

Reasoning about sequences of memory states[☆]

Rémi Brochenin^a, Stéphane Demri^a, Etienne Lozes^a

^a*LSV, ENS Cachan, CNRS, INRIA Saclay, France*

Abstract

Motivated by the verification of programs with pointer variables, we introduce a temporal logic LTL^{mem} whose underlying assertion language is the quantifier-free fragment of separation logic and the temporal logic on the top of it is the standard linear-time temporal logic LTL. We analyze the complexity of various model-checking and satisfiability problems for LTL^{mem} , considering various fragments of separation logic (including pointer arithmetic), various classes of models (with or without constant heap), and the influence of fixing the initial memory state. We provide a complete picture based on these criteria. Our main decidability result is PSPACE-completeness of the satisfiability problems on the record fragment and on a classical fragment allowing pointer arithmetic. Σ_1^0 -completeness or Σ_1^1 -completeness results are established for various problems by reducing standard problems for Minsky machines, and underline the tightness of our decidability results.

Key words: separation logic, temporal logic, Büchi automaton, computational complexity

1. Introduction

Verification of programs with pointers. Model-checking of infinite-state systems is a very active area of formal verification [1] even though in full generality, simple reachability questions are undecidable. Nevertheless, many classes of infinite-state systems can be analyzed, such as Petri nets, timed automata, etc. Programs with pointer variables suffer the same drawback since reachability problems are also undecidable, see e.g. [2, 3]. It is worth noting that specific properties need to be verified for such programs, such as the existence of memory leaks, memory violation, or shape analysis. Prominent logics for analyzing such programs are Separation Logic [4], pointer assertion logic PAL [5], TVLA [6] and alias logic [7], to quote a few examples.

Towards a temporal separation logic. Since [8], temporal logics are used as languages for formal specification of programs. General and powerful automata-based techniques for verification have been developed, (see for example the works [9, 10]). On the other hand, Separation Logic is a static logic for program annotation [4], and more recently for symbolic computation [11]. Extending the scope of application of Separation Logic to standard temporal logic-based verification techniques has many potential interests. First, it provides a rich underlying assertion language where properties more complex than accessibility can be stated. Second, this may open a new direction of investigation for the purely static Separation Logic extended with general recursion, in the same spirit as [12]. For instance, if we write Xx to denote the next value of x (also sometimes written x'), the formula $(x \hookrightarrow Xx)U(x \hookrightarrow \text{null})$, understood on a model with constant heap, characterises the existence of a simple flat list, which is usually written $\mu L(x). x \hookrightarrow \text{null} \vee \exists x'. x \hookrightarrow x' \wedge L(x')$. Third, temporal logics allow to work in the very convenient framework of "programs-as-formulae" and decision procedures for logical problems can be directly used for program verification. For instance, the previous formula can be seen as a program traversing a list and, more generally, programs without destructive updates can be expressed as formulae. Some programs with destructive updates that perform a simple pass

[☆]Work partially supported by the RNTL project "AVERILES". The first author is supported by a fellowship from CNRS/DGA.

on the heap, have an input-output relation that may be described by a formula. For instance, the formula $(x \mapsto_0 Xx \wedge Xx \mapsto_1 x)Ux \mapsto_0 \text{null}$ roughly expresses that the list in the initial heap h_0 is reversed in the final heap h_1 . Fourth, pointer arithmetic has been poorly studied until now, whereas arithmetical constraints in temporal logics are known to easily lead to undecidability, see e.g. [13, 14, 15]. Actually, there is a growing interest in understanding the interplay of pointer arithmetic, temporal reasoning, and non aliasing properties.

Our contribution. We introduce a linear-time temporal logic LTL^{mem} to specify sequences of memory states with underlying assertion language based on quantifier-free Separation Logic [4]. From a logical perspective, the logic LTL^{mem} can be viewed as a many-dimensional logic [16] since LTL^{mem} contains a temporal dimension and the spatial dimension for memory states. Other many-dimensional logics can be found in [17, 18, 16, 19]. Our logic addresses a very general notion of models, including the aspects of pointer arithmetic and recursive structures with records. We distinguish the satisfiability problems from the model-checking problems, as well as distinct subclasses of interesting programs, like for instance the programs without destructive update. The most promising result for future implementation is the PSPACE-completeness of the satisfiability problems $SAT(\text{CL})$ and $SAT(\text{RF})$ where CL is the classical fragment without separation connectives and RF is the record fragment with no pointer arithmetic but with separation connectives. This result is very tight, as both propositional LTL and static Separation Logic are already PSPACE-complete [20, 21]. These results are obtained by reduction to the nonemptiness problem for Büchi automata on an alphabet made of symbolic memory states obtained by an abstraction that we show sound and complete, see e.g. [22, 23]. Such abstractions are similar to resource graphs from [24, 25]. This is a variant of the automata-based approach introduced in [9] for plain LTL and further developed with concrete domains of interpretation in [19]. Surprisingly, the abstraction method used to establish these results does not scale to the whole logic, due to a subtle interplay between separation connectives and pointer arithmetic. Moreover, we provide new undecidability results for several problems, for instance $SAT^{ct}(\text{LF})$ (satisfiability with constant heap on the list fragment).

Related work. Previous temporal logics designed for pointer verification include Evolution Temporal Logic [26], based on the three-valued logic abstraction method that made the success of TVLA [6], and Navigation temporal logic [27], based on a tableau method quite similar to our automaton-based reduction. In these works, the assertion language for states is quite rich, as it includes, for instance, list predicate, quantification over addresses, and a freshness predicate. Because of this high expressive power, only incomplete abstractions are proposed, whereas we stick to exact methods. More importantly, our work addresses models with constant heaps and pointer arithmetic, which has not been done so far, and leads to a quite different perspective.

Structure of the paper. We define our logic LTL^{mem} and several fragments and problems in Section 2. Section 3 introduces the symbolic memory states (also useful in Section 4) and presents the PSPACE-completeness of the satisfiability and model-checking problems for SL with pointer arithmetic. Section 4 is dedicated to the decidability proof of satisfiability for various fragments and its consequences for other problems. In Section 5, we mention several seemingly optimal undecidability results by encoding computations of Minsky machines. Section 6 contains concluding remarks.

This paper is an extended version of [28].

2. Memory Model and Specification Language

In this section, we introduce a separation logic dealing with pointer arithmetic and record values, and a temporal logic LTL^{mem} . Unlike BI's pointer logic from [29], we allow pointer arithmetic.

Expressions	State Formulae
$e ::= \mathbf{x} \mid \mathbf{null}$	$\mathcal{A} ::= \pi$
Atomic formulae	$\mid \mathcal{A} * \mathcal{B} \mid \mathcal{A} \rightarrow \mathcal{B} \mid \mathbf{emp}$ (spatial fragment)
$\pi ::= e = e' \mid \mathbf{x} + i \stackrel{l}{\hookrightarrow} e$	$\mid \mathcal{A} \wedge \mathcal{B} \mid \mathcal{A} \rightarrow \mathcal{B} \mid \perp$ (classical fragment)
Satisfaction	
$(s, h) \models_{\text{SL}} e = e'$	iff $s(e) = s(e')$
$(s, h) \models_{\text{SL}} \mathbf{x} + i \stackrel{l}{\hookrightarrow} e$	iff $s(\mathbf{x}) \neq \mathit{nil}$ and $h(s(\mathbf{x}) + i)(l) = s(e)$
$(s, h) \models_{\text{SL}} \mathbf{emp}$	iff $\text{dom}(h) = \emptyset$
$(s, h) \models_{\text{SL}} \mathcal{A}_1 * \mathcal{A}_2$	iff $\exists h_1, h_2$ s.t. $h = h_1 * h_2$, $(s, h_1) \models_{\text{SL}} \mathcal{A}_1$ and $(s, h_2) \models_{\text{SL}} \mathcal{A}_2$
$(s, h) \models_{\text{SL}} \mathcal{A}_1 \rightarrow \mathcal{A}_2$	iff for all h' , if $h \perp h'$ and $(s, h') \models_{\text{SL}} \mathcal{A}_1$ then $(s, h * h') \models_{\text{SL}} \mathcal{A}_2$
$(s, h) \models_{\text{SL}} \mathcal{A}_1 \wedge \mathcal{A}_2$	iff $(s, h) \models_{\text{SL}} \mathcal{A}_1$ and $(s, h) \models_{\text{SL}} \mathcal{A}_2$
$(s, h) \models_{\text{SL}} \mathcal{A}_1 \rightarrow \mathcal{A}_2$	iff $(s, h) \models_{\text{SL}} \mathcal{A}_1$ implies $(s, h) \models_{\text{SL}} \mathcal{A}_2$
$(s, h) \models_{\text{SL}} \perp$	never

Table 1: The syntax and semantics of SL with pointer arithmetic and records

2.1. A separation logic with pointer arithmetic

Memory states. Let us introduce our model of memory. It captures features of programs with pointer variables that use pointer arithmetic and records. We assume a countably infinite set \mathbf{Var} of variables (as usual, for a fixed formula we need only a finite amount), and an infinite set \mathbf{Val} of values containing the set \mathbb{N} of naturals, thought as address indexes, and a special value nil . For simplicity, we assume that $\mathbf{Val} = \mathbb{N} \uplus \{\mathit{nil}\}$. In order to model field selectors, we consider an infinite set \mathbf{Lab} of labels. We will usually range over values with u, v , over naturals with i, j , over labels with $l, r, \mathit{next}, \mathit{prev}$, and over variables with \mathbf{x}, \mathbf{y} . In the remainder, we will assume a fixed injection $(\mathbf{x}, i) \in \mathbf{Var} \times \mathbb{N} \mapsto \langle \mathbf{x}, i \rangle \in \mathbf{Var}$.

We use the notation $E \rightarrow_{fin} F$ for the set of partial functions from E to F of finite domain, and \rightarrow_{fin+} the set of the ones of non-empty domain. The sets \mathcal{S} of stores and \mathcal{H} of heaps are then defined as follows:

$$\mathcal{S} \stackrel{\text{def}}{=} \mathbf{Var} \rightarrow \mathbf{Val} \quad \mathcal{H} \stackrel{\text{def}}{=} \mathbb{N} \rightarrow_{fin} (\mathbf{Lab} \rightarrow_{fin+} \mathbf{Val}).$$

We will range over a store with s, s' and over a heap with h, h', h_1, h_2 . We call *memory state* a couple $(s, h) \in \mathcal{S} \times \mathcal{H}$. Hence, a heap can be equivalently understood as a finite subset of $\mathbb{N} \times \mathbf{Lab} \times \mathbf{Val}$.

We will refer to the domain of a heap h by $\text{dom}(h) \subset \mathbb{N}$. Intuitively, in our memory model, each index is thought as an entry point on some record cell containing several fields. Cells are either not allocated, or allocated with some record stored in. In a memory state (s, h) , the memory cell at index i is *allocated* if $i \in \text{dom}(h)$; in this case the stored record is $h(i) = \{l_1 \mapsto v_1, \dots, l_n \mapsto v_n\}$.

Note that the size of the information held in a memory cell is neither fixed, nor bounded. Our models could be more concrete considering labels as offsets and relying on pointer arithmetic. However, for our classification of several problems, it will be useful to consider pointer arithmetic independently.

The size of the store s with respect to a finite set of variables $X \subseteq \mathbf{Var}$, written $\text{size}_X(s)$, is defined as $\text{card}(X) \times \max(1 + \log(1 + s(\mathbf{x})) : s(\mathbf{x}) \in \mathbb{N}, \mathbf{x} \in X)$. Similarly, the size of the heap h with respect to a finite set of labels $Y \subseteq \mathbf{Lab}$, written $\text{size}_Y(h)$, is defined as $\text{card}(\text{dom}(h)) \times \text{card}(Y) \times \max(1 + \log(1 + h(i)(l)) : i \in \text{dom}(h), h(i)(l) \text{ is defined and } h(i)(l) \in \mathbb{N})$. The size of the memory state (s, h) with respect to X and Y , written $\text{size}_{X,Y}((s, h))$, is $\text{size}_X(s) + \text{size}_Y(h)$.

Separation Logic. We now introduce the separation logic (SL) on top of which we will define our temporal logic. The syntax of the logic is given in Table 1.

In short, Separation logic is about reasoning on disjoint heaps, and we need to define what we mean by “disjoint heaps” in our model. Our level of granularity implies that a record cell cannot be decomposed in

disjoint parts. Let h_1 and h_2 be two heaps; we say that h_1 and h_2 are disjoint, noted $h_1 \perp h_2$, if $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$. The operation $h_1 * h_2$ is defined for disjoint heaps as the *disjoint union* of the two partial functions. Semantics of formulae is defined by the satisfaction relation \models_{SL} (see Table 1).

Formulae π are *atomic formulae*. The formula $\mathbf{x} + i \xrightarrow{l} e$ states that the value of the field l of the record stored at the address pointed by \mathbf{x} with offset i is equal to the value of the expression e . The formula $e = e'$ states the equality between the values of the two expressions, and emp means that the current heap has no memory cell allocated. In Table 1 (as well as in the rest of the paper), we assume that a store s is also defined for null with $s(\text{null}) = \text{nil}$. In some places, “ null ” is understood as a distinguished variable whose interpretation is fixed to the value $\text{nil} \in \text{Val} \setminus \mathbb{N}$.

Formulae \mathcal{A} of SL are called *state formulae*. A formula $\mathcal{A} * \mathcal{B}$ with the *separating conjunction* states that \mathcal{A} holds on some portion of the memory heap and \mathcal{B} holds on a disjoint portion. A formula $\mathcal{A} \text{--} * \mathcal{B}$ states that the current heap, when extended with any disjoint heap verifying \mathcal{A} , will verify \mathcal{B} . Boolean operators are understood as usual. Derivable connectives $\mathcal{A} \vee \mathcal{B}$ and $\neg \mathcal{A}$ are defined as usual. In the remainder, we focus on several specific fragments of this separation logic. We say that a formula is in the *record fragment* (RF) if all its subformulae of the form $\mathbf{x} + i \xrightarrow{l} e$ use $i = 0$. In that case, we write $\mathbf{x} \xrightarrow{l} e$. We say that a formula is in the *classical fragment* (CL) if it does not contain any of the connectives $*$ and $\text{--}*$. Finally, we say that a formula is in the *list fragment* (LF) if it is in the classical fragment and all subformulae $\mathbf{x} + i \xrightarrow{l} e$ use $i = 0$ and $l = \text{next}$, and we may simply write $\mathbf{x} \xrightarrow{l} e$. Clearly, the classical and record fragments are incomparable, while the list fragment is included in both of them.

Let us illustrate the expressive power on simple examples. The formula $\neg \text{emp} * \neg \text{emp}$ means that at least two memory cells are allocated. The formula $\mathbf{x} \xrightarrow{l} e$, defined as $\neg(\neg \text{emp} * \neg \text{emp}) \wedge \mathbf{x} \xrightarrow{l} e$, is the local version of $\mathbf{x} \xrightarrow{l} e$: $(s, h) \models_{\text{SL}} \mathbf{x} \xrightarrow{l} e$ iff $\text{dom}(h) = \{s(\mathbf{x})\}$ and $h(s(\mathbf{x}))(l) = s(e)$. The formula $(\mathbf{x} \xrightarrow{l} \text{null}) \text{--} * \perp$ is satisfied by (s_0, h_0) whenever there is no heap h_1 with $h_1 \perp h_0$ that allocates the variable \mathbf{x} to nil on l field. In other words, the variable \mathbf{x} is already allocated in the heap h_0 .

As usual, \mathcal{A} is valid iff for every memory state (s, h) , we have $(s, h) \models_{\text{SL}} \mathcal{A}$ (written $\models_{\text{SL}} \mathcal{A}$). Satisfiability is defined dually: \mathcal{A} is satisfiable iff there is a memory state (s, h) , such that $(s, h) \models_{\text{SL}} \mathcal{A}$.

The size of the state formula \mathcal{A} , written $|\mathcal{A}|$, is the length of the string \mathcal{A} for some reasonably succinct encoding of variables and integers with a binary representation. We will use the map $|\cdot|$ for other syntactic objects such as LTL^{mem} formulae.

2.2. Temporal extension

Memory states sequences. Models of the logic LTL^{mem} are ω -sequences of memory states, which means there are elements in $(\mathcal{S} \times \mathcal{H})^\omega$ and they are understood as infinite computations of programs with pointer variables. We range over ρ for a given model, and its i^{th} state $\rho(i)$ will be noted (s_i, h_i) . In order to analyze computations from programs without destructive update, we shall also consider models with constant heap, that is elements in $\mathcal{S}^\omega \times \mathcal{H}$.

The logic LTL^{mem} . Formulae of LTL^{mem} are defined in Table 2. Atomic formulae of LTL^{mem} are state formulae from SL except that variables can be prefixed by the symbol “ \mathbf{X} ”. For instance, $\mathbf{X}\mathbf{x}$ is interpreted

by the value of \mathbf{x} at the next memory state. We use the notation $\mathbf{X}^i \mathbf{x}$ for $\overbrace{\mathbf{X} \dots \mathbf{X}}^{i \text{ times}} \mathbf{x}$ (but keep in mind that encoding $\mathbf{X}^i \mathbf{x}$ requires memory space in $\mathcal{O}(i)$). The temporal operators are the standard next-time operator \mathbf{X} and until operator \mathbf{U} present in LTL, see e.g. [30, 20]. The satisfaction relation $\rho, t \models \phi$, where ρ is a model of LTL^{mem} , $t \in \mathbb{N}$ and ϕ is a formula, is also defined in Table 2. We use standard abbreviations such as $\mathbf{F}\phi$ for $\mathbf{T}\mathbf{U}\phi$ or $\mathbf{G}\phi$ for $\neg \mathbf{F}(\neg \phi)$. We freely use propositional variables p, q , having in mind that the propositional variable p should be understood as $\mathbf{x}_p = \mathbf{x}_\top$ for some fixed extra variables $\mathbf{x}_p, \mathbf{x}_q, \dots, \mathbf{x}_\top$. In the sequel, given an atomic formula \mathcal{A} , we write $\mathcal{A}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle]$ to denote the SL state formula in which every occurrence of a term of the form $\mathbf{X}^u \mathbf{x}$ is replaced by the variable $\langle \mathbf{x}, u \rangle$. Similarly, given a state formula \mathcal{A} , we write $\mathcal{A}[\mathbf{x} \leftarrow \langle \mathbf{x}, 0 \rangle]$ to denote the state formula in which every occurrence of a variable \mathbf{x} is replaced by $\langle \mathbf{x}, 0 \rangle$.

