge \ell) \wedge \langle E\rangle(\pi \wedge r)) \vee ((\pi \wedge \ell) \wedge (\pi \wedge r))$ and observe that every interval satisfying $\langle B\rangle\varphi$ must also*

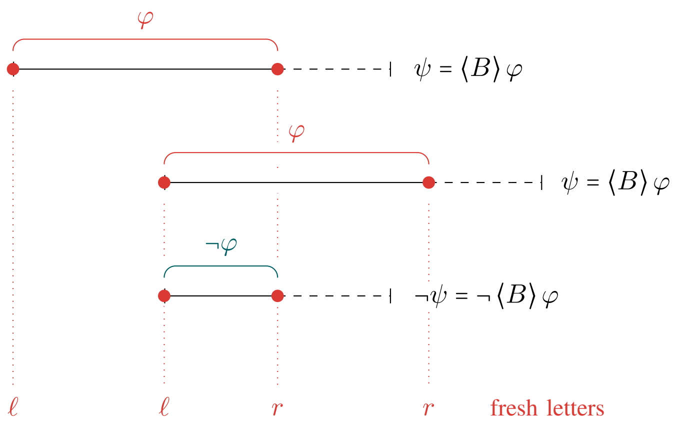

Fig. 2. Overlapping prefix-minimal intervals for φ, and their intersection.

satisfy $\langle B \rangle \varphi'$. So one might be tempted to replace φ with φ'. Unfortunately, while $\langle B \rangle \varphi$ entails $\langle B \rangle \varphi'$, the converse is not true, as the intersection of any two prefix-minimal intervals for φ does not always satisfy φ (see the blue bracket in Figure 2). In general, in order to mark the endpoints of minimal intervals without ambiguities, one could use different letters to mark the endpoints of any two overlapping intervals. More precisely, one should introduce as many copies of letters ℓ, r as the maximum number of overlapping prefix-minimal intervals for φ that have different right endpoints.

C. Encoding of minimal witnessing intervals

Example 10 brings up a third ingredient that is crucial for the normalization procedure, as it suggests that, in order to mark without ambiguities the endpoints of prefix-minimal (resp., suffix-minimal) intervals for a formula φ, one must first bound the number of distinct right (resp., left) endpoints of overlapping intervals. A bound will be shown precisely in Corollary 14 below.

Definition 11. *A set $\mathcal{I}$ of intervals is an* intersecting family *if there is a point x that is contained in every interval of $\mathcal{I}$.*

An example of an intersecting family of intervals is shown to the left of Figure 3.

Towards proving the desired bound, we shall first establish two auxiliary lemmas. The first lemma relates the maximum cardinality of a partially ordered set (e.g., an intersecting family of intervals, partially ordered by containment) to the maximum cardinality of its chains and anti-chains. Formally, a *chain* of a partially ordered set is a subset of pairwise comparable elements. An *anti-chain* is a subset of pairwise incomparable elements. The first lemma is in fact a rephrasing of Dilworth's theorem [8] (we give a proof here for self-containment):

Lemma 12. *Let X be a partially ordered set and suppose that all its chains and anti-chains have cardinality at most n. Then the cardinality of X is at most n^2.*

Proof. To begin with, notice that X is well-founded, due to the hypothesis that chains have cardinality at most n. Define the partition $Y_1, Y_2, \ldots$ of X, where each Y_i contains all and only the *minimal* elements of $X \setminus \bigcup_{j<i} Y_j$ — in particular, each Y_i is defined inductively on the basis of the previous sets $Y_1, \ldots, Y_{i-1}$. By construction, every subset Y_i is an anti-chain, and hence, by the hypotheses of the claim, it has cardinality at most n.

Let us now bound by n the number of subsets of the partition. Towards a contradiction, assume that $Y_1, Y_2, \ldots, Y_{n+1}$ belong to the partition of X. By construction, for every $1 < i \leq n+1$ and every $y \in Y_i$, there is $y' \in Y_{i-1}$ such that $y' < y$ (otherwise y should have been added to Y_{i-1}). Using this property and a simple induction, we can construct a chain of length $n+1$: we start by taking an arbitrary $y_{n+1} \in Y_{n+1}$ and then we repeatedly use the property to prepend to a chain $y_i < y_{i+1} < \cdots < y_{n+1}$, with $i > 1$, $y_i \in Y_i$, $y_{i+1} \in Y_{i+1}$, $\ldots$, $y_{n+1} \in Y_{n+1}$, a new element $y_{i-1} < y_i$, with $y_{i-1} \in Y_{i-1}$. Clearly, such a chain of length $n+1$ leads to a contradiction, and hence the partition $Y_1, Y_2, \ldots$ of X contains at most n elements. We conclude that $|X| = \sum_i |Y_i| \leq n^2$. ☐

Ultimately, we aim at applying Lemma 12 to bound the cardinality of every intersecting family of prefix-minimal (resp., suffix-minimal) intervals with pairwise distinct right (resp., left) endpoints, using the containment relation as partial order. To this end, it is crucial to bound the cardinalities of the chains and anti-chains of such an intersecting family. It will be also convenient to avoid singleton intervals when reasoning about intersecting families (note that there is at most one singleton interval in every intersecting family).

Lemma 13. *Let $\mathcal{S}$ be an interval structure, φ a BE_π formula, $\mathcal{I}$ an intersecting family of non-singleton prefix-minimal (resp., suffix-minimal) intervals for φ, with pairwise distinct right (resp., left) endpoints, and $\mathcal{I}'$ a chain or an anti-chain of $\mathcal{I}$, where the partial order is given by containment. We have that*

$$|\mathcal{I}'| \;\leq\; 2^{2|\varphi|}. \tag{1}$$

Proof. We present the proof for an intersecting family of non-singleton *prefix-minimal* intervals for φ (the case of suffix-minimal intervals uses symmetric arguments). Towards a contradiction, assume that there exist a BE_π formula φ, an intersecting family $\mathcal{I}$ of non-singleton prefix-minimal intervals for φ with pairwise distinct right endpoints, and a subset $\mathcal{I}'$ of $\mathcal{I}$ that is a chain or an anti-chain and that violates the bound (1), i.e., $\mathcal{I}'$ contains more than $2^{2|\varphi|}$ intervals. We also assume, without loss of generality, that φ is a smallest formula witnessing this violation of the bound (later we will exploit this assumption when considering families of prefix-minimal intervals for subformulas of φ).

Let $\boldsymbol{\partial}_B \varphi$ (resp., $\boldsymbol{\partial}_E \varphi$) be the set of formulas α such that $\langle B \rangle \alpha$ (resp., $\langle E \rangle \alpha$) is a subformula of φ with no other modal operator above it. For example, if $\varphi = \langle B \rangle \alpha_1 \wedge \langle B \rangle \langle B \rangle \alpha_2 \wedge \langle E \rangle \langle B \rangle \alpha_3$, then $\boldsymbol{\partial}_B \varphi = \{\alpha_1, \langle B \rangle \alpha_2\}$ and $\boldsymbol{\partial}_E = \{\langle B \rangle \alpha_3\}$. Note that $|\varphi| \geq |\boldsymbol{\partial}_B \varphi| + |\boldsymbol{\partial}_E \varphi|$.

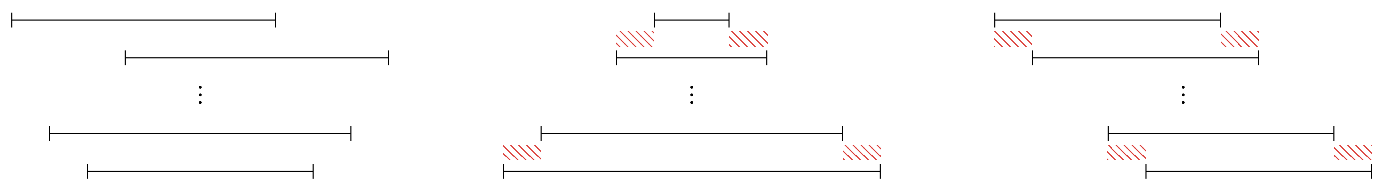

Fig. 3. From left to right: an intersecting family of intervals, a chain, and an anti-chain.

