Two-Variable Separation Logic and Its Inner Circle

Stéphane Demri, Morgan Deters

► To cite this version:

Stéphane Demri, Morgan Deters. Two-Variable Separation Logic and Its Inner Circle. ACM Transactions on Computational Logic, 2015, 16 (2:15), 10.1145/2724711 . hal-03192209

HAL Id: hal-03192209

<https://hal.science/hal-03192209>

Submitted on 8 Apr 2021

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

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

Two-Variable Separation Logic and Its Inner Circle

STEPHANE DEMRI, New York University, USA & CNRS, France
MORGAN DETERS, New York University, USA

Separation logic is a well-known assertion language for Hoare-style proof systems. We show that first-order separation logic with a unique record field restricted to two quantified variables and no program variables is undecidable. This is among the smallest fragments of separation logic known to be undecidable, and this contrasts with decidability of two-variable first-order logic. We also investigate its restriction by dropping the magic wand connective, known to be decidable with non-elementary complexity, and we show that the satisfiability problem with only two quantified variables is not yet elementary recursive. Furthermore, we establish insightful and concrete relationships between two-variable separation logic and propositional interval temporal logic (PITL), data logics, and modal logics, providing an inner circle of closely-related logics.

Categories and Subject Descriptors: F.3.1 [Specifying and Verifying and Reasoning about Programs]: Logics of Programs

General Terms: Theory, Verification

Additional Key Words and Phrases: Separation logic, two-variable logics, interval temporal logic, modal logic, data logic, decidability, complexity

ACM Reference Format:

Stéphane Demri, Morgan Deters, January 21st, 2015 [paper accepted], November 14th, 2014 [revised submission]. Initial submission: September 13th, 2013. Two-Variable Separation Logic and Its Inner Circle *ACM Trans. Comput. Logic* V, N, Article A (January YYYY), 37 pages.
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Bounding the number of variables. The satisfiability problem for the two-variable fragment of first-order logic (FO2) is decidable [Scott 1962; Mortimer 1975], and it is NEXPTIME-complete [Lewis 1980; Grädel et al. 1997a], even though the extension with three variables is known to be undecidable [Kahr et al. 1962]. Bounding the number of variables to find decidable fragments is usually motivated by identifying maximal fragments that are decidable or to determine a rough boundary between decidability and undecidability. It is often observed that the critical zone for the decidability border lies between two and three variables for first-order dialects; for instance, the two-variable fragment of first-order intuitionistic logic (with constant domain) has been shown undecidable in [Gabbay and Shehtman 1993], whereas FO2 extended with weak forms of recursion such as transitive closure operators is also undecidable [Grädel et al. 1999]. Similarly, monodic two-variable first-order linear-time temporal logic with equality is undecidable [Degtyarev et al. 2002]. By contrast, decidability results of FO2 interpreted over linear structures can be found in [Otto 2001]; FO2 on data words is de-

Work partially supported by the EU Seventh Framework Programme under grant agreement No. PEOF-GA-2011-301166 (DATAVERIF), the Air Force Office of Scientific Research (under award FA9550-09-1-0596), and the National Science Foundation (under grant 0644299).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1529-3785/YYYY/01-ARTA \$15.00
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

cidable, too, and somewhat equivalent to the reachability problem for Petri nets [Bojańczyk et al. 2011]. There are also other decidable extensions of FO2, see [Pacholski et al. 1997; Grädel et al. 1997b; Kieroński et al. 2012; Szwaast and Tendera 2013; Charatonik et al. 2014]. At the propositional level, bounding the number of propositional variables also makes sense, to restrict syntactic resources and to study its impact on the complexity of decision problems; see examples for modal and temporal logics in [Halpern 1995; Demri and Schnoebelen 2002]. Herein, we follow a similar path, and we consider first-order separation logic [Reynolds 2002].

Decidability and complexity issues for separation logics. Before going any further, let us recall briefly that the models (heaps) in separation logics can be understood as partial functions with profile $\mathfrak{h} : \mathbb{N} \rightarrow \mathbb{N}^k$ for some $k \geq 1$ and with finite domain (k corresponds to the number of record fields). The set of natural numbers behaves here as an abstraction for the sets of locations and values, respectively. Separating connectives are those in the logical formalism that operate on models by possibly updating them by addition or removal of elements in the domain of the heap (this is done in a controlled way, of course). Program variables are understood as first-order variables with a rigid interpretation (in contrast to quantified variables).

First-order separation logic is used as an assertion language for Hoare-style proof systems about programs with pointers [Reynolds 2002]. There is an ongoing quest to determine expressive fragments with relatively low complexity (see e.g. [Cook et al. 2011]) and to extend known decidability results (see e.g. [Iosif et al. 2013; Brotherston et al. 2014]). As far as decidability is concerned, first-order separation logic with two record fields (2SL) is shown undecidable in [Calcagno et al. 2001] (see also [Yang 2001, Section 8.1] or the undecidability results from [Brotherston and Kanovich 2014; Larchey-Wendling and Galmiche 2013], which are obtained in an alternative setting with propositional variables and no first-order quantification) and the proof does not require any use of separating connectives.

So, the result in [Calcagno et al. 2001] has been strengthened in [Brochenin et al. 2012] by showing that first-order separation logic with a unique record field (1SL) is also undecidable, which is obtained by showing that 1SL is equivalent to weak second-order logic (weakness refers to quantifications over finite sets only). Recently, in [Demri and Deters 2014], 1SL restricted to two quantified variables has also been shown to be equivalent to weak second-order logic, providing another undecidability proof but one more complex than what is presented below, since one needs to go through an encoding of weak second-order logic. If separating implication (also known as the “magic wand” operator) is dropped, decidability is regained but with non-elementary complexity as a consequence of [Rabin 1969; Stockmeyer 1974]. A natural continuation of this undecidability result, quite surprising considering that there is a single record field, would be to see how the undecidability is sensitive to the number of quantified variables, following the trend of works related to two-variable fragments of first-order dialects. Whereas 1SL restricted to one variable and augmented with program variables admits a decidable satisfiability problem and is the subject of another paper [Demri et al. 2014], we study herein the two-variable first-order separation logic without program variables; the atomic formulae are made of the equality predicate symbol and the points-to predicate only (no other predicate symbols are allowed), which is a considerably downgraded version of the full logic. Separating connectives include separating conjunction and separating implication.

Our contribution. The paper presents two main results and provides interesting relationships between separation logic [Reynolds 2002], interval temporal logic [Moszkowski 2004], data logics [Bojańczyk et al. 2011] and modal logics, see

e.g. [Blackburn et al. 2001]. In that way, we improve our understanding about the expressive power of separation logic fragments.

- (1) We show that two-variable first-order separation logic with a unique record field (1SL2) has an undecidable satisfiability problem by reduction from the halting problem for Minsky machines (Section 4). It is certainly minimal in terms of the number of quantified variables, considering that the case with one variable is decidable [Demri et al. 2014]. This concludes the classification of fragments of 1SL with respect to the number of variables. For first-order separation logic, undecidability already strikes with two variables (by contrast to the NEXPTIME-completeness of FO2), which could be explained by the fact that the magic wand can be used to simulate quantification of locations (see the proof in Section 4). In our proof, all the difficulties are concentrated on the use of only two variables: of course, we can take advantage of the recycling of variables as done for modal logics [Gabbay 1981], but this is not sufficient since we need to compare the neighborhood of locations that are not direct predecessors or successors. This is where the separating operators for separation logic are helpful, for instance to operate surgically on selections of the heap (see Section 4.4). Additionally, we provide evidence that an undecidable logic on data words from [Bojańczyk et al. 2011] can be reduced to two-variable first-order separation logic (Section 4.6). Logics on data words have already been used to get undecidability results for separation logic with data fields, see e.g. [Bansal et al. 2009]. Herein, we use a simple version of first-order separation logic without program variables, without data fields, and apart from equality, there is only one binary relation, and it is functional and finite. This marks a substantial difference with existing work, and it highlights just how few our syntactic resources are while still getting undecidability and complexity lower bounds.
- (2) Two-variable first-order separation logic with a unique record field and without the magic wand (1SL2(*)) is known to be decidable (as a consequence of [Brochenin et al. 2012, Corollary 3.3]), but we establish non-elementary complexity by a reduction from Moszkowski’s Propositional Interval Temporal Logic PITL (see e.g. [Moszkowski 2004]), despite the restriction on the number of variables. So, we are able to make an interesting bridge between interval temporal logics and separation logics (see Section 3). This is not completely surprising since there is a clear similarity between the “chop” operator in interval temporal logics and the separating conjunction in separation logic, but we are able to conclude an original result about a fragment of two-variable first-order separation logic. To our knowledge, this is the first time that the similarity has been turned into a concrete, interesting result. The possibility to relate separation logic and interval temporal logic has been already envisaged by Tony Hoare, see e.g. [Zhou Chaochen 2008].¹ Finally, we refine this result by proposing a new simple modal logic MLH interpreted on heaps that can be translated into two-variable separation logic with a unique record field via a standard translation schema (see Section 2.4 and Section 3.4). For instance, we are able to show that MLH restricted to only the separating conjunction (i.e. without the magic wand operator) has non-elementary complexity. So, we are also happy to provide a modal logic closely related to two-variable separation logic, strengthening relationships between modal logics and two-variable fragments (see also [Lutz and Sattler 2002] in a more classical setting).

¹We thank Ben Moszkowski for pointing us to this work.

Figure 1 diagrams the contributions of this paper and puts them in context with previously-known results. Logic definitions can be found in Sections 2, 3.1, and 4.6.

Fig. 1. A depiction of the contributions of this paper, in context, with both novel and previously-known results marked. Complexity and decidability results are shown for the satisfiability problem in each respective logic. Each arrow represents a known satisfiability-preserving translation of formulae. For example, any PITL formula ϕ can be translated into an MLH(*) formula ψ such that ϕ is satisfiable in PITL if and only if ψ is satisfiable in MLH(*); we abuse this notation to mean “encoding” in the case of our encoding of the halting problem for Minsky machines in 1SL2. 1SL2(*) is a syntactic fragment of 1SL(*), so certainly such a “translation” is possible.

Structure of the paper. Section 2 is primarily dedicated to the presentation of the separation logic 1SL, but it also explains how data words can be encoded in 1SL2(*). This encoding will be useful many times in succeeding sections. Moreover, we introduce a new modal logic on heaps (MLH) that is a fragment of 1SL2(*). Section 3 recalls the basics of propositional interval temporal logic PITL and we present an elementary satisfiability-preserving reduction from PITL to 1SL2(*), leading to a non-elementary lower bound for 1SL2(*). This result is strengthened by showing non-elementarity of MLH restricted to only the separating conjunction *. Section 4 contains a master reduction from the halting problem for Minsky machines into the satisfiability problem for 1SL2. To do so, several technical problems need to be solved and we dedicate one

subsection to each of them. We also explain how undecidability results for first-order data logics can be used to obtain a similar result for 1SL2, which provides an alternative to the master reduction. Finally, Section 5 contains concluding remarks.

2. PRELIMINARIES

2.1. First-order separation logic with one record field (1SL)

A *heap* h is a partial function $h : \mathbb{N} \rightarrow \mathbb{N}$ with finite domain. We write $\text{dom}(h)$ to denote its *domain* and $\text{ran}(h)$ to denote its *range*. Two heaps h_1, h_2 are said to be *disjoint*, denoted $h_1 \perp h_2$, if their domains are disjoint; when this holds, we write $h_1 \uplus h_2$ to denote the heap obtained from h_1 and h_2 by taking their disjoint union. We use l_i , with $i, l_i \in \mathbb{N}$ and $i \geq 0$, to represent *locations*. We write $l_1 \rightarrow l_2 \rightarrow \dots \rightarrow l_m$ (or, equivalently, $l_m \leftarrow l_{m-1} \leftarrow \dots \leftarrow l_1$) to mean that for every $i \in [1, m-1]$, $h(l_i) = l_{i+1}$. In that case $\{l_1, \dots, l_{m-1}\} \subseteq \text{dom}(h)$. We write $\#l$ to denote the cardinal of the set $\{l' : h(l') = l\}$ made of *predecessors* of l (heap h is implicit in the expression $\#l$). A location l is an *ancestor* of a location l' iff there exists $i \geq 0$ such that $h^i(l) = l'$ where $h^i(l)$ is shorthand for $h(h(\dots(h(l)\dots)))$ (i applications of h to l).