Enriched expressions	$\eta ::= \mathbf{x} \mid X\eta \quad \xi ::= \eta \mid \mathbf{null}$
Atomic formulae	$\pi ::= \xi = \xi' \mid \eta + i \xrightarrow{l} \xi$
State formulae	$\mathcal{A} ::= \pi \mid \mathbf{emp} \mid \mathcal{A} * \mathcal{B} \mid \mathcal{A} \text{-} * \mathcal{B} \mid \mathcal{A} \wedge \mathcal{B} \mid \mathcal{A} \rightarrow \mathcal{B} \mid \perp$
Temporal formulae	$\phi ::= \mathcal{A} \mid X\phi \mid \phi U \psi \mid \phi \wedge \psi \mid \neg\phi$
Semantics	
$\rho, t \models X\phi$	iff $\rho, t + 1 \models \phi$.
$\rho, t \models \phi U \psi$	iff there is $t_1 \geq t$ s.t. $\rho, t_1 \models \psi$ and $\rho, t' \models \phi$ for all $t' \in [t, t_1[$.
$\rho, t \models \phi \wedge \psi$	iff $\rho, t \models \phi$ and $\rho, t \models \psi$.
$\rho, t \models \phi \rightarrow \psi$	iff $\rho, t \models \phi$ implies $\rho, t \models \psi$.
$\rho, t \models \mathcal{A}$	iff $s'_t, h_t \models_{\text{SL}} \mathcal{A}[X^k \mathbf{x} \leftarrow \langle \mathbf{x}, k \rangle]$ where $\rho = (s_t, h_t)_{t \geq 0}$ and s'_t is defined by $s'_t(\langle \mathbf{x}, k \rangle) = s_{t+k}(\mathbf{x})$.

Table 2: The syntax and semantics of LTL^{mem}

Given a fragment *Frag* of SL, such as RF or LF, $\text{LTL}^{\text{mem}}(\text{Frag})$ is the restriction of LTL^{mem} to formulae in which occur only state formulae built over *Frag* (with extended variables $X^u \mathbf{x}$), and we write $\text{SAT}(\text{Frag})$ to denote the satisfiability problem for $\text{LTL}^{\text{mem}}(\text{Frag})$: given a temporal formula ϕ in $\text{LTL}^{\text{mem}}(\text{Frag})$, is there a model ρ such that $\rho, 0 \models \phi$? The variant problem in which we require that the model has a constant heap [resp. that the initial memory state is fixed] is denoted by $\text{SAT}^{ct}(\text{Frag})$ [resp. $\text{SAT}_{init}(\text{Frag})$]. The problem $\text{SAT}_{init}^{ct}(\text{Frag})$ is defined analogously.

2.3. Programs with pointer variables

In this section, we define the model-checking problems for programs with pointer variables over LTL^{mem} specifications. The set \mathbf{I} of *instructions* used in the programs is defined by the grammar below:

$$\begin{aligned} \mathbf{instr} ::= & \mathbf{x} := \mathbf{y} \mid \mathbf{skip} \\ & \mid \mathbf{x} := \mathbf{y} \rightarrow l \mid \mathbf{x} \rightarrow l := \mathbf{y} \mid \mathbf{x} := \mathbf{cons}(l_1 : \mathbf{x}_1, \dots, l_k : \mathbf{x}_k) \mid \mathbf{free} \mathbf{x} \\ & \mid \mathbf{x} := \mathbf{y}[i] \mid \mathbf{x}[i] := \mathbf{y} \mid \mathbf{x} := \mathbf{malloc}(i) \mid \mathbf{free} \mathbf{x}, i \end{aligned}$$

The denotational semantics of an instruction \mathbf{instr} is defined as a binary relation $(\mathcal{S} \times \mathcal{H}) \times (\mathcal{S} \times \mathcal{H}) \subseteq \llbracket \mathbf{instr} \rrbracket$ in order to deal with the nondeterministic allocation of new memory cells. We list in Table 3 the formal denotational semantics of our instruction set. Observe that the instructions $\mathbf{x} := \mathbf{y}[i]$, $\mathbf{x} := \mathbf{malloc}(i)$ and $\mathbf{x}[i] := \mathbf{y}$ deal with the specific label *next*. Boolean combinations of equalities between expressions are called guards and its set is denoted by G . A program is defined as a triple (Q, δ, q_I) such that Q is a finite set of control states, q_I is the initial state and δ is the transition relation, a subset of $Q \times G \times \mathbf{I} \times Q$. We use $q \xrightarrow{g, \mathbf{instr}} q'$ to denote a transition. We say that a program is *without destructive update* if transitions are labeled only with instructions of the form $\mathbf{x} := \mathbf{y}$, $\mathbf{x} := \mathbf{y} \rightarrow l$, and $\mathbf{x} := \mathbf{y}[i]$. We write \mathbf{P} to denote the set of programs and \mathbf{P}^{ct} to denote the set of programs without destructive update.

A program is a finite object whose interpretation can be viewed as an infinite-state system. More precisely, given a program $\mathbf{p} = (Q, \delta, q_I)$, the transition system $\mathcal{S}_{\mathbf{p}} = (S, \rightarrow)$ is defined as follows: $S = Q \times (\mathcal{S} \times \mathcal{H})$ (set of configurations) and $(q, (s, h)) \rightarrow (q', (s', h'))$ iff there is a transition $q \xrightarrow{g, \mathbf{instr}} q' \in \delta$ such that $(s, h) \models g$ and $(s', h') \in \llbracket \mathbf{instr} \rrbracket(s, h)$. Note that $\mathcal{S}_{\mathbf{p}}$ is not necessarily linear. A computation (or execution) of \mathbf{p} is defined as an infinite path in $\mathcal{S}_{\mathbf{p}}$ starting with control state q_I .

Computations of \mathbf{p} can be viewed as LTL^{mem} models, using propositional variables to encode the extra information about the control states (details are omitted herein).

Model-checking aims at checking properties expressible in LTL^{mem} along computations of programs. To a logical fragment (SL, CL, RF, or LF), we associate a set of programs: all programs for SL and CL, programs with instructions having $i = 0$ for RF, and moreover with only the label *next* for LF. Given one

$\llbracket \mathbf{x} := \mathbf{y} \rrbracket (s, h)$	$\ni (s[\mathbf{x} \mapsto s(\mathbf{y})], h).$
$\llbracket \mathbf{x} := \mathbf{y} \rightarrow l \rrbracket (s, h * \{i \mapsto \{l \mapsto v, \dots\}\})$	$\ni (s[\mathbf{x} \mapsto v], h * \{i \mapsto \{l \mapsto v, \dots\}\})$ with $s(\mathbf{y}) = i$
$\llbracket \mathbf{x} \rightarrow l := \mathbf{y} \rrbracket (s, h * \{i \mapsto \{l \mapsto v, \dots\}\})$	$\ni (s, h * \{i \mapsto \{l \mapsto s(\mathbf{y}), \dots\}\})$ with $s(\mathbf{x}) = i$
$\llbracket \mathbf{x} := \mathbf{cons}(l_1 : \mathbf{x}_1, \dots, l_k : \mathbf{x}_k) \rrbracket (s, h)$	$\ni (s[\mathbf{x} \mapsto i], h * \{i \mapsto \{l_1 \mapsto s(\mathbf{x}_1), \dots, l_k \mapsto s(\mathbf{x}_k)\}\})$ with $i \notin \mathbf{dom}(h)$
$\llbracket \mathbf{free} \mathbf{x} \rrbracket (s, h * \{i \mapsto \{\dots\}\})$	$\ni (s, h)$ with $s(\mathbf{x}) = i$
$\llbracket \mathbf{skip} \rrbracket (s, h)$	$\ni (s, h)$
$\llbracket \mathbf{x} := \mathbf{y}[i] \rrbracket (s, h * \{i + i' \mapsto \{next \mapsto v\}\})$	$\ni (s[\mathbf{x} \mapsto v], h * \{i + i' \mapsto \{next \mapsto v\}\})$ with $s(\mathbf{y}) = i'$
$\llbracket \mathbf{x}[i] := \mathbf{y} \rrbracket (s, h * \{i' + i \mapsto \{next \mapsto v\}\})$	$\ni (s, h * \{i' + i \mapsto \{next \mapsto s(\mathbf{y})\}\})$ with $s(\mathbf{x}) = i'$
$\llbracket \mathbf{x} := \mathbf{malloc}(i) \rrbracket (s, h)$	$\ni (s[\mathbf{x} \mapsto i'], h * \{i' \mapsto \{next \mapsto nil\}, \dots, i' + (i - 1) \mapsto \{next \mapsto nil\}\})$ with $i', \dots, i' + (i - 1) \notin \mathbf{dom}(h)$
$\llbracket \mathbf{free} \mathbf{x}, i \rrbracket (s, h * \{i' + i \mapsto f\})$	$\ni (s, h)$ with $s(\mathbf{x}) = i'$

Table 3: Semantics for instructions

of these fragments Frag of SL, we write $\text{MC}(\text{Frag})$ to denote the model-checking problem for Frag: given a temporal formula ϕ in LTL^{mem} with state formulae built over Frag and a program \mathbf{p} of the associated set of programs, is there an infinite computation ρ of \mathbf{p} such that $\rho, 0 \models \phi$ (which we write $\mathbf{p} \models \phi$)? This is the existential variant of the problem. The variant problem in which we require that the program is without destructive update [resp. that the initial memory state is fixed, say (s, h)] is denoted by $\text{MC}^{\text{ct}}(\text{Frag})$ [resp. $\text{MC}_{\text{init}}^{\text{ct}}(\text{Frag})$]. The problem $\text{MC}_{\text{init}}^{\text{ct}}(\text{Frag})$ is defined analogously. We may write $\mathbf{p}, (s, h) \models \phi$ to emphasize what is the initial memory state.

Basic results. Using extended variables \mathbf{Xx} , we may express some programs as formulae. This actually holds only for programs without destructive update, for the semantics with constant heap. Intuitively, we express the control of the program with propositional variables, and define a formula that encodes the transitions. As a consequence, the following result can be derived.

Lemma 1. *Let Frag be a fragment among SL, CL, RF, or LF. There is a logspace reduction from $\text{MC}^{\text{ct}}(\text{Frag})$ to $\text{SAT}^{\text{ct}}(\text{Frag})$ (resp. from $\text{MC}_{\text{init}}^{\text{ct}}(\text{Frag})$ to $\text{SAT}_{\text{init}}^{\text{ct}}(\text{Frag})$).*

Proof. We adapt the proof in [20] for reducing LTL model-checking to LTL satisfiability. To a program $\mathbf{p} = (Q, \delta, q_I)$, we associate the formula $\phi_{\mathbf{p}}$ below built over the propositional variables in Q :

$$q_I \wedge \mathbf{G} \bigwedge_{q \in Q} (q \Rightarrow (\bigwedge_{q' \in Q \setminus \{q\}} \neg q' \wedge \bigvee_{\tau \in \delta_q^+} \phi_{\tau}))$$

where ϕ_{τ} expresses that transition τ is fired between the current state and the next state and, δ_q^+ is the set of transitions starting at the state q . In order to define ϕ_{τ} , we need to translate instructions and guards into the logic (remember that there are limitations on the instructions). We translate instructions of the form

- $\mathbf{x} := \mathbf{y}$ into $\mathbf{Xx} = \mathbf{y}$,

- $x := y \rightarrow l$ into $y \xrightarrow{l} Xx$,
- $x := y[i]$ into $y + i \xrightarrow{next} Xx$.

Guards are translated accordingly. It is then standard to show that $p \models \phi$ iff $\phi \wedge \phi_p$ is satisfiable. \square

All the model-checking and satisfiability problems defined in this paper belong to Σ_1^1 in the analytical hierarchy. Indeed, the models and computations of programs can be viewed as functions $f : \mathbb{N} \rightarrow \mathbb{N}$ by encoding memory states and configurations by natural numbers (details are tedious). Then, the satisfaction relation between models and LTL^{mem} formulae and the transition relations obtained from programs can be encoded by a first-order formula. This guarantees that these problems are in Σ_1^1 . Additionally, all the problems can easily be shown PSPACE-hard since they all generalize LTL satisfiability and model-checking [20].

2.4. Discussion

Let us discuss few issues about the expressive power of this logical formalism. First, the interest of model-checking programs with heap updates stems from early works on automata-based verification. Decision procedures are obtained at the cost of limitations: to define approximations as done in [26, 27] or to restrict the programming language, see e.g. [31]. However, with this approach, compositionality principles are lost, which is a pity since they made the success of separation logic, as the frame rule and the composition rule.

Second, assuming that the heap is constant is subject to promising developments. Indeed, it is then possible to define spatial operators at the same syntactic level as temporal operators, and write formulae as e.g. $[(x \leftrightarrow Xx)U(x \leftrightarrow \text{null})] * (y \mapsto \text{null})$. Observe that this formula does not belong to LTL^{mem} . This might be a way to model modularity in model-checking programs without destructive updates, but there are other points of interest we will try to advocate now.

2.4.1. Recursion with local parameters

The constant heap semantics provides an original viewpoint for recursion with local parameters and local quantification. The design of decision procedures in the presence of general recursive predicates was introduced in [12], as well as incomplete methods of inference even though they are apparently good in practice. Complete methods have been proposed for some standard recursive structures like trees, lists, or doubly-linked lists [32]. But we are not aware of complete methods for a general form of recursive data structures defined on top of Separation Logic, and we believed that our logic could give an alternative way of specifying recursion, although we did not manage to characterize an interesting decidable fragment.

In order to be a bit more precise, let us consider the fragment of recursive separation logic where all recursive formulae are of the form:

$$(1) \quad \mu X(x_1, \dots, x_k). A(x_1, \dots, x_k) \vee \exists x'_1 \dots x'_k. B(x_1, \dots, x_k, x'_1, \dots, x'_k) \wedge X(x'_1, \dots, x'_k)$$

This fragment is rich enough to express single lists, cyclic lists, and doubly-linked lists. However, we conjecture that it is not expressive enough for trees and DAGs. We conjecture that deciding satisfiability in the fragment of recursive separation logic mentioned above reduces to $SAT^{ct}(\text{SL})$, and the model-checking problem reduces to SAT_{init}^{ct} , considering that (1) can be rewritten as:

$$(B(x_1, \dots, x_k, Xx_1, \dots, Xx_k)) \cup A(x_1, \dots, x_k).$$

In this perspective, our results could arise interesting decidability results for model-checking some of the recursive separation logic with local quantifiers. For satisfiability, we expect to define decidable fragments for $SAT^{ct}(\text{SL})$, for instance considering the techniques for checking temporal properties of so-called flat programs without destructive updates introduced in [33]. Another interesting fragment of recursive separation logic is probably the one where recursion is guarded by the separation operator $*$, but we do not currently see how to treat it in the temporal logic perspective.

2.4.2. Programs as formulae

Let us speculate a bit more. We may take advantage of expressing programs as formulae in order to reduce model-checking to satisfiability, a known approach since [20]. For programs without destructive update, we take advantage of Lemma 1. Moreover, we believe we can extend this result to programs with updates, but with a slightly different perspective. The constant heap semantics can be helpful to define the input-output relation of programs, even with destructive updates, provided some conditions on the way the program read and write over the memory are satisfied. To do so, we consider the extension of LTL^{mem} with two predicates \hookrightarrow_0 and \hookrightarrow_1 instead of the single \hookrightarrow , and models are couples of state sequences with constant heap, that is tuples $\langle (s_i)_{i \geq 0}, h_0, h_1 \rangle$. Let us define the input-output relation IO_p of a program p as : for all $(s_0, h_0), (s_1, h_1), (s_0, h_0) \text{IO}_p (s_1, h_1)$ if there is a run of p that starts with (s_0, h_0) and ends with (s_1, h_1) . Then we conjecture that for an interesting class of programs, this relation is definable in LTL^{mem} extended with \hookrightarrow_0 and \hookrightarrow_1 . Basically, the encoding of the control of the program will be the same as for programs without destructive updates, but the encoding of the instructions will be different. For instance, $\mathbf{x} \rightarrow l := \mathbf{y}$ would be encoded by $(\mathbf{Xx}) \xrightarrow{l} \mathbf{y}$ whereas $\mathbf{x} := \mathbf{y} \rightarrow l$ would be encoded as $\mathbf{y} \xrightarrow{l} \mathbf{Xx}$.

3. Separation Logic: Complexity and Abstraction

After defining an abstraction for the fragment RF of SL, which will be proved sound, we will be able to decide the complexity of model checking and satisfiability problems for SL.

3.1. Syntactic measures

The main approach to get decision procedures to verify infinite-state systems consists of introducing a symbolic representation for infinite sets of configurations. The symbolic representation defined below is motivated by a similar goal and it has similarities with symbolic heaps for Separation Logic in Smallfoot [11]. Let us start by some useful definitions. Following [22], we introduce the set of *test formulae* that are SL formulae of the forms below:

- $\text{alloc}(\mathbf{x} + i) \stackrel{\text{def}}{=} (\mathbf{x} + i \xrightarrow{\text{next}} \text{null}) * \perp$ ($\mathbf{x} + i$ is allocated).
- $\text{size} \geq k \stackrel{\text{def}}{=} \overbrace{\neg \text{emp} * \dots * \neg \text{emp}}^{k \text{ times}}$ (at least k indices are allocated).
- $\mathbf{x} + i \xrightarrow{l} e, e = e'$ (see Table 1 for notations).