Define the *φ-profile* of a non-singleton interval I as the pair (B, E), where B (resp., E) is the set of formulas $\alpha \in \boldsymbol{\partial}_B\varphi$ (resp., $\alpha \in \boldsymbol{\partial}_E\varphi$) that hold at prefixes (resp., suffixes) of I. Note that any two non-singleton intervals with the same φ-profile either both satisfy φ or both satisfy $\neg\varphi$; in particular, this holds thanks to the fact that φ is in homogeneous normal form.

We also observe that there are at most $2^{|\boldsymbol{\partial}_B\varphi|+|\boldsymbol{\partial}_E\varphi|}$ distinct φ-profiles. Therefore, by our assumption on $\mathcal{I}'$, there are

$$n \;>\; 2^{2|\varphi|-|\boldsymbol{\partial}_B\varphi|-|\boldsymbol{\partial}_E\varphi|}$$

intervals $I_1, \ldots, I_n \in \mathcal{I}'$ with the same φ-profile. Without loss of generality, assume that the intervals $I_1, \ldots, I_n$ are listed based on the natural ordering of their right endpoints, that is, $\max(I_1) < \cdots < \max(I_n)$. Depending on $\mathcal{I}'$ being a chain or an anti-chain, the left endpoints of these intervals are also ordered, in descending, resp., ascending order (see Figure 3).

For the rest of the proof, unless otherwise stated, i will denote a natural number from 1 to $n-1$, and will be used in particular to index pairs of consecutive intervals, say I_i and I_{i+1}. For every i, let $w_i < x_i \leq y_i < z_i$ be the four endpoints of I_i and I_{i+1}. Further, let left-$\Delta_i = [w_i + 1, x_i]$ and right-$\Delta_i = [y_i, z_i - 1]$ (these intervals are represented by the red dashed rectangles in Figure 3). Thanks to the fact that $\mathcal{I}'$ is a chain or an anti-chain, the left-Δ_i's and the right-Δ_i's are pairwise disjoint across all i (this property will be used later and is the main reason for restricting our attention to chains and anti-chains).

Given $\alpha \in \boldsymbol{\partial}_B\varphi$ and $1 \leq i < n$, a *special $\langle B\rangle$-witness of α at i* (if it exists) is the prefix-minimal interval for α that has the same left endpoint as I_{i+1} and whose right endpoint belongs to right-Δ_i. Figure 4 gives two examples of special $\langle B\rangle$-witnesses, represented by green brackets: one example is for the chain arrangement and the other is for the anti-chain arrangement. Symmetrically, given $\alpha \in \boldsymbol{\partial}_E\varphi$ and $1 \leq i < n$, a *special $\langle E\rangle$-witness of α at i* (if it exists) is the suffix-minimal interval for α that has the same right endpoint as I_i and whose left endpoint belongs to left-Δ_i. Special $\langle E\rangle$-witnesses are represented in Figure 4 by blue brackets.

Now, we tag an index $1 \leq i < n$ with a pair (B, α) (resp., (E, α)) whenever $\alpha \in \boldsymbol{\partial}_B\varphi$ (resp., $\alpha \in \boldsymbol{\partial}_E\varphi$) and there is a special $\langle B\rangle$-witness (resp., $\langle E\rangle$-witness) of α at i. If there is no special witness for any α, then we tag i with the symbol $\bot$. Let $K_i = [\min(I_{i+1}), \max(I_i)]$ and observe that K_i is a *proper prefix* of I_{i+1}. We will prove that, for some index i, the interval K_i satisfies φ, thus contradicting prefix-minimality of I_{i+1}. Towards this, it will be sufficient to find an index i tagged with $\bot$. Indeed, we claim that

Claim 13.1. *If index i is tagged with $\bot$, then every formula $\alpha \in \boldsymbol{\partial}_B\varphi$ (resp., $\alpha \in \boldsymbol{\partial}_E\varphi$) that holds at a prefix (resp., suffix) of I_i also holds at a prefix (resp., suffix) of K_i, and vice versa.*

The above claim would then imply that the φ-profile of K_i coincides with that of I_i, and hence $K_i \vDash \varphi$.

Proof of the claim. Assume that index i is tagged with $\bot$. Consider some $\alpha \in \boldsymbol{\partial}_B\varphi$. If α holds at a prefix of I_i, then α holds at some prefix of I_{i+1} as well, because I_i and I_{i+1} have the same φ-profile. Let J be the smallest prefix of I_{i+1} that satisfies α. Due to i being tagged with $\bot$, we have that $\max(J) < \max(I_i) = \max(K_i)$, meaning that J is also a prefix of K_i. Conversely, if α holds at a prefix of K_i, then it trivially holds at a prefix of I_{i+1} as well, and thus it holds at a prefix of I_i, too, because I_i and I_{i+1} have the same φ-profile. Next, consider some $\alpha \in \boldsymbol{\partial}_E\varphi$. If α holds at some suffix of I_i, then let J be the smallest suffix of I_i that satisfies α. Due to i being tagged with $\bot$, we have that $\min(J) > \min(I_{i+1}) = \min(K_i)$, meaning that J is also a suffix of K_i. Conversely, assume that α holds at some suffix of K_i and let J be the smallest suffix of K_i that satisfies α. Once again, since i is tagged with $\bot$, we have that $\min(J) > \min(I_i)$, meaning that J is a suffix of I_i, too. ☐

It remains to prove that at least one index i is tagged with $\bot$. For this, we bound the number of indices tagged with pairs of the form (X, α), with $X = B$ (resp., $X = E$) and $\alpha \in \boldsymbol{\partial}_X\varphi$. By construction, for each tag (X, α), the special $\langle X\rangle$-witnesses of α form an intersecting (anti-)chain $\mathcal{I}_{X,\alpha}$ of prefix-minimal (resp., suffix-minimal) intervals for α. Moreover, we know that:

- All intervals in $\mathcal{I}_{X,\alpha}$ are non-singleton.
  This is because the only scenario where a *singleton* special $\langle X\rangle$-witness arises is when $\mathcal{I}'$ is anti-chain, $n = 2$, and $\max(I_1) = \min(I_2)$. This scenario is however excluded by the fact that $n > 2^{2|\varphi|-|\boldsymbol{\partial}_B\varphi|-|\boldsymbol{\partial}_E\varphi|} \geq 2$.
- The intervals in $\mathcal{I}_{X,\alpha}$ have pairwise distinct right (resp., left) endpoints.
  This is because those endpoints belong to the intervals right-Δ_i (resp., left-Δ_i), which are pairwise disjoint across all i's.
- The cardinality of each (anti-)chain $\mathcal{I}_{X,\alpha}$ is at most $2^{2|\alpha|}$.

Fig. 4. Special witnesses in a chain (left) and in an anti-chain (right).

This is thanks to the previous properties and because α is a proper subformula of φ, which was assumed to be a smallest formula violating the bound (1).

