



The Lower Bound of Decidable Entailments in Separation Logic with Inductive Definitions

Mnacho Echenim, Radu Iosif, Nicolas Peltier

► To cite this version:

Mnacho Echenim, Radu Iosif, Nicolas Peltier. The Lower Bound of Decidable Entailments in Separation Logic with Inductive Definitions. ADSL 2020 - Workshop on Automated Deduction in Separation Logic, 2020, New Orleans, United States. hal-02388028

HAL Id: hal-02388028

<https://hal.science/hal-02388028>

Submitted on 30 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Lower Bound of Decidable Entailments in Separation Logic with Inductive Definitions

Mnacho Echenim¹, Radu Iosif² and Nicolas Peltier¹

¹ Univ. Grenoble Alpes, CNRS, LIG, F-38000 Grenoble France

² Univ. Grenoble Alpes, CNRS, VERIMAG, F-38000 Grenoble France

Abstract. The entailment between separation logic formulae with inductive predicates (also known as symbolic heaps) has been shown to be decidable for a large class of inductive definitions [4]. Recently, a 2EXPTIME algorithm has been proposed [7], however no precise lower bound is known (although a EXPTIME-hard bound for this problem has been established in [5]). In this paper, we show that deciding entailment between predicate atoms is 2EXPTIME-hard. The proof is based on a reduction from the membership problem for exponential-space bounded alternating Turing machines [3].

1 Introduction

In this paper, we consider the fragment of separation logic formulae known as *symbolic heaps*, defined as (separated) conjunctions of atoms. Such atoms may be either equational atoms, asserting equalities or disequalities between memory locations, points-to atoms asserting that some location refers to a given record, or may be built on additional predicates that assert that a part of the memory has a specific shape. For genericity, such predicates are associated with user-provided inductive definitions that allow one to describe custom data-structures. Satisfiability is EXPTIME-complete for such formulae [2], but entailment is not decidable in general³ [5,1]. However, the entailment problem was proven to be decidable for a large class of inductive definitions, whose syntactical restrictions ensure that the generated heap structures have a bounded-tree width [4], using a reduction to monadic second order logic interpreted over graphs. A EXPTIME-hard bound was established in [5], and very recently, a 2EXPTIME algorithm has been proposed [7]. However, personal communication with the authors of [7] revealed several problems related to quantifier instantiation that are potential sources of incompleteness. Thus we believe that their decision procedure might not cover the entire class of entailments between symbolic heaps and is complete for a strict subclass in which the existential quantifiers of the right-hand side of entailments can be instantiated by the free variables on the left. In this paper, we show that the problem is 2EXPTIME-hard, even if only entailment between predicate atoms is considered. The proof uses a reduction from the membership problem for alternating Turing machines [3] whose working tape is exponentially bounded in the size of the input. Proving a matching 2EXPTIME upper bound reduces essentially to fixing the issues of the algorithm proposed in [7] and is considered for future work.

³ Entailment does not reduce to satisfiability since the considered logic has no negation.

2 Separation Logic with Inductive Definitions

For a finite set S , we denote by $\|S\| \in \mathbb{N}$ its cardinality. For a partial mapping $f : A \rightarrow B$, let $\text{dom}(f) \stackrel{\text{def}}{=} \{x \in A \mid f(x) \in B\}$ and $\text{img}(f) \stackrel{\text{def}}{=} \{f(x) \mid x \in A\}$ be its domain and range, respectively, and we write $f : A \rightarrow_{\text{fin}} B$ if $\|\text{dom}(f)\| < \infty$. Given integers $n \leq m$, we denote by $\llbracket n \dots m \rrbracket$ the set $\{n, n+1, \dots, m\}$. For a tuple $t = (t_1, \dots, t_n)$ and $i \in \llbracket 1 \dots n \rrbracket$, we denote by t_i the element on the i -th position in t . By slight abuse of notation, we write $t \in t$ if $t = t_i$, for some $i \in \llbracket 1 \dots n \rrbracket$.

Let $\text{Var} = \{x, y, \dots\}$ be an infinite countable set of logical first-order variables and $\text{Pred} = \{p, q, \dots\}$ be an infinite countable set of uninterpreted relation symbols, called *predicates*. Each predicate p has an arity $\#p \geq 1$, denoting the number of arguments. In addition, we consider a special function symbol nil , of arity zero. A *term* is an element of the set $\text{Term} \stackrel{\text{def}}{=} \text{Var} \cup \{\text{nil}\}$. Let $k \geq 1$ be an integer constant fixed throughout this paper. The logic SL^k is the set of formulæ generated inductively as follows:

$$\phi := \text{emp} \mid t_0 \mapsto (t_1, \dots, t_k) \mid p(t_1, \dots, t_{\#p}) \mid t_1 \doteq t_2 \mid t_1 \not\doteq t_2 \mid \phi_1 * \phi_2 \mid \exists x. \phi_1$$

where $p \in \text{Pred}$, $t_1, t_2, \dots, t_{\#p} \in \text{Term}$ and $x \in \text{Var}$. We write \perp for $t \not\doteq t$, for any $t \in \text{Term}$.

A *ground formula* is a formula of SL^k in which no predicates occur. We write $\text{fv}(\phi)$ for the set of *free* variables, that occur in ϕ not within the scope of an existential quantifier. By writing $\phi(x_1, \dots, x_n)$ we mean $x_1, \dots, x_n \in \text{fv}(\phi)$ and $\phi[y_1/x_1, \dots, y_n/x_n]$ is the formula obtained from ϕ by simultaneously substituting each x_i with y_i , for $i \in \llbracket 1 \dots n \rrbracket$. A *substitution* is a mapping $\sigma : \text{Var} \rightarrow \text{Term}$ and we denote by $\phi\sigma$ the formula $\phi[\sigma(x_1)/x_1, \dots, \sigma(x_n)/x_n]$.

To interpret SL^k formulæ, we consider an infinite countable set Loc of *locations* and a designated location $\text{nil} \in \text{Loc}$. Let $\text{Loc}^i \stackrel{\text{def}}{=} \{(\ell_1, \dots, \ell_i) \mid \ell_1, \dots, \ell_i \in \text{Loc}\}$. The semantics of SL^k formulæ is defined in terms of *structures* (s, h, \mathfrak{I}) , where:

- $s : \text{Term} \rightarrow \text{Loc}$ is a total mapping of terms into locations, called *store*, such that always $s(\text{nil}) = \text{nil}$,
- $h : \text{Loc} \rightarrow_{\text{fin}} \text{Loc}^k$ is a finite partial mapping of locations into k -tuples of locations, called *heap*, such that $\text{nil} \notin \text{dom}(h)$. Let Heaps denotes the set of heaps, and
- $\mathfrak{I} : \text{Pred} \rightarrow \bigcup_{i \geq 1} 2^{\text{Loc}^i \times \text{Heaps}}$ is an *interpretation* associating each predicate p a set of pairs $\langle (\ell_1, \dots, \ell_{\#p}), h' \rangle \in \text{Loc}^{\#p} \times \text{Heaps}$.

Two heaps h_1 and h_2 are *disjoint* iff $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$, in which case their *disjoint union* is denoted $h_1 \uplus h_2$, undefined otherwise. The satisfaction relation \models between structures and SL^k formulæ is defined, as usual, recursively on the syntax of formulæ:

$$\begin{aligned} (s, h, \mathfrak{I}) \models t_1 \doteq t_2 & \Leftrightarrow h = \emptyset \text{ and } s(t_1) = s(t_2) \\ (s, h, \mathfrak{I}) \models t_1 \not\doteq t_2 & \Leftrightarrow h = \emptyset \text{ and } s(t_1) \neq s(t_2) \\ (s, h, \mathfrak{I}) \models \text{emp} & \Leftrightarrow h = \emptyset \\ (s, h, \mathfrak{I}) \models t_0 \mapsto (t_1, \dots, t_k) & \Leftrightarrow \text{dom}(h) = \{s(t_0)\} \text{ and } h(s(t_0)) = (s(t_1), \dots, s(t_k)) \\ (s, h, \mathfrak{I}) \models p(t_1, \dots, t_{\#p}) & \Leftrightarrow \langle (s(t_1), \dots, s(t_{\#p})), h \rangle \in \mathfrak{I}(p) \\ (s, h, \mathfrak{I}) \models \phi_1 * \phi_2 & \Leftrightarrow \text{there are disjoint heaps } h_1 \text{ and } h_2, \text{ such that } h = h_1 \uplus h_2 \\ & \text{and } (s, h_i, \mathfrak{I}) \models \phi_i, \text{ for each } i = 1, 2 \\ (s, h, \mathfrak{I}) \models \exists x. \phi & \Leftrightarrow (s[x \leftarrow \ell], h, \mathfrak{I}) \models \phi, \text{ for some } \ell \in \text{Loc} \end{aligned}$$

where $\mathfrak{s}[x \leftarrow \ell]$ is the store mapping x into ℓ and behaving like \mathfrak{s} for all $t \in \text{Term} \setminus \{x\}$.

Given formulæ ϕ and ψ , we write $\phi \stackrel{\text{ac}}{=} \psi$ if ψ is obtained from ϕ by a reordering of $*$ -connected subformulæ. Since $*$ is associative and commutative, any two such formulæ are equivalent. It is known that, for any formulæ ϕ and ψ , if $x \notin \text{fv}(\psi)$ then the formulæ $(\exists x . \phi) * \psi$ and $\exists x . \phi * \psi$ are equivalent [6]. Thus, each SL^k formula ϕ can be written as $\exists x_1 \dots \exists x_n . \psi * \bigstar_{j=1}^m p_j(u_1^j, \dots, u_{\#p_j}^j)$, where $\widehat{\phi} \stackrel{\text{def}}{=} \psi$ is a ground formula, called the *hat* of ϕ . Note that, moreover, $\widehat{\phi}$ is unique, up to the $\stackrel{\text{ac}}{=}$ relation.

In the rest of this paper, we consider that the predicates are interpreted by a set \mathcal{S} of rules $p(x_1, \dots, x_{\#p}) \Leftarrow \rho$, where ρ is an SL^k formula, such that $\text{fv}(\rho) \subseteq \{x_1, \dots, x_{\#p}\}$. We refer to $p(x_1, \dots, x_{\#p})$ as to the *head* and to ρ as to the *body* of the rule. A rule is a *base rule* if its body is a ground formula. W.l.o.g., we assume that any two heads with the predicate symbol p share the same tuple of variables \mathbf{x}_p (this can be achieved by variable renaming) and by writing $p(\mathbf{x}_p)$ we mean that \mathbf{x}_p is the unique tuple of variables associated with p . We write $p(\mathbf{x}_p) \Leftarrow_S \rho$ if the rule $p(\mathbf{x}_p) \Leftarrow \rho$ belongs to \mathcal{S} .

