

Reasoning about sequences of memory states (preliminary version) ^{*}

Rémi Brochenin, Stéphane Demri, and Etienne Lozes

LSV, ENS Cachan, CNRS, INRIA
61, av. Pdt. Wilson, 94235 Cachan Cedex, France
email: {brocheni,demri,lozes}@lsv.ens-cachan.fr

Abstract. In order to verify 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 state 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. 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, and underline the tightness of our decidability results. This paper is a preliminary version of [BDL07].

1 Introduction

Verification of programs with pointers. Programs with pointer variables are worth to be analyzed and certified, since they easily contain programming errors that are sometimes difficult to detect. Such errors include the existence of memory leaks, memory violation, or undesirable aliasing. Prominent logics for analyzing such programs are Separation Logic [Rey02], pointer assertion logic PAL [JJKS97], TVLA [LAS00] and alias logic [BIL04], to quote a few examples.

Temporal Separation Logic: what for? Since [Pnu77], temporal logics are also used as languages for formal specification of programs. General and powerful automata-based techniques for verification have been developed, see e.g. [VW94]. On the other hand, Separation Logic is a static logic having a great success for program annotation [Rey02], and more recently for symbolic computation [BCO05]. Extending the scope of application of Separation Logic to standard temporal logic-based verification technique has many potential interests, either for model-checking programs, or for defining restricted forms of recursive predicates. For instance, if we write Xx to denote the next value of x (also sometimes written

^{*} Work supported by the RNTL project “AVERILES”. The first author is supported by a fellowship from CNRS/DGA.

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 $\mu L(x). x \hookrightarrow \text{null} \vee \exists x'. x \hookrightarrow x' \wedge L(x')$.

Temporal logics also 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, see a standard reduction in [SC85]. For instance, the previous formula can be seen as a program walking on a list, and more generally programs without updates can be expressed as formulae. Some programs with update that perform a simple pass on the heap, have an input-output relation that may be described by a formula. For instance, the formula $(x \hookrightarrow_0 Xx \wedge Xx \hookrightarrow_1 x)Ux \hookrightarrow_0 \text{null}$ expresses broadly that the list in the initial heap h_0 is reversed in the final heap h_1 (see also Sect. 4.2).

As a side interest, up to our knowledge, few decision procedures for programs working with pointer arithmetic have been proposed up to now, whereas arithmetical constraints in temporal logics are known to lead to undecidability, see e.g. [CC00]. Actually, there is a growing interest in understanding the interplay of pointer arithmetic, temporal reasoning, and non aliasing properties.

Our contribution. In this note, 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 [Rey02]. 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. We have shown 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 [SC85,CYO01a]. We have obtained these results by reduction to the nonemptiness problem for Büchi automata on an alphabet of symbolic memory states obtained by an abstraction that we have shown sound and complete. This is a variant of the automata-based approach introduced in [VW94] for plain LTL and further developed with concrete domains of interpretation in [DD07]. This result is not a direct consequence of the decidability of the state logic used in this temporal logic. For instance, Presburger arithmetic is decidable, but LTL with Presburger constraints is not. As a matter of fact, the abstraction method we use based on [Loz04] does not scale to the whole temporal logic, due to a subtle interplay between separation connectives and pointer arithmetic. Similar techniques can be found in [GM05]. Observe that the satisfiability problem for the whole state logic (SL) is decidable. Moreover, we have obtained new undecidability results for several problems, for instance for $SAT^{ct}(\text{LF})$ (satisfiability with constant heap on the list fragment).

Related work. Previous temporal logics designed for pointer verification include Evolution Temporal Logic [YRSW03], based on the three-valued logic abstrac-

tion method that made the success of TVLA [LAS00], and Navigation temporal logic [DKR04], based on a tableau method for model-checking 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 adresses, and a freshness predicate. The price of this expressiveness is that 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.

This paper is a preliminary version of [BDL07] even though Figure 1 contains few new results, see also a discussion in Sect. 4.

2 Memory model and specification language

We introduce below a separation logic dealing with pointer arithmetic and record values, and a temporal logic LTL^{mem} . Model-checking programs with pointer variables over LTL^{mem} specifications is our main problem of interest.

2.1 A separation logic with pointer arithmetic

Memory states. First, 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 $\mathbf{Val} = \mathbb{N} \uplus \{nil\}$. In order to model field selectors, we also consider some infinite set \mathbf{Lab} of labels. In the remainder, we will assume some fixed injection $(x, i) \in \mathbf{Var} \times \mathbb{N} \mapsto \langle x, i \rangle \in \mathbf{Var}$.