In view of the last property, we derive that the number of indices that are *not* tagged with $\perp$ is

$$\begin{aligned} n' &\le \textstyle\sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2^{2|\alpha|} + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2^{2|\alpha|} \\ &\le 2^{\sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2|\alpha| + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2|\alpha|} , \end{aligned}$$

where the last inequality follows from majorating sums with products. Next, recall that $n > 2^{2|\varphi|-|\boldsymbol{\partial}_B\varphi|-|\boldsymbol{\partial}_E\varphi|}$, and hence the number of indices $1 \le i < n$ that are tagged with $\perp$ is

$$n - 1 - n' \ \ge\ 2^{2|\varphi|-|\boldsymbol{\partial}_B\varphi|-|\boldsymbol{\partial}_E\varphi|} \ - \ 2^{\sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2|\alpha| + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2|\alpha|} .$$

We prove that the right hand-side number is always positive by showing that $2|\varphi|-|\boldsymbol{\partial}_B\varphi|-|\boldsymbol{\partial}_E\varphi| > \sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2|\alpha| + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2|\alpha|$. We distinguish two cases, depending on whether or not φ contains modal operators. If φ contains no modal operators, then $2|\varphi| - |\boldsymbol{\partial}_B\varphi| - |\boldsymbol{\partial}_E\varphi| = 2|\varphi| > 0 = \sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2|\alpha| + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2|\alpha|$. Otherwise, if φ contains at least one modal operator, then we observe that *(i)* the size of φ is at least the sum of the sizes of the subformulas $\langle X\rangle\,\alpha$, for $X \in \{B,E\}$ and $\alpha \in \boldsymbol{\partial}_X\varphi$, which are $|\langle X\rangle\,\alpha| = |\alpha| + 1$, and *(ii)* $\sum_{\alpha\in\boldsymbol{\partial}_X\varphi}(|\alpha| + 1) = \big(\sum_{\alpha\in\boldsymbol{\partial}_X\varphi}|\alpha|\big) + |\boldsymbol{\partial}_X\varphi|$, for $X \in \{B,E\}$. From this we derive:

$$\begin{aligned} & 2|\varphi| - |\boldsymbol{\partial}_B\varphi| - |\boldsymbol{\partial}_E\varphi| \\ \ge\ & \textstyle\sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2|\alpha| + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2|\alpha| + |\boldsymbol{\partial}_B\varphi| + |\boldsymbol{\partial}_E\varphi| \\ >\ & \textstyle\sum_{\alpha\in\boldsymbol{\partial}_B\varphi} 2|\alpha| + \sum_{\alpha\in\boldsymbol{\partial}_E\varphi} 2|\alpha| . \end{aligned}$$

We have just shown that at least one index i must be tagged with $\perp$, which completes the proof of the lemma. □

Putting together Lemmas 12 and 13, we obtain the desired bound for an arbitrary intersecting family of non-singleton prefix/suffix-minimal intervals for φ:

Corollary 14. *Let $\mathcal{S}$ be an interval structure, φ a BE_π formula, and $\mathcal{I}$ an intersecting family of non-singleton prefix-minimal (resp., suffix-minimal) intervals for φ. Then the number of distinct right (resp., left) endpoints of intervals of $\mathcal{I}$ is at most $2^{4|\varphi|}$.*

We conclude this part by showing how prefix-minimal intervals for φ can be characterized using fresh propositional letters and suitable formulas $\mathrm{flat}(\varphi)$ and $\mathrm{enc}(\varphi)$ (a similar corollary can be stated for suffix-minimal intervals).

Corollary 15. *Consider a BE_π formula φ over a signature Σ and let $\Sigma' = \Sigma \uplus \{p_1,\dots,p_m,\ell,r,s\}$, where $p_1,\dots,p_m$, ℓ,r,s are fresh propositional letters and $m = 4|\varphi|$ (this m is precisely the exponent appearing in the bound of Corollary 14). Define the BE_π formulas*[2]

$$\begin{aligned} \mathrm{flat}(\varphi) \ &= \ \Big(\langle B\rangle(\pi\wedge\ell) \ \wedge \ \langle E\rangle(\pi\wedge r) \ \wedge \\ &\qquad \textstyle\bigwedge_{i=1,\dots,m} \big(\langle B\rangle(\pi\wedge p_i) \leftrightarrow \langle E\rangle(\pi\wedge p_i) \big) \Big) \\ &\quad \vee \ (\pi\wedge s) \\ \mathrm{enc}(\varphi) \ &= \ \underbrace{\big(\mathrm{flat}(\varphi)\wedge\neg\langle B\rangle\,\mathrm{flat}(\varphi) \ \to \ \varphi\big)}_{\text{prefix-minimial intervals for } \mathrm{flat}(\varphi) \text{ satisfy } \varphi} \\ &\quad \wedge \ \underbrace{\big(\varphi\wedge\neg\langle B\rangle\,\mathrm{flat}(\varphi) \ \to \ \mathrm{flat}(\varphi)\big)}_{\text{prefix-minimal intervals for } \varphi \text{ satisfy } \mathrm{flat}(\varphi)} \end{aligned}$$

We have that $\mathrm{enc}(\varphi)$ is an expander from Σ to Σ' and that $\langle B\rangle\,\varphi$ and $\langle B\rangle\,\mathrm{flat}(\varphi)$ are equivalent over all interval structures that make $\mathrm{enc}(\varphi)$ valid.

Proof. Let us first explain the intended use of the fresh propositional letters $p_1,\dots,p_m$, ℓ,r,s. The letters $p_1,\dots,p_m$ will annotate points of an interval structure with m-tuples of bits, thus enumerating an exponentially-large set (e.g., $\{1,\dots,2^m\}$). More precisely, the left and right endpoints of every non-singleton prefix-minimal interval for φ will be identified by having labels ℓ and r, respectively, and the same m-tuple of bits — this correlation between endpoints is checked by the first disjunct of $\mathrm{flat}(\varphi)$. Singleton prefix-minimal intervals for φ will instead be identified by the special label s — this is checked in the second disjunct of $\mathrm{flat}(\varphi)$. Another important constraint is that every two intersecting intervals that are non-singleton, prefix-minimal for φ, and not a suffix one of another will have their endpoints marked by different m-tuples of bits.

Corollary 14 guarantees the existence of an annotation satisfying all the above constraints. Such an annotation is enforced precisely by the formula $\mathrm{enc}(\varphi)$, which turns out to be an expander from Σ to Σ' (namely, every interval structure over Σ admits an expansion over Σ' that makes $\mathrm{enc}(\varphi)$ valid). Moreover, if the annotation is correct, namely, if $\mathrm{enc}(\varphi)$ is valid over an expanded interval structure, then every prefix-minimal interval for φ is also a prefix-minimal interval for $\mathrm{flat}(\varphi)$, and vice versa. Note that there may still exist intervals

[2]Note that, despite the notation, the formula $\mathrm{flat}(\varphi)$ only depends on the signature and the size of φ, whereas $\mathrm{enc}(\varphi)$ depends entirely on φ.

that satisfy φ but not $\mathrm{flat}(\varphi)$, or vice versa; however, those intervals will always contain proper prefixes that satisfy both φ and $\mathrm{flat}(\varphi)$. Overall, this proves that the two formulas $\langle B \rangle \varphi$ and $\langle B \rangle \mathrm{flat}(\varphi)$ are equivalent over expanded interval structures that make $\mathrm{enc}(\varphi)$ valid. □

D. Normalization procedure

We are now ready to describe the normalization procedure underlying Theorem 4. Let ψ be a BE_π formula. The normalization of ψ consists of repeatedly applying some rewriting steps that preserve satisfiability and progressively reduce the number of distinct subformulas of depth larger than 2, until a shallow normal form is eventually obtained.