Usually in models for separation logic(s), memory states have a heap and a store for interpreting program variables, see e.g. [Reynolds 2002]. Our work concentrates on hardness results, and we are able to obtain these results without using such program variables. For this reason, we do not introduce them. Observe also that heaps will be understood as first-order structures of the form $(\mathbb{N}, \mathfrak{R})$ where $\mathfrak{R}$ is a finite and functional binary relation. Indeed, $\mathfrak{R} = \{(l, h(l)) : l \in \text{dom}(h)\}$ for a heap h . A new modal logic with such frames is presented in Section 2.4. The locations l and l' are in the same *connected component* whenever $(l, l') \in (\mathfrak{R} \cup \mathfrak{R}^{-1})^*$. Usually, connected components are understood as non-singleton components. A finite functional graph $(\mathbb{N}, \mathfrak{R})$ can be made of several maximal connected subgraphs so that each connected subgraph is made of a cycle, possibly with trees attached to it. Figure 2 presents a heap, i.e. a finite functional graph on $\mathbb{N}$ with two maximal connected subgraphs.

Fig. 2. A heap/finite functional graph with two maximal connected subgraphs.

Let $\text{Var} = \{u_1, u_2, \dots\}$ be a countably infinite set of quantified variables. Formulae of 1SL are defined by the abstract grammar below:

$$\phi ::= u_i = u_j \mid u_i \hookrightarrow u_j \mid \text{emp} \mid \phi \wedge \phi \mid \neg \phi \mid \phi * \phi \mid \phi \multimap \phi \mid \exists u_i \phi$$

The connective $*$ is called the *separating conjunction* and the connective $\multimap$ is called the *separating implication* (also known as the *magic wand*). We will make use of standard notations for the derived connectives.

An *assignment* is a map f of the form $\text{Var} \rightarrow \mathbb{N}$. The satisfaction relation $\models$ is parameterized by assignments and defined as follows:

- $\mathfrak{h} \models_f \text{emp} \stackrel{\text{def}}{\iff} \text{dom}(\mathfrak{h}) = \emptyset$.
- $\mathfrak{h} \models_f u_i = u_j \stackrel{\text{def}}{\iff} f(u_i) = f(u_j)$.
- $\mathfrak{h} \models_f u_i \hookrightarrow u_j \stackrel{\text{def}}{\iff} f(u_i) \in \text{dom}(\mathfrak{h}) \text{ and } \mathfrak{h}(f(u_i)) = f(u_j)$.
- $\mathfrak{h} \models_f \phi_1 \wedge \phi_2 \stackrel{\text{def}}{\iff} \mathfrak{h} \models_f \phi_1 \text{ and } \mathfrak{h} \models_f \phi_2$.
- $\mathfrak{h} \models_f \neg \phi \stackrel{\text{def}}{\iff} \mathfrak{h} \not\models_f \phi$.
- $\mathfrak{h} \models_f \phi_1 * \phi_2 \stackrel{\text{def}}{\iff}$ there exist $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h}_1 \perp \mathfrak{h}_2$, $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$, $\mathfrak{h}_1 \models_f \phi_1$ and $\mathfrak{h}_2 \models_f \phi_2$.
- $\mathfrak{h} \models_f \phi_1 * \phi_2 \stackrel{\text{def}}{\iff}$ for all $\mathfrak{h}'$, if $\mathfrak{h} \perp \mathfrak{h}'$ and $\mathfrak{h}' \models_f \phi_1$ then $\mathfrak{h} \uplus \mathfrak{h}' \models_f \phi_2$.
- $\mathfrak{h} \models_f \exists u_i \phi \stackrel{\text{def}}{\iff}$ there is $l \in \mathbb{N}$ such that $\mathfrak{h} \models_{f[u_i \mapsto l]} \phi$ ($f[u_i \mapsto l]$ refers to a map equal to f except that u_i takes the value l).

Remark 2.1. Given a bijection $\sigma : \mathbb{N} \rightarrow \mathbb{N}$, we write $\mathfrak{h}' = \mathfrak{h} \circ \sigma$ to denote the heap whose graph is $\{(\sigma(l), \sigma(\mathfrak{h}(l))) : l \in \text{dom}(\mathfrak{h})\}$. Similarly, we write $f' = f \circ \sigma$ to denote the assignment such that $f'(u_i) = \sigma(f(u_i))$. Note that for all formulae ϕ , we have $\mathfrak{h} \models_f \phi$ iff $\mathfrak{h}' \models_{f'} \phi$. That is why heaps equipped with assignments are understood modulo bijections (i.e. concrete heaps are canonical elements of equivalence classes).

For a fixed $i \geq 0$, we write $1SLi$ to denote the fragment of $1SL$ restricted to i quantified variables and $1SLi(*)$ to denote its restriction when the magic wand operator is disallowed.

Let $\mathcal{L}$ be a logic among $1SL$, $1SLi$, $1SLi(*)$. The *satisfiability problem* for $\mathcal{L}$ takes as input a sentence ϕ from $\mathcal{L}$ and asks whether there is a heap $\mathfrak{h}$ such that $\mathfrak{h} \models \phi$ (regardless of assignment, as ϕ has no free variables).

THEOREM 2.2. [Rabin 1969; Brochenin et al. 2012] *The satisfiability problem for $1SL$ is undecidable, and the satisfiability problem for $1SL(*)$ is decidable with non-elementary complexity.*

Decidability of $1SL(*)$ is shown by (easy) reduction into weak monadic second-order logic of one unary (total) function with equality (referred to as MSO in Figure 1) that is shown decidable in [Börger et al. 1997, Corollary 7.2.11] by using [Rabin 1969].

2.2. Toolkit of formulae

In the following, let u and $\bar{u}$ be the variables u_1 and u_2 , in either order. Throughout this paper, we build formulae with the quantified variables u and $\bar{u}$. Note that any formula $\phi(u)$ with free variable u can be turned into an equivalent formula with free variable $\bar{u}$ by switching the two variables.

Below, we define (standard) formulae and explain which properties they express.

- The domain $\text{dom}(\mathfrak{h})$ has exactly one location:

$$\text{size} = 1 \stackrel{\text{def}}{=} \neg \text{emp} \wedge \neg(\neg \text{emp} * \neg \text{emp})$$

- The domain $\text{dom}(\mathfrak{h})$ has exactly two locations:

$$\text{size} = 2 \stackrel{\text{def}}{=} (\neg \text{emp} * \neg \text{emp}) \wedge \neg(\neg \text{emp} * \neg \text{emp} * \neg \text{emp})$$

- u has a successor: $\text{alloc}(u) \stackrel{\text{def}}{=} \exists \bar{u} u \hookrightarrow \bar{u}$

- u has at least k predecessors: $\sharp u \geq k \stackrel{\text{def}}{=} \overbrace{(\exists \bar{u} \bar{u} \hookrightarrow u) * \dots * (\exists \bar{u} \bar{u} \hookrightarrow u)}^{k \text{ times}}$

- u has at most k predecessors: $\sharp u \leq k \stackrel{\text{def}}{=} \neg(\sharp u \geq k + 1)$

- u has exactly k predecessors: $\sharp u = k \stackrel{\text{def}}{=} (\sharp u \geq k) \wedge \neg(\sharp u \geq k + 1)$

- For all $\sim \in \{\leq, \geq, =\}$ and $i \geq 0$, we define the following formulae:

$$\begin{aligned}
\#u^0 \sim k &\stackrel{\text{def}}{=} \#u \sim k \\
\#u^{i+1} \sim k &\stackrel{\text{def}}{=} \exists \bar{u} \ u \hookrightarrow \bar{u} \wedge \# \bar{u}^i \sim k \\
\#u^{-i-1} \sim k &\stackrel{\text{def}}{=} \exists \bar{u} \ \bar{u} \hookrightarrow u \wedge \# \bar{u}^{-i} \sim k
\end{aligned}$$

For instance, $\#u^6 \geq 5$ states that there is a (necessarily unique) location at distance 6 from u and its number of predecessors is greater than or equal to 5. The formula $\#u^{-5} \leq 2$ states that there is a (not necessarily unique) location at distance -5 from u and its number of predecessors is not strictly greater than 2. For instance, $\#u^1 \geq 1$ is logically equivalent to $\text{alloc}(u)$.

— There is a non-empty path from u to $\bar{u}$ and nothing else except loops that exclude $\bar{u}$:

$$\begin{aligned}
\text{ls}'(u, \bar{u}) &\stackrel{\text{def}}{=} \#u = 0 \wedge \text{alloc}(u) \wedge \neg \text{alloc}(\bar{u}) \wedge \\
&\quad \forall \bar{u} \ ((\text{alloc}(\bar{u}) \wedge \# \bar{u} = 0) \implies \bar{u} = u) \wedge \\
&\quad \forall u \ [(\#u \neq 0 \wedge u \neq \bar{u}) \implies (\#u = 1 \wedge \text{alloc}(u))]
\end{aligned}$$

— There is a (possibly empty) path from u to $\bar{u}$:

$$\text{ls}(u, \bar{u}) \stackrel{\text{def}}{=} u = \bar{u} \vee \left[\top * \text{ls}'(u, \bar{u}) \right]$$

One can show that $\mathfrak{h} \models_f \text{ls}(u, \bar{u})$ iff there is $i \in \mathbb{N}$ such that $\mathfrak{h}^i(f(u)) = f(\bar{u})$. The proof for this property can be found in [Brochenin et al. 2012, Lemma 2.4] (a similar property has been established for graph logics in [Dawar et al. 2007]).

— There is at most a single connected component (and nothing else):

$$\text{1comp} \stackrel{\text{def}}{=} \neg \text{emp} \wedge \exists u \forall \bar{u} \ \text{alloc}(\bar{u}) \implies \text{ls}(\bar{u}, u)$$

— There are exactly two components: $\text{2comps} \stackrel{\text{def}}{=} \text{1comp} * \text{1comp}$

Remark 2.3. The heap is a finite tree with at least two nodes can be expressed by the formula below:

$$\neg \text{emp} \wedge \exists u \neg \text{alloc}(u) \wedge (\forall \bar{u} \ \text{alloc}(\bar{u}) \implies \text{ls}(\bar{u}, u))$$

Complexity results about two-variable fragments of first-order logic over finite trees can be found in [Benaim et al. 2013] but we cannot really take advantage of them since we do not use predicate symbols apart from equality and the points-to relation. By contrast, we do admit separating connectives.

Remark 2.4. Observe that all formulae in our toolkit above are in the $\text{1SL2}(\ast)$ fragment. This is not by accident; we will study this fragment extensively in Section 3 and will need to employ several of the above formulae.

2.3. Encoding data words with multiple attributes

In this section, we present a simple encoding of data words with multiple attributes into heaps that will be useful in the rest of the paper. Finite data words [Bouyer 2002] are ubiquitous structures that include timed words, runs of Minsky machines, and runs of concurrent programs with an unbounded number of processes. These are finite words in which every position carries a label from a finite alphabet and a finite tuple of data values from some infinite alphabet. A wealth of specification formalisms for data words (and slight variants) has been introduced stemming from automata to

adequate logical languages such as first-order logic [Bojańczyk et al. 2011; Schwentick and Zeume 2012] and temporal logics [Figueira 2010; Decker et al. 2014].

A *data word* of dimension β is a finite non-empty sequence in $([1, \alpha] \times \mathbb{N}^\beta)^+$ for some $\alpha \geq 1$ and $\beta \geq 0$. The set $[1, \alpha]$ is understood as a finite alphabet of cardinal α whereas $\mathbb{N}$ is the infinite data domain. Data words of dimension zero are simply finite words over a finite alphabet whereas data words of dimension one correspond to data words in the sense introduced in [Bouyer 2002]. Finite runs of Minsky machines (with two counters) can be viewed as data words of dimension two over the alphabet $[1, \alpha]$ assuming that the Minsky machine has α distinct instructions (see also Section 4.1).

Let $\mathfrak{dw} = (\mathfrak{a}^1, \mathfrak{d}_1^1, \dots, \mathfrak{d}_\beta^1) \cdots (\mathfrak{a}^L, \mathfrak{d}_1^L, \dots, \mathfrak{d}_\beta^L)$ be a data word in $([1, \alpha] \times \mathbb{N}^\beta)^+$, i.e. $\mathfrak{dw}$ is of dimension β and its underlying alphabet has cardinal $\alpha \geq 1$. The data word $\mathfrak{dw}$ shall be encoded by the heap $\mathfrak{h}_{\mathfrak{dw}}$ containing a path of the form below:

$$\mathfrak{l}_0^1 \rightarrow \mathfrak{l}_1^1 \rightarrow \cdots \rightarrow \mathfrak{l}_\beta^1 \rightarrow \cdots \rightarrow \mathfrak{l}_0^L \rightarrow \mathfrak{l}_1^L \rightarrow \cdots \rightarrow \mathfrak{l}_\beta^L$$

where

- for every $i \in [1, L]$, $\mathfrak{l}_0^i$ has $\mathfrak{a}^i + 2$ predecessors,
- for all $i \in [1, L]$ and all $j \in [1, \beta]$, $\mathfrak{l}_j^i$ has $\mathfrak{d}_j^i + \alpha + 3$ predecessors,
- every location in the domain of the heap is either on that path or points to a location on that path.