A set \mathcal{S} of rules defines a transformer function $\mathfrak{T}_{\mathcal{S}}$ on interpretations of predicates. For an interpretation \mathfrak{I} , the interpretation $\mathfrak{T}_{\mathcal{S}}(\mathfrak{I})$ is defined as:

$$\mathfrak{T}_{\mathcal{S}}(\mathfrak{I}) \stackrel{\text{def}}{=} \lambda p . \left\{ \langle \ell, \mathfrak{h} \rangle \mid ([\mathbf{x}_p \leftarrow \ell], \mathfrak{h}, \mathfrak{I}) \models \rho, p(\mathbf{x}_p) \Leftarrow_S \rho \right\}$$

where $\mathbf{x}_p = (x_1, \dots, x_{\#p})$, $\ell = (\ell_1, \dots, \ell_{\#p})$ and $[\mathbf{x}_p \leftarrow \ell]$ is any store mapping x_i into ℓ_i , for all $i \in \llbracket 1 \dots \#p \rrbracket$. It is not hard to prove that the set of interpretations forms a complete lattice, equipped with the pointwise inclusion as a partial order, union as join and intersection as meet. Moreover, the transformer function is monotonic, because no predicate symbol occurs under negation in a rule body. Consequently, $\mathfrak{T}_{\mathcal{S}}$ has a unique least fixed point $\mathfrak{T}_{\mathcal{S}}^\mu$. In the rest of this paper, given a set of rules \mathcal{S} , we write $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} \phi$ for $(\mathfrak{s}, \mathfrak{h}, \mathfrak{T}_{\mathcal{S}}^\mu) \models \phi$ and refer also to the pairs $(\mathfrak{s}, \mathfrak{h})$ as structures. We are now ready to define the class of entailment problems, which are the concern of this paper:

Definition 1 (Entailment Problem). *Given a set of rules \mathcal{S} and two SL^k formulæ ϕ and ψ is it the case that for every store \mathfrak{s} and every heap \mathfrak{h} such that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} \phi$, we have $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} \psi$? Instances of the entailment problem are denoted $\phi \models_{\mathcal{S}} \psi$.*

2.1 Unfolding Trees

In many cases it is useful to think of heaps from the least fixed point interpretation of a predicate symbol as obtained by a finite unfolding of the predicate, which produces a ground formula in which all the predicate symbols have been successively replaced with the bodies of their defining rules. In this section, let \mathcal{S} be a given set of rules.

Definition 2 (Unfolding). *A formula ψ is a 1-unfolding of a formula ϕ , written $\phi \Longrightarrow_{\mathcal{S}} \psi$, if ψ is obtained by substituting an occurrence of an atom $p(y_1, \dots, y_{\#p})$ in ϕ with $\rho[y_1/x_1, \dots, y_{\#p}/x_{\#p}]$, where $(p(x_1, \dots, x_{\#p}), \rho) \in \mathcal{S}$ is a rule. An unfolding of ϕ is a formula ψ such that $\phi \Longrightarrow_{\mathcal{S}}^* \psi$, where $\Longrightarrow_{\mathcal{S}}^*$ is the reflexive and transitive closure of $\Longrightarrow_{\mathcal{S}}$. Moreover, ψ is a ground unfolding of ϕ if ψ is a ground formula.*

It is often convenient to place the steps of an unfolding sequence in a tree that records the partial order in which these steps occur. We assume the set of nodes in a tree t , denoted as $\text{nodes}(t)$, to be a finite prefix-closed subset of \mathbb{N}^* , where \mathbb{N}^* is the set of finite sequences of positive integers, such that if w and wi are nodes of t , for some $i \in \mathbb{N} \setminus \{0\}$, then so are wj , for all $j \in \llbracket 0 \dots i-1 \rrbracket$. We write $|w|$ for the length of the sequence w and λ for the empty sequence ($|\lambda| = 0$). The label of a node $w \in \text{nodes}(t)$ is denoted by $t(w)$ and the subtree of t rooted at w by $t|_w$. The height of a tree is $\text{height}(t) \stackrel{\text{def}}{=} \max_{w \in \text{nodes}(t)} |w|$. A *traversal* of a tree t is an enumeration of $\text{nodes}(t)$ such that the index of a node is smaller than all indices of its successors.

Definition 3 (Unfolding Tree). Given a formula ϕ , an unfolding tree for ϕ is a tree u , labeled with pairs (p, ψ) , where $p \in \text{Pred} \cup \{\perp\}$ and ψ is a formula, such that:

- $t(\lambda) = (\perp, \phi)$, and
- if $u(w) = (p, \psi)$, then there exists a bijective mapping from the set of occurrences of atoms $q(t_1, \dots, t_{\#q})$ in ψ and the children of w , such that if $q(t_1, \dots, t_{\#q})$ is mapped to wi , then $u(wi) = (q(t_1, \dots, t_{\#q}), \rho[t_1/x_1, \dots, t_{\#q}/x_{\#q}])$, where $q(x_1, \dots, x_{\#q}) \Leftarrow_S \rho$.

We denote by $\mathcal{T}_S(\phi)$ the set of unfolding trees for ϕ .

Clearly, every traversal of an unfolding tree $u \in \mathcal{T}_S(\phi_0)$ corresponds to a ground unfolding sequence $\sigma : \phi_0 \Rightarrow_S \phi_1 \Rightarrow_S \dots \Rightarrow_S \psi$ where, for each node w_i of label $(q(t), \psi)$ in σ , ϕ_i is obtained from ϕ_{i-1} by replacing $q(t)$ by ψ . It is easy to check that the outcome ψ only depends on u , not on σ . In other words, all traversals of u yield the same ground formula, denoted by $\mathcal{F}(u)$. By convention, we let $\mathcal{F}(u) \stackrel{\text{def}}{=} \phi$ if $\text{nodes}(u) = \{\lambda\}$. The following lemma provides an equivalent definition for the semantics of SL^k formulae:

Lemma 1. Given a SL^k formula ϕ , for any store s and any heap h , we have $(s, h) \models_S \phi$ if and only if $(s, h) \models \mathcal{F}(u)$, for some unfolding tree $u \in \mathcal{T}_S(\phi)$.

Proof: It is sufficient to give the proof only in the case $\phi = p(t_1, \dots, t_{\#p})$, for a predicate symbol $p \in \text{Pred}$ and terms $t_1, \dots, t_{\#p} \in \text{Term}$, the generalization to the case where ϕ is any SL^k formula being immediate. We start by proving the following fact:

Fact 1 Let $u \in \mathcal{T}_S(p(t_1, \dots, t_{\#p}))$ be an unfolding tree, such that (necessarily) $u(0) = (p(t_1, \dots, t_{\#p}), \rho\sigma)$, for some rule:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \underbrace{\exists y_1 \dots \exists y_n \cdot \hat{\rho} * \bigstar_{j=1}^m p_j(t_1^j, \dots, t_{\#p_j}^j)}_{\rho}$$

and some substitution σ such that $\sigma(x_k) = t_k$, for all $k \in \llbracket 1 \dots \#p \rrbracket$, then $\mathcal{F}(u) \stackrel{\text{ac}}{=} \exists y_1 \dots \exists y_n \cdot \hat{\rho}\sigma * \bigstar_{j=1}^m \mathcal{F}(u_j)$, where $u_j \in \mathcal{T}_S(p_j(t_1^j, \dots, t_{\#p_j}^j)\sigma)$, for all $j \in \llbracket 1 \dots m \rrbracket$.

Proof: Let $s : \phi_0 = p(t_1, \dots, t_{\#p}) \Rightarrow_S \phi_1 = \rho\sigma \Rightarrow_S \dots \Rightarrow_S \phi_r = \mathcal{F}(u)$ be a ground unfolding sequence. Then there exists $\ell \in \llbracket 1 \dots m \rrbracket$ and a position $j \in \llbracket 1 \dots r \rrbracket$ in this sequence such that the sequence $\phi_1 \Rightarrow_S \dots \Rightarrow_S \phi_j$ traverses all nodes from $u|_\ell$. W.l.o.g., let ℓ and j be the least such numbers, respectively. We re-organize the unfolding sequence s into s^1 by moving all steps between 1 and j , corresponding to visits of nodes

from $u_{|\ell'}$, for $\ell' \in \llbracket 1 \dots m \rrbracket \setminus \{\ell\}$, after j and let $k \leq j$ be the new position of ϕ_j in s^1 . Clearly this re-organization does not change the outcome ϕ_r , who is the same for s and s_1 . Because all nodes from $u_{|j}$ are visited by $s_1 : \phi_0^1 = \phi_0 \Rightarrow_S \phi_1^1 = \phi_1 \Rightarrow_S \dots \Rightarrow_S \phi_k^1 \Rightarrow_S \dots \Rightarrow_S \phi_r^1 = \phi_r$, we obtain:

$$\phi_k^1 \stackrel{\text{ac}}{=} \exists y_1 \dots \exists y_n \cdot \hat{\rho}\sigma * \underbrace{*}_{j \in \llbracket 1 \dots m \rrbracket \setminus \ell} p_j(t_1^j, \dots, t_{\#p_j}^j) * \mathcal{F}(u_\ell)$$

The unfolding sequence $\phi_k^1 \Rightarrow_S \dots \Rightarrow_S \phi_r$ is strictly smaller than s and we apply the same argument inductively until the ground formula $\mathcal{F}(u)$ is derived. \square

Let $\mathbf{t} \stackrel{\text{def}}{=} (t_1, \dots, t_{\#p})$ and $\mathbf{s}(\mathbf{t}) \stackrel{\text{def}}{=} (\mathbf{s}(t_1), \dots, \mathbf{s}(t_{\#p}))$. Let $\mathfrak{T}_S^0 \stackrel{\text{def}}{=} \lambda p. \emptyset$ be the interpretation mapping each predicate symbol into the empty set. We have $(\mathbf{s}, \mathbf{h}) \models_S p(\mathbf{t})$ if and only if $(\mathbf{s}(\mathbf{t}), \mathbf{h}) \in \mathfrak{T}_S^\mu(p) = (\bigcup_{i=0}^\infty \mathfrak{T}_S^i)(p)$. The last equality uses the fact that \mathfrak{T}_S is monotone and continuous which, by Kleene's Fixpoint Theorem, guarantees that its least fixpoint is the limit of the increasing sequence of approximants $\{\mathfrak{T}_S^i\}_{i \geq 0}$.