Every rewriting step is applied to a formula of the form $\psi_i \wedge [G]\,\xi_i$ over a signature Σ_i (initially, $\psi_0 = \psi$, $\xi_i = \mathrm{true}$, and $\Sigma_i = \Sigma$), and results in an equi-satisfiable formula $\psi_{i+1} \wedge [G]\,\xi_{i+1}$ over an extended signature Σ_{i+1}. To perform the rewriting step, we must choose a subformula $\langle X \rangle \varphi$ of $\psi_i \wedge [G]\,\xi_i$, for some $X \in \{B, E\}$, that has depth $d > 2$ and that does not occur under the scope of any other modal operator, except possibly the operator $[G]$ that has ξ_i as argument. We then use Corollary 15 to obtain an expander $\mathrm{enc}(\varphi)$ from Σ_i to Σ_{i+1} and a formula $\langle X \rangle \mathrm{flat}(\varphi)$ equivalent to $\langle X \rangle \varphi$ over every interval structure that makes $\mathrm{enc}(\varphi)$ valid. We then rewrite $\psi_i \wedge [G]\,\xi_i$ into the formula

$$\underbrace{\psi_i\big[\,\langle X \rangle \varphi \,/\, \langle X \rangle \mathrm{flat}(\varphi)\big]}_{\psi_{i+1}} \;\wedge\; [G]\Big(\underbrace{\xi_i\big[\,\langle X \rangle \varphi \,/\, \langle X \rangle \mathrm{flat}(\varphi)\big] \wedge \mathrm{enc}(\varphi)}_{\xi_{i+1}}\Big) \qquad (\dagger)$$

Thanks to distributivity of $[G]$ with respect to $\wedge$, the formula (†) is equivalent to $(\psi_i \wedge [G]\,\xi_i)[\langle X \rangle \varphi \,/\, \langle X \rangle \mathrm{flat}(\varphi)] \wedge [G]\,\mathrm{enc}(\varphi)$. Moreover, thanks to Lemma 8, the latter formula is equi-satisfiable as $\psi_i \wedge [G]\,\xi_i$. This completes the description of a rewriting step.

Let us now analyse the complexity of the normalization procedure. The procedure terminates when one cannot choose any subformula $\langle X \rangle \varphi$ with the desired properties: in this case the rewritten formula $\psi_i \wedge [G]\,\xi_i$ turns out to be in shallow normal form and we can let $\psi^\star = \psi_i \wedge [G]\,\xi_i$. To bound the number of rewriting steps, we study how a single rewriting step affects the number of distinct subformulas of depth larger than 2. As for $\langle X \rangle \varphi$, we observe that this subformula does not occur anymore in the rewritten formula $\psi_{i+1} \wedge [G]\,\xi_{i+1}$ (in particular, the inteded use of $\mathrm{enc}(\varphi)$ is to entail $\langle X \rangle \varphi \leftrightarrow \langle X \rangle \mathrm{flat}(\varphi)$, but the chosen writing in the statement of Corollary 15 avoids having φ under the scope of a modal operator, thus guaranteeing that $\mathrm{enc}(\varphi)$ has depth at most 2). On the other hand, new occurrences of subformulas may emerge in $\psi_{i+1} \wedge [G]\,\xi_{i+1}$: these are either formulas of depth at most 2 (e.g., $\langle X \rangle \mathrm{flat}(\varphi)$) or copies of formulas that already occur in $\psi_i \wedge [G]\,\xi_i$ (e.g., φ). Summing up, the effect of a rewriting step is to decrease the number of *distinct* subformulas of depth larger than 2. This implies that the number of rewriting steps is at most linear in the size of the original formula ψ. Finally, each rewriting step is purely syntactical and can be carried out efficiently on the involved formula $\psi_i \wedge [G]\,\xi_i$, whose size grows at most linearly with i. This shows that the entire normalization procedure can be performed in polynomial time w.r.t. $|\psi|$, and completes the proof of Theorem 4. □

IV. Complexity of the satisfiability problem

In this section, we build up on the previous normalization result to prove a tight complexity bound:

Theorem 16. *The satisfiability problem for* BE *logic restricted to homogeneous interval structures is* ExpSpace*-complete.*

An ExpSpace lowerbound for BE under homogeneity was already proven in [3], so we focus on the upperbound. In view of Proposition 2 and Theorem 4, given any BE formula ψ, one can compute in polynomial time a formula $\psi^\star$ in shallow normal form that is equi-satisfiable as ψ over homogeneous interval structures. This also means that $\psi^\star$ has size at most polynomial in $|\psi|$. We argue below that one can test satisfiability of a formula in shallow normal form in exponential space with respect to the size of the formula itself. Together with the previous observations, this proves Theorem 16.

A. Composition of logical types

We need to formalize a notion of logical type, similar to the notion of profile used in the proof of Lemma 13, that not only determines which formulas hold at a given interval, but also satisfies mild compositional properties, that is, under suitable conditions, one can compute the type of the sum of two adjacent intervals on the basis of the types of the original intervals. It will be convenient to define types separately for formulas of depth 0, 1, and 2 (there is no need to consider higher depths, as we assume to deal with formulas in shallow normal form). We will first present the rather simple definitions and properties of depth-0 and depth-1 types, and then focus on the more complex notion of depth-2 type.

Depth-0 *and depth-*1 *types.* We fix, once and for all, an interval structure $\mathcal{S} = (\mathbb{I}(N), \sigma)$ and we assume that all formulas are over the signature Σ of $\mathcal{S}$.

Definition 17. The *depth-*0 *type of an interval* I*, denoted* $\mathrm{type}^0(I)$*, is either the set* $\{\pi\} \cup \{p \in \Sigma : \mathcal{S}, I \vDash \pi \wedge p\}$ *or the empty set, depending on whether* I *is a singleton or not.*

The *depth-*1 *type of an interval* $I = [x, y]$ *is the quadruple* $\mathrm{type}^1(I) = (S, T, B, E)$*, where* S *is the symbol 1, 2, or 3, depending on whether* I *contains one point, two points, or more,* $T = \mathrm{type}^0(I)$*,* $B = \mathrm{type}^0([x, x])$*, and* $E = \mathrm{type}^0([y, y])$*.*

It is easy to see that depth-0 (resp., depth-1) types of adjacent intervals can be composed to form the depth-0 (resp., depth-1) type of the sum of the two intervals. One can also verify that the depth-0 (resp., depth-1) type of an interval determines which formulas of depth 0 (resp., depth at most 1) hold at that interval. These simple results are formalized in the next two lemmas below.

Lemma 18. *For both* $d = 0$ *and* $d = 1$*, there is a composition operator* $\cdot$ *on depth-*d *types that is computable in polynomial time and such that, for all pairs of adjacent intervals* I, J*, with* $\max(I) + 1 = \min(J)$*,* $\mathrm{type}^d(I) \cdot \mathrm{type}^d(J) = \mathrm{type}^d(I \cup J)$*.*

Proof. The composition of depth-0 types is trivial: for every pair of depth-0 types T, T', we simply let $T \cdot T' = \varnothing$. This is correct because the sum of two adjacent intervals always results in a non-singleton interval, whose depth-0 type is the empty set.