Given a formula ϕ of LTL^{mem} , we define its measure μ_ϕ , understood as pieces of information about the syntactic resources involved in ϕ . Indeed, forthcoming symbolic states are finite objects parameterized by such syntactic measures. For a state formula \mathcal{A} of LTL^{mem} , the size of memory potentially examined by \mathcal{A} , written $w_{\mathcal{A}}$, is inductively defined as follows: $w_{\mathcal{A}}$ is 1 for atomic formulae, $\max(w_{\mathcal{A}_1}, w_{\mathcal{A}_2})$ for $\mathcal{A}_1 \wedge \mathcal{A}_2$ or $\mathcal{A}_1 \rightarrow \mathcal{A}_2$ or $\mathcal{A}_1 * \mathcal{A}_2$, and $w_{\mathcal{A}_1} + w_{\mathcal{A}_2}$ for $\mathcal{A}_1 * \mathcal{A}_2$. Observe that $w_{\mathcal{A}} \leq |\mathcal{A}|$. Other simple sets about the syntactic resources of \mathcal{A} need to be defined: $\text{Lab}_{\mathcal{A}} (\in \mathcal{P}_f(\text{Lab}))$ is the set of labels from Lab occurring in \mathcal{A} , $\text{Var}_{\mathcal{A}} (\in \mathcal{P}_f(\text{Var}))$ is the set of variables from Var occurring in \mathcal{A} , $\text{offsets}_{\mathcal{A}} (\in \mathcal{P}_f(\mathbb{N}))$ is the set of natural numbers i such that $\mathbf{X}^u \mathbf{x} + i \xrightarrow{l} e$ occurs in \mathcal{A} , where $\mathcal{P}_f(X)$ denotes the set of finite subsets of some set X .

Definition 2. A measure μ is defined as a tuple $(\text{offsets}_{\mu}, w_{\mu}, \text{Lab}_{\mu}, \text{Var}_{\mu}) \in \mathcal{P}_f(\mathbb{N}) \times \mathbb{N} \times \mathcal{P}_f(\text{Lab}) \times \mathcal{P}_f(\text{Var})$.

The set of measures has a natural lattice structure for the pointwise order, noted below $\mu \leq \mu'$. We also write $\mu[w \leftarrow 0]$ to denote the measure μ except that the second component w is 0. The measure for \mathcal{A} , written $\mu_{\mathcal{A}}$, is the tuple $(\text{offsets}_{\mathcal{A}}, w_{\mathcal{A}}, \text{Lab}_{\mathcal{A}}, \text{Var}_{\mathcal{A}})$. The measure of some formula ϕ of LTL^{mem} , written μ_{ϕ} , is $\sup\{\mu_{\mathcal{A}}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle] : \mathcal{A} \text{ occurs in } \phi\}$. We write $\text{size}(\mu)$ to denote the size of the measure μ in some reasonable succinct encoding.

Definition 3. Given a measure $\mu = (\text{offsets}, w, \text{Lab}_\mu, \text{Var}_\mu)$, we write \mathcal{T}_μ to denote the finite set of test formulae ψ defined as follows:

$$\begin{aligned} e &::= \mathbf{x} \mid \text{null} & f &::= \mathbf{x} + i \\ \psi &::= f \xrightarrow{l} e \mid \text{alloc}(f) \mid e = e' \mid \text{size} \geq k \end{aligned}$$

with $i \in \text{offsets}$, $l \in \text{Lab}_\mu$, $k \in [0, w[$ and $\mathbf{x} \in \text{Var}_\mu$.

Observe that the cardinal of \mathcal{T}_{μ_ϕ} is polynomial in $|\phi|$. The variable $\langle \mathbf{x}, u \rangle$ will be used in subsequent developments to deal with the interpretation of the term $\mathbf{X}^u \mathbf{x}$ in the formulae of the temporal logic. Given a measure $\mu = (\text{offsets}, w, \text{Lab}_\mu, \text{Var}_\mu)$ and a memory state (s, h) , we write $\text{Abs}_\mu(s, h) = \{\mathcal{A} \in \mathcal{T}_\mu : (s, h) \models_{\text{SL}} \mathcal{A}\}$ to denote the abstraction of (s, h) with respect to μ . Given a measure μ and two memory states (s, h) and (s', h') , we write $(s, h) \simeq_\mu (s', h')$ iff $\text{Abs}_\mu(s, h) = \text{Abs}_\mu(s', h')$, that is, formulae in \mathcal{T}_μ cannot distinguish the two memory states.

The proof of Lemma 7 below is based on three technical lemmas. Before stating them and proving them, in Lemmas 4-6, we assume that the measure has $\text{offsets} = \{0\}$ since we are dealing with RF. Moreover, we introduce the following definition: $(\text{offsets}, w, \text{Lab}_0, \text{Var}_0) + (\text{offsets}, w', \text{Lab}_0, \text{Var}_0) = (\text{offsets}, w + w', \text{Lab}_0, \text{Var}_0)$.

Lemma 4 (Distributivity). Let μ, μ_1 and μ_2 be measures with all sets of offsets equal to $\{0\}$ and $\mu = \mu_1 + \mu_2$. Let (s, h) and (s', h') be memory states such that $(s, h) \simeq_\mu (s', h')$ and, h_1, h_2 be heaps such that $h = h_1 * h_2$. Then, there exist heaps h'_1 and h'_2 with $h' = h'_1 * h'_2$, $(s, h_1) \simeq_{\mu_1} (s', h'_1)$ and $(s, h_2) \simeq_{\mu_2} (s', h'_2)$.

Proof. Let $(s, h), (s', h'), h_1, h_2$ and $\mu = (\text{offsets}, w, \text{Lab}_0, \text{Var}_0)$, $\mu_1 = (\text{offsets}, w_1, \text{Lab}_0, \text{Var}_0)$ and $\mu_2 = (\text{offsets}, w_2, \text{Lab}_0, \text{Var}_0)$ satisfying the hypotheses of the lemma.

We shall define the disjoint heaps h'_1 and h'_2 by distinguishing the four disjoint sets of values S_1, S_2, A_1 and A_2 corresponding to the following sets:

- $S_1 = \text{dom}(h'_1) \cap \text{Im}(s'), S_2 = \text{dom}(h'_2) \cap \text{Im}(s')$,
- $A_1 = \text{dom}(h'_1) \setminus \text{Im}(s'), A_2 = \text{dom}(h'_2) \setminus \text{Im}(s')$.

Let us first separate $\text{dom}(h') \cap \text{Im}(s')$ into two parts S_1 and S_2 . We define $S_i = s'(s^{-1}(\text{dom}(h_i)) \cap \text{Var}_\mu)$, and we need to show that S_1, S_2 are disjoint. Let us assume by contradiction that they are not, thus there are some variables $\mathbf{x}, \mathbf{y} \in \text{Var}_\mu$ such that $s'(\mathbf{x}) = s'(\mathbf{y}) \in S_1 \cap S_2$, and $s(\mathbf{x}) \in \text{dom}(h_1)$ whereas $s(\mathbf{y}) \in \text{dom}(h_2)$. Since $h_1 \perp h_2$, $s(\mathbf{x}) \neq s(\mathbf{y})$, so $s, h \models_{\text{SL}} \mathbf{x} = \mathbf{y}$, but $s', h' \not\models_{\text{SL}} \mathbf{x} = \mathbf{y}$, hence the contradiction.

Now, we shall separate the set $\text{dom}(h') \setminus \text{Im}(s')$ into two parts A_1 and A_2 . Let $B = \text{dom}(h) \setminus \text{Im}(s)$, $B_1 = \text{dom}(h_1) \setminus \text{Im}(s)$ and $B_2 = \text{dom}(h_2) \setminus \text{Im}(s)$. We have $\text{card}(B_1) + \text{card}(S_1) = \text{card}(\text{dom}(h_1))$ and $\text{card}(B_2) + \text{card}(S_2) = \text{card}(\text{dom}(h_2))$. The sets A_1 and A_2 shall contain respectively $\text{card}(A_1)$ and $\text{card}(A_2)$ random elements of $\text{dom}(h') \setminus \text{Im}(s')$ so that $\text{card}(A_i) = \text{card}(\text{dom}(h_i)) - \text{card}(S_i)$ if $\text{card}(\text{dom}(h)) < w$; otherwise $\text{card}(A_i) = \min(w_i, \text{card}(\text{dom}(h_i))) - \text{card}(S_i)$. In order to select the elements of A_1 and A_2 , we distinguish different cases depending on $\text{card}(\text{dom}(h_1))$ and $\text{card}(\text{dom}(h_2))$.

Case 1: $\text{card}(\text{dom}(h)) < w$.

Since $(s, h) \simeq_\mu (s', h')$, we have $\text{card}(\text{dom}(h')) = \text{card}(\text{dom}(h))$. Hence, $\text{dom}(h') \setminus \text{Im}(s')$ can be divided into two parts A_1, A_2 such that $A_1 \uplus A_2 = \text{dom}(h') \setminus \text{Im}(s')$, $\text{card}(A_1) = \text{card}(\text{dom}(h_1)) - \text{card}(S_1)$ and $\text{card}(A_2) = \text{card}(\text{dom}(h_2)) - \text{card}(S_2)$.

Case 2: $\text{card}(\text{dom}(h)) \geq w$.

Consequently $\text{card}(\text{dom}(h')) \geq w$.

Case 2.1: $\text{card}(\text{dom}(h_1)) \geq w_1$ and $\text{card}(\text{dom}(h_2)) \geq w_2$.

There exist A_1, A_2 such that $A_1 \uplus A_2 = \text{dom}(h') \setminus \text{Im}(s')$, $\text{card}(A_1) = \text{card}(\text{dom}(h_1)) - \text{card}(S_1) \geq w_1 - \text{card}(S_1)$ and $\text{card}(A_2) = \text{card}(\text{dom}(h_2)) - \text{card}(S_2) \geq w_2 - \text{card}(S_2)$.

Case 2.2: For some $i \in \{1, 2\}$, $\text{card}(\text{dom}(h_i)) < w_i$ and $\text{card}(\text{dom}(h_{3-i})) \geq w_{3-i}$.

There exist A_1, A_2 such that $A_1 \uplus A_2 = \text{dom}(h') \setminus \text{Im}(s')$, $\text{card}(A_i) = \text{card}(\text{dom}(h_i)) - \text{card}(S_i)$.

Then $\text{card}(A_{3-i}) = \text{card}(\text{dom}(h_{3-i})) - \text{card}(S_{3-i}) \geq w_{3-i} - \text{card}(S_{3-i})$.

The heap h'_1 is defined as $h'_{|_{A_1 \cup S_1}}$ and the heap h'_2 is defined as $h'_{|_{A_2 \cup S_2}}$. Since A_1, A_2, S_1 and S_2 are disjoint sets, we have that $A_1 \cup S_1$ and $A_2 \cup S_2$ are disjoint. Moreover, $A_1 \cup A_2 \cup S_1 \cup S_2 = \text{dom}(h')$. So $h' = h'_1 * h'_2$. Observe that for $i \in \{1, 2\}$, we have $\text{card}(\text{dom}(h_i)) \geq w_i$ iff $\text{card}(\text{dom}(h'_i)) \geq w_i$ and, $\text{card}(\text{dom}(h_i)) < w_i$ implies $\text{card}(\text{dom}(h_i)) = \text{card}(\text{dom}(h'_i))$.

It remains to show that $(s, h_1) \simeq_{\mu_1} (s', h'_1)$ and $(s, h_2) \simeq_{\mu_2} (s', h'_2)$. The above considerations about cardinality entail that for all $i \in \{1, 2\}$ and $k < w_i$, we have $\text{size} \geq k \in \text{Abs}_{\mu_i}(s, h_i)$ iff $\text{size} \geq k \in \text{Abs}_{\mu_i}(s', h'_i)$. It is also easy to check that for all $e = e', \text{alloc}(x) \in \mathcal{T}_{\mu_i}$, $e = e' \in \text{Abs}_{\mu_i}(s, h_i)$ iff $e = e' \in \text{Abs}_{\mu_i}(s', h'_i)$ and $\text{alloc}(x) \in \text{Abs}_{\mu_i}(s, h_i)$ iff $\text{alloc}(x) \in \text{Abs}_{\mu_i}(s', h'_i)$.

It remains to consider test formulae of the form $\mathbf{x} \xrightarrow{l} e$. The following statements are equivalent:

- $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_{\mu_i}(s, h_i)$,
- $s(\mathbf{x}) \in \mathbb{N}$ and $h_i(s(\mathbf{x}))(l) = s(e)$,
- $s(\mathbf{x}) \in \mathbb{N} \cap \text{dom}(h_i)$ and $h(s(\mathbf{x}))(l) = s(e)$,
- $s(\mathbf{x}) \in \text{dom}(h_i) \cap \text{Im}(s)$ and $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_{\mu}(s, h)$,
- $s'(\mathbf{x}) \in \text{dom}(h'_i) \cap \text{Im}(s')$ and $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_{\mu}(s', h')$,
- $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_{\mu_i}(s', h'_i)$.

□

In the proof of Lemma 7, we need Lemma 5 below, which is indeed an instance of Lemma 9.

Lemma 5. *Let μ be a measure such that $\text{offsets}_{\mu} = \{0\}$. If $(s, h) \simeq_{\mu} (s', h')$, then for all $h_0 \perp h$, there is $h'_0 \perp h'$ such that $(s, h_0) \simeq_{\mu} (s', h'_0)$.*

Lemma 6 (Congruence). *Let $(s, h_0), (s', h'_0), (s, h_1), (s', h'_1)$ be memory states such that $h_0 \perp h_1, h'_0 \perp h'_1$. Let μ be a measure such that $\text{offsets}_{\mu} = \{0\}$, and assume that $(s, h_0) \simeq_{\mu} (s', h'_0)$ and $(s, h_1) \simeq_{\mu} (s', h'_1)$. Then, $(s, h_0 * h_1) \simeq_{\mu} (s', h'_0 * h'_1)$.*

Proof. Let μ be the measure $(\text{offsets}, w, \text{Lab}_{\mu}, \text{Var}_{\mu})$. We shall show that $(s, h_0 * h_1) \simeq_{\mu} (s', h'_0 * h'_1)$. By symmetry of \simeq_{μ} , it is sufficient to prove that $\text{Abs}_{\mu}(s, h_0 * h_1) \subseteq \text{Abs}_{\mu}(s', h'_0 * h'_1)$. Let $\mathcal{A} \in \text{Abs}_{\mu}(s, h_0 * h_1)$. We make a case analysis according to \mathcal{A} .

- If $\mathcal{A} = \text{size} \geq k$ for some $k < w$, then $k \leq \text{card}(\text{dom}(h_0 * h_1))$. We want to show that $k \leq \text{card}(\text{dom}(h'_0 * h'_1))$ which implies that $\mathcal{A} \in \text{Abs}_{\mu}(s', h'_0 * h'_1)$.
 - If $\text{card}(\text{dom}(h_1)) \geq w$ or $\text{card}(\text{dom}(h_0)) \geq w$, then $\text{card}(\text{dom}(h'_1)) \geq w$ or $\text{card}(\text{dom}(h'_0)) \geq w$, respectively. So $\text{card}(\text{dom}(h'_0 * h'_1)) \geq w$ and $\text{card}(\text{dom}(h'_0 * h'_1)) \geq k$ as $k < w$.
 - If $\text{card}(\text{dom}(h_1)) < w$ and $\text{card}(\text{dom}(h_0)) < w$, then $\text{card}(\text{dom}(h_0 * h_1)) = \text{card}(\text{dom}(h_1)) + \text{card}(\text{dom}(h_0)) = \text{card}(\text{dom}(h'_1)) + \text{card}(\text{dom}(h'_0)) = \text{card}(\text{dom}(h'_0 * h'_1))$. So $k \leq \text{card}(\text{dom}(h'_0 * h'_1))$.
- If \mathcal{A} is $e = e'$, then $s(e) = s(e')$. Moreover, $\mathcal{A} \in \text{Abs}_{\mu}(s, h_1)$ iff $\mathcal{A} \in \text{Abs}_{\mu}(s', h'_1)$. Therefore $s'(e) = s'(e')$ and $\mathcal{A} \in \text{Abs}_{\mu}(s', h'_0 * h'_1)$.
- If $\mathcal{A} = \mathbf{x} \xrightarrow{l} e$ then $(h_0 * h_1)(s(\mathbf{x}))(l) = s(e)$. Hence, there is $i \in \{0, 1\}$ such that $h_i(s(\mathbf{x}))(l) = s(e)$. Since $(s, h_i) \simeq_{\mu} (s', h'_i)$, $h'_i(s'(\mathbf{x}))(l) = s'(e)$ and $(h'_0 * h'_1)(s'(\mathbf{x}))(l) = s'(e)$. So $\mathcal{A} \in \text{Abs}_{\mu}(s', h'_0 * h'_1)$.
- If $\mathcal{A} = \text{alloc}(x)$ then $s(\mathbf{x}) \in \text{dom}(h_0 * h_1)$. Hence, there is $i \in \{0, 1\}$ such that $s(\mathbf{x}) \in \text{dom}(h_i)$. Since $(s, h_i) \simeq_{\mu} (s', h'_i)$, $s'(\mathbf{x}) \in \text{dom}(h'_i)$ and $s'(\mathbf{x}) \in \text{dom}(h'_0 * h'_1)$, which entails $\mathcal{A} \in \text{Abs}_{\mu}(s', h'_0 * h'_1)$. □

Lemma 7 below states that our abstraction is correct for the fragments CL and RF.

Lemma 7 (Soundness of abstraction for SL). *Let μ be a measure, (s, h) and (s', h') be two memory states such that $(s, h) \simeq_\mu (s', h')$ [resp. $(s, h) \simeq_{\mu[w \leftarrow 0]} (s', h')$]. For any SL formula \mathcal{A} such that $\mu_{\mathcal{A}} \leq \mu$ and \mathcal{A} belongs to RF [resp. CL], we have $(s, h) \models_{\text{SL}} \mathcal{A}$ iff $(s', h') \models_{\text{SL}} \mathcal{A}$.*

Proof. The proof of Lemma 7 for the classical fragment is rather straightforward. Indeed, any CL formula is a Boolean combination of test formulae. In order to deal with the record fragment, more efforts are needed. Suppose that $(s, h) \simeq_\mu (s', h')$ and $\mathcal{A} \in \text{RF}$ with $\mu_{\mathcal{A}} \leq \mu$. By structural induction, we show that $(s, h) \models_{\text{SL}} \mathcal{A}$ iff $(s', h') \models_{\text{SL}} \mathcal{A}$. The base case when \mathcal{A} has one of the forms $e = e'$, $f \xrightarrow{l} e$ and **emp** is by an easy verification. Similarly, in the induction step, the cases when the outermost connective is Boolean are straightforward.