Such a path from $\mathfrak{l}_0^1$ to $\mathfrak{l}_\beta^L$ is called the *main path*, and $\mathfrak{h}_{\mathfrak{dw}}^{(\beta+1)L-1}(\mathfrak{l}_0^1) = \mathfrak{l}_\beta^L$. Other simple encodings are possible (for instance without shifting the values from the finite alphabet or from the infinite domain) but the current one is well-suited for all the developments made in this paper. In particular, the encoding allows us to know easily whether a location encodes a letter from the finite alphabet or an element from the infinite domain. Note also that $\mathfrak{h}_{\mathfrak{dw}}$ is not uniquely specified, and we understand it modulo isomorphism as discussed in Remark 2.1.

Figure 3 presents the encoding of the data word $\mathfrak{dw}_0 = (2, 1)(1, 2)(2, 2)$ of dimension 1 with $\alpha = 2$ with its representation of the heap $\mathfrak{h}_{\mathfrak{dw}}$ in which the predecessors of the locations on the main path are provided schematically. We use this type of schema in Section 4 to illustrate a few constructions.

Fig. 3. The heap for data word $\mathfrak{dw}_0 = (2, 1)(1, 2)(2, 2)$ of dimension 1 and with $\alpha = 2$ (focus on the main path) with its schematic representation.

The heap $\mathfrak{h}_{\mathfrak{dw}}$ looks like a fishbone. Let us make this precise. A heap $\mathfrak{h}$ is a *fishbone* $\stackrel{\text{def}}{\iff}$

- (fb1). $\text{dom}(\mathfrak{h}) \neq \emptyset$,
- (fb2). there is a location reachable from all the locations of $\text{dom}(\mathfrak{h})$ that is not in $\text{dom}(\mathfrak{h})$, and
- (fb3). there are no distinct locations l_1, l_2, l_3, l_4, l_5 such that $l_1 \rightarrow l_2 \rightarrow l_3 \leftarrow l_4 \leftarrow l_5$ in the heap $\mathfrak{h}$.

When $\mathfrak{h}$ is a fishbone, it has a tree-like structure (when looking at the edges backward), equipped with a root (the unique location from (fb2)), but additionally, one can recognize the locations on the main path as those locations with at least one predecessor. The existence of such a main path is guaranteed by (fb3). The first location on the main path satisfies the formula

$$\text{first}(u) \stackrel{\text{def}}{=} (\#u \geq 1) \wedge \neg(\#u^{-1} \geq 1)$$

and the last location on the main path satisfies precisely the formula

$$\text{last}(u) \stackrel{\text{def}}{=} (\#u \geq 1) \wedge \neg \text{alloc}(u)$$

Let ϕ_{fb} be the formula below:

$$\overbrace{\neg \text{emp}}^{(\text{fb1})} \wedge \overbrace{(\exists u \neg \text{alloc}(u) \wedge (\forall \bar{u} \text{alloc}(\bar{u}) \Rightarrow \text{ls}(\bar{u}, u)))}^{(\text{fb2})} \wedge \overbrace{\neg(\exists u (\#u^{-2} \geq 0) * (\#u^{-2} \geq 0))}^{(\text{fb3})}$$

LEMMA 2.5. *Let $\mathfrak{h}$ be a heap. $\mathfrak{h} \models \phi_{\text{fb}}$ iff $\mathfrak{h}$ is a fishbone.*

The proof for Lemma 2.5 is by an easy verification.

Now, let us refine the notion of a fishbone heap so that it takes into account constraints on numbers of predecessors. An (α, β) -fishbone is a fishbone heap such that

- (C1). the first location on the main path has a number of predecessors in $[3, \alpha + 2]$,
- (C2). on the main path, a location with a number of predecessors in $[3, \alpha + 2]$, is followed by β locations with at least $\alpha + 3$ predecessors, and
- (C3). the number of locations on the main path is a multiple of $\beta + 1$.

It is easy to check that the formulae ϕ_{C1} , ϕ_{C2} and ϕ_{C3} in 1SL2(*) defined below are able to express the conditions (C1), (C2) and (C3), respectively. This assumes that the heap is already known to be a fishbone, which is equivalent to the satisfaction of ϕ_{fb} (by Lemma 2.5).

$$\phi_{(C1)} \stackrel{\text{def}}{=} \exists u \text{ first}(u) \wedge (3 \leq \#u \leq \alpha + 2)$$

$$\phi_{(C2)} \stackrel{\text{def}}{=} \forall u (3 \leq \#u \leq \alpha + 2) \Rightarrow \bigwedge_{i \in [1, \beta]} \#u^{+i} \geq \alpha + 3$$

$$\phi_{(C3)} \stackrel{\text{def}}{=} \forall u (3 \leq \#u \leq \alpha + 2) \Rightarrow ((-\#u^{+(\beta+1)} \geq 0) \vee (3 \leq \#u^{+(\beta+1)} \leq \alpha + 2))$$

We write $\text{dw}(\alpha, \beta)$ to denote the formula $\phi_{\text{fb}} \wedge \phi_{(C1)} \wedge \phi_{(C2)} \wedge \phi_{(C3)}$. It specifies the shape of the encoding of data words in $([1, \alpha] \times \mathbb{N}^\beta)^+$ as stated below.

LEMMA 2.6. *Let $\mathfrak{h}$ be a heap. We have $\mathfrak{h} \models \text{dw}(\alpha, \beta)$ iff $\mathfrak{h}$ is an (α, β) -fishbone.*

Again, the proof is by an easy verification by using Lemma 2.5 and the correspondence between condition (Ci) and the formula $\phi_{(Ci)}$.

Given a data word $\mathfrak{dw} = (a^1, \mathfrak{d}_1^1, \dots, \mathfrak{d}_\beta^1) \cdots (a^L, \mathfrak{d}_1^L, \dots, \mathfrak{d}_\beta^L)$, we can associate a (α, β) -fishbone $\mathfrak{h}_{\mathfrak{dw}}$ with $(1 + \beta) \times L$ locations on the main path, say

$$l_0^1 \rightarrow l_1^1 \rightarrow \cdots \rightarrow l_\beta^1 \rightarrow \cdots \rightarrow l_0^L \rightarrow l_1^L \rightarrow \cdots \rightarrow l_\beta^L$$

such that

- for every $i \in [1, L]$, $\widetilde{\#l_0^i} = a^i + 2$,
- for all $i \in [1, L]$ and all $j \in [1, \beta]$, $\widetilde{\#l_j^i} = d_j^i + \alpha + 3$.

The heap $\mathfrak{h}_{\text{dw}}$ is unique modulo isomorphism. This natural encoding generalizes the encoding of finite words by heaps in [Brochenin et al. 2012, Section 3] (see also [Bansal et al. 2009]) while providing a much more concise representation. Note also that the encoding by itself is of no use since it is essential to be able to operate on it with the logical language at hand.

Conversely, given a (α, β) -fishbone $\mathfrak{h}$ with $(1 + \beta) \times L$ locations on the main path, say

$$l_0^1 \rightarrow l_1^1 \rightarrow \dots \rightarrow l_\beta^1 \rightarrow \dots \rightarrow l_0^L \rightarrow l_1^L \rightarrow \dots \rightarrow l_\beta^L$$

we associate a (unique) data word $\text{dw}_{\mathfrak{h}} = (a^1, d_1^1, \dots, d_\beta^1) \dots (a^L, d_1^L, \dots, d_\beta^L)$ such that for every $i \in [1, L]$, $a^i \stackrel{\text{def}}{=} \widetilde{\#l_0^i} - 2$ and for all $i \in [1, L]$ and all $j \in [1, \beta]$, $d_j^i \stackrel{\text{def}}{=} \widetilde{\#l_j^i} - \alpha - 3$.

LEMMA 2.7. *There is a one-to-one map between data words in $([1, \alpha] \times \mathbb{N}^\beta)^+$ and (α, β) -fishbone heaps (modulo isomorphism).*

The proof is then by an easy verification.

So, we have seen that finite words can be encoded in $1\text{SL}2(*)$, which allows us to establish that $1\text{SL}2(*)$ is NEXPTIME-hard since first-order logic restricted to two quantified variables on finite words (written $\text{FO}2_{\alpha,0}(<, +1, =)$ herein) is NEXPTIME-complete [Etessami et al. 1997]. Indeed, consider a sentence ϕ in that fragment of first-order logic. Let us define $t(\phi)$ such that ϕ is satisfiable iff $\text{dw}(\alpha, 0) \wedge t(\phi)$ is satisfiable in $1\text{SL}2(*)$.

We define the logarithmic-space translation t as follows ($i, j \in \{1, 2\}$).

- t is homomorphic for Boolean connectives,
- $t(u_i = u_j) \stackrel{\text{def}}{=} u_i = u_j$,
- $t(a(u_i)) \stackrel{\text{def}}{=} (\#u_i = a + 2)$,
- $t(u_i = 1 + (u_j)) \stackrel{\text{def}}{=} u_j \hookrightarrow u_i$,
- $t(u_i < u_j) \stackrel{\text{def}}{=} \text{ls}(u_i, u_j) \wedge u_i \neq u_j$,
- $t(\exists u_i \phi) \stackrel{\text{def}}{=} \exists u_i (\#u_i \geq 1) \wedge t(\phi)$.

Note that $\text{FO}2_{\alpha,0}(<, +1, =)$ and $1\text{SL}2(*)$ share the same number of quantified variables and $\text{ls}(u_i, u_j)$ can be expressed in $1\text{SL}2(*)$ (see Section 2.2). We do not provide the correctness proof herein since we can do much better than NEXPTIME-hardness by making a strong connection with Moszkowski's Interval Temporal Logic ITL (with the locality condition), see Section 3. However, we shall use a similar type of reduction in Section 4.6 with $\beta > 0$.

2.4. A modal logic for heaps

Let us conclude this section about logics for heaps, by introducing a new modal logic. We introduce a modal logic that is closely related to $1\text{SL}2$. Modal Logic for Heaps (MLH) is a multimodal logic in which models are exactly heap graphs and it does not contain propositional variables (as 1SL does not contain unary predicate symbols). In a sense, it is similar to Hennessy-Milner logic HML [Hennessy and Milner 1980] in which the only atomic formulae are truth constants. However, the language contains modal operators and separating connectives, which is a feature shared with the logics defined in [Courtault and Galmiche 2013]. We define below the formulae of

the modal logic MLH.

$$\phi ::= \perp \mid \neg \phi \mid \phi \wedge \phi \mid \Diamond \phi \mid \Diamond^{-1} \phi \mid \langle \neq \rangle \phi \mid \langle \star \rangle \phi \mid \phi * \phi \mid \phi \multimap \phi$$

Note that there are no quantified variables involved in formulae, which is a feature shared with most known propositional modal logics, see e.g. [Blackburn et al. 2001]. We write $\text{MLH}(\star)$ to denote the fragment of MLH without the magic wand operator $\multimap$.