For any predicate symbol $p \in \text{Pred}$ and any tuple of terms $\mathbf{t} = (t_1, \dots, t_{\#p})$, we prove the following equivalence by induction on $i \geq 0$:

$$\langle \mathbf{s}(\mathbf{t}), \mathbf{h} \rangle \in \mathfrak{T}_S^i(p) \Leftrightarrow (\mathbf{s}, \mathbf{h}) \models \mathcal{F}(u), \text{ for some } u \in \mathcal{T}_S(p(\mathbf{t})), \text{ such that } \text{height}(u) \leq i$$

In the base case $i = 0$, the equivalence holds because both $\langle \mathbf{s}(\mathbf{t}), \mathbf{h} \rangle \in \mathfrak{T}_S^0(p)$ is false and $\mathcal{T}_S(p(\mathbf{t}))$ contains no unfolding tree of height 0. For the induction step $i \geq 1$, we consider two directions:

" \Rightarrow " If $\langle \mathbf{s}(\mathbf{t}), \mathbf{h} \rangle \in \mathfrak{T}_S^i(p)$ then $(\mathbf{s}, \mathbf{h}, \mathfrak{T}_S^{i-1}) \models \rho\sigma$, for some rule:

$$p(x_1, \dots, x_{\#p}) \Leftarrow_S \underbrace{\exists y_1 \dots \exists y_n \cdot \hat{\rho} * \underbrace{*}_{j=1}^m p_j(t_1^j, \dots, t_{\#p_j}^j)}_{\rho}$$

and substitution σ such that $\sigma(x_k) = t_k$, for all $k \in \llbracket 1 \dots \#p \rrbracket$. Then there exists locations $\ell_1, \dots, \ell_n \in \text{Loc}$ and heaps $\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_m$ such that $\mathbf{h} = \biguplus_{j=0}^m \mathbf{h}_j$ and the following hold:

- $(\mathbf{s}', \mathbf{h}_0) \models \hat{\rho}\sigma$, and
- $\langle \mathbf{s}'(\mathbf{t}_j\sigma), \mathbf{h}_j \rangle \in \mathfrak{T}_S^{i-1}(p_j)$, for all $j \in \llbracket 1 \dots m \rrbracket$,

where $\mathbf{s}' \stackrel{\text{def}}{=} \mathbf{s}[y_1 \leftarrow \ell_1, \dots, y_n \leftarrow \ell_n]$ and $\mathbf{t}_j \stackrel{\text{def}}{=} (t_1^j, \dots, t_{\#p_j}^j)$, for all $j \in \llbracket 1 \dots m \rrbracket$. By the induction hypothesis, we obtain unfolding trees $u_j \in \mathcal{T}_S(p_j(\mathbf{t}_j\sigma))$, of height at most $i-1$, such that $(\mathbf{s}', \mathbf{h}_j) \models \mathcal{F}(u_j)$, for all $j \in \llbracket 1 \dots m \rrbracket$. Note that, by Definition 3, the root of each u_j has a single child and let $u_{j|0}$ be the subtree of u_j rooted at the child of λ . Then we define an unfolding tree u , such that $u(\lambda) \stackrel{\text{def}}{=} (\perp, p(\mathbf{t}))$, $u(0) \stackrel{\text{def}}{=} (p, \rho\sigma)$ and $u_{|j} \stackrel{\text{def}}{=} u_{j|0}$, for all $j \in \llbracket 1 \dots m \rrbracket$. It is straightforward to check that $u \in \mathcal{T}_S(p(\mathbf{t}))$ and that the height of u is at most i . We are left with proving that $(\mathbf{s}, \mathbf{h}) \models \mathcal{F}(u)$, which follows from Fact 1 and the inductive hypothesis $(\mathbf{s}', \mathbf{h}_j) \models \mathcal{F}(u_j)$, for all $j \in \llbracket 1 \dots m \rrbracket$.

" \Leftarrow " If $(\mathbf{s}, \mathbf{h}) \models \mathcal{F}(u)$ then, by Definition 3, the root of u has a single child 0 whose label is $u(0) = (p(\mathbf{t}), \rho\sigma)$, for a rule:

$$p(x_1, \dots, x_{\#p}) \Leftarrow_S \underbrace{\exists y_1 \dots \exists y_n \cdot \hat{\rho} * \underbrace{*}_{j=1}^m p_j(t_1^j, \dots, t_{\#p_j}^j)}_{\rho}$$

and a substitution $\sigma(x_k) = t_k$, for all $k \in \llbracket 1 \dots \#p \rrbracket$. By Fact 1, we obtain:

$$\mathcal{F}(u) \stackrel{\text{ac}}{=} \exists y_1 \dots \exists y_n . \hat{\rho}\sigma * \bigstar_{j=1}^m \mathcal{F}(u_j)$$

for some unfolding trees $u_j \in \mathcal{T}_S(p_j(t_1^j, \dots, t_{\#p_j}^j))$, of height at most $i-1$, for $j \in \llbracket 1 \dots m \rrbracket$. Since $(s, h) \models \mathcal{F}(u)$, there exists locations $\ell_1, \dots, \ell_n \in \text{Loc}$, such that:

- $(s', h_0) \models \hat{\rho}\sigma$ and
- $(s', h_j) \models \mathcal{F}(u_j)$, for all $j \in \llbracket 1 \dots m \rrbracket$,

where $s' \stackrel{\text{def}}{=} s[y_1 \leftarrow \ell_1, \dots, y_n \leftarrow \ell_n]$ and $h = \biguplus_{j=0}^m h_j$. By the inductive hypothesis, we obtain:

$$\langle s'(t_j\sigma), h_j \rangle \in \mathfrak{T}_S^{i-1}(p_j) \subseteq \mathfrak{T}_S^\mu(p_j)$$

where $t_j \stackrel{\text{def}}{=} (t_1^j, \dots, t_{\#p}^j)$, for all $j \in \llbracket 1 \dots m \rrbracket$. But then $(s', h_j) \models_S p_j(t_j\sigma)$, for all $j \in \llbracket 1 \dots m \rrbracket$, thus $(s, h) \models_S \rho\sigma$ and $(s, h) \models_S p(t)$ follows. \square

3 A Decidable Class of Entailments

In general, the entailment problem is undecidable, proofs can be found in [5, 1]. Thus we consider a subclass of entailments for which decidability (with elementary recursive complexity) was proved in [4] and provide a 2EXPTIME lower bound for this problem. The decidable class is given by three restrictions of the systems used for the interpretation of predicates, namely *progress*, *connectivity* and *establishment*, defined in the rest of this section.

First, the *progress* condition requires that each rule adds to the heap exactly one location, namely the one associated to the first parameter of the head. Second, the *connectivity* condition requires that all locations added during an unfolding of a predicate atom $p(t)$ form a connected tree-like structure.

Definition 4 (Progress & Connectivity). A set of rules S is progressing if and only if, the body ρ of each rule $p(x_1, \dots, x_{\#p}) \Leftarrow_S \rho$ is of the form $\exists z_1 \dots \exists z_m . x_1 \mapsto (y_1, \dots, y_k) * \psi$ and ψ contains no occurrence of $a \mapsto \text{atom}$. If, moreover, each occurrence of a predicate atom in ψ is of the form $q(y_i, u_1, \dots, u_{\#q-1})$, for some $i \in \llbracket 1 \dots k \rrbracket$, then S is connected.

For upcoming developments, we make the connection between heaps and unfolding trees more precise, in the context of progressing connected sets of rules.

Lemma 2. Assume that S is a progressing and connected set of rules and that (s, h) is a structure such that $(s, h) \models_S p(t_1, \dots, t_{\#p})$, for some predicate symbol $p \in \text{Pred}$ and terms $t_1, \dots, t_{\#p} \in \text{Term}$. Then there exists an unfolding tree $u \in \mathcal{T}_S(p(t_1, \dots, t_{\#p}))$ and a bijection $\Lambda : \text{dom}(h) \rightarrow \text{nodes}(u) \setminus \{\lambda\}$ such that $(s, h) \models \mathcal{F}(u)$ and for each node $wi \in \text{nodes}(u) \setminus \{\lambda\}$, where $i \in \mathbb{N}$, we have $\Lambda^{-1}(wi) \in h(\Lambda^{-1}(w))$.

Proof: If $(s, h) \models_S p(t_1, \dots, t_{\#p})$, by Lemma 1, there exists $u \in \mathcal{T}_S(p(t_1, \dots, t_{\#p}))$, such that $(s, h) \models \mathcal{F}(u)$. Note that, by Definition 3, the root of u is labeled by $u(\lambda) = (\perp, p(t_1, \dots, t_{\#p}))$

and has a single child, namely 0. The bijection Λ is built inductively on the structure of the subtree $u|_0$ of u , rooted at 0, taking into account Definition 4. \square

Given a structure (s, h) such that $(s, h) \models_S p(t_1, \dots, t_{\#p})$, for any location $\ell \in \text{dom}(h)$, necessarily ℓ must be allocated by a 1-unfolding of some atom $q(\mathbf{t})$, which corresponds to a node in an unfolding tree $\Lambda(\ell) \in \mathcal{T}_S(p(t_1, \dots, t_{\#p}))$. In the following, we shall denote by $\text{Pr}_{s, h, \Lambda}(\ell)$ the predicate symbol q . Note that the existence of the unfolding tree u and bijection $\Lambda : \text{dom}(h) \rightarrow \text{nodes}(u) \setminus \{\lambda\}$ are guaranteed by Lemmas 1 and 2, respectively. The notation $\text{Pr}(\ell)$ stands for $\text{Pr}_{s, h, \Lambda}(\ell)$, for some Λ as in Lemma 2, where the structure (s, h) is clear from the context.

The third condition requires that all the existentially quantified variables introduced during a can only be associated with locations from the heap, in every ground unfolding of a formula. To formalize this condition, for a ground formula ϕ , we define the *partial satisfaction* relation as $(s, h) \models \phi$ iff $h = h_1 \uplus h_2$ and $(s, h_1) \models \phi$ and partial entailment as $\phi \models \psi$ iff $(s, h) \models \psi$ for each structure (s, h) , such that $(s, h) \models \phi$. If ϕ is ground, its set of allocated variables is $\text{alloc}(\phi) \stackrel{\text{def}}{=} \{x \in \text{fv}(\phi) \mid x' \mapsto (y_1, \dots, y_k) \text{ occurs in } \phi \text{ and } \phi \models x \dot{=} x'\}$. Note that $\text{alloc}(\phi) = \text{fv}(\phi)$ if ϕ is unsatisfiable. Extending this notion to formulae that are not necessarily ground, we define $\text{alloc}_S(\phi) \stackrel{\text{def}}{=} \bigcap_{u \in \mathcal{T}_S(\phi)} \text{alloc}(\mathcal{F}(u))$, i.e. a variable is allocated in ϕ iff it is allocated in every ground unfolding of ϕ w.r.t. a set of rules S .

Definition 5 (Establishment). *A set of rules S is established if and only if, for each rule $(p(x_p), \exists z_1 \dots \exists z_m . \psi) \in S$, where ψ is quantifier-free, we have $z_1, \dots, z_m \in \text{alloc}_S(\psi)$.*

In the following, we consider only sets of rules that are progressing, connected and established (PCE). The interest for PCE sets of rules is motivated by the following decidability result, proved in [4]:

Theorem 1 (Decidability). *Given a PCE set of rules S and two formulae ϕ and ψ such that the free variables of ψ are also free in ϕ , the problem $\phi \models_S \psi$ belongs to ELEMENTARY.*

The rest of this paper is concerned with proving that the entailment problem $\phi \models_S \psi$, for PCE sets of rules S , is 2EXPTIME-hard. Previously, a EXPTIME-hard bound for this problem was established in [5].

3.1 Syntactic Shorthands

Before proving 2EXPTIME-hardness by reduction from the membership problem of alternating Turing machines running in exponential space, we define several syntactical shorthands that simplify the presentation.