- Assume that $(s, h) \models_{\text{SL}} \mathcal{A}$ with $\mathcal{A} = \mathcal{B}_1 * \mathcal{B}_2$. There are heaps h_1 and h_2 such that $h = h_1 * h_2$, $(s_1, h_1) \models_{\text{SL}} \mathcal{B}_1$ and $(s_2, h_2) \models_{\text{SL}} \mathcal{B}_2$. As $\mu \geq \mu_{\mathcal{A}}$ and $\mu_{\mathcal{A}} \geq \mu_{\mathcal{B}_1} + \mu_{\mathcal{B}_2}$, there are μ_1 and μ_2 such that $\mu_1 \geq \mu_{\mathcal{B}_1}$, $\mu_2 \geq \mu_{\mathcal{B}_2}$ and $\mu_1 + \mu_2 = \mu$. By application of Lemma 4, there are heaps h'_1 and h'_2 verifying $h' = h'_1 * h'_2$, $(s, h_1) \simeq_{\mu_1} (s', h'_1)$ and $(s, h_2) \simeq_{\mu_2} (s', h'_2)$. By the induction hypothesis, we get $(s', h'_1) \models_{\text{SL}} \mathcal{B}_1$ and $(s', h'_2) \models_{\text{SL}} \mathcal{B}_2$. Consequently, $(s', h') \models_{\text{SL}} \mathcal{A}$ since $h' = h'_1 * h'_2$ and $\mathcal{A} = \mathcal{B}_1 * \mathcal{B}_2$.
- Finally, assume $\mathcal{A} = \mathcal{B}_1 * \mathcal{B}_2$. Let $h'_1 \perp h'$ be such that $(s', h'_1) \models_{\text{SL}} \mathcal{B}_1$. By Lemma 5, there is a heap h_1 such that $(s, h_1) \simeq_\mu (s', h'_1)$ and $h_1 \perp h$, and so $(s, h_1) \models_{\text{SL}} \mathcal{B}_1$ by the induction hypothesis. Then $(s, h * h_1) \models_{\text{SL}} \mathcal{B}_2$, and by Lemma 6, $(s', h' * h'_1) \models_{\text{SL}} \mathcal{B}_2$. Hence $(s', h') \models_{\text{SL}} \mathcal{B}_1 * \mathcal{B}_2$.

□

Note that we can extend this result to the whole SL by considering test formulae of the form $\mathbf{x} + i = \mathbf{x}' + i'$.

3.2. Complexity of reasoning tasks for SL

In this section, we show that model-checking, satisfiability, and validity, for SL, are PSPACE-complete. We use the abbreviations $\text{mc}(\text{SL})$, $\text{sat}(\text{SL})$ and $\text{val}(\text{SL})$ for the respective problems. These abbreviations are extended to any fragment of separation logic, for instance $\text{sat}(\text{RF})$ denotes the satisfiability problem for the record fragment.

PSPACE-hardness of $\text{mc}(\text{LF})$ and $\text{sat}(\text{LF})$ is a consequence of [21, Sect. 5.2]. As SL strictly contains LF, this entails the PSPACE-hardness of $\text{mc}(\text{SL})$ and $\text{sat}(\text{SL})$. Since SL is closed under negation, PSPACE-completeness of $\text{val}(\text{SL})$ will follow from PSPACE-completeness of $\text{sat}(\text{SL})$.

In order to show that $\text{mc}(\text{SL})$ and $\text{sat}(\text{SL})$ are in PSPACE, we establish the lemmas below. Lemma 8 establishes a reduction from $\text{mc}(\text{SL})$ to $\text{mc}(\text{RF})$, so that we only need to consider RF in order to find the complexity of model-checking. Then, in Lemma 9, we will provide a small model property for RF, leading to the PSPACE-easiness of $\text{mc}(\text{RF})$ (see Lemma 10). Finally, we characterize the computational complexity of the satisfiability problem thanks to Lemma 12, which entails a reduction from $\text{sat}(\text{SL})$ to $\text{mc}(\text{SL})$.

Lemma 8. *There is a logspace reduction from $\text{mc}(\text{SL})$ to $\text{mc}(\text{RF})$.*

Proof. Let $t(\mathcal{A})$ be the formula obtained from \mathcal{A} in SL by replacing each occurrence of $\mathbf{x} + i \xrightarrow{l} e$ by $\langle \mathbf{x}, i \rangle \xrightarrow{l} e$. The formula $t(\mathcal{A})$ belong to RF. Given a store s , we write $t(s)$ to denote the store such that $t(s)(\langle \mathbf{x}, i \rangle) = s(\mathbf{x}) + i$. One can show that for every heap h , we have $(s, h) \models_{\text{SL}} \mathcal{A}$ iff $(t(s), h) \models_{\text{SL}} t(\mathcal{A})$. The proof is by structural induction on \mathcal{A} . □

We need to establish a quite technical lemma. Given a heap h , let $\text{Im}^2(h)$ be the set of natural numbers i such that there are i' and l for which $h(i')(l) = i$.

Lemma 9. *Let $\mu = (\{0\}, w, \text{Lab}_\mu, \text{Var}_\mu)$ be a measure, and l_0 be a label that does not belong to the finite set of labels Lab_μ . If $(s, h) \simeq_\mu (s', h')$ and $h_0 \perp h$ is a heap, then there is a heap h'_0 such that*

- $h'_0 \perp h'$,

- $(s, h_0) \simeq_\mu (s', h'_0)$,
- $\text{card}(\text{dom}(h'_0)) \leq \max(w, \text{card}(\text{Var}_\mu))$,
- $\max(\text{dom}(h'_0) \cup \text{Im}^2(h'_0)) \leq \max((s'(\text{Var}_\mu) \cap \mathbb{N}) \cup \text{dom}(h')) + w$,
- for all $n \in \text{dom}(h'_0)$, $\{l : h'_0(n)(l) \text{ is defined}\} \subseteq \text{Lab}_\mu \uplus \{l_0\}$.

The heap h'_0 is said to be a small disjoint heap with respect to μ and (s', h') and it can be represented in polynomial space in $\text{size}(\mu) + \text{size}_{\text{Lab}_\mu}(h_0) + \text{size}_{\text{Lab}_\mu \text{ bis, Lab}_\mu}((s', h'))$.

Proof. Assume that $(s, h) \simeq_\mu (s', h')$ and $h_0 \perp h$. We introduce two disjoint heaps h_{01} and h_{02} such that $\text{dom}(h_{01}) = \text{Im}(s) \cap \text{dom}(h_0)$, $\text{dom}(h_{02}) = \text{dom}(h_0) \setminus \text{Im}(s)$ and $h_0 = h_{01} * h_{02}$. We define the heap h'_0 as the disjoint union $h'_{01} * h'_{02}$ where h'_{01} and h'_{02} are defined so as to satisfy $\text{dom}(h'_{01}) = \text{Im}(s') \cap \text{dom}(h'_0)$ and $\text{dom}(h'_{02}) = \text{dom}(h'_0) \setminus \text{Im}(s')$.

In the sequel, the constant `null` will be viewed as a distinguished variable always interpreted by *nil*. We write V to denote the set of “variables” $\text{Var}_\mu \cup \{\text{null}\}$.

- In order to define h'_{01} , let V_1, \dots, V_a be the equivalence classes over the set V for the relation \sim_s defined by $\mathbf{x} \sim_s \mathbf{y}$ if $s(\mathbf{x}) = s(\mathbf{y})$. Since $(s, h) \simeq_\mu (s', h')$, the relation $\sim_{s'}$ defines the same set of equivalence classes. For each class V_k , let i_k be the image of the variables of V_k through s , and i'_k through s' . Then, for each $k \in [1, a]$ and $l \in \text{dom}(h_{01}(i_k))$, the heap h'_{01} is defined as follows:

- if $l \notin \text{Lab}_\mu$, then $h'_{01}(i'_k)(l_0) = \text{nil}$ and $h'_{01}(i'_k)(l)$ is undefined,
- if $l \in \text{Lab}_\mu$ and $h_{01}(i_k)(l) = i_n$ for some n , then $h'_{01}(i'_k)(l) = i'_n$,
- if $l \in \text{Lab}_\mu$ and $h_{01}(i_k)(l) \neq i_n$ for all n , then $h'_{01}(i'_k)(l) = \text{nil}$ (in that case, $\text{nil} \notin \{i'_1, \dots, i'_a\}$).

The domain of h'_{01} is included in $\text{Im}(s')$, since $\text{Im}(s'_V) = \{i'_1, \dots, i'_a\}$.

- In order to define h'_{02} , let $b = \max(0, \min\{\text{card}(\text{dom}(h_{02})), w - \text{card}(\text{dom}(h_{01}))\})$ and j'_1, \dots, j'_b be the b smallest natural numbers disjoint from $\{i'_1, \dots, i'_a\} \cup \text{dom}(h')$. Hence, when $\text{card}(\text{dom}(h_{01})) \geq w$, $b = 0$ and therefore there are no such natural numbers. Otherwise, $\text{dom}(h'_{02}) = \{j'_1, \dots, j'_b\}$ and for each $k \in [1, b]$, we define $h'_{02}(j'_k)(l_0) = \text{nil}$ (which is an arbitrary value).

As announced, we define h'_0 as the heap $h'_{01} * h'_{02}$. Let us show that the heap h'_0 has all the desired properties.

- Let us check that $h' \perp h'_0$. First, $h' \perp h'_{01}$ since $h \perp h_{01}$. Second, $h' \perp h'_{02}$ by construction.
- Let us check that $(s, h_0) \simeq_\mu (s', h'_0)$. We proceed by a case analysis on the form of the test formulae.

$(e = e')$ Since $(s, h) \simeq_\mu (s', h')$, $s(e) = s(e')$ iff $s'(e) = s'(e')$.

$(\text{alloc}(\mathbf{x}))$ We have equivalences between the propositions below:

- $\text{alloc}(\mathbf{x}) \in \text{Abs}_\mu(s, h_0)$,
- $s(\mathbf{x}) \in \text{dom}(h_0)$,
- there is k such that $\mathbf{x} \in V_k$ and $i_k \in \text{dom}(h_0)$,
- there is k such that $\mathbf{x} \in V_k$ and $i'_k \in \text{dom}(h'_{01})$,
- $\text{alloc}(\mathbf{x}) \in \text{Abs}_\mu(s', h'_0)$.

$(\text{size} \geq k)$ First, observe that $\text{card}(\text{dom}(h_{01})) = \text{card}(\text{dom}(h'_{01}))$. Moreover, by construction, if $\text{card}(\text{dom}(h_0)) < w$, then $\text{card}(\text{dom}(h_0)) = \text{card}(\text{dom}(h'_0))$. When $\text{card}(\text{dom}(h_0)) \geq w$, the construction of h'_0 guarantees that $\text{card}(\text{dom}(h'_0)) \geq w$. So, for all formulae $\text{size} \geq k$ with $k < w$, $\text{size} \geq k \in \text{Abs}_\mu(s, h_0)$ iff $\text{size} \geq k \in \text{Abs}_\mu(s', h'_0)$.

$(\mathbf{x} \xrightarrow{l} e)$ We have the following implications:

- $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_\mu(s, h_0)$,
- there is k such that $\mathbf{x} \in V_k$ and $h_0(i_k)(l) = s(e)$,
- there are k, k' such that $h_0(i_k)(l) = i_{k'}$,

function $\text{MC}((s, h), \mathcal{A}, \mu)$

(base-cases) If \mathcal{A} is atomic, then return $(s, h) \models_{\text{SL}} \mathcal{A}$;

(Boolean-cases) If $\mathcal{A} = \mathcal{A}_1 \wedge \mathcal{A}_2$, then return $(\text{MC}((s, h), \mathcal{A}_1, \mu)$ and $\text{MC}((s, h), \mathcal{A}_2, \mu))$;
Other Boolean operators are treated analogously.

(* case) If $\mathcal{A} = \mathcal{A}_1 * \mathcal{A}_2$, then return \perp if there are no h_1, h_2 such that $h = h_1 * h_2$ and $\text{MC}((s, h_1), \mathcal{A}_1, \mu)$ and $\text{MC}((s, h_2), \mathcal{A}_2, \mu)$;

(\neg * case) If $\mathcal{A} = \mathcal{A}_1 \neg * \mathcal{A}_2$, then return \perp if for some small disjoint heap h' with respect to μ and (s, h) verifying $\text{MC}((s, h'), \mathcal{A}_1, \mu)$, we have not $\text{MC}((s, h * h'), \mathcal{A}_2, \mu)$;

Return \top ;

Figure 1: Model-checking algorithm

- there are k, k' such that $\mathbf{x} \in V_k$ and $h'_{01}(i'_k)(l) = i'_{k'}$,
- $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_\mu(s', h'_0)$.

Now suppose that not $\mathbf{x} \xrightarrow{l} e \in \text{Abs}_\mu(s, h_0)$. We distinguish 3 cases.

1. $s(\mathbf{x}) \notin \text{dom}(h_0)$.

From the above case with $\text{alloc}(\mathbf{x})$, $s'(\mathbf{x}) \notin \text{dom}(h'_0)$ and therefore $\mathbf{x} \xrightarrow{l} e \notin \text{Abs}_\mu(s', h'_0)$.

2. $s(\mathbf{x}) \in \text{dom}(h_0)$ (with $i_k = s(\mathbf{x})$), $l \in \text{dom}(h_0(i_k))$ and $h_0(i_k)(l) \neq s(e)$.

If $h_0(i_k)(l) = i_{k'}$ for some $k' \in [1, a]$, then $s'(e) \neq i'_{k'}$. If for all $k' \in [1, a]$, $i_{k'} \neq h_0(i_k)(l)$ (in particular $h_0(i_k)(l) \neq \text{nil}$), then by construction, $h'_{01}(i'_k)(l) \notin \{i'_1, \dots, i'_a\}$. In both cases, $\mathbf{x} \xrightarrow{l} e \notin \text{Abs}_\mu(s', h'_0)$.

3. $s(\mathbf{x}) \in \text{dom}(h_0)$ (with $i_k = s(\mathbf{x})$) and $l \notin \text{dom}(h_0(i_k))$. Consequently, $l \notin \text{dom}(h'_0(i'_k))$ and therefore $\mathbf{x} \xrightarrow{l} e \notin \text{Abs}_\mu(s', h'_0)$.

Therefore (s, h_0) and (s', h'_0) have the same abstraction.

- Let us check that $\text{card}(\text{dom}(h'_0)) \leq \max(w, \text{card}(\text{Var}_\mu))$. We already know that $a \leq \text{card}(\text{Var}_\mu)$. If $\text{card}(\text{dom}(h'_{01})) \geq w$, then h'_{02} is the empty heap and therefore $\text{card}(\text{dom}(h'_0)) \leq a$. Otherwise, by construction $\text{card}(\text{dom}(h'_{01})) + \text{card}(\text{dom}(h'_{02})) \leq w$. Consequently, $\text{card}(\text{dom}(h'_0)) \leq \max(w, \text{card}(\text{Var}_\mu))$.
- Let us check that $\max(\text{dom}(h'_0) \cup \text{Im}^2(h'_0)) \leq \max(\{s'(\mathbf{x}) : \mathbf{x} \in \text{Var}_\mu, s'(\mathbf{x}) \in \mathbb{N}\} \cup \text{dom}(h')) + w$. We have chosen the domain and image of h'_{01} to be included in the image of s' plus nil , and therefore $\text{dom}(h'_{01})$ satisfies the above condition. The image of h'_{02} is $\{\text{nil}\}$. The domain of h'_{02} is composed of the smallest natural numbers which neither belong to $s'(\text{Var}_\mu)$, nor to $\text{dom}(h')$. As $\text{dom}(h'_{02})$ has less than w elements, it is bounded by the w^{th} such natural number, which is bounded by $w + \max((s'(\text{Var}_\mu) \cap \mathbb{N}) \cup \text{dom}(h'))$.
- Let us check that for every $n \in \text{dom}(h'_0)$, $\{l : h'_0(n)(l) \text{ is defined}\} \subseteq \text{Lab}_\mu \uplus \{l_0\}$. This condition is satisfied by construction of h'_{01} and h'_{02} .

□

Lemma 10. $mc(\text{RF})$ is in PSPACE.