A *model for MLH* $\mathfrak{M}$ is a pair $(\mathbb{N}, \mathfrak{R})$ such that $\mathfrak{R}$ is a binary relation on $\mathbb{N}$ that is finite and functional. Otherwise said, the models for MLH are heap graphs. Models for MLH could be defined as heaps but we prefer to stick to the most usual presentation for modal logics with frames. The satisfaction relation $\models$ is defined below and it provides a standard semantics for the modal operators and separating connectives (we omit the clauses for Boolean connectives):

- never $\mathfrak{M}, l \models \perp$,
- $\mathfrak{M}, l \models \Diamond \phi \stackrel{\text{def}}{\iff}$ there is l' such that $(l, l') \in \mathfrak{R}$ and $\mathfrak{M}, l' \models \phi$,
- $\mathfrak{M}, l \models \Diamond^{-1} \phi \stackrel{\text{def}}{\iff}$ there is l' such that $(l', l) \in \mathfrak{R}$ and $\mathfrak{M}, l' \models \phi$,
- $\mathfrak{M}, l \models \langle \star \rangle \phi \stackrel{\text{def}}{\iff}$ there is l' such that $(l, l') \in \mathfrak{R}^*$ and $\mathfrak{M}, l' \models \phi$ where $\mathfrak{R}^*$ is the reflexive and transitive closure of $\mathfrak{R}$,
- $\mathfrak{M}, l \models \langle \neq \rangle \phi \stackrel{\text{def}}{\iff}$ there is $l' \neq l$ such that $\mathfrak{M}, l' \models \phi$,
- $\mathfrak{M}, l \models \phi_1 * \phi_2 \stackrel{\text{def}}{\iff} (\mathbb{N}, \mathfrak{R}_1), l \models \phi_1$ and $(\mathbb{N}, \mathfrak{R}_2), l \models \phi_2$ for some partition $\{\mathfrak{R}_1, \mathfrak{R}_2\}$ of $\mathfrak{R}$,
- $\mathfrak{M}, l \models \phi_1 \multimap \phi_2 \stackrel{\text{def}}{\iff}$ for all models $\mathfrak{M}' = (\mathbb{N}, \mathfrak{R}')$ such that $\mathfrak{R} \cap \mathfrak{R}' = \emptyset$ and $\mathfrak{R} \cup \mathfrak{R}'$ is functional, $\mathfrak{M}', l \models \phi_1$ implies $(\mathbb{N}, \mathfrak{R} \cup \mathfrak{R}'), l \models \phi_2$.

We use the following standard abbreviations:

- $\langle U \rangle \phi \stackrel{\text{def}}{=} \phi \vee \langle \neq \rangle \phi$, $[U] \phi \stackrel{\text{def}}{=} \neg \langle U \rangle \neg \phi$,
- $[\neq] \phi \stackrel{\text{def}}{=} \neg \langle \neq \rangle \neg \phi$,
- $\Diamond_{\geq k}^{-1} \top \stackrel{\text{def}}{=} \Diamond^{-1} \top * \dots * \Diamond^{-1} \top$ ($k \geq 1$ times), $\Diamond_{\leq k-1}^{-1} \top \stackrel{\text{def}}{=} \neg \Diamond_{\geq k}^{-1} \top$,
- $\Diamond_{[k_1, k_2]}^{-1} \top \stackrel{\text{def}}{=} \Diamond_{\geq k_1}^{-1} \top \wedge \Diamond_{\leq k_2}^{-1} \top$, $\Diamond_{=k}^{-1} \top \stackrel{\text{def}}{=} \Diamond_{\geq k}^{-1} \top \wedge \Diamond_{\leq k}^{-1} \top$.

A formula ϕ is satisfiable whenever there is a model $\mathfrak{M}$ and a location l such that $\mathfrak{M}, l \models \phi$. The satisfiability problem for MLH is therefore defined as any such problem for modal logics. Note that MLH has forward and backward modalities as in Prior's tense logic (see e.g. [Prior 1967]), the inequality modal operator (see e.g. [de Rijke 1992]) and the transitive closure operator as in PDL (see e.g. [Harel et al. 2000]). The most non-standard feature of MLH is certainly the presence of the separating connectives. It is possible to design a relational translation from MLH formulae into 1SL2 formulae by recycling variables (only u_1 and u_2 are used, so $i \in \{1, 2\}$):

- t is homomorphic for the connectives $\neg$, $\wedge$, $*$ and $\multimap$,
- $t(\perp, u_i) \stackrel{\text{def}}{=} \perp$,
- $t(\Diamond \phi, u_i) \stackrel{\text{def}}{=} \exists u_{3-i} (u_i \hookrightarrow u_{3-i}) \wedge t(\phi, u_{3-i})$,
- $t(\Diamond^{-1} \phi, u_i) \stackrel{\text{def}}{=} \exists u_{3-i} (u_{3-i} \hookrightarrow u_i) \wedge t(\phi, u_{3-i})$,
- $t(\langle \neq \rangle \phi, u_i) \stackrel{\text{def}}{=} \exists u_{3-i} (u_i \neq u_{3-i}) \wedge t(\phi, u_{3-i})$,
- $t(\langle \star \rangle \phi, u_i) \stackrel{\text{def}}{=} \exists u_{3-i} \text{ls}(u_i, u_{3-i}) \wedge t(\phi, u_{3-i})$.

PROPOSITION 2.8. *A formula ϕ in MLH is satisfiable iff $\exists u_1 t(\phi, u_1)$ is satisfiable in 1SL2. Moreover, if ϕ is in $\text{MLH}(\star)$, then $\exists u_1 t(\phi, u_1)$ is in 1SL2($\star$).*

PROOF. (sketch) The proof is obtained as an obvious adaptation of the proof for the relational translation from modal logic K into FO2, see e.g. [Morgan 1976; van Ben-

them 1976; Blackburn et al. 2001]. Indeed, the models $(\mathbb{N}, \mathfrak{R})$ for MLH are heap graphs and therefore formulae in 1SL2 can be equivalently interpreted on MLH models; for instance, we get $(\mathbb{N}, \mathfrak{R}) \models_f u_1 \hookrightarrow u_2$ iff $(f(u_1), f(u_2)) \in \mathfrak{R}$. Similarly, $(\mathbb{N}, \mathfrak{R}) \models_f \phi_1 * \phi_2$ iff for all MLH models $(\mathbb{N}, \mathfrak{R}')$ such that $(\mathbb{N}, \mathfrak{R} \cup \mathfrak{R}')$ is an MLH model too and $(\mathbb{N}, \mathfrak{R}') \models_f \phi_1$, we have $(\mathbb{N}, \mathfrak{R} \cup \mathfrak{R}') \models_f \phi_2$.

Note that u_j is the only free variable in $t(\phi, u_j)$. The standard translation t is semantically faithful in the following sense: for all MLH models $(\mathbb{N}, \mathfrak{R})$, $l \in \mathbb{N}$ and formulae ϕ in MLH, we have $(\mathbb{N}, \mathfrak{R}), l \models \phi$ iff $(\mathbb{N}, \mathfrak{R}) \models_{[u_1 \mapsto l]} t(\phi, u_1)$. This is sufficient to establish Proposition 2.8.

We show that for all $i \in \{1, 2\}$, for all formulae ψ in MLH, for all MLH models $\mathfrak{M} = (\mathbb{N}, \mathfrak{R})$ and for $l \in \mathbb{N}$, we have $\mathfrak{M}, l \models \psi$ iff $\mathfrak{M} \models_{[u_i \mapsto l]} t(\psi, u_i)$. The proof is by structural induction. The base case for $\perp$ and the cases in the induction step for the Boolean connectives are straightforward. By way of example, let us provide the cases in the induction step for $\psi = \langle \star \rangle \psi'$ and for $\psi = \psi_1 * \psi_2$. The proof for the other cases is similar and quite standard.

Case $\psi = \langle \star \rangle \psi'$. The following are equivalent:

- $\mathfrak{M}, l \models \psi$,
- $\mathfrak{M}, l' \models \psi'$ for some $l' \in \mathfrak{R}^*(l)$ (by definition of $\models$),
- $\mathfrak{M} \models_{[u_{3-i} \mapsto l']} t(\psi', x_{3-i})$ for some $l' \in \mathbb{N}$ such that $l' \in \mathfrak{R}^*(l)$ (by the induction hypothesis),
- $\mathfrak{M} \models_{[u_i \mapsto l]} \exists u_{3-i} \text{ ls}(u_i, u_{3-i}) \wedge t(\psi', x_{3-i})$ (by definition of $\models$ in 1SL2 and by the fact that ls is the reachability predicate),
- $\mathfrak{M} \models_{[u_i \mapsto l]} t(\psi, u_i)$ (by definition of t).

Case $\psi = \psi_1 * \psi_2$. The following are equivalent:

- $\mathfrak{M}, l \models \psi$,
- $(\mathbb{N}, \mathfrak{R}_1), l \models \psi_1$ and $(\mathbb{N}, \mathfrak{R}_2), l \models \psi_2$ for some partition $\{\mathfrak{R}_1, \mathfrak{R}_2\}$ of $\mathfrak{R}$, (by definition of $\models$ in MLH),
- $(\mathbb{N}, \mathfrak{R}_1) \models_{[u_i \mapsto l]} t(\psi_1, u_i)$ and $(\mathbb{N}, \mathfrak{R}_2) \models_{[u_i \mapsto l]} t(\psi_2, u_i)$ for some partition $\{\mathfrak{R}_1, \mathfrak{R}_2\}$ of $\mathfrak{R}$, (by the induction hypothesis),
- $\mathfrak{M} \models_{[u_i \mapsto l]} t(\psi_1, u_i) * t(\psi_2, u_i)$ (by definition of the satisfaction relation in 1SL2)
- $\mathfrak{M} \models_{[u_i \mapsto l]} t(\psi, u_i)$ (by definition of t).

□

Modal logic MLH can be viewed as a fragment of 1SL2. Any formula $\psi_1 * \psi_2$ [resp. $\psi_1 \rightarrow \psi_2$] in $t(\phi, u_1)$ has at most one free variable. A similar restriction can be found in monodic fragments for first-order temporal logics, see e.g. [Degtyarev et al. 2002].

Since $\text{MLH}(\star)$ can be translated into $1\text{SL2}(\star)$ and $1\text{SL}(\star)$ is decidable [Brochenin et al. 2012, Corollary 3.3], we get decidability of $\text{MLH}(\star)$ as a corollary.

COROLLARY 2.9. *The satisfiability problem for $\text{MLH}(\star)$ is decidable.*

Note that to be more uniform, we could have added to the modal language the converse operators $\langle \neq \rangle^{-1}$ and $\langle \star \rangle^{-1}$. However, since the inequality relation is symmetric, $\langle \neq \rangle^{-1}\phi$ is logically equivalent to $\langle \neq \rangle\phi$. The above translation can be obviously extended with the modal operator $\langle \star \rangle^{-1}$ and therefore decidability holds also for this extension. However, we have introduced MLH mainly to establish non-elementarity of $\text{MLH}(\star)$ (shown below), refining the result for $1\text{SL2}(\star)$. We did not include $\langle \star \rangle^{-1}$ because the proof of non-elementarity result does not require it. By contrast, we do not know whether the satisfiability problem for MLH is decidable. As far as we know, the characterization of the computational complexity of MLH without separating connec-

tives is open too. This corresponds to a fragment of deterministic PDL with (restricted) graded modalities and inequality modality.

3. WHEN INTERVAL TEMPORAL LOGIC MEETS SEPARATION LOGIC

Interval-based temporal logics admit time intervals as first-class objects (instead of time points), and an early and classical study for reasoning about intervals can be found in [Allen 1983]. One of the most prominent interval-based logics is Propositional Interval Temporal Logic (PITL), introduced by Ben Moszkowski in [Moszkowski 1983] for the verification of hardware components. It contains the so-called ‘chop’ operation that consists of chopping an interval into two subintervals. This is of course reminiscent of separating conjunction in separation logic, and in this section we make a formal statement about this correspondence, in addition to deriving new complexity results. Before doing so, it is worth noting that even though most standard point-based temporal logics used in computer science are decidable (CTL, CTL*, ECTL*, etc.), undecidability is much more common in the realm of interval-based temporal logics. Below, we consider PITL in which propositional variables are interpreted under the *locality condition* and for which decidability is guaranteed but computational complexity is very high. This will allow us to derive similar bounds for 1SL2(*).

Below, we recall the main definitions about PITL under the locality condition and we explain why formulae from PITL can be faithfully translated into formulae in 1SL2(*), leading to insights about both formalisms and new complexity results. A similar analysis is presented for MLH(*) making new bridges between modal logic, interval-based temporal logics and separation logic.

3.1. Logic PITL

Given $\alpha \geq 1$, we consider the finite alphabet $\Sigma = [1, \alpha]$ and we write PITL_Σ to denote propositional interval temporal logic in which the models are non-empty finite words in Σ^+ . We write PITL instead of PITL_Σ when the finite alphabet Σ is clear from the context. Formulae for PITL_Σ are defined according to the following abstract grammar:

$$\phi ::= a \mid \text{pt} \mid \neg\phi \mid \phi \wedge \phi \mid \phi \mathbf{C} \phi$$

with $a \in \Sigma$. Even though elements of Σ are natural numbers (for the sake of technical convenience), we write a to denote such an arbitrary element to emphasize that a is a letter from a finite alphabet. Roughly speaking, a holds true at word w when a is the first letter of w . Similarly, the atomic formula pt holds true at a word w when the word w is only a single letter. The connective $\mathbf{C}$ is the *chop* operator, which chops a word.