As for the composition of two depth-1 types, say $\mathcal{T} = (S, T, B, E)$ and $\mathcal{T}' = (S', T', B', E')$, we let $\mathcal{T} \cdot \mathcal{T}' = (S'', T \cdot T', B, E')$, where S'' is either 2 or 3 depending on whether $S = S' = 1$ or not, and $T \cdot T'$ is the composition of the depth-0 types T and T', as defined just above. It is immediate to check that if $\mathcal{T} = \mathrm{type}^1([x, y])$ and $\mathcal{T}' = \mathrm{type}^1([y+1, z])$, then $\mathrm{type}^1([x, z]) = \mathcal{T} \cdot \mathcal{T}'$. □

Lemma 19. *For both* $d = 0$ *and* $d = 1$*, for every* BE_π *formula* φ *of depth at most* d*, and for all intervals* I, J *such that* $\mathrm{type}^d(I) = \mathrm{type}^d(J)$*, we have* $\mathcal{S}, I \vDash \varphi$ *iff* $\mathcal{S}, J \vDash \varphi$*. Moreover, whether* $\mathcal{S}, I \vDash \varphi$ *holds or not can be decided in polynomial time given* φ *and* $\mathrm{type}^d(I)$*.*

Proof. We first prove the claim for $d = 0$. For the case $\varphi = \pi$, we have $\mathcal{S}, I \vDash \varphi$ if and only if I is a singleton, or, equally, $\pi \in \mathrm{type}^0(I)$. The case $\varphi = \pi \wedge p$ is trivial as well, as we have $\mathcal{S}, I \vDash \varphi$ if and only if $p \in \mathrm{type}^0(I)$. It remains to consider the case where φ is a Boolean combination of the previous atomic formulas. In this case, we determine the evaluation of φ at I "homomorphically" on the basis of the evaluations of the atomic formulas.

Let us now prove the claim for $d = 1$. The interesting cases are when φ has depth 0 or it is of the form $\langle B \rangle \alpha$ or $\langle E \rangle \alpha$, with α again of depth 0. Once the claim is proved for these cases, it can be generalized to Boolean combinations of those formulas using the same arguments as before. Let $I = [x, y]$ and $\mathrm{type}^1(I) = (S, T, B, E)$, and recall that $T = \mathrm{type}^0(I)$, $B = \mathrm{type}^0([x, x])$, and $E = \mathrm{type}^0([y, y])$.

If φ has depth 0, then we know that the component T (= $\mathrm{type}^0(I)$) already determines whether or not $\mathcal{S}, I \vDash \varphi$.

If $\varphi = \langle B \rangle \alpha$, we further distinguish three subcases, depending on S. If $S = 1$, then I is a singleton and hence $\mathcal{S}, I \nvDash \langle B \rangle \alpha$. If $S = 2$, then the only prefix of I is the singleton interval $[x, x]$, hence $\mathcal{S}, I \vDash \langle B \rangle \alpha$ iff $\mathcal{S}, [x, x] \vDash \alpha$. Since α has depth 0, the latter condition can be decided using the type $B = \mathrm{type}^0([x, x])$. If $S = 3$, then since α is a Boolean combination of formulas of the form π or $\pi \wedge p$, with $p \in \Sigma$, it suffices to consider only two prefixes of I: the singleton interval $J_0 = [x, x]$ and the interval $J_1 = [x, x+1]$. In particular, we have $\mathcal{S}, I \vDash \langle B \rangle \alpha$ if and only if $\mathcal{S}, J_0 \vDash \alpha$ or $\mathcal{S}, J_1 \vDash \alpha$. Again, the latter two conditions are determined by the depth-0 types of J_0 and J_1, which are B and $\varnothing$, respectively. This shows how to determine whether $\mathcal{S}, I \vDash \langle B \rangle \alpha$ using the type $\mathrm{type}^1(I) = (S, T, B, E)$.

The remaining case is that of a formula $\varphi = \langle E \rangle \alpha$, which can be handled by symmetric arguments, using the component E instead of B. □

Depth-2 types. We now introduce types for depth-2 formulas. The machinery here is not as neat as one could hope, as there is a trade-off between the desired compositional properties and the number of possible depth-2 types. As an example, full compositionality of types for depth-2 formulas can only hold if we allow doubly exponentially many types with respect to the size of the underlying signature — this can be shown formally using arguments based on communication complexity and the fact that a depth-2 formula can describe a Stockmeyer's counter of level 2 [15]. In order to ease compositional properties while maintaining the number of types as low as possible, we will parameterise depth-2 types by a formula and some contexts.

We first discuss a couple of tentative definitions, with their drawbacks. Following the same principle used to define depth-1 types, one may define the depth-2 type of an interval I as $(\mathcal{T}, \mathscr{B}, \mathscr{E})$, where $\mathcal{T} = \mathrm{type}^1(I)$, $\mathscr{B} = \{\mathrm{type}^1(J) : J <_B I\}$, and $\mathscr{E} = \{\mathrm{type}^1(J) : J <_E I\}$. This notion of depth-2 type would be fully compositional and would determine the evaluation of every depth-2 formula in homogeneous normal form (proofs omitted). Unfortunately, there could be doubly exponentially many such types with respect to the size of the signature, and this would not be compatible with the intended use that we will make in the satisfiability procedure. Another option would be to parameterise the depth-2 type of I by a formula φ and define it as the triple $(\mathcal{T}, \mathscr{B}, \mathscr{E})$, where $\mathcal{T} = \mathrm{type}^1(I)$ as before, and $\mathscr{B}$ (resp., $\mathscr{E}$) is the set of subformulas α of φ that hold at proper prefixes (resp., suffixes) of I. Of course, the resulting type would determine the evaluation of φ at the interval I. This second attempt would also generate at most exponentially many depth-2 type with respect to the size of φ. On the other hand, the resulting types would not carry enough information to be composable, the reason being that it is not sufficient to know which depth-1 subformulas hold at two adjacent intervals in order to derive which depth-1 subformulas hold at the union interval. The appropriate notion of depth-2 type is somehow a blend of the two attempts that we have just discussed.

Let us now fix some other useful notation and terminology:

- Given a BE_π formula φ, we denote by $\mathrm{Depth}^{\leq 1}(\varphi)$ the set of subformulas of φ of depth at most 1.
- Lemma 19 states that the depth-1 type of an interval I effectively determines which formulas of depth at most 1 hold at I. This motivates the following notation: given a depth-1 type $\mathcal{T}$ and a formula $\alpha \in \mathrm{Depth}^{\leq 1}(\varphi)$, we write $\mathcal{T} \vdash \alpha$ to state that $\mathcal{S}, I \vDash \alpha$ for some (or, equally, for every) interval I such that $\mathrm{type}^1(I) = \mathcal{T}$ (this latter property can be tested efficiently given $\mathcal{T}$ and α).
- By Lemma 18, depth-1 types are equipped with a composition operation $\cdot$ that forms a semigroup structure. We complete the structure into a monoid by introducing the

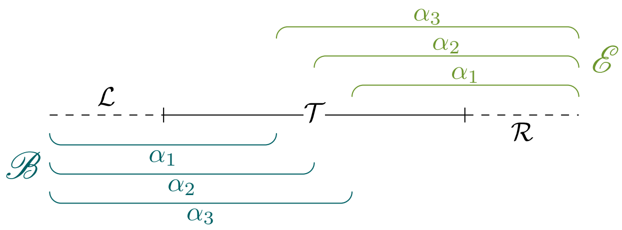

Fig. 5. Components of a depth-2 type.

*dummy depth-*1 *type* ε and by assuming that $\varepsilon \cdot \mathcal{T} = \mathcal{T} \cdot \varepsilon = \mathcal{T}$ for every depth-1 type $\mathcal{T}$.

Definition 20. *Let $\mathcal{L}, \mathcal{R}$ be some (possibly dummy) depth-*1 *types. The depth-*2 *φ-type of an interval I with left and right contexts $\mathcal{L}, \mathcal{R}$ is the tuple* $\text{type}^2_{\varphi,\mathcal{L},\mathcal{R}}(I) = (\mathcal{L}, \mathcal{R}, \mathcal{T}, \mathscr{B}, \mathscr{E})$, *where*

- $\mathcal{T} = \text{type}^1(I)$,
- $\mathscr{B} = \{\alpha \in \text{Depth}^{\leq 1}(\varphi) \,:\, \exists J <_B I \ \ \mathcal{L} \cdot \text{type}^1(J) \vdash \alpha\}$,
- $\mathscr{E} = \{\alpha \in \text{Depth}^{\leq 1}(\varphi) \,:\, \exists J <_E I \ \ \text{type}^1(J) \cdot \mathcal{R} \vdash \alpha\}$.