Proof. The algorithm is described in Figure 1. First of all, the algorithm can be implemented in polynomial space since the quantifications are over sets of exponential size in $|\mathcal{A}| + \text{size}_{\text{Var}_\mu, \text{Lab}_\mu}((s, h))$ where $\mu_{\mathcal{A}} = (\dots, \text{Lab}_\mu, \text{Var}_\mu)$, and the recursion depth is linear in $|\mathcal{A}|$. Hence, all the heaps considered in the algorithm are of polynomial-size in $|\mathcal{A}| + \text{size}_{\text{Var}_\mu, \text{Lab}_\mu}((s, h))$. It remains to be shown that the algorithm is correct: given \mathcal{A} with $\mu_{\mathcal{A}} \leq \mu$ and $(s, h) \models_{\text{SL}} \mathcal{A}$ iff $\text{MC}((s, h), \mathcal{A}, \mu)$ returns \top . The only point to

check in the proof by structural induction is the case when the outermost connective is the operator \ast . Whenever $(s, h) \not\models_{\text{SL}} \mathcal{A}_1 \ast \mathcal{A}_2$, there is a heap $h_0 \perp h$ such that $(s, h_0) \models_{\text{SL}} \mathcal{A}_1$ and $(s, h \ast h_0) \not\models_{\text{SL}} \mathcal{A}_2$. By Lemma 9 with $(s', h') = (s, h)$, there is a small disjoint heap h'_0 with respect to μ and (s, h) such that $(s, h'_0) \simeq_\mu (s, h_0)$. Since the measure of \mathcal{A}_1 is less than μ , Lemma 7 entails $(s, h'_0) \models_{\text{SL}} \mathcal{A}_1$. Moreover, by Lemma 6, $(s, h \ast h'_0) \not\models_{\text{SL}} \mathcal{A}_2$. Consequently, $(s, h) \not\models_{\text{SL}} \mathcal{A}_1 \ast \mathcal{A}_2$ iff there is a *small* heap h'_0 such that $(s, h'_0) \models_{\text{SL}} \mathcal{A}_1$ and $(s, h \ast h'_0) \not\models_{\text{SL}} \mathcal{A}_2$. \square

The rest of the section is dedicated to the characterization of the complexity of decision problems for SL. To do so, we need another technical lemma. Given a permutation $\sigma : \text{Val} \rightarrow \text{Val}$ with $\sigma(\text{nil}) = \text{nil}$ and a heap h , we write $\sigma \cdot h$ to denote the partial function which maps i to the partial function $\sigma \circ (h(i))$. When viewing heaps as finite subsets of $\mathbb{N} \times \text{Lab} \times \text{Val}$, $\sigma \cdot h$ is equal to $\{(i, l, \sigma(j)) : (i, l, j) \in h\}$. We write $\sigma \bullet h$ to denote the heap $\sigma \cdot (h \circ \sigma^{-1})$, which corresponds to $\{(\sigma(i), l, \sigma(j)) : (i, l, j) \in h\}$. For instance, given a label l and an address i , we have $(\sigma \bullet h)(i)(l) = \sigma(h(\sigma^{-1}(i))(l))$. This operation allows us to rename all the addresses according to the permutation: the memory graph keeps the same shape, but vertices are placed on different addresses. We shall use the properties below that can be easily checked:

- For all permutations σ and disjoint heaps h_1 and h_2 , $\sigma \bullet (h_1 \ast h_2) = (\sigma \bullet h_1) \ast (\sigma \bullet h_2)$.
- For all permutations σ and heaps h , we have $\sigma^{-1} \bullet (\sigma \bullet h) = h$.

Lemma 11. *Let \mathcal{A} be a state formula of SL with measure $\mu = (0, \text{offsets}, w, \text{Lab}_\mu, \text{Var}_\mu)$ and (s, h) be a memory state. For all permutations $\sigma : \text{Val} \rightarrow \text{Val}$ with $\sigma(\text{nil}) = \text{nil}$ such that for all $\mathbf{x} \in \text{Var}_\mu$ and $i \in \text{offsets}$, $\sigma(s(\mathbf{x}) + i) = \sigma(s(\mathbf{x})) + i$, we have $(s, h) \models_{\text{SL}} \mathcal{A}$ iff $(\sigma \circ s, \sigma \bullet h) \models_{\text{SL}} \mathcal{A}$.*

Proof. Let \mathcal{A} be an SL formula, μ be a measure greater than $\mu_{\mathcal{A}}$, s be a store and h be a heap. It is sufficient to show one direction of the equivalence since the other direction is obtained by application of the first one with the store $\sigma \circ s$ and the well-defined inverse bijection σ^{-1} . Indeed, for all $\mathbf{x} \in \text{Var}_\mu$, $\sigma^{-1}((\sigma \circ s)(\mathbf{x}) + u) = \sigma^{-1}((\sigma \circ s)(\mathbf{x})) + u$. Assume that $(s, h) \models_{\text{SL}} \mathcal{A}$. We show that $(\sigma \circ s, \sigma \bullet h) \models_{\text{SL}} \mathcal{A}$. We are going to prove this by induction on \mathcal{A} . The cases with boolean operators are trivial and are omitted herein. If \mathcal{A} is an atomic formula, then we proceed by a case analysis.

$e = e' \ s(e) = s(e')$ iff $\sigma(s(e)) = \sigma(s(e'))$ since σ is a bijection on Val .

\mathcal{A} is $\mathbf{x} + i \xrightarrow{l} e$: $h(s(\mathbf{x}) + i)(l) = s(e)$, and $\sigma \bullet h(\sigma \circ s(\mathbf{x}) + i)(l) = \sigma \bullet h(\sigma^{-1}(\sigma(s(\mathbf{x}) + i)))(l) = \sigma \bullet h(\sigma^{-1}(\sigma(s(\mathbf{x}) + i)))(l) = \sigma \bullet h(s(\mathbf{x}) + i)(l) = \sigma(h(s(\mathbf{x}) + i)(l)) = \sigma(s(e)) = \sigma \circ s(e)$,

\mathcal{A} is **emp**: $\text{dom}(\sigma \bullet h)$ is empty iff $\text{dom}(h)$ is empty.

If $\mathcal{A} = \mathcal{A}_1 \ast \mathcal{A}_2$, then there are h_1 and h_2 such that $h = h_1 \ast h_2$ and $(s, h_1) \models_{\text{SL}} \mathcal{A}_1$ and $(s, h_2) \models_{\text{SL}} \mathcal{A}_2$. For each measure $\mu_{\mathcal{A}_i}$, we have $\mu_{\mathcal{A}_i} \leq \mu_{\mathcal{A}} \leq \mu$. Then, by induction, $(\sigma \circ s, \sigma \bullet h_i) \models_{\text{SL}} \mathcal{A}_i$. Since $\sigma \bullet h = \sigma \bullet (h_1 \ast h_2) = (\sigma \bullet h_1) \ast (\sigma \bullet h_2)$, we can conclude that $(\sigma \circ s, \sigma \bullet h) \models_{\text{SL}} \mathcal{A}$.

If $\mathcal{A} = \mathcal{A}_1 \ast \mathcal{A}_2$, then let h_0 be a heap which is orthogonal to $\sigma \bullet h$. Assume that $(\sigma \circ s, h_0) \models_{\text{SL}} \mathcal{A}_1$. Then by induction, $(\sigma^{-1} \circ (\sigma \circ s), \sigma^{-1} \bullet h_0) \models_{\text{SL}} \mathcal{A}_1$, that is $(s, \sigma^{-1} \bullet h_0) \models_{\text{SL}} \mathcal{A}_1$. So $(s, h \ast (\sigma^{-1} \bullet h_0)) \models_{\text{SL}} \mathcal{A}_2$, and by induction $(\sigma \circ s, \sigma \bullet (h \ast (\sigma^{-1} \bullet h_0))) \models_{\text{SL}} \mathcal{A}_2$, that is $(\sigma \circ s, (\sigma \bullet h) \ast (\sigma \bullet (\sigma^{-1} \bullet h_0))) \models_{\text{SL}} \mathcal{A}_2$, and finally $(\sigma \circ s, (\sigma \bullet h) \ast h_0) \models_{\text{SL}} \mathcal{A}_2$. So, $(\sigma \circ s, \sigma \bullet h) \models_{\text{SL}} \mathcal{A}$. \square

We state below a small memory state property that happens to be central to establish the results about the forthcoming PSPACE upper bounds.

Lemma 12 (Small memory state property). *A state formula \mathcal{A} in SL is satisfiable iff there is a store s such that $(s, \emptyset) \models_{\text{SL}} \neg(\mathcal{A} \ast \perp)$ and for each variable $\mathbf{x} \in \text{Var}_{\mathcal{A}}$, $s(\mathbf{x}) \leq (\text{card}(\text{Var}_{\mathcal{A}}) + 1) \times (1 + \max(\text{offsets}_{\mathcal{A}}))$, where \emptyset stands for the heap with empty domain, $\text{Var}_{\mathcal{A}}$ is the set of variables occurring in \mathcal{A} , and $\text{offsets}_{\mathcal{A}}$ is the set of indices i such that $\mathbf{x} + i$ occurs in \mathcal{A} for some variable \mathbf{x} . If $\text{offsets}_{\mathcal{A}}$ is empty, we can replace $\max(\text{offsets}_{\mathcal{A}})$ by 0.*

Proof. First, it is straightforward to show that \mathcal{A} in SL is satisfiable iff there is a store s such that $(s, \emptyset) \models_{\text{SL}} \neg(\mathcal{A} * \perp)$, where \emptyset is the heap with empty domain. So, we only have to prove that given an SL state formula \mathcal{A} and a store s such that $(s, \emptyset) \models_{\text{SL}} \mathcal{A}$, there is a store s' such that $(s', \emptyset) \models_{\text{SL}} \mathcal{A}$ and for each $\mathbf{x} \in \text{Var}_{\mathcal{A}}$, $s'(\mathbf{x}) \leq (\text{card}(\text{Var}_{\mathcal{A}}) + 1) \times (1 + \max(\text{offsets}_{\mathcal{A}}))$ (the interpretation of other variables is irrelevant). In order to obtain this small store, we are going to decrease the value of the variables in several steps. Each step consists of applying a permutation to the memory graph.

Assume that $(s, \emptyset) \models_{\text{SL}} \mathcal{A}$ and let $\text{MAX} = 1 + \max(\text{offsets}_{\mathcal{A}})$. Let \mathbf{x}_0 be a dummy variable such that $s(\mathbf{x}_0) = 0$, and $\mathbf{x}_1, \dots, \mathbf{x}_n$ be an ordering of the variables occurring in \mathcal{A} such that for $j \in [0, n-1]$, $s(\mathbf{x}_j) \leq s(\mathbf{x}_{j+1})$. If there is no k such that $s(\mathbf{x}_{k+1}) \geq s(\mathbf{x}_k) + \text{MAX}$, then for all $\mathbf{x} \in \text{Var}_{\mathcal{A}}$, $s(\mathbf{x}) \leq (n+1) \times (1 + \text{MAX})$.

Otherwise, let k be the smallest index such that $s(\mathbf{x}_{k+1}) \geq s(\mathbf{x}_k) + \text{MAX}$. Let $\alpha = s(\mathbf{x}_{k+1}) - (s(\mathbf{x}_k) + \text{MAX})$. Let us define the permutation σ based on α :

- If $j \leq s(\mathbf{x}_k) + \text{MAX}$ then $\sigma(j) = j$;
- If $s(\mathbf{x}_{k+1}) \leq j \leq s(\mathbf{x}_n) + \text{MAX}$, then $\sigma(j) = j - \alpha$;
- If $j \geq s(\mathbf{x}_n) + \text{MAX}$ then $\sigma(j) = j$;
- If $s(\mathbf{x}_k) + \text{MAX} < j < s(\mathbf{x}_{k+1})$ then we have to complete this function so as to obtain a bijection, $\sigma(j) = j - (s(\mathbf{x}_k) + \text{MAX}) + (s(\mathbf{x}_n) + \text{MAX} - \alpha)$.

Observe that for all $\mathbf{x} \in \text{Var}_{\mathcal{A}}$ and $i \in \text{offsets}_{\mathcal{A}}$, $\sigma(s(\mathbf{x}) + i) = \sigma(s(\mathbf{x})) + i$. This permutation satisfies the hypotheses of Lemma 11, and thus may be applied to (s, \emptyset) , which then still satisfies \mathcal{A} . We apply this type of permutation until there is no k such that $s(\mathbf{x}_{k+1}) \geq s(\mathbf{x}_k) + \text{MAX}$. So, by simple multiplication, for all $\mathbf{x} \in \text{Var}_{\mathcal{A}}$, $s(\mathbf{x}) \leq (n+1) \times \text{MAX}$. □

Proposition 13. *The model-checking, satisfiability, and validity problems for SL are PSPACE-complete.*

Proof. PSPACE-hardness results are consequences of [21, Sect. 5.2]. The PSPACE upper bound for $\text{mc}(\text{SL})$ is a consequence of Lemmas 8 and 10. The PSPACE upper bound for $\text{sat}(\text{SL})$ is obtained by enumerating the small memory states of $\neg(\mathcal{A} * \perp)$ with empty heap (see Lemma 12) and then using Lemma 10. □

Decidable fragments of first-order SL can be found in [25, 34, 35].

4. Decidable Satisfiability Problems by Abstracting Computations

In this section, we establish the PSPACE-completeness of the problems $\text{SAT}(\text{CL})$ and $\text{SAT}(\text{RF})$. To do so, we abstract memory states whose size is a priori unbounded by finite symbolic memory states. As usual with linear temporal logic, temporal infinity in models is handled by Büchi automata recognizing ω -sequences. We propose below an abstraction that is correct for CL (allowing pointer arithmetic) and for RF (allowing all operators from Separation Logic) taken separately but that is not exact for the full language SL.

4.1. Symbolic models

Definitions. An important aspect of the method consists in defining a symbolic satisfiability relation. Here are the details. Given a measure μ , we write Σ_μ to denote the powerset of \mathcal{T}_μ ; Σ_μ is thought of as an alphabet, and elements $a \in \Sigma_\mu$ are called *letters*. A symbolic model with respect to μ is defined as an infinite sequence $\sigma \in \Sigma_\mu^\omega$. Symbolic models are abstractions of models from LTL^{mem} : given a model $\rho : \mathbb{N} \rightarrow \mathcal{S} \times \mathcal{H}$ and a measure μ , we write $\text{Abs}_\mu(\rho) : \mathbb{N} \rightarrow \Sigma_\mu$ to denote the symbolic model with respect to μ such that for every $t \in \mathbb{N}$, $\text{Abs}_\mu(\rho)(t) \stackrel{\text{def}}{=} \{\mathcal{A} \in \mathcal{T}_\mu : \rho, t \models \mathcal{A}[\langle \mathbf{x}, u \rangle \leftarrow X^u \mathbf{x}]\}$.

To a letter a , we associate the formula $TF[a] = \bigwedge_{\mathcal{A} \in a} \mathcal{A} \wedge \bigwedge_{\mathcal{A} \in (\mathcal{T}_\mu \setminus a)} \neg \mathcal{A}$. For all symbolic models σ and formulae ϕ such that $\mu_\phi \leq \mu$, we define the symbolic satisfaction relation $\sigma, t \models_\mu \phi$ as the satisfaction relation for models except for the clause about atomic subformulae is updated as follows: $\sigma, t \models_\mu \mathcal{A}$ iff $\models_{\text{SL}} TF[\sigma(t)] \Rightarrow \mathcal{A}[X^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle]$. We write $\text{L}^\mu(\phi)$ to denote the set of symbolic models σ with respect to μ such that $\sigma, 0 \models_\mu \phi$. As a corollary of Lemma 7, we get a soundness result for our abstraction:

Proposition 14. *Let ϕ be a formula of $\text{LTL}^{\text{mem}}(\text{RF})$ [resp. of $\text{LTL}^{\text{mem}}(\text{CL})$] and $\mu_\phi \leq \mu$. For any model ρ , we have $\rho \models \phi$ iff $\text{Abs}_\mu(\rho) \models_\mu \phi$ [resp. $\text{Abs}_{\mu[w \leftarrow 0]}(\rho) \models_\mu \phi$].*

Proof. We treat the case $\phi \in \text{LTL}^{\text{mem}}(\text{RF})$ (for the case $\phi \in \text{LTL}^{\text{mem}}(\text{CL})$, replace below μ by $\mu[w \leftarrow 0]$). The induction step for the cases with Boolean and temporal operators is by an easy verification. Let us check the base case, for a state formula. Suppose that $\rho, t \models \mathcal{B}$ for an atomic formula \mathcal{B} of LTL^{mem} . By definition, $\text{Abs}_\mu(\rho)(t) \stackrel{\text{def}}{=} \{\mathcal{A} \in \mathcal{T}_\mu : \rho, t \models \mathcal{A}[\langle \mathbf{x}, u \rangle \leftarrow \mathbf{X}^u \mathbf{x}]\}$. Let us show that $\models_{\text{SL}} \text{TF}[\text{Abs}_\mu(\rho)(t)] \Rightarrow \mathcal{B}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle]$. If for some memory state $(s, h) \models_{\text{SL}} \text{TF}[\text{Abs}_\mu(\rho)(t)]$, then by Lemma 7, $(s, h) \models_{\text{SL}} \mathcal{B}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle]$.

Suppose now that $\text{Abs}_\mu(\rho), t \models_\mu \mathcal{B}$. Hence, $\models_{\text{SL}} \text{TF}[\text{Abs}_\mu(\rho)(t)] \Rightarrow \mathcal{B}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle]$. Since $\rho, t \models_{\text{SL}} \text{TF}[\text{Abs}_\mu(\rho)(t)][\langle \mathbf{x}, u \rangle \leftarrow \mathbf{X}^u \mathbf{x}]$, we have $\rho, t \models (\mathcal{B}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle])[\langle \mathbf{x}, u \rangle \leftarrow \mathbf{X}^u \mathbf{x}]$. This entails that $\rho, t \models \mathcal{B}$. \square

Note that Abs_μ is not surjective; we note L_{sat}^μ the set of symbolic models with respect to μ that are abstractions of some model for LTL^{mem} . Consequently, ϕ in $\text{LTL}^{\text{mem}}(\text{RF})$ is satisfiable iff $L^{\mu_\phi}(\phi) \cap L_{\text{sat}}^{\mu_\phi}$ is nonempty.