We use the notation $E \rightarrow_{fin} F$ for the set of partial functions from E to F of finite domain; and $E \rightarrow_{fin+} F$ for the set of partial functions from E to F of finite and nonempty domain. The sets \mathcal{S} of stores and \mathcal{H} of heaps are then defined as follows: $\mathcal{S} \stackrel{\text{def}}{=} \mathbf{Var} \rightarrow \mathbf{Val}$ and $\mathcal{H} \stackrel{\text{def}}{=} \mathbb{N} \rightarrow_{fin} (\mathbf{Lab} \rightarrow_{fin+} \mathbf{Val})$. We call *memory state* a couple $(s, h) \in \mathcal{S} \times \mathcal{H}$.

We will refer to the domain of a heap h by $\text{dom}(h) \subseteq \mathbb{N}$. 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 hold in a memory cell is not fixed, nor bounded. Our models could be more concrete considering labels as offsets and relying on pointer arithmetic. But for our purpose, it will be convenient to consider pointer arithmetic and recursive structures independently.

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

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

Separation Logic. We now introduce the separation logic (SL), see e.g. [Rey02], on top of which we will define our temporal logic. The syntax of the logic is given in Table 1. As separation logic is about reasoning on disjoint heaps, and we need to define what we mean by “disjoint heaps” in our model. We choose to allow to reason at the granularity of record cells, so 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 $\mathbf{dom}(h_1) \cap \mathbf{dom}(h_2) = \emptyset$. The operation $h_1 * h_2$ is defined for disjoint heap 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*, in which $\mathbf{x} + i \xrightarrow{l} e$ states that the value of the l field 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 two expressions, and \mathbf{emp} means that the current heap has no memory cell allocated (empty heap).

Formulae \mathcal{A} of SL are called *state formulae*. A formula $\mathcal{A} * \mathcal{B}$ with the *separation conjunction* states that \mathcal{A} holds on some portion of the memory heap and \mathcal{B} holds on a disjoint portion. A formula $\mathcal{A} -* \mathcal{B}$ states that the current heap, when extended with any disjoint heap verifying \mathcal{A} , will verify \mathcal{B} .

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 subformulae $\mathbf{x} + i \xrightarrow{l} e$ use $i = 0$ (we then write $\mathbf{x} \xrightarrow{l} e$). We say that a formula is in the *classical fragment* (CL) if it does not use the connectives $*$, $-*$. 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 $l = \mathit{next}$ (we simply write $\mathbf{x} \xrightarrow{} e$). Clearly, the classical and record fragments are incomparable, while the list fragment is included in both of them.

Let us illustrate the expressive power of SL on examples. The formula $\neg\text{emp}*$ 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 $h(s(\mathbf{x}))(l) = \llbracket e \rrbracket_s$. The formula $(\mathbf{x} \xrightarrow{l} \text{null})-*\perp$ is satisfied at (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, that is \mathbf{x} is allocated in h_0 .

\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.

Proposition 1. *The model-checking, satisfiability and validity problems for SL are PSPACE-complete.*

PSPACE-hardness results are consequences of [CYO01b, Sect. 5.2].

2.2 Temporal extension

Memory states sequences Models of the logic LTL^{mem} are ω -sequences of memory states, that is elements in $(\mathcal{S} \times \mathcal{H})^\omega$ and they are understood as infinite computations of programs with pointer variables. 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 “X”. For instance, $\text{X}\mathbf{x}$ is interpreted by the value of \mathbf{x} at the next memory state. The temporal operators are the standard next-time operator X and until operator U present in LTL, see e.g. [SC85]. 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 $\diamond\Phi$, $\square\Phi$. . .

Given a fragment F of SL, we write $\text{SAT}(\text{F})$ to denote the satisfiability problem for F: given a temporal formula ϕ in LTL^{mem} with state formulae built over F, 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, say (s, h)] is denoted by $\text{SAT}^{\text{ct}}(\text{F})$ [resp. $\text{SAT}_{\text{init}}(\text{F})$]. The problem $\text{SAT}_{\text{init}}^{\text{ct}}(\text{F})$ 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 I of *instructions* used in the programs is defined by the grammar below:

$$\begin{aligned} \text{instr} ::= & \mathbf{x} := \mathbf{y} \mid \text{skip} \\ & \mid \mathbf{x} := \mathbf{y} \rightarrow l \mid \mathbf{x} \rightarrow l := \mathbf{y} \mid \mathbf{x} := \text{cons}(l_1 : x_1, \dots, l_k : x_k) \mid \text{free } \mathbf{x} \\ & \mid \mathbf{x} := \mathbf{y}[i] \mid \mathbf{x}[i] := \mathbf{y} \mid \mathbf{x} = \text{malloc}(i) \mid \text{free } \mathbf{x}, i \end{aligned}$$