Given a term t and $n \geq 0$, we denote by t^n the tuple consisting of n occurrences of t . In the rest of the paper we silently assume that any heap entry of the form $h(\ell) = (\ell_1, \dots, \ell_k)$ is represented by a binary heap $h_2 : \text{Loc} \rightarrow_{\text{fin}} \text{Loc}^2$ such that $\text{dom}(h_2) = \{\ell\} \cup \{\ell_0^i \mid i \in \llbracket 1 \dots k-1 \rrbracket\}$, $h_2(\ell) = (\ell_0^1, \ell_1)$ and $h_2(\ell_0^i) = (\ell_0^{i+1}, \ell_{i+1})$, for all $i \in \llbracket 1 \dots k-1 \rrbracket$, with $\ell_0^k = \text{nil}$. This allows one to encode records of non constant length (for instance tuples of length \mathfrak{N}) by using only a constant number of record fields (i.e., $k = 2$).

Deep Hats For a tuple $\mathbf{t} = (t_1, \dots, t_j) \in \text{Term}^j$, $j \geq 1$ and $n \geq 0$, by writing:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \exists y_1 \dots \exists y_r . x_1 \mapsto (u_1, \dots, u_{i-1}, [\mathbf{t}]^n, u_{i+1}, \dots, u_m) * \psi$$

we denote the rules:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \exists y_1 \dots \exists y_r . x_1 \mapsto (u_1, \dots, u_{i-1}, t_1, \dots, t_j, u_{i+1}, \dots, u_m) * \psi, \text{ if } n = 0$$

and, recursively, for all $n > 0$:

$$\begin{aligned} p(x_1, \dots, x_{\#p}) &\Leftarrow \exists y_1 \dots \exists y_r \exists z . x_1 \mapsto (u_1, \dots, u_{i-1}, z, u_{i+1}, \dots, u_m) * \widetilde{p}_{n-1}(z, x_1, \dots, x_{\#p}, y_1, \dots, y_r) * \psi' \\ \widetilde{p}_\ell(z_1, \dots, z_{\#p+r+1}) &\Leftarrow \exists z'_1 . z_1 \mapsto (z'_1) * \widetilde{p}_{\ell-1}(z'_1, z_2, \dots, z_{\#p+r+1}), \text{ for } \ell \in \llbracket 1 \dots n-1 \rrbracket \\ \widetilde{p}_0(z_1, \dots, z_{\#p+r+1}) &\Leftarrow z_1 \mapsto (t_1, \dots, t_j) * \psi'' \end{aligned}$$

where \widetilde{p}_ℓ , for $\ell \in \llbracket 0 \dots n-1 \rrbracket$ are fresh predicate symbols of arity $\# \widetilde{p}_\ell = \#p + r + 1$ and ψ'' (resp. ψ') is the separating conjunction of all atoms of ψ whose first argument is in \mathbf{t} (resp. is not in \mathbf{t}). Clearly, $\psi' * \psi'' \stackrel{\text{ac}}{=} \psi$, and the obtained rules are connected.

Special Variables We assume the existence of the following special variables that occur free in each formula: $\mathbf{0}, \mathbf{1}, \gamma_1, \dots, \gamma_N$, for some fixed constant $N \geq 2$, used throughout the paper. These variables will always be allocated and denote separate locations, as required by the following rule:

$$\text{Const}(x) \Leftarrow x \mapsto (\mathbf{0}, \mathbf{1}, \gamma_1, \dots, \gamma_N) * \mathbf{0} \mapsto (\text{nil}) * \mathbf{1} \mapsto (\text{nil}) * *_{i=1}^N \gamma_i \mapsto (\text{nil}) \quad (1)$$

Considering special variables is without loss of generality in the following, because these variables can be added to the parameter list of each head in the system, at the expense of cluttering the presentation.

Binary Choices The symbol \bullet occurring in the body of a rule ranges over the special variables $\mathbf{0}$ and $\mathbf{1}$. Thus any rule of the form:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \exists z_1 \dots \exists z_n . x_1 \mapsto (y_1, \dots, y_{i-1}, \bullet, y_{i+1}, \dots, y_m) * \psi$$

stands for the following two rules:

$$\begin{aligned} p(x_1, \dots, x_{\#p}) &\Leftarrow \exists z_1 \dots \exists z_n . x_1 \mapsto (y_1, \dots, y_{i-1}, \mathbf{0}, y_{i+1}, \dots, y_m) * \psi \\ p(x_1, \dots, x_{\#p}) &\Leftarrow \exists z_1 \dots \exists z_n . x_1 \mapsto (y_1, \dots, y_{i-1}, \mathbf{1}, y_{i+1}, \dots, y_m) * \psi \end{aligned}$$

Note that eliminating the occurrences of \bullet will increase the number of rules in \mathcal{S} by a constant at most 2^k , because k is assumed to be constant (i.e. independent of the input of a decision procedure).

Binary Variables A binary variable b is understood as ranging over the domain of the interpretation of $\mathbf{0}$ and $\mathbf{1}$, namely the locations assigned to $\mathbf{0}$ and $\mathbf{1}$ by the formula Const (1). Additionally, for each binary variable b , we consider the associated variable

\bar{b} , intended to denote the complement of b . More precisely, the formula $\exists b . \psi$ is understood as $\psi[0/b, 1/\bar{b}] \vee \psi[1/b, 0/\bar{b}]$. However, this direct substitution of the (existentially quantified) binary variables by **0** and **1** within the rules of an established system would break the establishment condition (Definition 5), because **0** and **1** are not necessarily allocated within the body of the rule⁴. This problem can be overcome by passing **0** and **1** as parameters to a fresh predicate. More precisely, a rule of the form:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \exists b_1 \dots \exists b_i \exists y_1 \dots \exists y_n . x_1 \mapsto ([t]^m) * \psi \quad (2)$$

where $1 \leq i \leq m$, is a shorthand for the following set of rules:

$$\begin{aligned} p(x_1, \dots, x_{\#p}) &\Leftarrow \exists y . x_1 \mapsto (y) * p'(y, x_1, \dots, x_{\#p}, \mathbf{0}, \mathbf{1}) \\ p(x_1, \dots, x_{\#p}) &\Leftarrow \exists y . x_1 \mapsto (y) * p'(y, x_1, \dots, x_{\#p}, \mathbf{1}, \mathbf{0}) \\ p'(y, x_1, \dots, x_{\#p}, b_1, \bar{b}_1) &\Leftarrow \exists b_2 \dots \exists b_i \exists y_1 \dots \exists y_n . y \mapsto ([t]^{m-1}) * \psi \end{aligned}$$

Clearly, the elimination of the binary existential quantifiers from the rule (2) will add $2i$ rules to the set. Note that the hat $[t]^m$, of height $m \geq i$ decreases at each step of the elimination. This guarantees that rules resulting from the elimination of $\exists b_1, \dots, \exists b_i$ respectively, are progressing (Definition 4).

Next, for a vector $\mathbf{b} = (b_1, \dots, b_n)$ of binary variables, we denote by $\bar{\mathbf{b}}$ the vector $(\bar{b}_1, \dots, \bar{b}_n)$. The following rule:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \exists y_1 \dots \exists y_m . x_1 \mapsto t * \psi \mid (b_1, \dots, b_n) \neq (\bar{c}_1, \dots, \bar{c}_n) \quad (3)$$

where $b_1, \dots, b_n \in \{x_2, \dots, x_{\#p}, y_1, \dots, y_m\}$ occur only once in t and do not occur in ψ and $c_1, \dots, c_n \in \{x_2, \dots, x_{\#p}, y_1, \dots, y_m\}$, is a shorthand for the following set of rules:

$$p(x_1, \dots, x_{\#p}) \Leftarrow \exists y_1 \dots \exists y_m . x_1 \mapsto (t[c_i/b_i])[\bullet / b_j]_{j \in \llbracket 1..n \rrbracket \setminus \{i\}} * \psi, i \in \llbracket 1..n \rrbracket \quad (4)$$

Intuitively, the rule (3) introduces new binary variables b_1, \dots, b_n , such that not all of them are equal to the complement of c_1, \dots, c_n , respectively. In other words, at least one b_i must be equal to c_i , for some $i \in \llbracket 1..n \rrbracket$. Note that expanding the rule (3) as described above results in at most n rules of the form (2), hence the full elimination of binary variables from the system is possible in polynomial time.

4 Alternating Turing Machines

An *Alternating Turing Machine* (ATM) is a tuple $M = (Q, \Gamma, \delta, q_0, g)$ where:

- Q is a finite set of control *states*,
- $\Gamma = \{\gamma_1, \dots, \gamma_N, \mathbf{b}\}$ is a finite alphabet, \mathbf{b} is the *blank* symbol and each symbol from $\Gamma \setminus \{\mathbf{b}\}$ is named after a special variable,
- $\delta \subseteq Q \times \Gamma \times Q \times (\Gamma \setminus \{\mathbf{b}\}) \times \{\leftarrow, \rightarrow\}$ is the *transition relation*, $(q, a, q', b, \mu) \in \delta$ meaning that, in state q , upon reading symbol a , the machine moves to state q' , writes $b \neq \mathbf{b}$ to the tape⁵ and moves the head by one to the left (resp. right) if $\mu = \leftarrow$ (resp. $\mu = \rightarrow$),

⁴ In fact they are allocated by the side condition Const.

⁵ A machine never writes blank symbols, that are used only for the initially empty tape cells.

- $q_0 \in Q$ is the *initial state*, and
- $g : Q \rightarrow \{\vee, \wedge\}$ partitions the set of states into *existential* ($g(q) = \vee$) and *universal* ($g(q) = \wedge$) states.

A *configuration* of M is a tuple (q, w, i) where $q \in Q$ is the current state, $w : \mathbb{N} \rightarrow \Gamma$ is (the contents of) the *tape*, such that $|\{j \in \mathbb{N} \mid w(j) \neq b\}| < \infty$ and $i \in [0 \dots \max\{j \in \mathbb{N} \mid w(j) \neq b\} + 1]$ is the current position of the head on the tape. We denote by ϵ the empty word over Γ . For any tape w and integer i , we denote by $w[i \leftarrow a]$ the tape w' such that $w'(i) = a$ and $w'(j) = w(j)$ for all $j \neq i$. In the following, we write $w_i \stackrel{\text{def}}{=} w(i)$, $i^{\leftarrow} \stackrel{\text{def}}{=} i - 1$ and $i^{\rightarrow} \stackrel{\text{def}}{=} i + 1$.

The step relation of M is the following relation between configurations: we write $(q, w, i) \xrightarrow{(q, a, q', b, \mu)} (q', w', j)$ if and only if there exists a transition $(q, a, q', b, \mu) \in \delta$ such that $w_i = a$, $w' = w[i \leftarrow b]$ and $j = i^\mu$. We write $(q, w, i) \rightarrow (q', w', j)$ when the applied transition is not important. An *execution* is a sequence $(q_0, w_0, 0) \xrightarrow{(q_0, a_0, q_1, b_0, \mu_0)} (q_1, w_1, i_1) \xrightarrow{(q_1, a_1, q_2, b_1, \mu_1)} \dots$. Note that an execution is entirely determined by the initial configuration $(q_0, w_0, 0)$ and the sequence $(q_0, a_0, q_1, b_0, \mu_0), (q_1, a_1, q_2, b_1, \mu_1), \dots$ of transition rules applied to it.