4.2. ω -regularity and PSPACE upper bound

In order to show that $\text{SAT}(\text{RF})$ and $\text{SAT}(\text{CL})$ are in PSPACE we shall explain why testing the nonemptiness of $L^{\mu_\phi}(\phi) \cap L_{\text{sat}}^{\mu_\phi}$ can be done in PSPACE. Below we always treat the case for RF. For CL, replace every occurrence of μ_ϕ by $\mu_\phi[w \leftarrow 0]$ and every occurrence of μ by $\mu[w \leftarrow 0]$. To do so, we show that each language can be recognized by an exponential-size Büchi automaton satisfying the good properties to establish the PSPACE upper bound. If \mathbb{A} is a Büchi automaton, we note $L(\mathbb{A})$ the language recognized by \mathbb{A} . Following [9, 19], let \mathbb{A} be the generalized Büchi automaton defined by the structure $(\Sigma, Q, \delta, I, \mathcal{F})$ such that $(\mu \geq \mu_\phi)$:

- Q is the set of so-called atoms of ϕ , that are sets of temporal formulae included in the so-called closure set $cl(\phi)$ (see [9]). Let us briefly recall that the closure set $cl(\phi)$ is the smallest set containing ϕ , closed under subformulae, negations (double negations are eliminated) and such that if $\psi \mathbf{U} \psi' \in cl(\phi)$, then $\mathbf{X}(\psi \mathbf{U} \psi') \in cl(\phi)$. A set $X \subseteq cl(\phi)$ is an atom whenever it satisfies the usual conditions for subformulae whose outermost connective is Boolean and, we have $\psi \mathbf{U} \psi' \in X$ iff $(\psi' \in X \text{ or } (\psi, (\psi \mathbf{U} \psi') \in X))$ whenever $\psi \mathbf{U} \psi' \in cl(\phi)$.
- $I = \{X \in Q : \phi \in X\}$.
- $\Sigma = \Sigma_\mu$.
- $X \xrightarrow{a} Y$ iff
 1. for every atomic formula \mathcal{A} of X , $\models_{\text{SL}} \text{TF}[a] \Rightarrow \mathcal{A}[\mathbf{X}^u \mathbf{x} \leftarrow \langle \mathbf{x}, u \rangle]$.
 2. for every $\mathbf{X}\phi' \in cl(\phi)$, $\mathbf{X}\phi' \in X$ iff $\phi' \in Y$.
- Let $\{\phi_1 \mathbf{U} \phi'_1, \dots, \phi_n \mathbf{U} \phi'_n\}$ be the set of until formulae in $cl(\phi)$. Let \mathcal{F} be equal to $\{F_1, \dots, F_n\}$ where $F_i = \{X \in Q : \phi_i \mathbf{U} \phi'_i \notin X \text{ or } \phi'_i \in X\}$ for $i \in \{1, \dots, n\}$.

Let \mathbb{A}_ϕ^μ be the Büchi automaton equivalent to the generalized Büchi automaton \mathbb{A} . It is easy to observe that \mathbb{A}_ϕ^μ has an exponential amount of states in the size of ϕ and its transition relation can be checked in polynomial space in the size of ϕ . Moreover,

Lemma 15. *Let ϕ in $\text{LTL}^{\text{mem}}(\text{RF})$ [resp. $\text{LTL}^{\text{mem}}(\text{CL})$] and $\mu \geq \mu_\phi$ [resp. $\mu[w \leftarrow 0] \geq \mu_\phi[w \leftarrow 0]$]. Then, $L(\mathbb{A}_\phi^\mu) = L^\mu(\phi)$ [resp. $L(\mathbb{A}_\phi^{\mu[w \leftarrow 0]}) = L^{\mu[w \leftarrow 0]}(\phi)$].*

We can also build a Büchi automaton $\mathbb{A}_{\text{sat}}^\mu$ such that $L(\mathbb{A}_{\text{sat}}^\mu) = L_{\text{sat}}^\mu$. $\mathbb{A}_{\text{sat}}^\mu$ is defined as $(\Sigma, Q, \delta, I, F)$, where $\Sigma = \Sigma_\mu$, $Q = \Sigma_\mu$, $F = I = Q$ and $a \xrightarrow{a'} a''$ iff:

1. $\text{TF}[a], \text{TF}[a'']$ are satisfiable, and $a = a'$,
2. for every formula $\langle \mathbf{x}, u \rangle = \langle \mathbf{x}', u' \rangle \in \mathcal{T}_\mu$ with $u, u' \geq 1$, $\langle \mathbf{x}, u \rangle = \langle \mathbf{x}', u' \rangle \in a$ iff $\langle \mathbf{x}, u-1 \rangle = \langle \mathbf{x}', u'-1 \rangle \in a''$,
3. for every formula $\langle \mathbf{x}, u \rangle = \text{null} \in \mathcal{T}_\mu$ with $u \geq 1$, $\langle \mathbf{x}, u \rangle = \text{null} \in a$ iff $\langle \mathbf{x}, u-1 \rangle = \text{null} \in a''$,

If $\mu = \mu_\phi$, then \mathbb{A}_{sat}^μ is of exponential-size in the size of ϕ and the transition relation can be checked in polynomial space in the size of ϕ . More importantly, this automaton recognizes satisfiable symbolic models.

Lemma 16. *Let ϕ in $LTL^{\text{mem}}(\text{RF})$ [resp. $LTL^{\text{mem}}(\text{CL})$] and $\mu = \mu_\phi$ [resp. $\mu = \mu_\phi[w \leftarrow 0]$]. Then, $L(\mathbb{A}_{sat}^\mu) = L_{sat}^\mu$.*

Proof. It is immediate that the abstraction of any model with respect to μ belongs to $L(\mathbb{A}_{sat}^\mu)$. Therefore, the set of abstractions of memory states with respect to μ is included in $L(\mathbb{A}_{sat}^\mu)$.

The other inclusion is shown by induction. Let $\mu = (\text{offsets}, w, \text{Lab}_0, \text{Var}_0)$ be the measure μ_ϕ , m be $\max\{u : \text{there is } \mathbf{x} \in \text{Var} \text{ such that } X^u \mathbf{x} \text{ occurs in } \phi\}$ and MAXI be $\max(\text{offsets}) + 1$. Let $(a_i)_{i \in \mathbb{N}}$ be an ω -sequence of symbolic memory states in $L(\mathbb{A}_{sat}^\mu)$. We shall build a sequence $(s_i, h_i)_{i \in \mathbb{N}}$ such that $\text{Abs}_\mu((s_i, h_i)_{i \in \mathbb{N}}) = (a_i)_{i \in \mathbb{N}}$. So, for $t \in \mathbb{N}$, $a_t = \{\mathcal{A} \in \mathcal{T}_\mu : \rho, t \models \mathcal{A}[\langle \mathbf{x}, u \rangle \leftarrow X^u \mathbf{x}]\}$. The construction is by induction on the position $t \in \mathbb{N}$.

Let us study the base case of the induction that will provide a value for s_0, \dots, s_m, h_0 . Since $(a_i)_{i \in \mathbb{N}} \in L(\mathbb{A}_{sat}^\mu)$, $TF[a_0]$ is satisfiable. There are s'_0 and h'_0 satisfying $(s'_0, h'_0) \models_{\text{SL}} TF[a_0]$. When dealing with the record fragment ($\text{offsets} = \{0\}$), the objects are appropriate for the initialization: $h_0 = h'_0$ and for $u \in [0, m]$ and $\mathbf{x} \in \text{Var}_0$, we set $s_u(\mathbf{x}) = s'_0(\langle \mathbf{x}, u \rangle)$. When $\text{offsets} \neq \{0\}$ ($w = 0$ and we are dealing with the fragment CL), there is no constraint on the size of the heap. We apply a permutation σ which maps all the images of variables to multiples of MAXI . For $u \in [0, m]$, we consider the store s_u such that for $\mathbf{x} \in \text{Var}_0$, $s_u(\mathbf{x}) = \sigma(s'_0(\langle \mathbf{x}, u \rangle))$. The heap h_0 is defined by enumerating the test formulae $\langle \mathbf{x}, u \rangle + j \xrightarrow{l} \langle \mathbf{x}', u' \rangle$, $\langle \mathbf{x}, u \rangle + j \xrightarrow{l} \text{null}$ and $\text{alloc}(\langle \mathbf{x}, u \rangle + j)$ of a_0 , and by defining the heap accordingly. When $\langle \mathbf{x}, u \rangle + j \xrightarrow{l} \langle \mathbf{x}', u' \rangle \in a_0$, $h_0(s_u(\mathbf{x}) + j)(l) = s_{u'}(\mathbf{x}')$; when $\langle \mathbf{x}, u \rangle + j \xrightarrow{l} \text{null} \in a_0$, $h_0(s_u(\mathbf{x}) + j)(l) = \text{nil}$; when $\text{alloc}(\langle \mathbf{x}, u \rangle + j) \in a_0$, we define $h_0(s_u(\mathbf{x}) + j)(l_0) = \text{nil}$, for some $l_0 \notin \text{Lab}_0$. Thanks to the distance MAXI imposed between the values of variables, test formulae about the heap which are not in a_0 are not satisfied. Equalities $e = e'$ are preserved since the store has only been modified by a permutation.

For the inductive step, suppose that we have already defined the stores s_0, \dots, s_{k+m} and heaps h_0, \dots, h_k for some position $k \geq 0$ satisfying the conditions below: for every $t \leq k$,

- for all $\mathcal{A} \in \mathcal{T}_\mu$, $(s_t^*, h_t) \models_{\text{SL}} \mathcal{A}$ iff $\mathcal{A} \in a_t$, where $s_t^* : \langle \mathbf{x}, u \rangle \mapsto s_{t+u}(\mathbf{x})$;
- $\text{Im}(s_t^*) \subseteq \text{MAXIN} \cup \{\text{nil}\}$.

Let us build the store s_{k+m+1} and the heap h_{k+1} . Since $(a_i)_{i \in \mathbb{N}} \in L(\mathbb{A}_{sat}^\mu)$, $TF[a_{k+1}]$ is satisfiable. There exist a memory state (s', h') satisfying $(s', h') \models_{\text{SL}} TF[a_{k+1}]$ and for all $\mathbf{x} \in \text{Var}$ and $u \in [0, m-1]$, $s'(\langle \mathbf{x}, u \rangle) = s_{k+1+u}(\mathbf{x})$. By definition of \mathbb{A}_{sat}^μ , $(\langle \mathbf{x}, u+1 \rangle = \langle \mathbf{x}', u'+1 \rangle \in a_k$ iff $\langle \mathbf{x}, u \rangle = \langle \mathbf{x}', u' \rangle \in a_{k+1}$) and $(\langle \mathbf{x}, u+1 \rangle = \text{null} \in a_k$ iff $\langle \mathbf{x}, u \rangle = \text{null} \in a_{k+1}$) for all $u, u' \in [0, m-1]$. Consequently, for all $u, u' \in [0, m-1]$, $s_{k+1+u}(\mathbf{x}) = s_{k+1+u'}(\mathbf{x}')$ iff $s'(\langle \mathbf{x}, u \rangle) = s'(\langle \mathbf{x}', u' \rangle)$. So, there is a permutation σ identical for the variables $\langle \mathbf{x}, u \rangle$ with $u \in [0, m-1]$ such that $\text{Im}(\sigma \circ s') \subseteq \text{MAXIN} \cup \{\text{nil}\}$. By construction, for $\langle \mathbf{x}, u \rangle \in V_\mu$, $\sigma(s'(\langle \mathbf{x}, u \rangle)) \in \text{MAXIN} \cup \{\text{nil}\}$. For $\mathbf{x} \in \text{Var}_0$, we set $s_{k+1+m}(\mathbf{x}) = \sigma(s'(\langle \mathbf{x}, m \rangle))$.

If we consider RF, this permutation satisfies the prerequisites of Lemma 11, since $\text{offsets} = \{0\}$. We can define $h_{k+1} = \sigma \bullet h'$. Thanks to Lemma 11, we know that both of these models satisfy the same test formulae, which are exactly those in a_{k+1} .

If we are dealing with CL, then the definition of s_{k+m+1} ensures that the equalities satisfied are exactly those of a_{k+1} . This time the prerequisites of Lemma 11 are not satisfied unless $\text{offsets} = \{0\}$. We know that $w = 0$, which means that the only test formula about size in a_{k+1} is $\text{size} \geq 0$; therefore there is no constraint on the size of the heap. The heap is defined by enumerating the test formulae of the form $\langle \mathbf{x}, u \rangle + j \xrightarrow{l} \langle \mathbf{x}', u' \rangle$ of a_{k+1} , and defining for each of them $h_{k+1}(s_{k+1+u}(\mathbf{x}) + j)(l) = s_{k+1+u'}(\mathbf{x}')$ (idem when $\langle \mathbf{x}, u \rangle + j \xrightarrow{l} \text{null} \in a_{k+1}$, $h_{k+1}(s_{k+1+u}(\mathbf{x}) + j)(l) = \text{nil}$); and then for each of the test formulae of the form $\text{alloc}(\langle \mathbf{x}, u \rangle + j)$ of a_{k+1} , we define $h_{k+1}(s_{k+1+u}(\mathbf{x}) + j)(l_0) = \text{nil}$, for some $l_0 \notin \text{Lab}_0$. Thanks to the distance MAXI between variables, the test formulae about the heap which are not in a_{k+1} are not satisfied. Equalities $e = e'$ are preserved since the store has only been modified by a permutation.

□

This lemma is essential and it is not possible to extend it to the whole logic LTL^{mem} even by allowing test formulae of the form $\mathbf{x} + i = \mathbf{y} + j$: we would then need automata with counters. Now, we can state our main complexity result.

Theorem 17. *SAT(RF) and SAT(CL) are PSPACE-complete.*

Proof. The lower bound is from LTL [20]. Let ϕ be an instance formula of SAT(RF) (for SAT(CL) replace below μ_ϕ by $\mu_\phi[w \leftarrow 0]$). As seen earlier, ϕ is satisfiable iff $L^{\mu_\phi}(\phi) \cap L_{\text{sat}}^{\mu_\phi}$ is nonempty. Hence, ϕ is satisfiable iff $L(\mathbb{A}_\phi^{\mu_\phi}) \cap L(\mathbb{A}_{\text{sat}}^{\mu_\phi}) \neq \emptyset$. The intersection automaton is of exponential size in the size of ϕ and can be checked nonempty by a nondeterministic on-the-fly algorithm. Since this algorithm, for the nonemptiness problem of Büchi automata, is in NLOGSPACE and the transition relation in the intersection automaton can be checked in polynomial space in the size of ϕ , we obtain a nondeterministic polynomial space algorithm for testing satisfiability of ϕ . As usual, by Savitch's theorem, we get the PSPACE upper bound. □

4.3. Other problems in PSPACE

Let Frag be either the classical fragment or the record fragment. Lemma 1 provides a reduction from $MC_{\text{init}}^{\text{ct}}(\text{Frag})$ to $\text{SAT}_{\text{init}}^{\text{ct}}(\text{Frag})$ based on a program-as-formula encoding. As we will see now, we may also reduce $\text{SAT}_{\text{init}}^{\text{ct}}(\text{Frag})$ to $\text{SAT}(\text{Frag})$ internalizing an approximation of the initial memory state which the logical language cannot distinguish from the initial memory state. As a consequence, the PSPACE upper bound for $\text{SAT}(\text{Frag})$ entails the PSPACE upper bound for both $\text{SAT}_{\text{init}}^{\text{ct}}(\text{Frag})$ and $MC_{\text{init}}^{\text{ct}}(\text{Frag})$.

Proposition 18. *Each of the problems $\text{SAT}_{\text{init}}^{\text{ct}}(\text{RF})$, $MC_{\text{init}}^{\text{ct}}(\text{RF})$, $\text{SAT}_{\text{init}}^{\text{ct}}(\text{CL})$ and $MC_{\text{init}}^{\text{ct}}(\text{CL})$ is PSPACE-complete.*

Proof. We begin with the fragment RF. By Lemma 1 and since $\text{SAT}_{\text{init}}^{\text{ct}}(\text{RF})$ is known to be PSPACE-hard, it remains to establish the PSPACE upper bound for $\text{SAT}_{\text{init}}^{\text{ct}}(\text{RF})$.

Given a formula ϕ and an initial memory state (s, h) , we shall build in polynomial-time a formula $\phi_{s,h}^{\text{ct}}$ in $\text{SAT}(\text{RF})$ such that ϕ is satisfiable in a model with initial memory state (s, h) and constant heap iff $\phi_{s,h}^{\text{ct}}$ is satisfiable by a general model. Since we have shown that $\text{SAT}(\text{RF})$ is in PSPACE, this guarantees that $\text{SAT}_{\text{init}}^{\text{ct}}(\text{RF})$ is in PSPACE. The idea of the proof is to internalize the initial memory state and the fact that the heap is constant in the logic $\text{SAT}(\text{RF})$. Actually, we will not exactly express that the heap is constant but the approximation we use will be sufficient for our purpose.

Apart from the variables of ϕ , the formula $\phi_{s,h}^{\text{ct}}$ is built over additional variables in $V = \{\mathbf{x}_i : i \in \text{dom}(h) \cup \text{Im}(s)\} \cup \{\mathbf{x}_{i,l} : i \in \text{dom}(h), l \in \text{dom}(h(i))\}$. The formula $\phi_{s,h}^{\text{ct}}$ is of the form $G(\psi_1 \wedge \psi_2 \wedge \psi_3) \wedge \psi_s \wedge \psi'$, where the subformulae are defined as follows.

- ψ_1 states that the heap is almost equal to h since we cannot forbid additional labels in the logical language ($\text{dom}(h) = \{i_1, \dots, i_k\}$):

$$\psi_1 \stackrel{\text{def}}{=} (\bigwedge_{l \in \text{dom}(h(i_1))} \mathbf{x}_{i_1} \stackrel{l}{\mapsto} \mathbf{x}_{i_1,l}) * \dots * (\bigwedge_{l \in \text{dom}(h(i_k))} \mathbf{x}_{i_k} \stackrel{l}{\mapsto} \mathbf{x}_{i_k,l}).$$
- ψ_2 states which variables are equal and which ones are not, depending on the initial memory state. It is a conjunction of simple formulae. As an example, for $i \neq j \in \text{dom}(h)$, a simple formula of ψ_2 is $\mathbf{x}_i \neq \mathbf{x}_j$. Similarly, if $h(i)(l) = j$ and $j \in \text{dom}(h)$, then $\mathbf{x}_{i,l} = \mathbf{x}_j$ is a simple formula of ψ_2 . Details are omitted.
- ψ_3 states that the additional variables remain constant: $\bigwedge_{\mathbf{x} \in V} \mathbf{x} = \mathbf{X}\mathbf{x}$.
- The formula ψ' is obtained from ϕ by replacing each occurrence of $\mathbf{x} \stackrel{l}{\mapsto} e$ by

$$\mathbf{x} \stackrel{l}{\mapsto} e \wedge \bigwedge_{i \in \text{dom}(h), l \notin \text{dom}(h(i))} \mathbf{x} \neq \mathbf{x}_i.$$

The additional conjunction is useful because our logical language cannot state that a label is not in the domain of some allocated address.