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

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

The denotational semantics of an instruction `instr` is defined as a partial function $\llbracket \text{instr} \rrbracket : \mathcal{S} \times \mathcal{H} \rightarrow \mathcal{S} \times \mathcal{H}$, undefined when the instruction would cause a memory violation. We list in Table 3 the formal denotational semantics of our instruction set. Boolean combinations of equalities between expressions are called guards and their 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 $q \xrightarrow{g, \text{instr}} q'$ to denote a transition. We say that a program is *without destructive update* if transitions are labeled only with instructions of the form $\mathbf{x} := y$, $\mathbf{x} := y \rightarrow l$, and $\mathbf{x} := y[i]$.

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),
- $(q, (s, h)) \rightarrow (q', (s', h'))$ iff there is a transition $q \xrightarrow{g, \text{instr}} q' \in \delta$ such that $(s, h) \models g$ and $(s', h') = \llbracket \text{instr} \rrbracket(s, h)$.

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.

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 of these fragments F of SL, we write $\text{MC}(F)$ to denote the model-checking problem for F: given a temporal formula ϕ in LTL^{mem} with state formulae built over F and a program \mathbf{p} of the associated fragment, is there an infinite computation ρ

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

Table 3. Semantics for instructions

of \mathbf{p} such that $\rho, 0 \models \phi$ (which we write $\mathbf{p} \models \phi$)? 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 $\text{MC}^{ct}(\mathbf{F})$ [resp. $\text{MC}_{init}(\mathbf{F})$]. The problem $\text{MC}_{init}^{ct}(\mathbf{F})$ is defined analogously. We may write $\mathbf{p}, (s, h) \models \phi$ to emphasize what is the initial memory state.

All the model-checking and satisfiability problems defined above can be placed in Σ_1^1 in the analytical hierarchy. Additionally, all the above problems can easily be shown PSPACE-hard since they all generalize LTL satisfiability and model-checking [SC85].

3 Complexity results

Figure 1 contains a summary of the complexity results about fragments of LTL^{mem} , more details can be found in [BDL07]. However, let us present some ideas below. The decidability results for SAT(CL) and SAT(RF) are obtained thanks to a polynomial size abstraction of the memory states similar to symbolic heaps. A notable difference is that values of variables at successive states can be compared. This abstraction is used to define a finite alphabet on which is based a Büchi automaton recognizing a language of abstract runs. Unlike the standard

| | MC | MC ^{ct} | MC ^{ct} _{init} | MC _{init} | SAT | SAT ^{ct} | SAT ^{ct} _{init} |
|-----------|------------------------|------------------------|----------------------------------|------------------------|------------------------|------------------------|-----------------------------------|
| LF | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^0\text{-c.}$ | PSPACE-c. | $\Sigma_1^1\text{-c.}$ | PSPACE-c. | $\Sigma_1^0\text{-c.}$ | PSPACE-c. |
| CL and RF | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^0\text{-c.}$ | PSPACE-c. | $\Sigma_1^1\text{-c.}$ | PSPACE-c. | $\Sigma_1^0\text{-c.}$ | PSPACE-c. |
| SL\{-*\} | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^0\text{-c.}$ | PSPACE-c. | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^0\text{-c.}$ | PSPACE-c. |
| SL | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^0\text{-c.}$ | PSPACE-c. | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^1\text{-c.}$ | $\Sigma_1^1\text{-c.}$ |

Fig. 1. Complexity of reasoning about program with pointer variables

construction for LTL, we need to recognize sequences of abstract memory states that admit a concrete sequence of memory states. Decidability is shown when the set of abstract sequences is ω -regular.

By way of example, let us also consider the undecidability proof for SAT(SL). The recurrence problem for non-deterministic Minsky machines can be reduced to it. The main difficulty is to be able to encode incrementations and decrements of a variable x . Observe that expressions of the form $x = y + 1$ do not belong to the logical language. The encoding will use in some essential way the interplay between the separation connectives and the temporal operators. We found two different ways to encode increment and decrement: using non-aliasing expressed by the separating conjunction, and using the precise pointing assertion $x \xrightarrow{next} \eta$ stating that the heap contains only one cell, in conjunction with the -* operator.