Formally, we have a nonempty word $w \in \Sigma^+$, its length $|w|$, extractions of the i th letter w_i where $1 \leq i \leq |w|$, and extractions of nonempty subwords $w_{i..j} = w_i w_{i+1} \dots w_j$, where $1 \leq i \leq j \leq |w|$. We define a ternary relation *chops* on words:

$$\text{chops} \stackrel{\text{def}}{=} \{(w_1, w_2, w_3) \mid \exists a, w', w'' \text{ such that } w_1 = w'aw'', w_2 = w'a, w_3 = aw''\}$$

Observe that when a word w_1 is chopped into two subwords w_2 and w_3 , there is an overlap between the last letter of w_2 and the first letter of w_3 . For instance, $(abb, ab, bb) \in \text{chops}$ but $(ab, a, b) \notin \text{chops}$.

Let us define the satisfaction relation $\models$ for PITL_Σ between a word $w \in \Sigma^+$ and a formula ϕ :

- $w \models a \stackrel{\text{def}}{\iff} w_1 = a$ (here, w_1 denotes the first letter of w).
- $w \models \text{pt} \stackrel{\text{def}}{\iff} |w| = 1$.
- $w \models \neg\phi \stackrel{\text{def}}{\iff} w \not\models \phi$.
- $w \models \phi \wedge \psi \stackrel{\text{def}}{\iff} w \models \phi \text{ and } w \models \psi$.

— $w \models \phi \mathbf{C} \psi \stackrel{\text{def}}{\iff}$ there exist words w_1, w_2 such that $\text{chops}(w, w_1, w_2)$, $w_1 \models \phi$ and $w_2 \models \psi$.

The satisfiability problem for PITL_Σ consists in checking whether a PITL_Σ formula admits a model satisfying it. Note that the models are *nonempty, finite* words and the satisfaction of a letter on a word depends only on its first letter (the locality condition).

Two examples. Consider the alphabet Σ with two distinct letters a and b and the PITL_Σ formula below:

$$(b \mathbf{C} a) \mathbf{C} \neg \text{pt}$$

This formula is satisfiable; many words satisfy this formula, for example the word “bab”—the top-level chop is satisfied since $ba \models b \mathbf{C} a$ and $ab \models \neg \text{pt}$. This gives insight on how to specify a lower-bound on word length, by applying sufficiently many chops and $\neg \text{pt}$ to force a particular (minimum) length. Of course, $b \mathbf{C} a$ also enforces a minimum word length (of 2), but constrains also the word content.

Consider another example:

$$\text{pt} \wedge (a \mathbf{C} b)$$

For this formula to be satisfiable, there must exist a word w for which both $w \models \text{pt}$ and $w \models a \mathbf{C} b$. This is impossible, as the first implies $|w| = 1$, and there is no way to chop a single-letter word into subwords that satisfy both a and b ; the formula is unsatisfiable.

THEOREM 3.1. (see e.g. [Moszkowski 1983; 2004]) *Given $\alpha \geq 1$ and $\Sigma = [1, \alpha]$, the satisfiability problem for PITL_Σ is decidable, but with $\alpha \geq 2$ is not elementary recursive.*

3.2. Correspondence between words and heaps

From now on, we use the data word representation of Section 2.3. From Lemma 2.7, we know there is a fishbone heap h_w corresponding to each nonempty word $w \in \Sigma^+$. Let us define a relation $\sim$ that establishes this correspondence between words and their fishbone representations, adding also a correspondence between the empty word and the empty heap:

$$\sim \stackrel{\text{def}}{=} \{(w, h_w) \mid w \in \Sigma^+\} \cup \{(\epsilon, \emptyset)\}$$

Here, observe that:

- (1) $\sim$ is a bijection between the set of finite words in Σ^* and the set of (equivalence classes of isomorphic) $(\alpha, 0)$ -fishbone heaps augmented with the empty heap;
- (2) so, every word w is in $\text{dom}(\sim)$;
- (3) so, every $(\alpha, 0)$ -fishbone heap is in $\text{ran}(\sim)$;
- (4) so, if $w \sim h$, h is either empty or an $(\alpha, 0)$ -fishbone heap; and
- (5) if $w \sim h$, then w is empty iff $\text{dom}(h)$ is empty.

In this section, we will only employ $(\alpha, 0)$ -fishbone heaps, with $\alpha = \text{card}(\Sigma)$.

The correspondence between finite words in Σ^+ and $(\alpha, 0)$ -fishbone heaps satisfies a nice property as far as splitting a word into two disjoint subwords is concerned (which is a slight variant of chopping). Before making a formal statement, let us introduce the following notion.

A *clean cut* of a $(\alpha, 0)$ -fishbone heap h is a pair of $(\alpha, 0)$ -fishbone heaps (h_1, h_2) such that $h = h_1 \uplus h_2$, and for some words $w_1 \sim h_1$ and $w_2 \sim h_2$, we have $w_1 w_2 \sim h$. That is, a clean cut is one that neatly cleaves a heap representation of a word into two subheaps in correspondence with two subwords. Figure 4 illustrates examples of a clean cut and a non-clean cut on a fishbone heap. Informally, a non-clean cut is one that either results in one subheap (or both) no longer satisfying the $(\alpha, 0)$ -fishbone conditions, or that

results in subheaps that don't preserve predecessor counts and thus don't represent subwords of the original.

Fig. 4. A visual depiction of *clean* and *non-clean cuts*. Clockwise from left: the original $(\alpha, 0)$ -fishbone heap; a clean cut of the original heap; a non-clean cut of the original heap. Note that clean cuts must result in two $(\alpha, 0)$ -fishbone heaps. A non-clean cut may or may not do so; the figure depicts a non-clean cut that does result in two $(\alpha, 0)$ -fishbone heaps.

LEMMA 3.2. *Let $w \sim h$ with $w = w_1 w_2 \in \Sigma^*$. There exist heaps h_1 and h_2 such that $h = h_1 \uplus h_2$, $w_1 \sim h_1$, and $w_2 \sim h_2$.*

PROOF. Suppose that $w \sim h$ and $w = w_1 w_2 \in \Sigma^*$. If $w_1 = \varepsilon$ or $w_2 = \varepsilon$, then the proof is by an easy verification with h equal to h_1 or h_2 respectively. In particular, if $w = \varepsilon$, then h is the empty heap and therefore $w_1 = w_2 = \varepsilon$ and $h_1 = h_2 = \emptyset$, which satisfies the statement.

Otherwise suppose that $w = a_1 \cdots a_K \in \Sigma^+$, $w_1 = a_1 \cdots a_{K'} \in \Sigma^+$, $w_2 = a_{K'+1} \cdots a_K \in \Sigma^+$ ($K > K'$). Since w is nonempty and $w \sim h$, h is a fishbone heap and the main path of h is of the form $l_1 \rightarrow l_2 \rightarrow \cdots \rightarrow l_K$ and for every $i \in [1, K]$, $\sharp l_i = a_i + 2$. Let h_1 be the subheap of h whose domain is $\{l' \in \mathbb{N} : l' \text{ is an ancestor of } l_{K'} \text{ in } h\}$, and let h_2 be the unique heap such that $h = h_1 \uplus h_2$. It is easy to show that $w_1 \sim h_1$ and $w_2 \sim h_2$. Moreover, it is not difficult to see that (h_1, h_2) is a clean cut of h . $\square$

Lemma 3.2 entails the following lemma, that will be useful to show the correctness of our reduction from PITL_Σ into $\text{1SL2}(\ast)$. It is tailored to the semantics of the chop operator in PITL_Σ .

LEMMA 3.3. *For all letters $a, b \in \Sigma$, words $w \in \Sigma^+$ and $w', w'' \in \Sigma^*$, and heaps h such that $w \sim h$ and chops $(aw, aw'b, bw'')$, there exist heaps h_1, h_2 such that $w'b \sim h_1$, $w'' \sim h_2$, and $h = h_1 \uplus h_2$.*

3.3. A reduction and its three ways to chop

In this section, we present a satisfiability-preserving translation of PITL_Σ into $\text{ISL2}(\ast)$. This translation hinges on the insight that the chop operation is very similar to the separating conjunction in separation logic. However, the correspondence is not an exact one: the connective $\mathbf{C}$ of PITL_Σ does not cut into disjoint pieces, but rather preserves one letter on both sides, in a sense “duplicating” the letter upon which the chop operates.

Fig. 5. The correspondence between PITL’s chop ‘ $\mathbf{C}$ ’ and separation logic’s separating conjunction ‘ $\ast$ ’ (before and after).

To handle this discrepancy, our translation uses the standard separating conjunction on heaps, but internally carries a “ghost letter” (a parameter to the translation) on one side to represent this “lost” letter. In the translation, we denote this ghost letter parameter $a \in \Sigma$. Figure 5 illustrates how a chop operation on words is translated into a separation on heaps. It is worth noting that we must always obtain a clean cut from the original heap.

Before presenting the formal definition of the translation, let us present a formula that allows us to perform a clean cut for which one of the subheaps contains all the ancestors of $f(u)$. Such a formula will be used in the translation and this is the purpose of Lemma 3.4.

LEMMA 3.4. *Given a fishbone heap h and a word w such that $w \sim h$, and an assignment f such that $f(u)$ is a location on the main path of h with $h \models_f \text{alloc}(u)$, any pair of heaps (h_1, h_2) such that $h = h_1 \uplus h_2$, $h_1 \models_f \text{dw}(\alpha, 0) \wedge \neg \text{alloc}(u)$, and $h_2 \models_f \text{dw}(\alpha, 0) \wedge \#u = 0$, is a clean cut of h .*

PROOF. Since $h_1 \models_f \text{dw}(\alpha, 0)$ and $h_2 \models_f \text{dw}(\alpha, 0)$, we know the heaps h_1 and h_2 are $(\alpha, 0)$ -fishbones. Fishbones are single components, so we know that h must be separated into exactly two connected components. It remains to analyze precisely how h can be separated into two fishbones, and to show that it must be a clean cut.

We know $l = f(u)$ is on the main path of h , so that means $h \models_f \#u > 0$. Since $h_2 \models_f \#u = 0$, that means l must have the same number of predecessors in h as it does in h_1 . We know $h_1 \models_f \neg \text{alloc}(u)$, and we know l has at least one predecessor in h_1 . Therefore, l is on the main path of h_1 . We know $h_2 \models_f \#u = 0$, so l is not on the main path of h_2 . However, it is allocated (since $h \models_f \text{alloc}(u)$ and $h_1 \models_f \neg \text{alloc}(u)$), so its successor (call the location l') is on the main path of h_2 . Let $f' = f[\bar{u} \mapsto l']$. Now, note that $h \models_{f'} u \hookrightarrow \bar{u}$ and $h_1 \not\models_{f'} u \hookrightarrow \bar{u}$, and recall that that on a fishbone, no two predecessors of an element can both have predecessors (fb3). Therefore, l' must have the same number of predecessors in h_2 as it did in h , and none of these predecessors can be on the main path.

Thus l is the final location on the main path of h_1 , and l' is the first location on the main path of h_2 . Further, l has 0 predecessors in h_2 and the same number of predecessors in h and h_1 . l' has 0 predecessors in h_1 and the same number of predecessors in h and h_2 .

Putting the above together, since l and l' are positions on the main path of h such that $h(l) = l'$, and since $l \notin \text{dom}(h_1)$, $l \in \text{dom}(h_2)$, $l \in \text{ran}(h_1)$, $l \notin \text{ran}(h_2)$, we must have a clean cut. $\square$

We reduce a PITL_Σ formula ϕ to a $1\text{SL}2(*)$ formula $t(\phi)$ with the help of the main translation $t(\cdot)$. We use the auxiliary translation map $t_a(\cdot)$ parameterized by a ghost letter a . The three disjuncts in the translation of $\phi \mathbf{C} \psi$ correspond to three types of chopping of w that leads to three ways of separating the heap h (assuming that $w \sim h$):

- (1) When $(w, aw_1b, bw_2) \in \text{chops}$ and the ghost letter is a , the heap h is separated into the heap h_1 with $w_1b \sim h_1$ (with ghost letter a) and into the heap h_2 with $w_2 \sim h_2$ (with ghost letter b).
- (2) When $(w, w, b) \in \text{chops}$ and the ghost letter is a , the heap h is separated into itself (again with ghost letter a) and into the empty heap (with ghost letter b).
- (3) When $(w, a, w) \in \text{chops}$ and the ghost letter is a , the heap h is separated into the empty heap (with ghost letter a) and into itself (again with ghost letter a).