Given a function $f : \mathbb{N} \rightarrow \mathbb{N}$, an execution is *f-space bounded* if and only if there exists a constant $c > 0$ such that $|w_i| \leq c \cdot f(|w_0|)$, for all $i > 0$. The ATM M is *exponential-space bounded* if it admits only *f-space bounded* executions, where $f(x) = 2^{g(x)}$ and g is a univariate polynomial function.

Definition 6. A derivation of an ATM $M = (Q, \Gamma, \delta, q_0, g)$, starting from a configuration $(q_0, w_0, 0)$, is a tree t , whose nodes are either:

1. branching nodes labeled with configurations $(q, w, i) \in Q \times \Gamma^* \times \mathbb{N}$, or
 2. action nodes labeled with tuples $(a, b, \mu) \in \Gamma \times \Gamma \setminus \{b\} \times \{\leftarrow, \rightarrow\}$, where a is the symbol read, b is the symbol written and μ is the move of the head at that step,
- such that the root of t is a branching node, $t(\lambda) = (q_0, w_0, 0)$ and, moreover:
- a. each branching node labeled by (q, w, i) , such that $g(q) = \vee$, has exactly one successor, that is an action node labeled by (a, b, μ) , where $(q, a, q', b, \mu) \in \delta$, whose successor is a branching node labeled by (q', w', j) such that $(q, w, i) \xrightarrow{(q, a, q', b, \mu)} (q', w', j)$,
 - b. each branching node labeled by (q, w, i) , such that $g(q) = \wedge$ has exactly one successor for each tuple $(q, a, q', b, \mu) \in \delta$ such that $a = w(i)$, and the successor associated with a transition (q, a, q', b, μ) is a branching node, labeled by (q', w', j) with $(q, w, i) \xrightarrow{(q, a, q', b, \mu)} (q', w', j)$.

We say that M accepts (q_0, w_0, i_0) iff it has a derivation starting from (q_0, w_0, i_0) .

Definition 7. The membership problem (M, w) asks the following: given an ATM $M = (Q, \Gamma, \delta, q_0, g)$ and an input word $w \in (\Gamma \setminus \{b\})^*$ does M accept $(q_0, w, 0)$?

The complexity class **AEXPSPACE** is the class of membership problems where M is exponential-space bounded. It is known that $\text{AEXPSPACE} = \text{co-AEXPSPACE} = \text{2EXPTIME}$ [3], where **co-AEXPSPACE** is the complement class of **AEXPSPACE**⁶. In the following, we shall w.l.o.g. consider only the membership problem (M, ϵ) . Indeed, let (M, w)

⁶ Every ATM can be complemented in linear time, by interchanging the existential with the universal states, thus all alternating classes are closed under complement.

be any instance of the membership problem, and let c and g be the constant and polynomial function witnessing the fact that M is exponential-space bounded. Let M_w be the ATM that produces w starting from input ϵ . Clearly, M_w uses at most $|w|$ working space, thus the machine $M_w; M$ which runs M_w on the empty word and then continues with M runs in space $c \cdot 2^{g(|w|)}$ and accepts $(q_0, \epsilon, 0)$ if and only if M accepts $(q_0, w, 0)$. If $\mathfrak{N} \geq \log_2(c) + g(w)$, then $M_w; M$ runs in space $2^{\mathfrak{N}}$. Therefore, we assume from now on that $M = (Q, \Gamma, \delta, q_0, g)$ is an ATM started in the configuration $(q_0, \epsilon, 0)$ and that M runs in space at most $2^{\mathfrak{N}}$ on $(q_0, \epsilon, 0)$, where \mathfrak{N} is polynomial w.r.t. the length of w .

4.1 Pseudo-derivations as Heaps

Given an ATM $M = (Q, \Gamma, \delta, q_0, g)$, we consider its *pseudo-derivations*, which are the trees of Definition 6 relaxed so that any sequence of three consecutive branching-action-branching node labels (q, w, i) , (a, b, μ) and (q', w', i') only needs to ensure the existence of a transition $(q, a, q', b, \mu) \in \delta$, i.e. we drop the requirements $w_i = a$, $w' = w[i \leftarrow b]$ and $i' = i^\mu$ from the condition $(q, w, i) \xrightarrow{(q, a, q', b, \mu)} (q', w', i')$ in Definition 6. Clearly, the leaves of a pseudo-derivation are all labeled by universal states.

We represent the pseudo-derivations of M as tree-shaped heaps generated by a set of rules where, intuitively, each predicate $q(x)$ allocates a branching node labeled by a configuration (q, w, i) , for some $w \in \Gamma^*$ and some $i \in \llbracket 0 \dots 2^{\mathfrak{N}} - 1 \rrbracket$, and each predicate $\bar{q}(x, a, b, \mu)$ allocates an action node labeled (a, b, μ) . Importantly, since M starts on the empty word ϵ , the tape contents in a branching node can be derived from the sequence of actions along the path from the root to that node. For this reason, we shall not explicitly represent tape contents within the configurations and label branching nodes with pairs $(q, i) \in Q \times \llbracket 0 \dots 2^{\mathfrak{N}} - 1 \rrbracket$.

We represent each position $i \in \llbracket 0 \dots 2^{\mathfrak{N}} - 1 \rrbracket$ on the tape succinctly, by an \mathfrak{N} -tuple of binary digits $\text{bin}(i) \in \{0, 1\}^{\mathfrak{N}}$ and encode the left and right moves as $\leftarrow \stackrel{\text{def}}{=} \mathbf{0}$ and $\rightarrow \stackrel{\text{def}}{=} \mathbf{1}$. Let $\tau(q, a) \stackrel{\text{def}}{=} \delta \cap (\{q\} \times \{a\} \times Q \times \Gamma \setminus \{b\} \times \{\leftarrow, \rightarrow\})$ be the set of transitions of M with source state q , reading symbol a from the tape.

As previously mentioned, we simplify the presentation by considering atoms of the form $x \mapsto (y_1, \dots, y_m)$ for an arbitrary $m \geq 1$, with the understanding that the heap corresponding to this atom is uniquely encoded by a binary heap. This is without loss of generality, because a rule containing such atoms can be transformed into a finite set of rules containing only atoms of the form $x \mapsto (y_1, y_2)$, resulting in a progressing, connected and established set of rules. With this in mind, we consider the following rules, for each state $q \in Q$:

$$q(x) \Leftarrow \exists x' . x \mapsto (\bullet^{\mathfrak{N}}, x') * \bar{q}'(x', a, b, \mu) \quad \text{if } g(q) = \vee \text{ and } (q, a, q', b, \mu) \in \tau(q, a) \quad (5)$$

$$q(x) \Leftarrow \exists y_1 \dots \exists y_n . x \mapsto (\bullet^{\mathfrak{N}}, y_1, \dots, y_m) * \bigstar_{j=1}^m \bar{q}_j(y_j, a, b_j, \mu_j) \quad \text{if } g(q) = \wedge \text{ and } \tau(q, a) = \{(q, a, q_1, b_1, \mu_1), \dots, (q, a, q_m, b_m, \mu_m)\} \quad (6)$$

$$\bar{q}(x, y, z, u) \Leftarrow \exists x' . x \mapsto (y, z, u, x') * q(x') \quad (7)$$

The heaps defined by the above rules ensure only that the control structure of a derivation of M is respected, namely that the branching and action nodes alternate correctly,

and that the sequence of control states labeling the branching nodes on any path is consistent with the transition relation of M . In other words, these trees encode pseudo-derivations of M . Further, we introduce a top-level predicate $p_M(x)$ that allocates the special variables $\mathbf{0}, \mathbf{1}, \gamma_1, \dots, \gamma_N$ and ensures that the initial state q_0 of M is the first control state that occurs on a path of a pseudo-derivation:

$$p_M(x) \Leftarrow \exists y_0 \exists z_0 . x \mapsto ([y_0]^{\mathfrak{N}}, z_0) * q_0(y_0) * \text{Const}(z_0) \quad (8)$$

Note that the hat $[y_0]^{\mathfrak{N}}$ above ensures that every heap in the least fixed point interpretation of p begins with a tree of height \mathfrak{N} . The use of this tree will be made clear below. For now, let \mathcal{S} be the set consisting of the rules above. We formalize the encoding of a pseudo-derivation by a structure:

Definition 8. A structure $(\mathfrak{s}, \mathfrak{h})$ such that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} p_M(x)$ encodes a pseudo-derivation t of M , written as $(\mathfrak{s}, \mathfrak{h}) \triangleright t$, if and only if there exists a injective mapping $f : \text{nodes}(t) \rightarrow \text{dom}(\mathfrak{h}) \setminus \mathfrak{s}(\Gamma \cup \{\mathbf{0}, \mathbf{1}\})$ such that, for all $w \in \text{nodes}(t)$, the following hold:

1. for all $i \in \mathbb{N}$, if $wi \in \text{nodes}(t)$ then $f(wi) \in \mathfrak{h}(f(w))$,
2. if w is a branching node and $t(w) = (q, i)$ then $\text{Pr}(f(w)) = q$ and $\mathfrak{h}(f(w)) = (\ell_1, \dots, \ell_{\mathfrak{N}+1})$, where $\ell_j = \mathfrak{s}(\text{bin}(i)_j)$, for all $j \in \llbracket 1 \dots \mathfrak{N} \rrbracket$ and $\ell_{\mathfrak{N}+1} \in \text{dom}(\mathfrak{h})$,
3. if w is an action node and $t(w) = (a, b, \mu)$ then $\mathfrak{h}(f(w)) = (\mathfrak{s}(a), \mathfrak{s}(b), \mathfrak{s}(\widetilde{\mu}), \ell_4)$, where $\ell_4 \in \text{dom}(\mathfrak{h})$,

We write $(\mathfrak{s}, \mathfrak{h}) > t$ instead of $(\mathfrak{s}, \mathfrak{h}) \triangleright t$ if condition (2) above is replaced by the weaker:

- 2'.** if w is a branching node and $t(w) = (q, i)$ then $\mathfrak{h}(f(w)) = (\ell_1, \dots, \ell_{\mathfrak{N}+1})$, where $\ell_j = \mathfrak{s}(\text{bin}(i)_j)$, for all $j \in \llbracket 1 \dots \mathfrak{N} \rrbracket$ and $\ell_{\mathfrak{N}+1} \in \text{dom}(\mathfrak{h})$.