We give some intuition about the components of a depth-2 φ-type (the reader can also refer to Figure 5). The component $\mathcal{T}$ is nothing but the depth-1 type of the reference interval I, thus determining which formulas of depth at most 1 hold at I. The components $\mathcal{L}$ and $\mathcal{R}$ represent the depth-1 types of some intervals adjacent to I, to the left and to the right respectively, and will be used as contexes for an operation of composition. The set $\mathscr{B}$ represents which subformulas of φ of depth at most 1 hold at some intervals I' that overlap I to the left (i.e., such that $\min(I') \leq \min(I) \leq \max(I') < \max(I)$), provided that the depth-1 type of $K = I' \smallsetminus I$ coincides with the left context $\mathcal{L}$. The set $\mathscr{E}$ provides similar information for the intervals I' that overlap I to the right and such that $\text{type}^1(I' \smallsetminus I) = \mathcal{R}$. As a special case, we observe that when $\mathcal{L} = \mathcal{R} = \varepsilon$, one could let I' range over prefixes or suffixes of I, thus determining which subformulas hold at prefixes and suffixes of the reference interval I. In particular, this can be used to determine the evaluation of φ at I, and generalizes the second attempt of definition of type that we discussed earlier.

Below, we prove the analogous of Lemmas 18 and 19 for depth-2 types.

Lemma 21. *There is a composition operator $\cdot$ on depth-*2 *φ-types that is computable in polynomial time and such that, for all contexts $\mathcal{L}, \mathcal{L}', \mathcal{R}, \mathcal{R}'$ and for all pairs of adjacent intervals I, I', if $\mathcal{L} \cdot \text{type}^1(I) = \mathcal{L}'$ and $\text{type}^1(I') \cdot \mathcal{R}' = \mathcal{R}$, then*

$$\text{type}^2_{\varphi,\mathcal{L},\mathcal{R}}(I) \cdot \text{type}^2_{\varphi,\mathcal{L}',\mathcal{R}'}(I') \;=\; \text{type}^2_{\varphi,\mathcal{L},\mathcal{R}'}(I \cup I') \;.$$

Proof. For the sake of brevity, let $\mathscr{T} = \text{type}^2_{\varphi,\mathcal{L},\mathcal{R}}(I)$ and $\mathscr{T}' = \text{type}^2_{\varphi,\mathcal{L}',\mathcal{R}'}(I')$, where $\mathscr{T} = (\mathcal{L}, \mathcal{R}, \mathcal{T}, \mathscr{B}, \mathscr{E})$, $\mathscr{T}' = (\mathcal{L}', \mathcal{R}', \mathcal{T}', \mathscr{B}', \mathscr{E}')$, $\mathcal{L} \cdot \mathcal{T} = \mathcal{L}'$, and $\mathcal{T}' \cdot \mathcal{R}' = \mathcal{R}$. We define the composition as

$$\mathscr{T} \cdot \mathscr{T}' \;=\; (\mathcal{L},\, \mathcal{R}',\, \mathcal{T} \cdot \mathcal{T}',\, \mathscr{B} \cup \mathscr{B}' \cup \mathscr{B}_\star,\, \mathscr{E} \cup \mathscr{E}' \cup \mathscr{E}_\star)$$

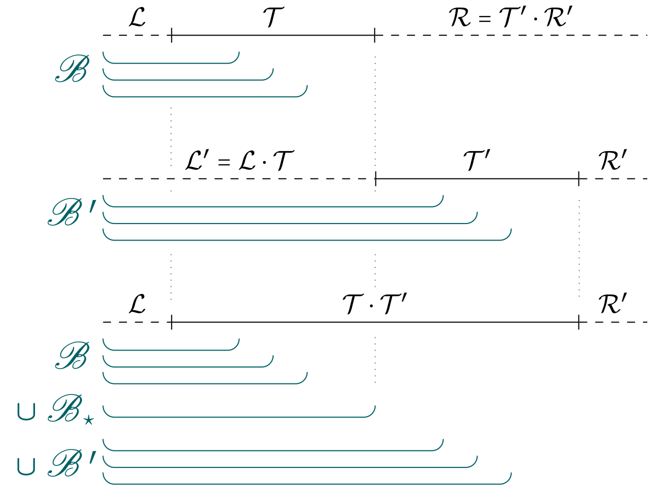

Fig. 6. Composition of depth-2 types.

where

$$\begin{aligned} \mathcal{B}_\star &= \{\alpha \in \text{Depth}^{\leq 1}(\varphi) \,:\, \mathcal{L}' \vdash \alpha\} \\ \mathcal{E}_\star &= \{\alpha \in \text{Depth}^{\leq 1}(\varphi) \,:\, \mathcal{R} \vdash \alpha\} \end{aligned}$$

(see Figure 6).

Note that, thanks to Lemma 19, the composition $\mathscr{T} \cdot \mathscr{T}'$ can be computed in polynomial time given the types $\mathscr{T}$ and $\mathscr{T}'$.

Below, we prove that the defined composition $\mathscr{T} \cdot \mathscr{T}'$ is correct, namely, it coincides with $\text{type}^2_{\varphi,\mathcal{L},\mathcal{R}'}(I \cup I')$. The latter type is of the form $(\mathcal{L}, \mathcal{R}', \mathcal{T}'', \mathscr{B}'', \mathscr{E}'')$, so the first two components of $\mathscr{T} \cdot \mathscr{T}'$ are clearly correct. It remains to prove that $\mathcal{T}'' = \mathcal{T} \cdot \mathcal{T}'$, $\mathscr{B}'' = \mathscr{B} \cup \mathscr{B}' \cup \mathscr{B}_\star$, and $\mathscr{E}'' = \mathscr{E} \cup \mathscr{E}' \cup \mathscr{E}_\star$. By Lemma 18 we have $\mathcal{T}'' = \text{type}^1(I \cup I') = \text{type}^1(I) \cdot \text{type}^1(I') = \mathcal{T} \cdot \mathcal{T}'$. Moreover, by Definition 20, $\mathscr{B}''$ contains the formulas $\alpha \in \text{Depth}^{\leq 1}(\varphi)$ that satisfy one of the following conditions:

1) $I'' \vDash \alpha$, for some interval I'' that overlaps I to the left (i.e., $\min(I'') \leq \min(I) \leq \max(I'') < \max(I)$) and such that $\text{type}^1(I'' \smallsetminus I) = \mathcal{L}$.
   Letting $K = I'' \smallsetminus I$ and $J = I'' \cap I$ and using Lemma 19, this condition is equivalent to
   $$\text{type}^1(I'') \;=\; \text{type}^1(K) \cdot \text{type}^1(J) \;=\; \mathcal{L} \cdot \text{type}^1(J) \;\vdash\; \alpha$$
   and hence to $\alpha \in \mathscr{B}$.
2) $I'' \vDash \alpha$, for some interval I'' that has I as a suffix (i.e., $\min(I'') \leq \min(I) \leq \max(I'') = \max(I)$) and such that $\text{type}^1(I'' \smallsetminus I) = \mathcal{L}$.
   Letting $K = I'' \smallsetminus I$ and using Lemma 19, together with the assumptions about the contexts $\mathcal{L}$ and $\mathcal{L}'$, this condition turns out to be equivalent to
   $$\text{type}^1(I'') \;=\; \text{type}^1(K) \cdot \mathcal{T} \;=\; \mathcal{L} \cdot \mathcal{T} \;=\; \mathcal{L}' \;\vdash\; \alpha$$
   and hence to $\alpha \in \mathscr{B}_\star$.
3) $I'' \vDash \alpha$, for some interval I'' that contains I, overlaps I' to the left (i.e., $\min(I'') \leq \min(I) \leq \max(I) < \max(I'') < \max(I')$), and such that $\text{type}^1(I'' \smallsetminus I') = \mathcal{L}$.