These are the three possible cases and the rest of the translation is quite straightforward.

$$t(\phi) \stackrel{\text{def}}{=} (\mathbf{dw}(\alpha, 0) \vee \mathbf{emp}) \wedge \bigvee_{a \in \Sigma} t_a(\phi)$$

$$t_a(b) \stackrel{\text{def}}{=} \begin{cases} \top & \text{if } b = a \\ \perp & \text{if } b \neq a \end{cases}$$

$$t_a(\mathbf{pt}) \stackrel{\text{def}}{=} \mathbf{emp}$$

$$t_a(\neg\phi) \stackrel{\text{def}}{=} \neg t_a(\phi)$$

$$t_a(\phi \wedge \psi) \stackrel{\text{def}}{=} t_a(\phi) \wedge t_a(\psi)$$

$$t_a(\phi \mathbf{C} \psi) \stackrel{\text{def}}{=} \mathbf{chop1}_a \vee \mathbf{chop2}_a \vee \mathbf{chop3}_a \quad (\text{with the three formulae as defined below})$$

$$\mathbf{chop1}_a \stackrel{\text{def}}{=} \bigvee_{b \in \Sigma} \exists u \left(\#u = b + 2 \wedge \right.$$

$$\left. \left[\mathbf{dw}(\alpha, 0) \wedge \neg \mathbf{alloc}(u) \wedge t_a(\phi) * \mathbf{dw}(\alpha, 0) \wedge \#u = 0 \wedge t_b(\psi) \right] \right)$$

$$\mathbf{chop2}_a \stackrel{\text{def}}{=} \bigvee_{b \in \Sigma} (\exists u \mathbf{last}(u) \wedge \#u = b + 2) \wedge \left[t_a(\phi) * \mathbf{emp} \wedge t_b(\psi) \right]$$

$$\mathbf{chop3}_a \stackrel{\text{def}}{=} \mathbf{emp} \wedge t_a(\phi) * t_a(\psi)$$

In full generality, $t_a(\cdot)$ is also parameterized by the alphabet Σ (see the clause for formulae with outermost chop operator $\mathbf{C}$) and the formulae $\mathbf{chop1}_a$, $\mathbf{chop2}_a$, and $\mathbf{chop3}_a$ are parameterized by $\phi \mathbf{C} \psi$. Clearly the translation $t(\cdot)$ can only produce 1SL2(*) formulae, as the right-hand side of each translation step above is in 1SL2(*). Note also that $t_a(\phi)$ always produces a closed formula (i.e., without free occurrences of individual variables).

The correctness of the translation is stated below, making completely explicit the role of the ghost letter in the translation process.

LEMMA 3.5. *Let $a \in \Sigma$, $w \in \Sigma^*$, and h be a heap such that $w \sim h$. For every PITLE $_{\Sigma}$ formula ϕ , we have $aw \models \phi$ iff $h \models t_a(\phi)$.*

Remark 3.6. Note that since $t_a(\phi)$ has no free occurrences of individual variables, in Lemma 3.5, there is no need to specify what the assignments are.

PROOF. The proof is by structural induction.
The base cases are:

- $t_a(b)$. ($\Rightarrow$) If $aw \models b$, then $a = b$. Clearly $h \models \top$, so $h \models t_a(b)$.
($\Leftarrow$) $h \models t_a(b)$, so it must be that $a = b$ (since $h \not\models \perp$). Thus $aw \models b$.
- $t_a(\mathbf{pt})$. ($\Rightarrow$) If $aw \models \mathbf{pt}$, then $w = \epsilon$. Thus $h = \emptyset$. Since $t_a(\mathbf{pt}) = \mathbf{emp}$, we have $h \models t_a(\mathbf{pt})$.
($\Leftarrow$) $h \models t_a(\mathbf{pt})$, and since $t_a(\mathbf{pt}) = \mathbf{emp}$, we know $h = \emptyset$. Thus $w = \epsilon$, so clearly $aw \models \mathbf{pt}$.

The inductive cases are:

- $t_a(\neg\phi)$. ($\Rightarrow$) $aw \models \neg\phi$, so $aw \not\models \phi$. By IH, $h \not\models t_a(\phi)$, so $h \models \neg t_a(\phi)$ and precisely $\neg t_a(\phi) = t_a(\neg\phi)$.
($\Leftarrow$) $h \models \neg t_a(\phi)$, so $h \not\models t_a(\phi)$. By IH, $aw \not\models \phi$, so $aw \models \neg\phi$.

- $t_a(\phi \wedge \psi)$. ($\Rightarrow$) $aw \models \phi \wedge \psi$, so $aw \models \phi$ and $aw \models \psi$. By IH, $h \models t_a(\phi)$ and $h \models t_a(\psi)$, so then $h \models t_a(\phi) \wedge t_a(\psi)$. We have $h \models t_a(\phi \wedge \psi)$.
- ($\Leftarrow$) $h \models t_a(\phi \wedge \psi)$, so $h \models t_a(\phi) \wedge t_a(\psi)$. Then $h \models t_a(\phi)$ and $h \models t_a(\psi)$. By IH, $aw \models \phi$ and $aw \models \psi$, so $aw \models \phi \wedge \psi$.
- $t_a(\phi \mathbf{C} \psi)$. ($\Rightarrow$) If $aw \models \phi \mathbf{C} \psi$, then there are w_1, w_2 such that $\text{chops}(aw, w_1, w_2)$, $w_1 \models \phi$, and $w_2 \models \psi$. From the definition of chops , we have, further, that there are b, w', w'' such that $aw = w'b w''$, $w_1 = w'b$, and $w_2 = b w''$.

Case 1. $w' = \epsilon$ (the chop occurs at the first position). Then $aw = b w''$ and $w_1 = b = a$, so $w = w''$. Since $a \models \phi$ and $aw \models \psi$, we have by IH that $\emptyset \models t_a(\phi)$ and $h \models t_a(\psi)$. Then $h \models \text{emp} \wedge t_a(\phi) * t_a(\psi)$, so $h \models \text{chop3}_a$, and $h \models t_a(\phi \mathbf{C} \psi)$.

Case 2. $w' \neq \epsilon$, $w'' = \epsilon$ (the chop occurs at the last position). $aw = w_1$, so $aw \models \phi$. $w_2 = b$, so $b \models \psi$. By IH, we then have $h \models t_a(\phi)$ and $\emptyset \models t_b(\psi)$. Thus $h \models t_a(\phi) * \text{emp} \wedge t_b(\psi)$. Further, since the last letter of w is b and $w \sim h$, the last location on the main path of h must have $b + 2$ predecessors. Therefore, $h \models \exists u \text{ last}(u) \wedge \#u = b + 2$. So $h \models \text{chop2}_a$, and thus we have $h \models t_a(\phi \mathbf{C} \psi)$.

Case 3. $w' \neq \epsilon$, $w'' \neq \epsilon$ (the chop occurs in the middle of the word). We know w_1 is nonempty and starts with a , so let w'_1 be a word such that $w_1 = aw'_1$. From Lemma 3.3, we have heaps h_1, h_2 such that $h = h_1 * h_2$, $w'_1 \sim h_1$, and $w'' \sim h_2$. Observe that, since they are in $\text{ran}(\sim)$, h_1 and h_2 are $(\alpha, 0)$ -fishbones. Since $aw'_1 \models \phi$ and $bw'' \models \psi$, then by IH, we have $h_1 \models t_a(\phi)$ and $h_2 \models t_b(\psi)$. Now, let l be the location of the chop, so that l is the last position on the main path of h_1 and a predecessor of the first position on the main path of h_2 . Let f be an assignment such that $f(u) = l$. Since $h_1 \models_f \text{dw}(\alpha, 0) \wedge \neg \text{alloc}(u) \wedge t_a(\phi)$ and $h_2 \models_f \text{dw}(\alpha, 0) \wedge \#u = 0 \wedge t_b(\psi)$, we have $h \models_f [\text{dw}(\alpha, 0) \wedge \neg \text{alloc}(u) \wedge t_a(\phi) * \text{dw}(\alpha, 0) \wedge \#u = 0 \wedge t_b(\psi)]$. Further, since in h , location l encodes the letter b , it has $b + 2$ predecessors, so we have $h \models_f \#u = b + 2$, so $h \models \text{chop1}_a$. Then $h \models t_a(\phi \mathbf{C} \psi)$.

($\Leftarrow$) If we have that $h \models t_a(\phi \mathbf{C} \psi)$, then one of the disjuncts in the translation holds, depending on where the separation applies in the heap and the content at that position. We consider three cases based on which of the constraints (chop1_a , chop2_a , or chop3_a) applies.

Case 1. The chop1_a constraint applies (neither subheap is empty). We know it holds for some $b \in \Sigma$, and that the existentially quantified formulae holds for some location; w.l.o.g., assume that the case $b \in \Sigma$ holds, with location l . Let f be an assignment such that $f(u) = l$. We know that $h \models_f \#u = b + 2$, and that h can be separated into two disjoint subheaps h_1, h_2 such that $h_1 \models_f \text{dw}(\alpha, 0) \wedge \neg \text{alloc}(u) \wedge t_a(\phi)$ and $h_2 \models_f \text{dw}(\alpha, 0) \wedge \#u = 0 \wedge t_b(\psi)$. This cut must be a *clean cut* by Lemma 3.4. By IH, then, we have words w_1, w_2 such that $w_1 \sim h_1$, $w_2 \sim h_2$, $aw_1 \models \phi$, and $bw_2 \models \psi$. We know w_1 ends with a letter b , since u is the last position on the main path of h_1 and has $b + 2$ predecessors. So then $\text{chops}(aw_1 w_2, aw_1, bw_2)$ holds. Thus $aw_1 w_2 \models \phi \mathbf{C} \psi$.

Case 2. The chop2_a constraint applies (the right subheap is empty). We know it holds for some $b \in \Sigma$, and that the existentially quantified formula holds for some location; w.l.o.g., then, assume that the case $b \in \Sigma$ holds, at some location l . We know that heap h can be separated into two disjoint subheaps h_1, h_2 such that $h_1 \models t_a(\phi)$, $h_2 = \emptyset$, and $h_2 \models t_b(\psi)$. Thus $h = h_1$, and by IH we have $aw \models \phi$ and $b \models \psi$. Next, let f be an assignment such that $f(u) = l$. We know $h \models_f \text{last}(u)$ and that $h \models_f \#u = b + 2$.

Since $w \sim h$, the last letter of w is b . So $\text{chops}(aw, aw, b)$, and therefore $aw \models \phi \mathbf{C} \psi$.

Case 3. The chop3_a constraint applies (the left subheap is empty). We then have $h_1 = \emptyset$, $\emptyset \models t_a(\phi)$, and $h_2 \models t_a(\psi)$. Since $h = h_2$, by IH we have $a \models \phi$ and $aw \models \psi$. Then $\text{chops}(aw, a, aw)$; thus $aw \models \phi \mathbf{C} \psi$.

Therefore, $aw \models \phi$ iff $h \models t_a(\phi)$ $\square$

As a result, we obtain a reduction between the satisfiability problems, as stated below.

LEMMA 3.7. *Given $\alpha \geq 1$ and $\Sigma = [1, \alpha]$, a PITL_Σ formula ϕ is satisfiable if and only if the $1\text{SL2}(\ast)$ formula $t(\phi)$ is satisfiable.*

PROOF. ($\Rightarrow$) Suppose that ϕ is satisfiable. This means that there exists a nonempty word w such that $w \models \phi$. The word w can be written in the form $w = aw'$ for some letter a . If $w' = \epsilon$, we have $w' \sim \emptyset$ and by Lemma 3.5, we have $\emptyset \models \text{emp} \wedge t_a(\phi)$. So $\emptyset \models t(\phi)$ and therefore $t(\phi)$ is satisfiable. If $w' \neq \epsilon$, then there is a $(\alpha, 0)$ -fishbone heap h' such that $w' \sim h'$. By Lemma 3.5, we have $h' \models \text{dw}(\alpha, 0) \wedge t_a(\phi)$. So $h' \models t(\phi)$ and therefore $t(\phi)$ is satisfiable.