- ψ_s states constraints about the initial store s : $\psi_s \stackrel{\text{def}}{=} \bigwedge_{\mathbf{x} \in \phi} \mathbf{x} = \mathbf{x}_{s(\mathbf{x})}$.

It is then easy to check that ϕ is satisfiable by a model with initial memory state (s, h) and constant heap iff $\phi_{s,h}^{ct}$ is satisfiable by a general model.

As far as the results for the classical fragment are concerned, by Lemma 1, there is a logspace reduction from $\text{MC}_{init}^{ct}(\text{CL})$ to $\text{SAT}_{init}^{ct}(\text{CL})$ and as done above, one can reduce $\text{SAT}_{init}^{ct}(\text{CL})$ to $\text{SAT}(\text{CL})$. \square

Proposition 19. $\text{MC}_{init}^{ct}(\text{SL})$ is PSPACE-complete.

Proof. As $\text{MC}_{init}^{ct}(\text{RF})$ is a subproblem of $\text{MC}_{init}^{ct}(\text{SL})$, Proposition 18 entails the PSPACE-hardness. It remains to prove the PSPACE upper bound. The proof goes by designing a polynomial space reduction to the model-checking problem for propositional LTL. Let $(\mathbf{p}, s_0, h_0, \phi)$ be an instance of $\text{MC}_{init}^{ct}(\text{RF})$, where $\mathbf{p} = (Q, \delta, q_{init})$ is a program without destructive updates, (s_0, h_0) is an initial memory state, and ϕ is a temporal formula in $\text{LTL}^{\text{mem}}(\text{SL})$. Let Σ be the finite set of stores $\{s : \text{Im}(s) \subseteq \text{Im}(s_0) \cup \text{Im}(h_0)\}$ restricted to variables occurring in \mathbf{p} and ϕ . Its cardinality is bounded by $(\text{card}(\text{Im}(s_0) \cup \text{Im}(h_0)))^{|\phi|+|\mathbf{p}|}$. All the memory states in the transition system $\mathcal{S}_{\mathbf{p}}$ restricted to the configurations reachable from the initial memory state (s_0, h_0) are in $\Sigma \times \{h_0\}$, since \mathbf{p} is without destructive updates.

Let wdw be one plus the maximal natural number j such that $X^j \mathbf{x}$ appears in ϕ (size of the window made of consecutive states that need to be considered simultaneously). We define the transition graph $G = (Q_G, \rightarrow, Q_{init})$ such that: $Q_G = Q \times \Sigma^{wdw}$, Q_{init} is the set of tuples $(q_{init}, s_1, s_2, \dots, s_{wdw})$ such that $(s_1, h_0), \dots, (s_{wdw}, h_0)$ is a prefix of a run of \mathbf{p} with initial memory state (s_0, h_0) , and the transition relation \rightarrow is defined as follows:

$$(q, s_1, \dots, s_{wdw}) \rightarrow (q', s'_1, \dots, s'_{wdw}) \text{ iff } \begin{cases} s_{k+1} = s'_k, k = 1, \dots, wdw - 1, \text{ and } \exists q \xrightarrow{g, \text{instr}} q' \in \delta \\ \text{such that } (s_1, h_0) \models g \text{ and } (s_2, h_0) \in \llbracket \text{instr} \rrbracket (s_1, h_0). \end{cases}$$

We now define the propositional LTL model by associating to each vertex of the transition graph a set of propositional variables that are true. We define $Prop$ to be the set of atomic formulae occurring in ϕ , so that ϕ can be seen as a propositional LTL formula over $Prop$. Then the LTL model is the vertex-labeled transition graph $M = (G, \lambda)$, with

$$\lambda : Q_G \rightarrow \mathcal{P}(Prop), (q, s_1, \dots, s_{wdw}) \mapsto \{\mathcal{A} \in Prop : s_1, \dots, s_{wdw}, h_0 \models_{\text{SL}} \mathcal{A}\}.$$

By construction, $M, (q_{init}, s_1, s_2, \dots, s_{wdw}) \models \phi$ in LTL for some $(q_{init}, s_1, s_2, \dots, s_{wdw}) \in Q_{init}$ (existential version) if and only if $\mathbf{p}, (s_0, h_0) \models \phi$. The model M can be computed in polynomial space in the size of $(\mathbf{p}, s_0, h_0, \phi)$ in the sense that the (nondeterministic) transition function and the labelling function are computable in polynomial space. M has an exponential size in the size of $(\mathbf{p}, s_0, h_0, \phi)$, but let us explain now why the existence of $(q_{init}, s_1, s_2, \dots, s_{wdw}) \in Q_{init}$ such that $M, (q_{init}, s_1, s_2, \dots, s_{wdw}) \models \phi$ can be checked in polynomial space. Let \mathbb{A}_{ϕ} be the automaton recognizing the models of ϕ over the set $Prop$ of propositions: it has an exponential size in the size of $(\mathbf{p}, s_0, h_0, \phi)$, and so is the product with M . Now the existence of $(q_{init}, s_1, s_2, \dots, s_{wdw}) \in Q_{init}$ such that $M, (q_{init}, s_1, s_2, \dots, s_{wdw}) \models \phi$ reduces to check the non-emptiness of $\mathbb{A}_{\phi} \cap M$, which is decidable in space $\mathcal{O}(\log(|\mathbb{A}_{\phi}|) + \log(|M|))$ by a nondeterministic on-the-fly algorithm. The problem can therefore be solved in polynomial space in the size of $(\mathbf{p}, s_0, h_0, \phi)$ by a non-deterministic algorithm, and by Savitch's theorem this can be turned into a deterministic polynomial space algorithm. \square

Proposition 20. $\text{SAT}_{init}^{ct}(\text{SL} \setminus \{-*\})$ is PSPACE-complete.

Proof. PSPACE-hardness is a consequence of the PSPACE-hardness of $\text{SAT}_{init}^{ct}(\text{CL})$ since CL is a fragment of $\text{SL} \setminus \{-*\}$. In order to get the PSPACE upper bound, we are going to reduce the problem $\text{SAT}_{init}^{ct}(\text{SL} \setminus \{-*\})$ to $\text{SAT}_{init}^{ct}(\text{RF})$. Let $(s_0, h), \phi$ be an instance of $\text{SAT}_{init}^{ct}(\text{SL} \setminus \{-*\})$. We shall build an instance $(s'_0, h), \phi'$ of $\text{SAT}_{init}^{ct}(\text{RF})$.

Let $E = \text{dom}(h) \cup \{k - i \in \mathbb{N} : k \in \text{dom}(h) \text{ and } X^u \mathbf{x} + i \text{ occurs in } \phi\}$. We modify the set Var by adding a new variable $var(k)$ for each $k \in E$, and a variable $var(\mathbf{x}, i)$ for all \mathbf{x} and i occurring in ϕ in an expression of the form $X^u \mathbf{x} + i$ (possibly u or i is equal to zero). These variables do not occur in ϕ .

The initial store s'_0 is the extension of s_0 which maps $var(k)$ to k , and $var(\mathbf{x}, i)$ to $s_0(\mathbf{x}) + i$. Finally:

$$\begin{aligned} \phi' &= \phi[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)] \\ &\quad \wedge \mathbf{G} \bigwedge_{k \in E} (var(k) = \mathbf{X} var(k)) \\ &\quad \wedge \mathbf{G} \bigwedge_{\mathbf{x} + i \in \phi} \bigwedge_{(k+i) \in \text{dom}(h)} (\mathbf{x} = var(k) \Leftrightarrow var(\mathbf{x}, i) = var(k + i)) \end{aligned}$$

s'_0 and ϕ' have a polynomial size in the size of the instance $(s_0, h), \phi$.

Assume that $(s_0, h), \phi$ is accepted by $\text{SAT}_{init}^{ct}(\text{SL} \setminus \{-*\})$. Then there is $(s_i)_{i \in \mathbb{N}}$ such that $(s_i, h)_{i \in \mathbb{N}} \models \phi$. Let s'_i be s_i extended so as to map $var(k)$ to k and $var(\mathbf{x}, j)$ to $s_i(\mathbf{x}) + j$. Clearly $(s'_i, h)_{i \in \mathbb{N}} \models \phi[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)]$. Our definition of each s'_i also ensures that $(s'_i, h)_{i \in \mathbb{N}} \models \mathbf{G} \bigwedge_{k \in E} (var(k) = \mathbf{X} var(k))$ since the value of a variable $var(k)$ is constantly equal to k , and that $(s'_i, h)_{i \in \mathbb{N}} \models \mathbf{G} \bigwedge_{\mathbf{x} + i \in \phi} \bigwedge_{(k+i) \in \text{dom}(h)} (\mathbf{x} = var(k) \Leftrightarrow var(\mathbf{x}, i) = var(k + i))$ since for all positions, the value of $var(k + i)$ plus i and the value of $var(\mathbf{x}, i)$ is that of \mathbf{x} plus i . So $(s'_i, h)_{i \in \mathbb{N}} \models \phi'$, and therefore $(s'_0, h), \phi'$ is accepted by $\text{SAT}_{init}^{ct}(\text{RF})$.

Now, assume that $(s'_0, h), \phi'$ is accepted by $\text{SAT}_{init}^{ct}(\text{RF})$. Then there is a sequence $(s'_i)_{i \in \mathbb{N}}$ such that $(s'_i, h)_{i \in \mathbb{N}} \models \phi'$. Then $(s'_i, h)_{i \in \mathbb{N}} \models \mathbf{G} \bigwedge_{k \in E} (var(k) = \mathbf{X} var(k))$, and so, at each time state t , we have $s'_t(var(k)) = s'_0(var(k)) = k$. Moreover, $(s'_i, h)_{i \in \mathbb{N}} \models \mathbf{G} \bigwedge_{\mathbf{x} + i \in \phi} \bigwedge_{(k+i) \in \text{dom}(h)} (\mathbf{x} = var(k) \Leftrightarrow var(\mathbf{x}, i) = var(k + i))$, and so, if $k \in \text{dom}(h)$ and $\mathbf{X}^u \mathbf{x} + i$ occurs in ϕ , we have $s'_{t+u}(\mathbf{x}) = k - i$ iff $s'_{t+u}(var(\mathbf{x}, i)) = k$ (I).

We write $h' \leq h$ when there is another heap h'' for which $h = h' * h''$. Let us prove by induction on subformulae ϕ_0 of ϕ that for all $t \in \mathbb{N}$ and $h' \leq h$, we have $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0$ iff $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)]$. This will ensure that $(s'_i, h)_{i \in \mathbb{N}}, 0 \models \phi$, so that $(s'_0, h), \phi$ is accepted by $\text{SAT}_{init}^{ct}(\text{SL} \setminus \{-*\})$, from which we will conclude that $(s_0, h), \phi$ is also accepted. Indeed, if V_0 is the set of variables occurring in ϕ , the restriction $s'_0|_{V_0}$ is $s_0|_{V_0}$. Here is the proof by induction:

- If ϕ_0 is $\mathbf{X}^u \mathbf{x} + i \xrightarrow{l} \mathbf{X}^{u'} \mathbf{y}$, let $k = s'_{t+u}(var(\mathbf{x}, i))$ (the proof with $\mathbf{X}^u \mathbf{x} + i \xrightarrow{l} \text{null}$ is analogous and omitted below).
 - Suppose that $k \notin \text{dom}(h)$. We are going to prove that neither $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)]$, nor $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0$. First, it is clear that $(s'_i, h')_{i \in \mathbb{N}}, t \not\models \phi_0[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)]$. Second, assume there is $k' \in \text{dom}(h)$ such that $k' = s'_{t+u}(\mathbf{x}) + i$. Thanks to the property (I), from $s'_{t+u}(\mathbf{x}) = k' - i$, we get $s'_{t+u}(var(\mathbf{x}, i)) = k'$, and so $k = k' \in \text{dom}(h)$, which leads to a contradiction. So there is no such k' , and not $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0$.
 - Now suppose that $k \in \text{dom}(h)$. We have $s'_{t+u}(\mathbf{x}) = k = s'_{t+u}(var(\mathbf{x}, i)) - i$ thanks to property (I). Consequently, $h'(s'_{t+u}(\mathbf{x}) + i)(l) = s'_{t+u'}(\mathbf{y})$ iff $h'(s'_{t+u}(var(\mathbf{x}, i)))(l) = s'_{t+u}(var(\mathbf{y}, 0))$, and $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0$ iff $(s'_i, h')_{i \in \mathbb{N}}, t \models \phi_0[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)]$.
- If $\phi_0 = \mathcal{A}_1 * \mathcal{A}_2$, then there are h'_1 and h'_2 such that $(s'_i, h'_1)_{i \in \mathbb{N}}, t \models \mathcal{A}_1$ and $(s'_i, h'_2)_{i \in \mathbb{N}}, t \models \mathcal{A}_2$. By the induction hypothesis, since $h = (h'_1 * h'_2) * h'' = h'_1 * (h'_2 * h'')$, we have $(s'_i, h'_1)_{i \in \mathbb{N}}, t \models \mathcal{A}_1[\mathbf{X}^u \mathbf{x} + i \leftarrow \mathbf{X}^u var(\mathbf{x}, i)]$ iff $(s'_i, h'_1)_{i \in \mathbb{N}}, t \models \mathcal{A}_1$; and the same equivalence is true for h'_2 . From the two equivalences for h'_1 and h'_2 , we can conclude the same equivalence for $h' = h'_1 * h'_2$.
- Other cases are straightforward.

□

If we allow the operator $*$ in the above proposition, the current proof may not be adapted, since we would have to deal with heaps which are not sub-heaps of h in the induction step.

5. Undecidability Results

In this section, we show several undecidability results by using reduction from problems for Minsky machines. So, first we recall that a Minsky machine M consists of two counters C_1 and C_2 , and a sequence of $n \geq 1$ instructions of one of the forms below:

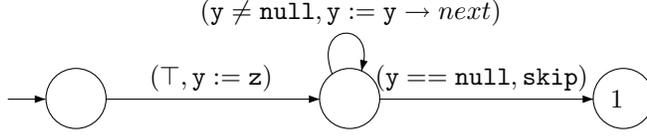


Figure 2: Checking that z points to a list

l : $C_i := C_i + 1$; goto l' l : if $C_i = 0$ then goto l' else $C_i := C_i - 1$; goto l'' .

In a nondeterministic machine, after an incrementation or a decrementation, a nondeterministic choice of the form “goto l_1 or goto l_2 ” is performed. The configurations of M are triples (l, c_1, c_2) , where $l \in [1, n]$ and $c_1, c_2 \geq 0$ are the current values of the location counter and the two counters C_1 and C_2 , respectively. The consecution relation on configurations is defined in the obvious way. A computation of M is a sequence of related configurations, starting with the initial configuration $(1, 0, 0)$.

Different encodings of counters are used here as for instance the one in [31] for which a counter C with value l is represented by a list of length l pointed to by an x dedicated to C . The same idea is used in the proof of Proposition 21 below. In order to show undecidability of $\text{SAT}(\text{SL})$, we alternatively encode counters by relying on pointer arithmetic and properties of heaps. Programs without destructive updates can simulate finite computations of Minsky machines on counters bounded by the size of some parts of the heap (the length of a list). In problems with an existential quantification on the initial heap, the maximal value of the counters can be guessed as shown to prove the results below.

Proposition 21. $\text{SAT}^{ct}(\text{LF})$ and $\text{MC}^{ct}(\text{LF})$ are Σ_1^0 -complete.

Proof. By Proposition 18, $\text{SAT}_{init}^{ct}(\text{LF})$ is decidable in polynomial space using a finite abstraction argument. Hence, $\text{SAT}^{ct}(\text{LF})$ is in Σ_1^0 by adding an existential quantification over the initial memory state. Similarly, by Proposition 18, $\text{MC}_{init}^{ct}(\text{LF})$ is decidable in polynomial space. Hence, $\text{MC}^{ct}(\text{LF})$ is also in Σ_1^0 .

By Lemma 1, we only need to show that $\text{MC}^{ct}(\text{LF})$ is Σ_1^0 -hard. We reduce the Σ_1^0 -complete halting problem for Minsky machines to it. The halting problem consists of determining whether M can reach a configuration with location counter n .

Let us build a formula ϕ and a program p in P^{ct} such that the existence of some memory state (s_0, h_0) for which $p, (s_0, h_0) \models \phi$ is equivalent to the fact that the machine M reaches a configuration with location counter n . In order to encode the values of counters, we consider a variable z pointing to a list (as shown below) in the initial memory state (s_0, h_0) :

$$z \square \xrightarrow{\text{next}} \square \xrightarrow{\text{next}} \dots \square \xrightarrow{\text{next}} \square \xrightarrow{\text{next}} \text{nil}$$

The variable z remains constant along any execution of p and the length of the list is greater than the maximal value of the counters in some finite computation (hopefully ending at the instruction corresponding to location n). We consider also the variables x_1 and x_2 and along any execution of p , each variable x_i points to a cell of the above sequence: the length of the list starting at x_i encodes the value of the counter C_i . Hence, in p , each x_i is initialized to **null**.

The program p is made of the following stages:

1. Check that z points to a list;
2. Initialize the variables;
3. Simulate M .

Figure 2 shows how to perform stage 1 with a simple “while” loop. Observe that checking whether a counter is equal to zero corresponds in p to an equality test with **null**. In order to simulate M , its structure can be embedded in the control graph of p . For instance, a decrementation instruction is encoded in p by

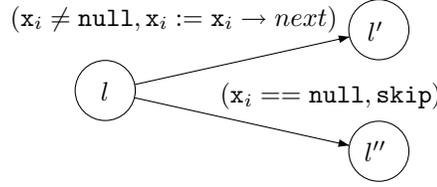


Figure 3: Simulating a decrementation

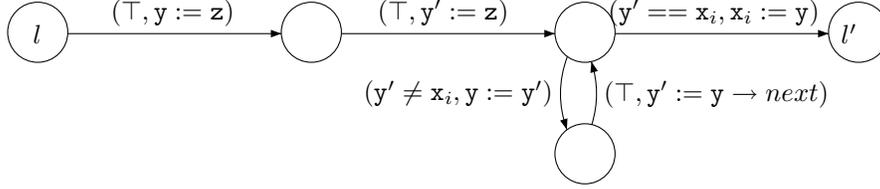


Figure 4: Simulating an incrementation

the transitions shown in Figure 3. An incrementation instruction requires a bit more care and its encoding in \mathbf{p} is presented in Figure 4. Indeed, auxiliary variables y and y' initialized to z visits the list until it meets x_i .