Letting $K = I'' \setminus I'$ and $J = I'' \cap I'$, and using again Lemma 19 and the assumptions about the contexts $\mathcal{L}$ and $\mathcal{L}'$, this condition turns out to be equivalent to

$$\begin{aligned}\mathrm{type}^1(I'') &= \mathrm{type}^1(K) \cdot \mathcal{T} \cdot \mathrm{type}^1(J) \\ &= \mathcal{L} \cdot \mathcal{T} \cdot \mathrm{type}^1(J) = \mathcal{L}' \cdot \mathrm{type}^1(J) \vdash \alpha\end{aligned}$$

and hence to $\alpha \in \mathscr{B}'$.

We have just shown that $\mathscr{B}'' = \mathscr{B} \cup \mathscr{B}' \cup \mathscr{B}_\star$. One proves $\mathscr{E}'' = \mathscr{E} \cup \mathscr{E}' \cup \mathscr{E}_\star$ using symmetric arguments. □

Lemma 22. *For all intervals I and J such that $\mathrm{type}^2_{\varphi,\varepsilon,\varepsilon}(I) = \mathrm{type}^2_{\varphi,\varepsilon,\varepsilon}(J)$ and for every BE_π formula φ of depth at most 2, we have $\mathcal{S}, I \vDash \varphi$ iff $\mathcal{S}, J \vDash \varphi$. Moreover, whether $\mathcal{S}, I \vDash \varphi$ holds can be decided in polynomial time from the given type $\mathrm{type}^2_{\varphi,\varepsilon,\varepsilon}(I)$.*

Proof. Let φ be a formula of depth at most 2 and let I be an interval with depth-2 type $(\varepsilon, \varepsilon, \mathcal{T}, \mathscr{B}, \mathscr{E})$, where both left and right contexts are ε.

If φ has depth smaller than 2, then by Lemma 19 the component $\mathcal{T} = \mathrm{type}^1(I)$ already determines (effectively in polynomial time) whether $\mathcal{S}, I \vDash \varphi$.

Otherwise, if φ has depth 2 and is of the form $\langle B \rangle \alpha$, then $\mathcal{S}, I \vDash \langle B \rangle \alpha$ iff there is a proper prefix J of I such that $\mathcal{S}, J \vDash \alpha$. Since $\alpha \in \mathrm{Depth}^{\leq 1}(\varphi)$, the latter condition is equivalent to $\varepsilon \cdot \mathrm{type}^1(J) \vdash \alpha$, and hence $\mathcal{S}, I \vDash \langle B \rangle \alpha$ iff $\alpha \in \mathscr{B}$. The case of $\varphi = \langle E \rangle \alpha$ is similar, but uses the component $\mathscr{E}$.

Finally, Boolean combinations of the previous formulas are evaluated homomorphically. □

B. Satisfiability procedure

As a warm-up, let us first describe the satisfiability procedure for a formula of depth at most 2; later we will generalize this to a formula in shallow normal form.

Let us fix a BE_π formula ψ of depth at most 2. Deciding satisfiability of ψ can be done in polynomial space, by reducing to non-emptiness of a language recognized by a suitable finite state automaton. To formalize the construction of the automaton from the given formula ψ, it is convenient to encode an interval structure $\mathcal{S} = (\mathbb{I}(N), \sigma)$ over the signature Σ by the finite word $w_{\mathcal{S}} = a_0 \dots a_{\max(N)}$ over the alphabet $\wp(\Sigma)$, where $a_i = \sigma(i)$ for all $i \in N$ (recall that N is a finite prefix of the natural numbers).

Lemma 23. *Given a BE_π formula ψ of depth at most 2, one can compute in polynomial space*[3] *a finite state automaton $\mathcal{A}_\psi$ that accepts all and only the encodings $w_{\mathcal{S}}$ of the interval structures $\mathcal{S}$ such that $\mathcal{S}, I \vDash \psi$, where I is the largest interval of $\mathcal{S}$.*

Proof sketch. The construction of $\mathcal{A}_\psi$ is quite standard, as it is the cascade product of three automata:

1) a deterministic automaton that computes in its states the depth-1 types of intervals corresponding to prefixes of the input,
2) a co-deterministic automaton that computes in its states the depth-1 types of intervals corresponding to suffixes of the input,
3) a deterministic automaton that computes the depth-2 ψ-type of prefixes of the input, with a constant dummy left context and right contexts given by the states of the previous automaton.

Transitions of these automata are defined using compositional properties of depth-1 and depth-2 types (Lemmas 18 and 21).

Below, we provide full details for the construction of $\mathcal{A}_\psi$. Like we have done for depth-1 types, we introduce *dummy depth-2 types* for abstracting an empty interval: these are tuples of the form $(\mathcal{L}, \mathcal{R}, \varepsilon, \varnothing, \varnothing)$, where $\mathcal{L}$ and $\mathcal{R}$ are left and right contexts and ε is the dummy depth-1 type (of course, there is exactly one dummy depth-2 type for each choice of the left and right contexts). As usual, a dummy type behaves as an identity w.r.t. composition with a depth-2 type, provided the contexts are compatible. We shall also use a generalization of the relation $\vdash$ that works with depth-2 types. Precisely, given a depth-2 type $\mathscr{T}$, we write $\mathscr{T} \vdash \psi$ whenever $\mathcal{S}, I \vDash \psi$ for some (or, equally, for every) interval I such that $\mathrm{type}^2_{\varphi,\varepsilon,\varepsilon}(I) = \mathscr{T}$.

- the alphabet A consists of subsets of the signature Σ;
- the state space Q consists of triples $q = (\mathcal{L}, \mathcal{R}, \mathscr{T})$, where $\mathcal{L}, \mathcal{R}$ are a depth-1 types and $\mathscr{T}$ is a depth-2 ψ-type with ε as left context and $\mathcal{R}$ as right context;
- the set I of initial states consists of triples $q = (\mathcal{L}, \mathcal{R}, \mathscr{T})$, where $\mathcal{L} = \varepsilon$ is the dummy depth-1 type and $\mathscr{T} = (\mathcal{L}, \mathcal{R}, \varepsilon, \varnothing, \varnothing)$ is a dummy depth-2 type;
- the set F of final sates consists of triples $q = (\mathcal{L}, \mathcal{R}, \mathscr{T})$, with $\mathcal{R} = \varepsilon$ and $\mathscr{T} \vdash \psi$;
- the set T of transition rules consists of the triples (q, a, q'), with $q = (\mathcal{L}, \mathcal{R}, \mathscr{T})$, $a \subseteq \Sigma$, and $q' = (\mathcal{L}', \mathcal{R}', \mathscr{T}')$, such that $\mathcal{L}' = \mathcal{L} \cdot \mathrm{type}^1(I_a)$, $\mathcal{R} = \mathrm{type}^1(I_a) \cdot \mathcal{R}$, and $\mathcal{T}' = \mathscr{T} \cdot \mathrm{type}^2_{\psi,\mathcal{L},\mathcal{R}'}(I_a)$, where I_a denotes the singleton interval labelled by the set a of propositional letters.

It is worth noting that the automaton $\mathcal{A}_\psi$ is unambiguous, namely, it admits at most one successful run on each input.