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

The formulae based on the separating conjunction correctly express incrementation and decrementation when the cells at index $x, x + 1, x - 1$ are allocated, whereas formulae based on the operator -* do not need the same assumption.

So, when the heap is constant, only the second way to encode increment applies. Moreover, in the absence of the operator -* , we can show that the problem $\text{SAT}_{init}^{ct}(\text{SL} \setminus \{-*\})$ is in PSPACE, which contrasts with the Σ_1^1 -hardness of $\text{SAT}_{init}^{ct}(\text{SL})$. This PSPACE result is obtained by reduction into $\text{SAT}_{init}^{ct}(\text{RF})$. In some essential way, we take into account that the heap is constant and that only subheaps can be considered thanks to the absence of the operator -* . This decidability result also implies that $\text{SAT}^{ct}(\text{SL})$ is in Σ_1^0 .

As far as model-checking problems are concerned, the complexity results do not depend on fragments. For the problems of the form MC_{init}^{ct} , the program can be abstracted as a finite-state automaton. Using standard results for LTL and Proposition 1, we get the PSPACE upper bound. For the problems of the form MC^{ct} , halting problem for Minsky machines can be encoded by guessing the maximal value of counters for reaching the halting state, whence the Σ_1^0 -hardness results. Finally, the problems of the form MC can encode infinite runs

of non-deterministic Minsky machines since the memory is not bounded, whence the Σ_1^1 -hardness results.

4 Discussion

We provided above complexity results for reasoning tasks about LTL^{mem} . We are currently investigating issues about the expressive power of this logical formalism. Let us discuss few issues below.

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 [YRSW03,DKR04] or to restrict the programming language, see e.g. [BFLS06]. However, with this approach, compositionality principles are lost which is a pity since they made the success of separation logic, as frame rule and composition rule.

Second, assuming that the heap is constant is subject to promising development. Indeed, it is then possible to define spatial operators at the same syntactic level as temporal operators, and write formulae as e.g. $(\mathbf{x} \hookrightarrow \mathbf{X}\mathbf{x}\mathbf{U}\mathbf{x} \hookrightarrow \text{null}) * y \mapsto \text{null}$. 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.

4.1 Recursion with local parameters

The constant heap semantics provides an original viewpoint for recursion with local parameters and local quantification. The problem of decision procedures in presence of recursive predicates has not yet completely satisfactory answers, as particular axiomatizations have been proposed for some standard recursive structures [CBO05], or incomplete, though apparently good in practice, methods of inference.

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 $\text{SAT}^{ct}(\text{SL})$, and the model-checking problem reduces to SAT_{init}^{ct} , considering that (1) can be rewritten as:

$$(B(x_1, \dots, x_k, \mathbf{X}x_1, \dots, \mathbf{X}x_k)) \mathbf{U} A(x_1, \dots, x_k).$$

In this perspective, our results could rise interesting decidability results for model-checking some of the recursive separation logic with local quantifiers. For satisfiability, we expect to define decidable fragments for $\text{SAT}^{ct}(\text{SL})$, for

instance considering the techniques for checking temporal properties of flat programs without destructive updates introduced in [FLS07]. 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.

4.2 Programs as formulae

Let us speculate some more. We may take advantage of expressing programs as formulae in order to reduce model-checking to satisfiability, a known approach from [SC85]. For programs without destructive update, we have the following result.

Proposition 2. *Let F be a fragment of SL among SL, CL, RF, or LF. Then there is a logspace reduction from $MC^{ct}(F)$ to $SAT^{ct}(F)$.*

Intuitively, we translate instructions of the form $x := y$ into $Xx = y, x := y \rightarrow l$ into $y \xrightarrow{l} Xx$, and $x := y[i]$ into $y + i \xrightarrow{l} Xx$. Guards are translated accordingly. To translate the control of the program, we use special variables to encode the control state and define a formula that expresses the transition relation.