Note that, by Lemma 1 $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} p_M(x)$ iff there exists an unfolding tree $u \in \mathcal{T}_{\mathcal{S}}(p_M(x))$ such that $(\mathfrak{s}, \mathfrak{h}) \models \mathcal{F}(u)$ and, by Lemma 2, there exists a bijection $\Lambda : \text{dom}(\mathfrak{h}) \rightarrow \text{nodes}(u) \setminus \{\lambda\}$, such that $\Lambda^{-1}(wi) \in \mathfrak{h}(\Lambda^{-1}(w))$, for all $wi \in \text{nodes}(u)$. Then the occurrence of $\text{Pr}(\cdot)$ at point (2) of Definition 8 stands for $\text{Pr}_{\mathfrak{s}, \mathfrak{h}, \Lambda}(\cdot)$, the structure $(\mathfrak{s}, \mathfrak{h})$ being clear from the context and the existence of Λ being guaranteed by Lemma 2. The weaker relation $(\mathfrak{s}, \mathfrak{h}) > t$ shall be used next to encode pseudo-derivations by structures that are not necessarily models of the above rules (more rules will be added later).

Lemma 3. (A) For each pseudo-derivation t of an ATM $M = (Q, \Gamma, \delta, q_0, g)$, there exists a structure $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} p_M(x)$ such that $(\mathfrak{s}, \mathfrak{h}) \triangleright t$. (B) Dually, for each structure $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} p_M(x)$, there exists an accepting pseudo-derivation t of M such that $(\mathfrak{s}, \mathfrak{h}) \triangleright t$.

Proof: (A) Let t be a pseudo-derivation of M . We build an isomorphic tree u , such that $\text{nodes}(t) = \text{nodes}(u)$ and the labels of u are of the form (q, ϕ) , where $q \in \text{Pred}$ and ϕ is a formula. The definition of u is top-down on the structure of t , for each node $w \in \text{nodes}(t)$, with label $t(w) = (q, i)$:

- if $g(q) = \vee$, where $t(w0) = (a, b, \mu)$ and $t(w00) = (q', i')$ (we have necessarily $w0, w00 \in \text{nodes}(t)$, where $w0$ and $w00$ are the only children of w and $w0$, respectively), then we define:

$$\begin{aligned} u(w) &\stackrel{\text{def}}{=} (q, \exists x' . x \mapsto (\text{bin}(i), x') * \bar{q}(x', a, b, \widetilde{\mu})) \quad (\text{see rule 5}) \\ u(w0) &\stackrel{\text{def}}{=} (\bar{q}, \exists x'' . x' \mapsto (a, b, \widetilde{\mu}, x'') * q(x'')) \quad (\text{see rule 7}) \end{aligned}$$

- otherwise, $g(q) = \wedge$ and $t(wj) = (a_j, b_j, \mu_j)$ for $j \in \llbracket 1 \dots n \rrbracket$ are the children of w , and $t(wj0) = (q'_j, i'_j)$, for all $j \in \llbracket 1 \dots n \rrbracket$, in which case we define:

$$\begin{aligned} u(w) &\stackrel{\text{def}}{=} \left(q, \exists y_1 \dots \exists y_n . x \mapsto (\text{bin}(i), y_1, \dots, y_n) * \bigstar_{j=1}^n \bar{q}_j(y_j, a_j, b_j, \bar{\mu}_j) \right) \quad (\text{see rule } \color{red}{6}) \\ u(wj) &\stackrel{\text{def}}{=} \left(\bar{q}_j, \exists x' . y_j \mapsto (a_j, b_j, \bar{\mu}_j, x') * q'_j(x') \right) \quad (\text{see rule } \color{red}{7}) \end{aligned}$$

Finally, either one of the following holds:

- $w = vi$, for some $v \in \text{nodes}(t)$, $i \in \mathbb{N}$ and the atom $q(x)$ occurs exactly once in $u(v)$,
- $w = \lambda$ and the variable x in the definition of $u(w)$ is the same as y_0 , in which case necessarily $t(\lambda) = (q_0, 0)$.

Next, we extend u to an unfolding tree $\hat{u} \in \mathcal{T}_S(p_M(x))$ by adding to it a hat of height \aleph and a sibling tree $u' \in \mathcal{T}_S(\text{Const}(z_0))$. It is not hard to check that $\hat{u} \in \mathcal{T}_S(p_M(x))$ and that $\mathcal{F}(\hat{u})$ is satisfiable, since there are no equality or disequality atoms and no variable occurs allocated in two different subtrees of \hat{u} . Then, let (s, h) be a structure such that $(s, h) \models \mathcal{F}(\hat{u})$. To check that $(s, h) \triangleright t$, we need to exhibit an injective mapping $f : \text{nodes}(t) \rightarrow \text{dom}(h)$ that meets the conditions (1), (2) and (3) from Definition 8. Because \mathcal{S} is a progressing and connected set of rules and $(s, h) \models \mathcal{F}(\hat{u})$, by Lemma 2, there exists a bijective mapping $\Lambda : \text{dom}(h) \rightarrow \text{nodes}(\hat{u}) \setminus \{\lambda\}$ such that $\Lambda^{-1}(wi) \in h(\Lambda^{-1}(w))$, for all $w \in \text{nodes}(\hat{u}) \setminus \{\lambda\}$ and all $i \in \mathbb{N}$, such that $wi \in \text{nodes}(\hat{u})$. Let f be the restriction of Λ^{-1} to $\text{nodes}(u)$. Point (1) follows by the definition of Λ , whereas points (2) and (3) are simple checks.

(B) Conversely, if $(s, h) \models_S p_M(x)$ then, by Lemma 1, there exists an unfolding tree $u \in \mathcal{T}_S(p_M(x))$ such that $(s, h) \models \mathcal{F}(u)$. Since \mathcal{S} is progressing and connected, by Lemma 2, there exists a bijective mapping $\Lambda : \text{dom}(h) \rightarrow \text{nodes}(u) \setminus \{\lambda\}$, as before. By the definition of \mathcal{S} , all nodes $w \in \text{nodes}(u) \setminus \{\lambda\}$ that are labeled with predicates q and \bar{q} , for some $q \in \mathcal{Q}$, occur below a unique node $w_0 \in \text{nodes}(u)$, such that $t(w_0) = (q_0, \phi)$, for some formula ϕ . We build a pseudo-derivation t of M such that $\text{nodes}(t) = \text{nodes}(u|_{w_0})$, by induction on the structure of $u|_{w_0}$. Namely, for each $w \in \text{nodes}(u|_{w_0})$:

- If $u|_{w_0}(w) = (q, \phi)$ then ϕ is the body of a rule (5) or (6). In the case (5) (the other case is similar and left to the reader) we have $h(\Lambda^{-1}(w)) = (\ell_1, \dots, \ell_{\aleph+1})$, with $\ell_1, \dots, \ell_{\aleph} \in \{s(\mathbf{0}), s(\mathbf{1})\}$ and let i be the integer such that $\text{bin}(i) = (s^{-1}(\ell_1), \dots, s^{-1}(\ell_{\aleph}))$. Note that, since $\mathbf{0}$ and $\mathbf{1}$ are always allocated separately in $\mathcal{F}(u)$, the restriction of s to the set $\{\mathbf{0}, \mathbf{1}\}$ is a bijection. In this case, we define $t(w) \stackrel{\text{def}}{=} (q, i)$.
- Otherwise, $u|_{w_0}(w) = (\bar{q}, \phi)$ and ϕ is the body of a rule (7). In this case, we have $h(\Lambda^{-1}(w)) = (\ell_1, \ell_2, \ell_3, \ell_4)$, with $\ell_1 \in s(\Gamma)$, $\ell_2 \in s(\Gamma \setminus \{b\})$ and $\ell_3 \in \{s(\mathbf{0}), s(\mathbf{1})\}$. Note that, since each $\gamma \in \Gamma \cup \{\mathbf{0}, \mathbf{1}\}$ is allocated separately in $\mathcal{F}(u)$, the restriction of s to the set $\Gamma \cup \{\mathbf{0}, \mathbf{1}\}$ is a bijection. In this case, we define $t(w) \stackrel{\text{def}}{=} (s^{-1}(\ell_1), s^{-1}(\ell_2), \mu)$, where $\mu = \leftarrow$ if $\ell_3 = s(\mathbf{0})$ and $\mu = \rightarrow$ if $\ell_3 = s(\mathbf{1})$.

It is easy to check that, indeed t is a pseudo-derivation of M . To check that $(s, h) \triangleright t$, we take $f : \text{nodes}(t) \rightarrow \text{dom}(h)$ as the restriction of Λ^{-1} to the nodes of $u|_{w_0}$. Clearly, f is injective and the conditions (1), (2) and (3) of Definition 8 are easy checks. \square

4.2 Encoding Complement Membership as Entailment Problems

Given an ATM $M = (Q, \Gamma, \delta, q_0, g)$, a pseudo-derivation of M is a derivation of M if the contents of the tape is consistent with the sequence of the actions applied, in particular the following conditions must hold:

- I. If a branching node labeled (q, i) is followed by an action node labeled (a, b, \rightarrow) [resp. (a, b, \leftarrow)], itself followed by a branching node labeled (q', i') then necessarily $i' = i + 1$ [resp. $i' = i - 1$], i.e. the position of the head changes according to the action executed between the adjacent configurations.
- II. For every $i \in \llbracket 0 \dots 2^{\mathfrak{N}} - 1 \rrbracket$, if along a path from a branching node labeled (q, i) , followed by an action node labeled (a, b, μ) , to another branching node labeled (q', i) , followed by an action node labeled (a', b', μ') , there is no branching node labeled (q'', i) , then necessarily $a' = b$. Indeed, the symbol read on position i must be the one previously written, since it was not changed in the meantime.
- III. For every $i \in \llbracket 0 \dots 2^{\mathfrak{N}} - 1 \rrbracket$, if along a path from the root to a branching node labeled (q, i) , followed by an action node labeled (a, b, μ) , there is no branching node labeled (q', i) , then necessarily $a = b$, i.e. the tape is initially empty.

In the following, we shall not check that the above conditions hold for some derivation of M , but rather the opposite: that for each derivation of M , at least one of the above conditions is broken. In other words, we reduce from the complement of the membership problem (M, ϵ) to an entailment problem, defined next. This does not change the final 2EXPTIME-hardness result, because $2EXPTIME = AEXPSPACE = co-AEXPSPACE$, as previously mentioned.

In the following, we shall give the rules for a predicate $c_M(x)$ such that the entailment $p_M(x) \models_{\mathcal{S}} c_M(x)$ holds, for a suitable set of rules \mathcal{S} , containing the rules for $p_M(x)$ and $c_M(x)$, iff every pseudo-derivation of M breaks at least one of the conditions (I), (II) or (III), in other words, that M , started on input ϵ , has no derivation.