In the above encoding, every instruction l in M corresponds to a control state of \mathbf{p} . Hence, the formula ϕ is simply F_n (remember that we may encode propositional variable n by additional variables dedicated only for this purpose, as stated earlier).

It is then easy to show that there is an initial memory state (s_0, h_0) such that \mathbf{p} reaches the control n starting with (s_0, h_0) iff the machine M reaches the location counter n . Observe that both \mathbf{p} and M are deterministic. \square

By contrast, programs with destructive update can work with unbounded heaps, and by using the representation of counters as above, they can faithfully simulate a Minsky machine, even if the initial heap is an empty heap, without any bound on the counters. Then, as LTL can express repeated accessibility, Σ_1^1 -hardness can be obtained.

Proposition 22. *The problems $\text{MC}(\text{LF})$ and $\text{MC}_{\text{init}}(\text{LF})$ are Σ_1^1 -complete.*

Proof. It is possible to reduce the recurring problem for nondeterministic Minsky machines to $\text{MC}(\text{LF})$ and to $\text{MC}_{\text{init}}(\text{LF})$. This problem is Σ_1^1 -hard [36]. The question is whether the machine has a computation with the location counter n repeated infinitely often; and this can be expressed by $\text{GF}n$ in LTL^{mem} .

The proof is quite similar to the proof of Proposition 21 except that there is no maximal value of the counters, the initial heap is empty (which can be expressed in LTL^{mem}), and the behavior of counters is encoded by updating the memory states. For instance, incrementing C_i amounts to execute $x_i := \text{cons}(\text{next} : x_i)$ (the length of the list pointed by x_i is incremented), decrementing C_i amounts to execute $x_i := x_i \rightarrow \text{next}$. Zero tests are encoded by null tests and the initial values of the variables is null . Details are omitted since there are no technical difficulties. \square

Now, let us explain how to encode incrementation and decrementation with separating connectives and pointer arithmetic. Observe that expressions of the form $x = y + 1$ are not allowed in the logical language. We repair this “defect” in two different ways: using non-aliasing expressed by the separating conjunction, and using the precise pointing assertion $x \xrightarrow{\text{next}} \eta$ stating that the heap contains only one cell, in conjunction with the \ast operator.

$$\begin{aligned}
\phi_{x++}^* &= (\text{X}x \xleftrightarrow{\text{next}} \text{null} \wedge x + 1 \xleftrightarrow{\text{next}} \text{null}) \wedge \neg(\text{X}x \xleftrightarrow{\text{next}} \text{null} \ast x + 1 \xleftrightarrow{\text{next}} \text{null}) \\
\phi_{x--}^* &= (\text{X}x + 1 \xleftrightarrow{\text{next}} \text{null} \wedge x \xleftrightarrow{\text{next}} \text{null}) \wedge \neg(\text{X}x + 1 \xleftrightarrow{\text{next}} \text{null} \ast x \xleftrightarrow{\text{next}} \text{null}) \\
\phi_{x++}^{-\ast} &= \text{emp} \wedge ((\text{X}x \xrightarrow{\text{next}} \text{null}) \ast x + 1 \xrightarrow{\text{next}} \text{null}) \\
\phi_{x--}^{-\ast} &= \text{emp} \wedge ((x \xrightarrow{\text{next}} \text{null}) \ast \text{X}x + 1 \xrightarrow{\text{next}} \text{null})
\end{aligned}$$

The formulae based on the separating conjunction correctly express incrementation and decrementation when the cells at indices $\mathbf{x} + 1$ and $\mathbf{X}\mathbf{x}$ are allocated, whereas formulae based on the operator --^* work when the heap is empty.

Let $\text{SAT}_?^2(\text{SL})$ be any satisfiability problem among the four variants.

Proposition 23. $\text{SAT}_?^2(\text{SL})$ is Σ_1^1 -complete.

Proof. We reduce the recurrence problem for nondeterministic Minsky machines [36] to $\text{SAT}_?^2(\text{SL})$. Let ϕ_0 be the formula $\mathbf{G}(\text{emp} \wedge \bigwedge_{i=1}^2 (\mathbf{x}_i \neq \text{null}))$. Incrementation and decrementation are performed thanks to the formulae $\phi_{\mathbf{x}++}^{\text{--}^*}$ and $\phi_{\mathbf{x}--}^{\text{--}^*}$, respectively. For any model ρ such that $\rho, 0 \models \phi_0$, and for any t , we have $\rho, t \models \phi_{\mathbf{x}++}^{\text{--}^*}$ iff $s_t(\mathbf{x}_i) + 1 = s_{t+1}(\mathbf{x}_i)$. Hence, we have a means to encode incrementation. Similarly, $\rho, t \models \phi_{\mathbf{x}--}^{\text{--}^*}$ and $s_t(\mathbf{x}_i) > 0$ iff $s_t(\mathbf{x}_i) - 1 = s_{t+1}(\mathbf{x}_i)$. The fact that a counter does not change is encoded by $\mathbf{x}_i = \mathbf{X}\mathbf{x}_i$. Given that $\phi_1 = \mathbf{G}(\mathbf{x}_{\text{zero}} = \mathbf{X}\mathbf{x}_{\text{zero}} \wedge \mathbf{x}_{\text{zero}} \neq \text{null})$ holds, zero tests are encoded by $\mathbf{x}_i = \mathbf{x}_{\text{zero}}$.

Given a nondeterministic Minsky machine M , we write ψ_l to denote the formula encoding instruction l . For instance for the instruction “ l : if $\mathbf{C}_1 = 0$ then goto l' else $\mathbf{C}_1 := \mathbf{C}_1 - 1$; goto l'_1 or goto l'_2 ,” ψ_l is equal to the formula below:

$$\begin{aligned} & \mathbf{G}((l \wedge \mathbf{x}_1 \neq \mathbf{x}_{\text{zero}}) \Rightarrow (\mathbf{x}_2 = \mathbf{X}\mathbf{x}_2 \wedge (\mathbf{X}l'_1 \vee \mathbf{X}l'_2) \wedge \phi_{\mathbf{x}_1--}^{\text{--}^*})) \wedge \\ & \mathbf{G}((l \wedge \mathbf{x}_1 = \mathbf{x}_{\text{zero}}) \Rightarrow (\mathbf{x}_1 = \mathbf{X}\mathbf{x}_1 \wedge \mathbf{x}_2 = \mathbf{X}\mathbf{x}_2 \wedge \mathbf{X}l')). \end{aligned}$$

Finally, let ϕ_2 be a formula stating that each position corresponds to a unique configuration and the first instruction is 1: $\phi_2 = \mathbf{G}(\bigwedge_l (\bigwedge_{l' \neq l} (l \rightarrow \neg l')))) \wedge 1$.

Hence, $(\mathbf{x}_1 = \mathbf{x}_2 = \mathbf{x}_{\text{zero}}) \wedge \phi_0 \wedge \phi_1 \wedge \bigwedge_l \psi_l \wedge \mathbf{G}\mathbf{F}n$ is satisfiable iff M has a computation with instruction n repeated infinitely often. \square

Proposition 24. The problem $\text{SAT}(\text{SL} \setminus \{\text{--}^*\})$ is Σ_1^1 -complete.

The proof of Proposition 24 is similar to the proof of Theorem 23 except that incrementation and decrementation are performed with the formulae $\phi_{\mathbf{x}++}^*$ and $\phi_{\mathbf{x}--}^*$ respectively, and the heap is not always empty: at each increment or decrement, it has size precisely 1.

6. Conclusion

In the paper, we have introduced a temporal logic LTL^{mem} for which assertion language is quantifier-free separation logic. Figure 5 shows the reductions between problems. Curved lines represent reductions for proving hardness in a class. Straight lines represent reductions for showing that a problem belongs to its class. Figure 6 contains a summary of the complexity results about fragments of LTL^{mem} .

Finally, extending LTL^{mem} with a special propositional variable $\text{heap}^=$ stating that the current heap is equal to the next one, can lead to undecidability (look at the problems of the form $\text{SAT}_?^{ct}(\text{Frag})$). However, it is open whether satisfiability becomes decidable if we restrict the interplay between the “until” operator \mathbf{U} and $\text{heap}^=$, for instance to forbid subformulae of the form $\mathbf{G} \text{heap}^=$ with positive polarity.

References

- [1] O. Burkart, D. Caucal, F. Moller, B. Steffen, Verification of infinite structures., in: Handbook of Process Algebra, Elsevier, 2001, pp. 545–623.
- [2] S. Bardin, A. Finkel, D. Nowak, Toward symbolic verification of programs handling pointers, in: 3rd International Workshop on Automated Verification of Infinite-State Systems (AVIS’04), 2004.
- [3] A. Bouajjani, M. Bozga, P. Habermehl, R. Iosif, P. Moro, T. Vojnar, Programs with lists are counter automata, in: CAV’06, Vol. 4144 of Lecture Notes in Computer Science, Springer, 2006, pp. 517–531.
- [4] J. Reynolds, Separation logic: a logic for shared mutable data structures, in: LICS’02, IEEE, 2002, pp. 55–74.
- [5] J. Jensen, M. Jorgensen, N. Klarlund, M. Schwartzbach, Automatic verification of pointer programs using monadic second-order logic, in: PLDI’97, ACM, 1997, pp. 226–236.
- [6] T. Lev-Ami, M. Sagiv, TVLA: A system for implementing static analyses, in: SAS’00, 2000, pp. 280–301.

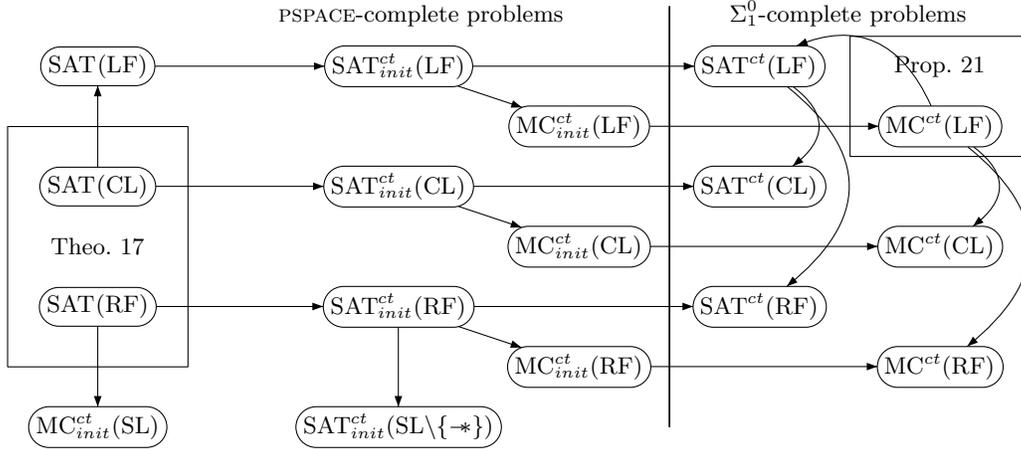


Figure 5: Reductions

	MC	MC ^{ct}	MC ^{ct} _{init}	MC _{init}	SAT	SAT ^{ct}	SAT ^{ct} _{init}
LF	Σ_1^1 -c.	Σ_1^0 -c.	PSPACE-c.	Σ_1^1 -c.	PSPACE-c.	Σ_1^0 -c.	PSPACE-c.
CL and RF	Σ_1^1 -c.	Σ_1^0 -c.	PSPACE-c.	Σ_1^1 -c.	PSPACE-c.	Σ_1^0 -c.	PSPACE-c.
SL \ {-*}	Σ_1^1 -c.	Σ_1^0 -c.	PSPACE-c.	Σ_1^1 -c.	Σ_1^1 -c.	Σ_1^0 -c.	PSPACE-c.
SL	Σ_1^1 -c.	Σ_1^0 -c.	PSPACE-c.	Σ_1^1 -c.	Σ_1^1 -c.	Σ_1^1 -c.	Σ_1^1 -c.

Figure 6: Complexity of reasoning about program with pointer variables

- [7] M. Bozga, R. Iosif, Y. Lakhnech, On logics of aliasing, in: SAS'04, Vol. 3148 of Lecture Notes in Computer Science, Springer, 2004, pp. 344–360.
- [8] A. Pnueli, The temporal logic of programs, in: FOCS'77, IEEE, 1977, pp. 46–57.
- [9] M. Vardi, P. Wolper, Reasoning about infinite computations, Information and Computation 115 (1994) 1–37.
- [10] O. Kupferman, M. Y. Vardi, P. Wolper, An automata-theoretic approach to branching-time model checking, Journal of the Association for Computing Machinery 47 (2) (2000) 312–360.
- [11] J. Berdine, C. Calcagno, P. W. O'Hearn, Symbolic execution with separation logic, in: APLAS'05, Vol. 3780 of Lecture Notes in Computer Science, Springer, 2005, pp. 52–68.
- [12] J. Berdine, C. Calcagno, B. Cook, D. Distefano, P. W. O'Hearn, T. Wies, H. Yang, Shape analysis for composite data structures, in: W. Damm, H. Hermanns (Eds.), CAV, Vol. 4590 of Lecture Notes in Computer Science, Springer, 2007, pp. 178–192.
- [13] A. Bouajjani, R. Echahed, P. Habermehl, On the verification problem of nonregular properties for nonregular processes, in: LICS'95, 1995, pp. 123–133.
- [14] H. Comon, V. Cortier, Flatness is not a weakness, in: CSL'00, Vol. 1862 of Lecture Notes in Computer Science, Springer, 2000, pp. 262–276.
- [15] S. Demri, R. Gascon, The effects of bounding syntactic resources on Presburger LTL, Journal of Logic and Computation to appear.
- [16] D. Gabbay, A. Kurucz, F. Wolter, M. Zakharyashev, Many-dimensional modal logics: theory and applications, CUP, 2003.
- [17] B. Bennett, F. Wolter, M. Zakharyashev, Multi-dimensional modal logic as a framework for spatio-temporal reasoning, Applied Intelligence 17 (3) (2002) 239–251.
- [18] P. Balbiani, J. Condotta, Computational complexity of propositional linear temporal logics based on qualitative spatial or temporal reasoning, in: FROCOs'02, Vol. 2309 of Lecture Notes in Artificial Intelligence, Springer, 2002, pp. 162–173.
- [19] S. Demri, D. D'Souza, An automata-theoretic approach to constraint LTL, Information and Computation 205 (3) (2007) 380–415.
- [20] A. Sistla, E. Clarke, The complexity of propositional linear temporal logic, Journal of the Association for Computing Machinery 32 (3) (1985) 733–749.
- [21] C. Calcagno, H. Yang, P. O'Hearn, Computability and complexity results for a spatial assertion language for data structures, in: FST&TCS'01, Vol. 2245 of Lecture Notes in Computer Science, Springer, 2001, pp. 108–119.
- [22] E. Lozes, Expressivité des logiques spatiales, Ph.D. thesis, Laboratoire de l'Informatique du Parallélisme, ENS Lyon,

- France (2004).
- [23] C. Calcagno, P. Gardner, M. Hague, From separation logic to first-order logic, in: FOSSACS'05, Vol. 3441 of Lecture Notes in Computer Science, Springer, 2005, pp. 395–409.
 - [24] D. Galmiche, D. Mery, Characterizing provability in BI's pointer logic through resource graphs, in: LPAR'05, Vol. 3835 of Lecture Notes in Computer Science, Springer, 2005, pp. 459–473.
 - [25] D. Galmiche, D. Méry, Tableaux and resource graphs for separation logic, *Journal of Logic and Computation* to appear.
 - [26] E. Yahav, T. Reps, M. Sagiv, R. Wilhelm, Verifying temporal heap properties specified via evolution logic, in: ESOP'03, Vol. 2618 of Lecture Notes in Computer Science, Springer, 2003, pp. 204–22.
 - [27] D. Distefano, J.-P. Katoen, A. Rensink, Who is pointing when to whom? on the automated verification of linked list structures, in: FST&TCS'04, Vol. 3328 of Lecture Notes in Computer Science, Springer, 2004, pp. 250–262.
 - [28] R. Brochenin, S. Demri, E. Lozes, Reasoning about sequences of memory states, in: LFCS'07, Vol. 4514 of Lecture Notes in Computer Science, Springer, 2007, pp. 100–114.
 - [29] S. Ishtiaq, P. O'Hearn, BI as an assertion language for mutable data structures, in: POPL'01, 2001, pp. 14–26.
 - [30] D. Gabbay, A. Pnueli, S. Shelah, J. Stavi, On the temporal analysis of fairness, in: 7th Annual ACM Symposium on Principles of Programming Languages, ACM Press, 1980, pp. 163–173.
 - [31] S. Bardin, A. Finkel, E. Lozes, A. Sangnier, From pointer systems to counter systems using shape analysis, 5th International Workshop on Automated Verification of Infinite-State Systems (AVIS'06).
 - [32] J. Berdine, C. Calcagno, P. O'Hearn, Smallfoot: Modular automatic assertion checking with separation logic, in: FMCO'05, Vol. 4111 of Lecture Notes in Computer Science, Springer, 2005, pp. 115–137.
 - [33] A. Finkel, E. Lozes, A. Sangnier, Towards model-checking programs with lists, in: *Infinity in Logic and Computation*, Vol. 5489 of Lecture Notes in Artificial Intelligence, Springer, 2009, to appear.
 - [34] R. Brochenin, S. Demri, E. Lozes, On the almighty wand, in: CSL'08, Vol. 5213 of Lecture Notes in Computer Science, Springer, 2008, pp. 322–337.
 - [35] K. Bansal, R. Brochenin, E. Lozes, Beyond shapes: Lists with ordered data, in: FOSSACS'09, Vol. 5504 of Lecture Notes in Computer Science, Springer, 2009, pp. 425–439.
 - [36] R. Alur, T. Henzinger, A really temporal logic, *Journal of the Association for Computing Machinery* 41 (1994) 181–204.