We now claim that, on every input $w_{\mathcal{S}} = a_0 \dots a_{n-1}$, the only possible runs of $\mathcal{A}_\psi$ that start and end in arbitrary states (not necessarily initial or final ones) are of the form

$$q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} \dots\dots \xrightarrow{a_{n-1}} q_n$$

with $q_i = (\mathcal{L}_i, \mathcal{R}_i, \mathscr{T}_i)$ such that, for all $i = 0, \dots, n$,

1) $\mathcal{L}_i = \mathcal{L}_0 \cdot \mathrm{type}^1([0, i-1])$,
2) $\mathcal{R}_i = \mathrm{type}^1([i, n-1]) \cdot \mathcal{R}_n$,
3) $\mathscr{T}_i = \mathscr{T}_0 \cdot \mathrm{type}^2_{\psi,\varepsilon,\mathcal{R}_i}([0, i-1])$.

Each of the above properties can be verified using a simple induction, either from smaller to larger i's or vice versa (we omit the tedious details).

From the properties stated in items 1., 2., 3. and the definitions of initial and final states, it immediately follows that $\mathcal{A}_\psi$ admits a successful run on $w_{\mathcal{S}}$ if and only if $\mathcal{S}, [0, n-1] \vDash \psi$.

[3]By computing an automaton in polynomial space we mean that its initial states, final states, and transitions can be enumerated in polynomial space. The enumeration procedures can be used within other algorithms of similar complexity, e.g., to test emptiness of the recognized language.

Finally, as for the complexity of constructing $\mathcal{A}_\psi$, we recall from Lemmas 18 and 21 that depth-1 and depth-2 types can be enumerated in polynomial space, and can be composed in polynomial time. This implies that the initial states and the transitions of $\mathcal{A}_\psi$ can be enumerated in polynomial space. To enumerate the final states, it suffices to test properties like $\mathscr{T} \vdash \psi$, for a given depth-2 type $\mathscr{T}$. This can be done in polynomial time thanks to Lemma 22. □

The fact that the automaton $\mathcal{A}_\psi$ above can be constructed from ψ in polynomial space, implies that (non-)emptiness of the recognized language can also be decided in polynomial space w.r.t. $|\psi|$. In its turn, this shows that the satisfiability of a BE_π formula ψ of depth at most 2 can be decided in polynomial space.

To conclude the proof of Theorem 16 it remains to reduce the satisfiability problem for a BE_π formula ψ in shallow normal form to the non-emptiness problem of an automaton $\mathcal{A}_\psi$ that is computable from ψ in exponential space. For this, it suffices to recall that ψ must be of the form $\varphi \wedge [G]\,\xi$, where both φ and ξ are BE_π formulas of depth at most 2. One uses Lemma 23 to construct the automata $\mathcal{A}_\varphi$ and $\mathcal{A}_{\neg\xi}$, whose languages contain encodings of models of φ and $\neg\xi$, respectively. From $\mathcal{A}_{\neg\xi}$, one can efficiently construct an automaton $\mathcal{A}_{\langle G\rangle\neg\xi}$ recognizing the language of words with infixes accepted by $\mathcal{A}_{\neg\xi}$, thus encoding models of $\langle G\rangle\neg\xi$ ($= \neg\,[G]\,\xi$). One then complements the latter automaton to obtain an automaton $\mathcal{A}_{[G]\,\xi}$ accepting the encodings of models of $[G]\,\xi$. Note that the latter step can be performed in exponential space in the size of ξ, by using an online version of the classical subset construction. Finally, one computes the product of the automata $\mathcal{A}_\varphi$ and $\mathcal{A}_{[G]\,\xi}$, so as to recognize the language of encodings of models of $\psi = \varphi \wedge [G]\,\xi$. It follows that non-emptiness of the latter language can be decided in exponential space w.r.t. the size of the original formula ψ. □

V. Conclusions

We have settled the question of whether the logic BE, interpreted over homogeneous interval structures, admits an elementary satisfiability problem. We have actually answered the question by giving an optimal ExpSpace decision procedure (ExpSpace-hardness was shown in [3]). As a by-product result, we have also devised a normal form for BE formulas that enforces a small bound to the number of nested modalities, while preserving satisfiability. Quite suprisingly such a normal form can be computed in polynomial time from arbitrary BE formulas, using a series of rewriting steps reminiscent of a quantifier elimination technique a-la Scott.

As for future work, one could try to see whether similar techniques are applicable to extensions of BE with modalities based on other Allen's interval relations (e.g., overlap, meet, the inverses of the prefix and suffix relations, etc.).

References

[1] J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983. doi:10.1145/182.358434.

[2] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, and P. Sala. Interval vs. point temporal logic model checking: An expressiveness comparison. *ACM Trans. Comput. Log.*, 20(1):4:1–4:31, 2019. URL: https://doi.org/10.1145/3281028, doi:10.1145/3281028.

[3] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, and P. Sala. Which fragments of the interval temporal logic HS are tractable in model checking? *Theor. Comput. Sci.*, 764:125–144, 2019. URL: https://doi.org/10.1016/j.tcs.2018.04.011, doi:10.1016/j.tcs.2018.04.011.

[4] L. Bozzelli, A. Molinari, A. Montanari, A. Peron, and P. Sala. Satisfiability and model checking for the logic of sub-intervals under the homogeneity assumption. *Log. Methods Comput. Sci.*, 18(1), 2022.

[5] L. Bozzelli, A. Montanari, A. Peron, and P. Sala. Adding the relation meets to the temporal logic of prefixes and infixes makes it EXPSPACE-complete. In *Proceedings 12th International Symposium on Games, Automata, Logics, and Formal Verification*, volume 346 of *EPTCS*, pages 179–194, 2021. URL: https://doi.org/10.4204/EPTCS.346.12, doi:10.4204/EPTCS.346.12.

[6] L. Bozzelli, A. Montanari, A. Peron, and P. Sala. PSPACE-completeness of the temporal logic of sub-intervals and suffixes. In *28th International Symposium on Temporal Representation and Reasoning*, volume 206 of *LIPIcs*, pages 9:1–9:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. URL: https://doi.org/10.4230/LIPIcs.TIME.2021.9, doi:10.4230/LIPIcs.TIME.2021.9.

[7] D. Bresolin, D. Della Monica, V. Goranko, A. Montanari, and G. Sciavicco. The dark side of interval temporal logic: marking the undecidability border. *Annals of Mathematics and Artificial Intelligence*, 71(1-3):41–83, 2014. doi:10.1007/s10472-013-9376-4.

[8] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, 51(1):161–166, 1950.

[9] J. Y. Halpern and Y. Shoham. A propositional modal logic of time intervals. *Journal of the ACM*, 38(4):935–962, 1991. doi:10.1145/115234.115351.

[10] K. Lodaya. Sharpening the undecidability of interval temporal logic. In *ASIAN*, LNCS 1961, pages 290–298. Springer, 2000.

[11] J. Marcinkowski and J. Michaliszyn. The undecidability of the logic of subintervals. *Fundamenta Informaticae*, 131(2):217–240, 2014. doi:10.3233/FI-2014-1011.

[12] A. Molinari, A. Montanari, A. Murano, G. Perelli, and A. Peron. Checking interval properties of computations. *Acta Informatica*, 53(6-8):587–619, 2016. doi:10.1007/s00236-015-0250-1.

[13] S. Schmitz. Complexity hierarchies beyond elementary. *ACM Trans. Comput. Theory*, 8(1):3:1–3:36, 2016. doi:10.1145/2858784.

[14] D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:377, 1962.

[15] L. J. Stockmeyer. *The complexity of decision problems in automata theory and logic*. PhD thesis, Massachusetts Institute of Technology. Dept. of Electrical Engineering, 1974.

[16] Y. Venema. A modal logic for chopping intervals. *Journal of Logic and Computation*, 1(4):453–476, 1991. doi:10.1093/logcom/1.4.453.