Let $\mathfrak{B} \stackrel{\text{def}}{=} \max_{q \in Q, a \in \Gamma} \|\tau(q, a)\|$ be the maximum branching degree of a derivation of M . First, we define an auxiliary predicate $r(x)$ that generates all tree-shaped heaps in which branching nodes correctly alternate with action nodes, with no regard to the labels of those nodes. In the following, we stick to the convention that predicate symbols p represent branching nodes, whereas \bar{p} represent action nodes:

$$\begin{aligned} r(x) &\Leftarrow \exists y_1 \dots \exists y_n . x \mapsto (\bullet^{\mathfrak{N}}, y_1, \dots, y_n) * \bigstar_{j=1}^n \bar{r}(y_j), \text{ for each } n \in \llbracket 0 \dots \mathfrak{B} \rrbracket \\ \bar{r}(x) &\Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * r(y), \text{ for each } a \in \Gamma \text{ and } b \in \Gamma \setminus \{b\} \end{aligned}$$

Next, we define the heap encodings of those derivation trees that violate condition (I). To this end, we guess a vector \mathbf{b} in $\{0, 1\}^{\mathfrak{N}}$, encoding a position on the tape $i \in \llbracket 0 \dots 2^{\mathfrak{N}} - 1 \rrbracket$, a shift $\mu \in \{\leftarrow, \rightarrow\}$, encoded by $\bar{\mu} \in \{0, 1\}$ and get the binary complement of the (encoding of the) position reached from \mathbf{b} by applying μ . Here we distinguish two cases, depending on the choice of μ :

- (a) if μ is \rightarrow then $\text{bin}(i) = \mathbf{b} \stackrel{\text{def}}{=} (b_1, \dots, b_n, \mathbf{0}, \mathbf{1}^{\mathfrak{N}-1-n})$ for some $n \in \llbracket 0 \dots \mathfrak{N} - 1 \rrbracket$ and let $\mathbf{c} \stackrel{\text{def}}{=} (\bar{b}_1, \dots, \bar{b}_n, \mathbf{0}, \mathbf{1}^{\mathfrak{N}-1-n})$ be the complement of $\text{bin}(i+1) = (b_1, \dots, b_n, \mathbf{1}, \mathbf{0}^{\mathfrak{N}-1-n})$.
- (b) otherwise, $\text{bin}(i) = \mathbf{b} \stackrel{\text{def}}{=} (b_1, \dots, b_n, \mathbf{1}, \mathbf{0}^{\mathfrak{N}-1-n})$ and let $\mathbf{c} \stackrel{\text{def}}{=} (\bar{b}_1, \dots, \bar{b}_n, \mathbf{1}, \mathbf{0}^{\mathfrak{N}-1-n})$ be the complement of $\text{bin}(i-1) = (b_1, \dots, b_n, \mathbf{0}, \mathbf{1}^{\mathfrak{N}-1-n})$.

For every $n \in \llbracket 0 \dots \aleph - 1 \rrbracket$ and every $m \in \llbracket 0 \dots \aleph \rrbracket$ and $i \in \llbracket 1 \dots m \rrbracket$, we consider the rules:

$$c_1(x) \Leftarrow \exists b_1 \dots \exists b_n \exists y . x \mapsto ([y]^{\aleph}) * d_1(y, \underbrace{1, b_1 \dots b_n, 0, 1^{\aleph-n-1}}_b, \underbrace{\bar{b}_1 \dots \bar{b}_n, 0, 1^{\aleph-n-1}}_c) \quad (9)$$

$$c_1(x) \Leftarrow \exists b_1 \dots \exists b_n \exists y . x \mapsto ([y]^{\aleph}) * d_1(y, \underbrace{0, b_1 \dots b_n, 1, 0^{\aleph-n-1}}_b, \underbrace{\bar{b}_1 \dots \bar{b}_n, 1, 0^{\aleph-n-1}}_c) \quad (10)$$

$$d_1(x, u, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\bullet^{\aleph}, y_1, \dots, y_m) * \bigstar_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} \bar{r}(y_j) * \bar{d}_1(y_i, u, \mathbf{b}, \mathbf{c}) \quad (11)$$

$$d_1(x, u, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\mathbf{b}, y_1, \dots, y_m) * \bigstar_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} \bar{r}(y_j) * \bar{e}_1(y_i, u, \mathbf{b}, \mathbf{c}) \quad (12)$$

$$\bar{d}_1(x, u, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * d_1(y, u, \mathbf{b}, \mathbf{c}), \text{ for each } a \in \Gamma, b \in \Gamma \setminus \{b\} \quad (13)$$

$$\bar{e}_1(x, u, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * f_1(y, u, \mathbf{b}, \mathbf{c}), \text{ for each } a \in \Gamma, b \in \Gamma \setminus \{b\} \quad (14)$$

$$f_1(x, u, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (e, y_1, \dots, y_m) * \bigstar_{j=1}^m \bar{r}(y_j) \mid e \neq \bar{c} \quad (15)$$

Intuitively, rules (9) and (10) choose the move $u \in \{0, 1\}$ and the binary vectors $\mathbf{b}, \mathbf{c} \in \{0, 1\}^{\aleph}$, according to the cases (a) and (b) above, respectively. Note that we use the hat $[y]^{\aleph}$ to eliminate the binary variables b_1, \dots, b_n , as $n \leq \aleph$, according to the elimination procedure described in §3.1. Then a path to the branching node, labeled (q', i') , that violates condition (I) is chosen, by alternating the branching and action nodes allocated by rules (11) and (13), respectively. The offending branching node is allocated by rule (15) and its predecessors are the branching and the action nodes, labeled with (q, i) and (i, a, b, μ) , such that $i' \neq i^{\mu}$. These latter nodes are allocated by rules (12) and (14), respectively.

The pseudo-derivations of M that violate condition (II) are encoded by the tree-structured heaps defined by the rules below. To this end, we guess a binary vector $\mathbf{b} \in \{0, 1\}^{\aleph}$ denoting the position of a write action that has an inconsistent read descendant and let \mathbf{c} be its binary complement. Then, for every $m \in \llbracket 0 \dots \aleph \rrbracket$ and $i \in \llbracket 1 \dots m \rrbracket$, we consider the rules below:

$$c_2(x) \Leftarrow \exists b_1 \dots \exists b_{\aleph} \exists y . x \mapsto ([y]^{\aleph}) * d_2(y, \underbrace{b_1, \dots, b_{\aleph}}_b, \underbrace{\bar{b}_1, \dots, \bar{b}_{\aleph}}_c) \quad (16)$$

$$d_2(x, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\bullet^{\aleph}, y_1, \dots, y_m) * \bigstar_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} \bar{r}(y_j) * \bar{d}_2(y_i, \mathbf{b}, \mathbf{c}) \quad (17)$$

$$\bar{d}_2(x, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * d_2(y, \mathbf{b}, \mathbf{c}), \text{ for each } a \in \Gamma, b \in \Gamma \setminus \{b\} \quad (18)$$

$$\bar{d}_2(x, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * e_2(y, \mathbf{b}, \mathbf{c}), \text{ for each } a \in \Gamma, b \in \Gamma \setminus \{b\} \quad (19)$$

$$e_2(x, \gamma, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\mathbf{b}, y_1, \dots, y_m) * \bigstar_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} \bar{r}(y_j) * \bar{e}_2(y_i, \gamma, \mathbf{b}, \mathbf{c}) \quad (20)$$

$$\bar{e}_2(x, \gamma, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * f_2(y, \gamma, \mathbf{b}, \mathbf{c}), \text{ for each } a \in \Gamma, b \in \Gamma \setminus \{b\} \quad (21)$$

$$\bar{e}_2(x, \gamma, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * g_2(y, \gamma, \mathbf{b}, \mathbf{c}), \text{ for each } a \in \Gamma, b \in \Gamma \setminus \{b\} \quad (22)$$

$$f_2(x, \gamma, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (e, y_1, \dots, y_m) * \bigstar_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} \bar{r}(y_j) * \bar{e}_2(y_i, \gamma, \mathbf{b}, \mathbf{c}) \mid e \neq \bar{c} \quad (23)$$

$$g_2(x, \gamma, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\mathbf{b}, y_1, \dots, y_m) * \bigstar_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} \bar{r}(y_j) * \bar{g}_2(y_i, \gamma) \quad (24)$$

$$\bar{g}_2(x, \gamma) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * r(y), \text{ for each } a \in \Gamma \setminus \{\gamma\}, b \in \Gamma \setminus \{b\} \quad (25)$$

Intuitively, rule (16) uses the hat $[y]^{\mathfrak{y}}$ to choose the tuple of binary variables $\mathbf{b} = (b_1, \dots, b_{\mathfrak{y}})$ and their complements $\mathbf{c} = (\bar{b}_1, \dots, \bar{b}_{\mathfrak{y}})$. First, the path to a branching node labeled by the binary position \mathbf{b} is non-deterministically chosen by an alternation of branching and action nodes allocated by the rules (17) and (18), respectively, until the node and its predecessor are allocated by rules (20) and (19), respectively. The symbol written on the tape by this node is stored in the second parameter of $e_2(x, \gamma, \mathbf{b}, \mathbf{c})$. Next, a path to a second branching node labeled by the binary position \mathbf{b} is non-deterministically chosen by an alternation of branching and action nodes allocated by the rules (21) and (23) respectively, while checking that no branching node with the same position \mathbf{b} occurs on this second path (23). At the end, we reach the offending branching node (24), whose predecessor is allocated by rule (22). At this point, we check that the symbol read by the last action node is different than the symbol previously written at position \mathbf{b} , by rule (19). This check is done by rules (24) and (25), thus ensuring that condition (II) is indeed violated.

Next, we define the tree-structured heap encodings of the derivation trees that violate condition (III). To this end, we guess a binary vector $\mathbf{b} \in \{0, 1\}^{\mathfrak{y}}$ denoting the position where a symbol different from \mathbf{b} has been read, with no previous write action at that position and let \mathbf{c} be its complement. We consider the rules below, for every $m \in \llbracket 0 \dots \mathfrak{B} \rrbracket$ and $i \in \llbracket 1 \dots m \rrbracket$:

$$c_3(x) \Leftarrow \exists b_1 \dots \exists b_{\mathfrak{y}} \exists y . x \mapsto ([y]^{\mathfrak{y}}) * d_3(y, \underbrace{b_1, \dots, b_{\mathfrak{y}}}_{\mathbf{b}}, \underbrace{\bar{b}_1, \dots, \bar{b}_{\mathfrak{y}}}_{\mathbf{c}}) \quad (26)$$

$$d_3(x, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\mathbf{e}, y_1, \dots, y_m) * \underbrace{* \bar{r}(y_j)}_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} * \bar{d}_3(y_i, \mathbf{b}, \mathbf{c}) \mid \mathbf{e} \neq \bar{\mathbf{c}} \quad (27)$$

$$\bar{d}_3(x, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * d_3(y, \mathbf{b}, \mathbf{c}), \text{ for all } a \in \Gamma, b \in \Gamma \setminus \{\mathbf{b}\} \quad (28)$$

$$\bar{d}_3(x, \mathbf{b}, \mathbf{c}) \Leftarrow x \mapsto (a, b, \bullet, y) * e_3(y, \mathbf{b}, \mathbf{c}), \text{ for all } a \in \Gamma, b \in \Gamma \setminus \{\mathbf{b}\} \quad (29)$$

$$e_3(x, \mathbf{b}, \mathbf{c}) \Leftarrow \exists y_1 \dots \exists y_m . x \mapsto (\mathbf{b}, y_1, \dots, y_m) * \underbrace{* \bar{r}(y_j)}_{j \in \llbracket 1..m \rrbracket \setminus \{i\}} * \bar{f}_3(y_i) \quad (30)$$

$$\bar{f}_3(x) \Leftarrow \exists y . x \mapsto (a, b, \bullet, y) * r(y), \text{ for all } a, b \in \Gamma \setminus \{\mathbf{b}\} \quad (31)$$

After the initial guess of the binary position \mathbf{b} , by rule (26), a path to a branching node labeled by \mathbf{b} is nondeterministically guessed, by an alternation of branching and action nodes corresponding to the rules (27) and (28), respectively, while checking that no branching node labeled with position \mathbf{b} occurs on this path. Once this node is reached, by rule (29), we check that its action node successor reads a symbol different than \mathbf{b} , by rules (30) and (31), which is in violation of condition (III).