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 $\xrightarrow{0}$ and $\xrightarrow{1}$ instead of $\xrightarrow{\cdot}$, 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 R_P of a program P as : for all $(s_0, h_0), (s_1, h_1), (s_0, h_0)R_P(s_1, h_1)$ 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 $\xrightarrow{0}$ and $\xrightarrow{1}$. 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, $x \rightarrow l := y$ would be encoded by $y \xrightarrow{l} Xx$, whereas $x := y \rightarrow l$ would be encoded as $(Xx) \xrightarrow{l} y$

5 Conclusion

We have introduced a logic that combines both aspects of temporal logic and separation logic, and permits to express constraints between values at different instants. We defined several decision problems, and studied their decidability and complexity. One of the most important complexity results states the decidability of the satisfiability problem without pointer arithmetic or with pointer arithmetic but without separation operators. In those cases, the problems are in PSPACE, which is really tight considering that both SL and LTL satisfiability problems are PSPACE-complete. This is not completely expected since LTL with

simple Presburger constraints is undecidable even though both LTL and expressive fragments of Presburger arithmetic can be solved in PSPACE. Our result strongly relies on a faithful abstraction of memory states. As a future work, we plan to investigate decidable fragments of separation logic with recursion.

References

- [BCO05] J. Berdine, C. Calcagno, and P. W. O’Hearn. Symbolic execution with separation logic. *APLAS’05*, 3780:52–68, 2005.
- [BDL07] R. Brochenin, S. Demri, and E. Lozes. Reasoning about sequences of memory states. In *LFCS’07*, LNCS. Springer, 2007. to appear.
- [BFLS06] S. Bardin, A. Finkel, E. Lozes, and A. Sangnier. From pointer systems to counter systems using shape analysis. *5th International Workshop on Automated Verification of Infinite-State Systems (AVIS’06)*, 2006.
- [BIL04] M. Bozga, R. Iosif, and Y. Lakhnech. On logics of aliasing. In *SAS’04*, volume 3148 of *LNCS*, pages 344–360. Springer, 2004.
- [CBO05] C. Calcagno, J. Berdine, and P. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. 4111:115–137, 2005.
- [CC00] H. Comon and V. Cortier. Flatness is not a weakness. *CSL’00*, 1862:262–276, 2000.
- [CYO01a] C. Calcagno, H. Yang, and P. O’Hearn. Computability and complexity results for a spatial assertion language. In *APLAS’01*, pages 289–300, 2001.
- [CYO01b] C. Calcagno, H. Yang, and P. O’Hearn. Computability and complexity results for a spatial assertion language for data structures. In *FST&TCS’01*, volume 2245 of *LNCS*, pages 108–119. Springer, 2001.
- [DD07] S. Demri and D. D’Souza. An automata-theoretic approach to constraint LTL. *Information and Computation*, 205(3):380–415, 2007.
- [DKR04] D. Distefano, J.-P. Katoen, and A. Rensink. Who is pointing when to whom? on the automated verification of linked list structures. In *FSTTCS’04*, volume 3328 of *LNCS*, pages 250–262. Springer, 2004.
- [FLS07] A. Finkel, E. Lozes, and A. Sangnier. Towards model-checking pointer systems without destructive update. 2007. Under submission.
- [GM05] D. Galmiche and D. Mery. Characterizing provability in BI’s pointer logic through resource graphs. In *LPAR’05*, volume 3835 of *LNCS*, pages 459–473. Springer, 2005.
- [JJKS97] J. Jensen, M. Jorgensen, N. Klarlund, and M. Schwartzbach. Automatic verification of pointer programs using monadic second-order logic. In *PLDI’97*, pages 226–236. ACM, 1997.
- [LAS00] T. Lev-Ami and M. Sagiv. TVLA: A system for implementing static analyses. In *SAS’00*, pages 280–301, 2000.
- [Loz04] E. Lozes. Separation logic preserves the expressive power of classical logic. In *2nd Workshop on Semantics, Program Analysis, and Computing Environments for Memory Management (SPACE’04)*, 2004.
- [Pnu77] A. Pnueli. The temporal logic of programs. In *FOCS’77*, pages 46–57. IEEE Computer Society Press, 1977.
- [Rey02] J.C. Reynolds. Separation logic: a logic for shared mutable data structures. In *LICS’02*, pages 55–74. IEEE, 2002.
- [SC85] A. Sistla and E. Clarke. The complexity of propositional linear temporal logic. *JACM*, 32(3):733–749, 1985.

- [VW94] M. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:1–37, 1994.
- [YRSW03] E. Yahav, Th. Reps, M. Sagiv, and R. Wilhelm. Verifying temporal heap properties specified via evolution logic. In *ESOP'03*, volume 2618 of *LNCS*, pages 204–22. Springer, 2003.