($\Leftarrow$) If $t(\phi)$ is satisfiable, then there exists a heap h such that $h \models (\text{dw}(\alpha, 0) \vee \text{emp}) \wedge \bigvee_{a \in \Sigma} t_a(\phi)$. If $h \models \text{emp} \wedge t_a(\phi)$ for some letter a , then by Lemma 3.5, we have $a \models \phi$. Otherwise, if $h \models \text{dw}(\alpha, 0) \wedge t_a(\phi)$, then h is an $(\alpha, 0)$ -fishbone heap by Lemma 2.6 and then there is a word w such that $w \sim h$ such that by Lemma 3.5, we have $aw \models \phi$. In both cases, ϕ is a satisfiable formula in PITL_Σ . $\square$

THEOREM 3.8. *The satisfiability problem for $1\text{SL2}(\ast)$ is decidable but not elementary recursive.*

PROOF. Satisfiability for PITL_Σ is known to be decidable with non-elementary complexity when Σ has at least two elements, see e.g. [Moszkowski 1983; 2004], and $1\text{SL}(\ast)$ is decidable [Brochenin et al. 2012]. From the correctness of our translation $t(\cdot)$ of PITL_Σ to $1\text{SL2}(\ast)$ (Lemma 3.7), we then conclude that $1\text{SL2}(\ast)$ is decidable but not elementary recursive. Note that the map $t(\cdot)$ may require exponential time and space in the size of the input formula in the worst-case but this is still fine to establish that $1\text{SL2}(\ast)$ is not elementary recursive, since this adds only a single exponential. $\square$

As mentioned earlier, Theorem 3.8 refines the non-elementarity result for $1\text{SL}(\ast)$ established in [Brochenin et al. 2012].

Remark 3.9. The reduction from PITL to $1\text{SL2}(\ast)$ provided in this section allows us to underline the common features of both formalisms. However, non-elementarity of $1\text{SL2}(\ast)$ can be established in a slightly different way as explained below. First, non-elementarity of PITL is due to Dexter Kozen (see e.g. [Moszkowski 2004, Appendix A.3])², and the proof is by reduction from the nonemptiness problem of regular expressions built over a binary alphabet with union, concatenation and complement [Stockmeyer 1974]. Nonelementarity of $1\text{SL2}(\ast)$ can be obtained by defining a similar reduction, but this is of course less insightful to understand the relationships between interval temporal logic and separation logic. Alternatively, it is also possible to consider the variant of PITL in which the chop operator does not share a letter, since this variant is of identical expressive power and complexity. In that way, we may avoid the introduction of the ghost letter but at the cost of introducing empty models (which may occur when chopping has no sharing) and of using a less standard interval temporal

²We thank Ben Moszkowski for pointing us to this fact.

logic. So, the current reduction from PITL is quite an attractive option to relate the logics. Finally, as noted in [Moszkowski 2004, Appendix A], complexity results about PITL presented in [Moszkowski 1983] were obtained in collaboration with Joseph Halpern.

In Section 3.4 below, we establish an even stronger result (see Theorem 3.12). The proof uses the same principles as for the proof of Theorem 3.8 and we only need to express the properties in *modal lingua*.

3.4. A refinement with the modal fragment of 1SL2(*)

In this section, we show that the satisfiability problem for MLH(*) is decidable but it is not elementary recursive. Decidability is due to the fact that the standard translation leads to formulae in 1SL2(*), see Section 2.4. In order to establish the lower bound, we express in MLH(*) all the properties that were useful to translate PITL_Σ formulae into 1SL2(*). For instance, note that the empty heap is the only heap validating the formula $([U] \neg \Diamond \top)$. Similarly, a location with at least one predecessor and with no successor (for instance, last location on the main path in a fishbone heap) satisfies the formula $(\Diamond^{-1} \top \wedge \neg \Diamond \top)$.

More interestingly, the formula in 1SL2(*) characterizing the (α, β) -fishbone heaps has a modal counterpart. Let us consider the following formulae.

- The formula $\phi_{\text{fb}}^{\square}$ defined below is designed exactly as the formula ϕ_{fb} (see Section 2.3).

$$\begin{aligned} & (\langle U \rangle \Diamond \top) \wedge \\ & \langle U \rangle ((\Diamond^{-1} \top \wedge \neg \Diamond \top) \wedge [\neq] \neg (\Diamond^{-1} \top \wedge \neg \Diamond \top)) \wedge [U] (\Diamond \top \Rightarrow \langle \star \rangle (\Diamond^{-1} \top \wedge \neg \Diamond \top)) \wedge \\ & (\neg \langle U \rangle (\Diamond^{-1} \Diamond^{-1} \top * \Diamond^{-1} \Diamond^{-1} \top)) \end{aligned}$$

This is a faithful translation except that we use the specification language MLH(*).

- The formula $\phi_{(C1)}^{\square}$ defined below is also designed exactly as the formula $\phi_{(C1)}$ (see Section 2.3).

$$\langle U \rangle ((\Diamond^{-1} \top) \wedge (\neg \Diamond^{-1} \Diamond^{-1} \top) \wedge \Diamond_{[3, \alpha+3]}^{-1} \top)$$

- The formula $\phi_{(C2)}^{\square}$ is equal to $[U] (\Diamond_{[3, \alpha+3]}^{-1} \top \Rightarrow \bigwedge_{i \in [1, \beta]} \overbrace{\Diamond \dots \Diamond}^{i \text{ times}} \Diamond_{\geq \alpha+3}^{-1} \top)$
- The formula $\phi_{(C3)}^{\square}$ is defined below:

$$[U] (\Diamond_{[3, \alpha+3]}^{-1} \top \Rightarrow (\neg \overbrace{\Diamond \dots \Diamond}^{\beta+1 \text{ times}} \top) \vee \overbrace{\Diamond \dots \Diamond}^{\beta+1 \text{ times}} (\Diamond_{[3, \alpha+3]}^{-1} \top))$$

We write $\text{dw}^{\square}(\alpha, \beta)$ to denote the formula $\phi_{\text{fb}}^{\square} \wedge \phi_{(C1)}^{\square} \wedge \phi_{(C2)}^{\square} \wedge \phi_{(C3)}^{\square}$. It specifies the shape of the encoding of data words in $([1, \alpha] \times \mathbb{N}^{\beta})^{+}$ as stated below. Note that since $\text{dw}^{\square}(\alpha, \beta)$ is a Boolean combination of formulae whose outermost connectives are $[U]$ or $\langle U \rangle$, then $\text{dw}^{\square}(\alpha, \beta)$ holds true at some location iff $\text{dw}^{\square}(\alpha, \beta)$ holds true at any location.

LEMMA 3.10. *Let $\mathfrak{M} = (\mathbb{N}, \mathfrak{R})$ be a model for MLH. $\mathfrak{M}, \mathfrak{l} \models \text{dw}^{\square}(\alpha, \beta)$ for some location $\mathfrak{l}$ iff $\mathfrak{M}$ is the graph of an (α, β) -fishbone heap.*

Again, the proof is by an easy verification by using Lemma 2.5 and the correspondence between condition (Ci) and the formula $\phi_{(Ci)}^{\square}$. In the rest of this section we are back to the case $\beta = 0$.

Given a formula ϕ in PITL_Σ with $\Sigma = [1, \alpha]$, we define a modal formula $t^{\square}(\phi)$ such that ϕ is satisfiable iff $t^{\square}(\phi)$ is satisfiable. Actually, the modal formula $t^{\square}(\phi)$ will express exactly the same properties as in the translation into 1SL2(*). For instance, $t^{\square}(\phi)$

is precisely the formula below:

$$(\mathbf{dw}^\square(\alpha, 0) \vee ([U] \neg \Diamond \top)) \wedge (\bigvee_{a \in \Sigma} t_a^\square(\phi))$$

The formula $t_a^\square(\phi)$ is defined inductively as follows.

- $t_a^\square(a) \stackrel{\text{def}}{=} \top$ and $t_a^\square(b) \stackrel{\text{def}}{=} \perp$ for every letter $b \in \Sigma \setminus \{a\}$.
- $t_a^\square(\cdot)$ is homomorphic for Boolean connectives.
- $t_a^\square(\text{pt}) \stackrel{\text{def}}{=} ([U] \neg \Diamond \top)$.
- The formula $t_a^\square(\phi \mathbf{C} \psi)$ is defined as $\text{chop1}_a^\square \vee \text{chop2}_a^\square \vee \text{chop3}_a^\square$ where:
 - $\text{chop1}_a^\square \stackrel{\text{def}}{=} \bigvee_{b \in \Sigma} \langle U \rangle ((\Diamond_{=b+2}^{-1} \top \wedge \mathbf{dw}^\square(\alpha, 0) \wedge \neg \Diamond \top \wedge t_a^\square(\phi)) * (\mathbf{dw}^\square(\alpha, 0) \wedge \neg \Diamond^{-1} \top \wedge t_b^\square(\psi)))$,
 - $\text{chop2}_a^\square \stackrel{\text{def}}{=} (\bigvee_{b \in \Sigma} \langle U \rangle ((\Diamond^{-1} \top \wedge \neg \Diamond \top) \wedge \Diamond_{=b+2}^{-1} \top) \wedge (t_a^\square(\phi) * (t_b^\square(\psi) \wedge ([U] \neg \Diamond \top))))$,
 - $\text{chop3}_a^\square \stackrel{\text{def}}{=} ((t_a^\square(\phi) \wedge ([U] \neg \Diamond \top)) * t_a^\square(\psi))$.

LEMMA 3.11. *Let $\alpha \geq 1$, $\Sigma = [1, \alpha]$, ϕ be a PITL_Σ formula and $t^\square(\phi)$ be its translation in MLH . We have ϕ is satisfiable iff $t^\square(\phi)$ is satisfiable.*

The proof goes as in the case for the direct translation into $\text{1SL2}(\ast)$ since the modal subformulae express exactly the same properties. Therefore, we can refine Theorem 3.8 as follows.

THEOREM 3.12. *The satisfiability problem for $\text{MLH}(\ast)$ is decidable but not elementary recursive.*

Interestingly, we do not know the decidability status for full MLH (i.e., with the magic wand operator).

4. HOW TWO VARIABLES WITH THE MAGIC WAND ENCODE RUNS

In this section, we consider the logic 1SL2 (i.e., equipped with both the separating conjunction *and* the separating implication), and we prove its undecidability. In order to show that the 1SL2 satisfiability problem is undecidable, we reduce the halting problem for Minsky machines [Minsky 1967].

4.1. Minsky machines in a nutshell

Let M be a Minsky machine with $\alpha \geq 1$ instructions, where 1 is the initial instruction and α is the halting instruction. Machine M has two counters C_1 and C_2 and the instructions are of the following types:

- (a) $I: C_j := C_j + 1; \text{goto } J$.
- (b) $I: \text{if } C_j = 0 \text{ then goto } J_1 \text{ else } (C_j := C_j - 1; \text{goto } J_2)$.
- (c) $\alpha: \text{halt}$.

($j \in [1, 2]$, $I \in [1, \alpha - 1]$, $J, J_1, J_2 \in [1, \alpha]$)

Machine M halts if there is a run of the form

$$(I_0, c_0^1, c_0^2), (I_1, c_1^1, c_1^2), \dots, (I_L, c_L^1, c_L^2)$$

such that $(I_i, c_i^1, c_i^2) \in [1, \alpha] \times \mathbb{N}^2$ ($i \in [1, L]$), the succession of configurations respects the instructions (in the obvious way), $I_0 = 1$, $I_L = \alpha$, and $c_0^1 = c_0^2 = 0$. The halting problem consists in checking whether a machine halts and it is known to be undecidable, see

e.g. [Minsky 1967]. Clearly, a halting run is a data word $\mathfrak{dw}$ of dimension 2 such that the first letter is 1 and the last letter is α .

When a Minsky machine M has $\alpha \geq 1$ instructions, any run starting from the initial instruction 1 and ending by the halting instruction α (there is a single such run since M is deterministic) is a data word of dimension two over the finite alphabet $[1, \alpha]$. The main and obvious idea to get undecidability is to show how 1SL2 can characterize heaps encoding data words of dimension two corresponding to halting runs of M .

4.2. Roadmap

Let us start by explaining how the rest of the section is structured. The next paragraph describes how initial and final conditions on the run are encoded in 1SL2; typically the run starts by instruction 1 and possibly ends by the halting instruction α . Then, the next two paragraphs deal with the description of two problems that we need to tackle. (Section 4.3 and Section 4.4 contain the technical developments.) In a nutshell, the first problem consists of being able to compare the numbers of predecessors of two locations (which corresponds to comparison of two counter values in Minsky machines), whereas the second problem is related to the fact that these two locations may be separated by distance three along the main path. Both issues stem from the fact that we have only two individual variables at hand. (It is already known how to solve these problems with an unbounded number of variables [Brochenin et al. 2012].) Section 4.5 provides the final definition of the reduction from the halting problem for Minsky machines to the satisfiability problem for 1SL2 whereas Section 4.6 presents the idea of an alternative undecidability proof by reduction from a first-order logic on data words (see Section 2.3). These proofs share essential building blocks, for instance the constructions allowing to compare numbers of predecessors. However, this provides different points of view as well.