Finally, the predicate $c_M(x)$ that choses the condition (I), (II) or (III) to be violated, is defined by the following rules:

$$c_M(x) \Leftarrow \exists y_0 \exists z_0 . x \mapsto (y_0, z_0) * c_i(y_0) * \text{Const}(z_0), \text{ for all } i \in \{1, 2, 3, \} \quad (32)$$

Let \mathcal{S} denote the set of rules introduced so far. The following lemma states the property of the models of $c_M(x)$:

Lemma 4. *Given a pseudo-derivation t of M and a structure $(\mathfrak{s}, \mathfrak{h})$, such that $(\mathfrak{s}, \mathfrak{h}) > t$, we have $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{S}} c_M(x)$ if and only if t is not a derivation of M .*

Proof: (sketch)

Fact 2 Assume that $(s, h) \models_S c_M(x)$. For all $\ell, \ell' \in \text{dom}(h)$ and $i, j \in \llbracket 1 \dots k \rrbracket$, $h(\ell)_i = h(\ell')_j$ iff either (i) $\ell = h(\ell)_i$, (ii) $h(\ell)_i \in s(\Gamma \cup \{\mathbf{0}, \mathbf{1}\})$, or (iii) $\ell = \ell'$ and $i = j$.

Proof: By close analysis of allocation within the rules of \mathcal{S} . \square

“ \Rightarrow ” If $(s, h) \models_S c_M(x)$ then, by Lemma 1, there exists an unfolding tree $u \in \mathcal{T}_S(c_M(x))$ such that $(s, h) \models \mathcal{F}(u)$. Since \mathcal{S} is a progressing and connected set of rules and, moreover $(s, h) \models \mathcal{F}(u)$, by Lemma 2, there exists a bijection $\Lambda : \text{dom}(h) \rightarrow \text{nodes}(u) \setminus \{\lambda\}$ such that, for all $wi \in \text{nodes}(u)$, where $i \in \mathbb{N}$, we have $\Lambda^{-1}(wi) \in h(\Lambda^{-1}(w))$. Moreover, since $(s, h) > t$, there exists an injective mapping $f : \text{nodes}(t) \rightarrow \text{dom}(h) \setminus s(\Gamma \cup \{\mathbf{0}, \mathbf{1}\})$ such that, for all $w \in \text{nodes}(t)$ such that $wi \in \text{nodes}(t)$, we have $f(wi) \in h(f(w))$. Consequently, we obtain the bijection $\Lambda \circ f : \text{nodes}(t) \rightarrow \text{nodes}(u) \setminus \{\lambda\}$ where, for all $wi \in \text{nodes}(t)$, such that $i \in \mathbb{N}$, there exists $j \in \mathbb{N}$ such that $f(wi) = h(f(w))_j$. Since Λ is bijective and $\lambda \notin \text{img}(\Lambda)$, there exists a node $vk \in \text{nodes}(u)$, for some $k \in \mathbb{N}$, such that $h(f(w))_j = \Lambda^{-1}(vk)$. By Lemma 2, there exists $m \in \mathbb{N}$ such that $\Lambda^{-1}(vk) = h(\Lambda^{-1}(v))_m$. Since f is injective, $h(f(w))_j = f(wi) \neq f(w)$. Moreover, $h(f(w))_j \notin s(\Gamma \cup \{\mathbf{0}, \mathbf{1}\})$ thus, by Fact 2, we obtain that $f(w) = \Lambda^{-1}(v)$, as a consequence of $h(f(w))_j = h(\Lambda^{-1}(v))_m$. Then we obtain that, for all $wi \in \text{nodes}(t)$, such that $i \in \mathbb{N}$, there exists $k \in \mathbb{N}$, such that:

$$\Lambda(f(wi)) = \Lambda(h(f(w))_j) = vk = \Lambda(f(w))_k$$

We can conclude that t and $u|_0$ are isomorphic, i.e. $\text{nodes}(t) = \text{nodes}(u) \setminus \{\lambda\}$.

By the definition of \mathcal{S} , namely rule (32), there exists a unique node $w_0 \in \text{nodes}(u) \setminus \{\lambda\}$ such that $u(w_0) = (c_i, \phi_i)$, for some formula ϕ_i , where $i \in \{1, 2, 3\}$. Distinguishing the cases $i = 1, 2, 3$ and using the fact that $(s, h) > t$, one shows that t breaks the one of the conditions (I), (II) or (III), respectively, thus t is not a derivation of M . The proof is along the lines of the second point of Lemma 3.

“ \Leftarrow ” If t is an accepting pseudo-derivation but not a derivation of M , then t violates one of the conditions (I), (II) or (III). In each case, we build an unfolding tree $u \in \mathcal{T}_S(c_M(x))$, along the lines of the proof of the first point of Lemma 3. Using the fact that $(s, h) > t$, we show that $(s, h) \models \mathcal{F}(u)$, leading to $(s, h) \models_S c_M(x)$. \square

Note that, akin to the rule for $p_M(x)$ (8), the rules (32) contain an occurrence of $\text{Const}(z_0)$ as a sibling to a hat of height \mathfrak{N} , that occurs in $c_i(y_0)$, for all $i = 1, 2, 3$. Then the entailment $p_M(x) \models c_M(x)$ holds if and only if, for each structure (s, h) such that $(s, h) \models p_M(x)$ and each extension $s[y_0 \leftarrow \ell_0]$, for some location $\ell_0 \in \text{Loc}$, the heap h is matched by the unfolding of one of the rules with head $c_1(y_0)$, $c_2(y_0)$ or $c_3(y_0)$. This is possible because each such rule uses a hat $[y]^{\mathfrak{N}}$, matching the one from the rule with head $p_M(x)$ (8).

Lemma 5. The entailment $p_M(x) \models_S c_M(x)$ holds if and only if the membership problem (M, ϵ) has a negative answer.

Proof: “ \Rightarrow ” Suppose, for a contradiction, that M accepts $(q_0, \epsilon, 0)$. By Definition 6 it has a derivation t . Since t is a derivation, it is also a pseudo-derivation of M and, by Lemma 3, there exists a structure (s, h) such that $(s, h) \models_S p_M(x)$ and $(s, h) \triangleright t$. Moreover, $(s, h) > t$ follows from $(s, h) \triangleright t$ and, by Lemma 4, we obtain $(s, h) \not\models_S c_M(x)$, thus $p_M(x) \not\models_S c_M(x)$, contradiction.

” \Leftarrow ” Suppose, for a contradiction, that $p_M(x) \not\models_S c_M(x)$, hence there exists a structure (s, h) such that $(s, h) \models_S p_M(x)$ and $(s, h) \not\models_S c_M(x)$. By Lemma 3, there exists a pseudo-derivation t of M such that $(s, h) \triangleright t$, hence $(s, h) > t$. By Lemma 4, t is a derivation of M , hence (M, ϵ) has a positive answer, contradiction. \square

We state the main result of this paper below:

Theorem 2. *The entailment problem $p(x) \models_S q(x)$, where S is a progressing, connected and established set of rules and p, q are predicate symbols in Pred that occur as heads in S , is 2EXPTIME-hard.*

Proof: Given an exponential-space bounded ATM M we define a set of rules S , based on the description of M , such that $p_M(x) \models_S c_M(x)$ if and only if (M, ϵ) has a negative answer (Lemma 5). Moreover, the set of rules is easy shown to be progressing, connected and established. The reduction is possible in time polynomial in the size of the standard encoding of M . Indeed, the number of rules in S is $O(\|Q\| \cdot \aleph \cdot \aleph)$ and the succinct representation of each rule, using deep hats, binary choices and binary variables can be generated in time $O(\|I\| \cdot \aleph \cdot \aleph)$. Finally, the complete elimination of binary variables is possible in polynomial time. Since we reduce from the complement of a AEXPSpace-complete membership problem and $\text{co-AEXPSpace} = \text{AEXPSpace} = \text{2EXPTIME}$, we obtain the 2EXPTIME-hardness result. \square

5 Conclusions

We show that the entailment problem, for symbolic heaps with inductively defined predicates, showed to be decidable (with elementary recursive time complexity) in [4] has an actual 2EXPTIME-hard lower bound. In the light of the recent results of [7], we expect 2EXPTIME to be the tight complexity for what is currently the most general decidable class of entailments for Separation Logic with inductive definitions.

References

1. Timos Antonopoulos, Nikos Gorogiannis, Christoph Haase, Max I. Kanovich, and Joël Ouaknine. Foundations for decision problems in separation logic with general inductive predicates. In Anca Muscholl, editor, *Foundations of Software Science and Computation Structures - 17th International Conference, FOSSACS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8412 of *Lecture Notes in Computer Science*, pages 411–425, 2014.
2. James Brotherston, Carsten Fuhs, Juan Antonio Navarro Pérez, and Nikos Gorogiannis. A decision procedure for satisfiability in separation logic with inductive predicates. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 25:1–25:10. ACM, 2014.
3. Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981. URL: <https://doi.org/10.1145/322234.322243>, doi:10.1145/322234.322243.

4. Radu Iosif, Adam Rogalewicz, and Jiri Simacek. The tree width of separation logic with recursive definitions. In *Proc. of CADE-24*, volume 7898 of *LNCS*, 2013.
5. Radu Iosif, Adam Rogalewicz, and Tomás Vojnar. Deciding entailments in inductive separation logic with tree automata. In Franck Cassez and Jean-François Raskin, editors, *Automated Technology for Verification and Analysis - 12th International Symposium, ATVA 2014, Sydney, NSW, Australia, November 3-7, 2014, Proceedings*, volume 8837 of *Lecture Notes in Computer Science*, pages 201–218. Springer, 2014.
6. Samin S Ishtiaq and Peter W O’Hearn. Bi as an assertion language for mutable data structures. In *ACM SIGPLAN Notices*, volume 36, pages 14–26, 2001.
7. Jens Katelaan, Christoph Matheja, and Florian Zuleger. Effective entailment checking for separation logic with inductive definitions. In Tomás Vojnar and Lijun Zhang, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part II*, volume 11428 of *Lecture Notes in Computer Science*, pages 319–336. Springer, 2019.