Limit conditions. We have seen that $(\alpha, 2)$ -fishbone heaps can be characterized thanks to the formula $\mathfrak{dw}(\alpha, 2)$. Obviously, more constraints need to be expressed, typically those related to the first instruction and those related to the halting instruction. Let us start by specifying the limit conditions thanks to the formulae ϕ_{first} and ϕ_{last} below.

- The first three locations on the main path have 3, $\alpha + 3$, and $\alpha + 3$ predecessors respectively:

$$\phi_{\text{first}} \stackrel{\text{def}}{=} \exists u \text{ first}(u) \wedge (\#u = 3) \wedge (\#u^{+1} = \alpha + 3) \wedge (\#u^{+2} = \alpha + 3)$$

- The main path encoding the run ends by a configuration with the halting instruction:

$$\phi_{\text{last}} \stackrel{\text{def}}{=} \exists u ((\#u = \alpha + 2) \wedge (\#u^{+2} \geq 0) \wedge \neg(\#u^{+3} \geq 0))$$

Let us call ϕ^* the conjunction of $\mathfrak{dw}(\alpha, 2) \wedge \phi_{\text{first}} \wedge \phi_{\text{last}}$. It specifies the shape of the encoding of the run without taking care of the constraints about counter values and instruction counter.

LEMMA 4.1. *Let $\mathfrak{h}$ be a heap. $\mathfrak{h} \models \phi^*$ iff $\mathfrak{h}$ encodes a data word*

$$\mathfrak{dw} = (\mathfrak{a}^1, \mathfrak{d}_1^1, \mathfrak{d}_2^1) \cdots (\mathfrak{a}^L, \mathfrak{d}_1^L, \mathfrak{d}_2^L)$$

such that $\mathfrak{a}^1 = 1$, $\mathfrak{a}^L = \alpha$, and $\mathfrak{d}_1^1 = \mathfrak{d}_2^1 = 0$.

We have provided formulae for basic properties about the encoding of the runs, but this is insufficient. Indeed, three consecutive locations on the main path encode a configuration of the Minsky machine M . In order to check that two consecutive configurations correspond to a step that is valid for M , we need to compare numbers of

predecessors for locations on the main path at distance three from each other. For instance, considering locations l and l' on the main path such that $l' = h^3(l)$, we plan to build formulae to express constraints between $\#l$ and $\#l'$. There are two problems there.

Distance three. Firstly, with two variables, one can explore the heap but is obliged to “forget” previously visited locations (in fact, it is possible to store only a finite amount of information). Since [Gabbay 1981], it is known how to visit a graph with only two quantified variables. In the current encoding of runs, we will need to compare the numbers of predecessors of locations at distance three on the main path. Let us consider the formula below:

$$\exists u [\forall \bar{u} (\bar{u} \hookrightarrow u) \Rightarrow (\# \bar{u} = 3 \vee \# \bar{u} = 7)] \wedge [\exists \bar{u} (u \hookrightarrow \bar{u}) \wedge \exists u ((\bar{u} \hookrightarrow u) \wedge \# u = 11)]$$

This formula states that there are two locations in the model for which there is a path of length 2 between them, the second location has exactly 11 predecessors and for the first one, every predecessor has either 3 or 7 predecessors.

In MLH (without the use of $*$ or $-*$), this property can be written as follows:

$$\langle U \rangle (\Box^{-1} (\Diamond_{=3}^{-1} \top \vee \Diamond_{=7}^{-1} \top) \wedge \Diamond \Diamond_{=11}^{-1} \top)$$

So, it is possible to state properties between locations that are not direct successors, but note that we can only enforce properties while we are visiting the nodes: once we move forward or backward we have no more access to previously visited locations. This becomes a problem when we need to compare number of predecessors for locations at distance three. Observe that if we had proposed another encoding of the runs, at some stage we would have to deal with locations that are not direct successors and for which we need to compare some potentially unbounded amount of information (since we need anyhow to encode counter values that are potentially unbounded). Section 4.4 provides a solution to this problem by presenting a selective chopping of the heap that preserves the numbers of predecessors we wish to compare while being able to access easily those locations for which the number of predecessors are compared.

It is worth observing that for locations on the main path related to counter instructions, we are not in big trouble because it is possible with two quantified variables to visit two successive locations related to counter instructions and to check that they respect the instructions of the Minsky machine (because only a finite amount of information needs to be encoded). While moving along the model, we may “forget” where we start, but this is fine since we can remember the value of the counter instruction in the formula. By contrast, when comparing two locations corresponding to two successive values of the same counter, this does not work anymore.

Arithmetical constraints. Secondly, we have also to be able to compare numbers of predecessors between locations. For instance, given two locations l and l' , we wish to be able to check whether $\#l = \#l'$ or $\#l = \#l' + 1$. Such a need should not come as a surprise, since in our encoding of data words, data values are represented by numbers of predecessors. Such arithmetical constraints can be expressed in 1SL (i.e. without limiting the number of quantified variables), and this has been a key step to establish 1SL’s equivalence to weak second-order logic on heaps [Brochenin et al. 2012]. In Section 4.3, we provide a fine-tuned adaptation with only two quantified variables (instead of an unbounded number of variables) and with substantial simplifications. The developments in Section 4.3 are not consequences of developments from [Brochenin et al. 2012, Section 5.2] but rather refinements using similar principles. Actually, we can take advantage of the fact that we do not work on an arbitrary heap but rather on an $(\alpha, 2)$ -fishbone heap. Moreover, we do not seek expressive completeness but rather we aim at expressing a sufficient set of arithmetical constraints to allow us to characterize $(\alpha, 2)$ -fishbone heaps that encode halting runs of M .

4.3. Expressing arithmetical constraints

Below, we show how to express in 1SL2 the constraints $\#u = \#u$, $\#u = \#u + 1$ and $\#u = \#u + 1$, which is not at all obvious. We explain why this can be done in 1SL2 by a careful recycling of variables; along the way, we also take advantage of the properties of heaps satisfying ϕ^* .

We shall use the fact that $N \leq N'$ iff for every $n \geq 0$, we have $N' \leq n$ implies $N \leq n$. Quantification over the set of natural numbers will be simulated by quantification over disjoint heaps in which n is related to the cardinal of their domains. Such quantification is performed thanks to the magic wand operator.

A *fork* in $\mathfrak{h}$ is a sequence of distinct locations l, l_0, l_1, l_2 such that $\mathfrak{h}(l_0) = l$, $\widetilde{\#}l_0 = 2$, $\mathfrak{h}(l_1) = l_0$, $\mathfrak{h}(l_2) = l_0$ and $\widetilde{\#}l_1 = \widetilde{\#}l_2 = 0$. The *endpoint* of the fork is l . Similarly, a *knife* in $\mathfrak{h}$ is a sequence of distinct locations l, l_0, l_1 such that $\mathfrak{h}(l_0) = l$, $\widetilde{\#}l_0 = 1$, $\mathfrak{h}(l_1) = l_0$ and $\widetilde{\#}l_1 = 0$. The *endpoint* of the knife is l . By way of example, the heap of Figure 6 contains three knives, two forks and four endpoints (identified by ‘*’).

Fig. 6. A heap with three knives, two forks and four endpoints.

LEMMA 4.2. *Let $\mathfrak{h}$ be a (α, β) -fishbone heap with $\alpha \geq 1$ and $\beta \geq 0$. Then, $\mathfrak{h}$ has no knife and no fork.*

Indeed, in such heaps, any allocated location has no predecessor or at least three predecessors.

A heap $\mathfrak{h}$ is a *collection of knives* $\stackrel{\text{def}}{\iff}$ there is no location in $\text{dom}(\mathfrak{h})$ that does not belong to a knife and no distinct knives share the same endpoint. A heap $\mathfrak{h}$ is *segmented* whenever $\text{dom}(\mathfrak{h}) \cap \text{ran}(\mathfrak{h}) = \emptyset$ and no location has strictly more than one predecessor.

LEMMA 4.3. *Let $\mathfrak{h}$ be a (α, β) -fishbone heap with $\alpha \geq 1$, $\beta \geq 0$ and $\mathfrak{h}'$ be a segmented heap disjoint from $\mathfrak{h}$. Then, $\mathfrak{h} \uplus \mathfrak{h}'$ has no fork.*

Being segmented can be naturally expressed in 1SL2:

$$\text{seg} \stackrel{\text{def}}{=} \forall u \bar{u} (u \hookrightarrow \bar{u} \Rightarrow ((\# \bar{u} = 1) \wedge (\# u = 0) \wedge \neg \text{alloc}(\bar{u})))$$

The statement below is counterpart to [Brochenin et al. 2012, Lemma 5.2] with simplified properties and with simpler formulae but using only two quantified variables.

LEMMA 4.4. *There are formulae $\text{forky}(u)$, KS and KS1F in 1SL2 such that for every heap $\mathfrak{h}$,*

- (I) $\mathfrak{h} \models_f \text{forky}(u)$ *iff all the predecessors of $f(u)$ are endpoints of forks,*
- (II) $\mathfrak{h} \models \text{KS}$ *iff $\mathfrak{h}$ is a collection of knives,*

(III) $\mathfrak{h} \models \text{KS1F}$ iff there are $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$, $\mathfrak{h}_1$ is a collection of knives and $\mathfrak{h}_2$ is made of a unique fork such that its unique endpoint is not in the range of $\mathfrak{h}_1$.

PROOF. $\text{forky}(\mathfrak{u})$ is equal to:

$$\forall \bar{u} (\bar{u} \hookrightarrow \mathfrak{u}) \Rightarrow (\exists u (\mathfrak{u} \hookrightarrow \bar{u}) \wedge (\#u = 2) \wedge \neg(\#u^{-1} \geq 1))$$

A knife is made of two consecutive memory cells that can be respectively called part 1 and part 2 as shown in $\mathfrak{l} \xrightarrow{\text{part 1}} \mathfrak{l}' \xrightarrow{\text{part 2}} \mathfrak{l}''$.

$$\text{KS} \stackrel{\text{def}}{=} \forall u \text{ alloc}(\mathfrak{u}) \Rightarrow (\phi_{\text{part1}}(\mathfrak{u}) \vee \phi_{\text{part2}}(\mathfrak{u}))$$

where

$$\phi_{\text{part1}}(\mathfrak{u}) \stackrel{\text{def}}{=} (\#u = 0) \wedge (\#u^{+1} = 1) \wedge (\#u^{+2} = 1) \wedge \neg(\#u^{+3} \geq 0)$$

$$\phi_{\text{part2}}(\mathfrak{u}) \stackrel{\text{def}}{=} (\#u = 1) \wedge (\#u^{-1} = 0) \wedge (\#u^{+1} = 1) \wedge \neg(\#u^{+2} \geq 0)$$

$$\begin{aligned} \text{KS1F} &\stackrel{\text{def}}{=} \\ &\overbrace{[\exists u (\#u = 2) \wedge (\#u^{+1} = 1) \wedge \neg(\#u^{+2} \geq 0) \wedge \neg(\#u^{-1} \geq 1) \wedge \neg(\exists \bar{u} (\mathfrak{u} \neq \bar{u}) \wedge (\#\bar{u} = 2))]}^{\text{unique fork}} \wedge \\ &[\forall u \text{ alloc}(\mathfrak{u}) \Rightarrow \\ &\underbrace{(\phi_{\text{part1}}(\mathfrak{u}) \vee \phi_{\text{part2}}(\mathfrak{u}))}_{\text{part with knives}} \vee \underbrace{((\#u = 0) \wedge (\#u^{+1} = 2)) \vee (\#u = 2)}_{\text{part with one fork}})] \end{aligned}$$

□

In our proof, we use the idea of augmenting the heap with a segmented heap, then augmenting it further with knives to form forks whose endpoints are predecessors of $\mathfrak{u}$; this is borrowed from [Brochenin et al. 2012]. As it is, this would not be sufficient to express arithmetical constraints on fishbone heaps since only two quantified variables are allowed. This restriction is not considered in [Brochenin et al. 2012]—the formulae there use strictly more than two quantified variables. This is why we had to provide specific developments that are well-tailored to fishbone heaps while taking into account our limited amount of syntactic resources. Simplifications have also been made in order to focus on undecidability rather than on questions of expressive power.

LEMMA 4.5. *Let $\mathfrak{h}$ be a heap with $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$ and f be an assignment such that $\mathfrak{h}_1 \models_f \phi^*$, $f(\mathfrak{u})$ is on the main path of the $(\alpha, 2)$ -fishbone heap $\mathfrak{h}$, $\mathfrak{h}_2 \models_f \text{seg} \wedge \#u = 0$, $n = \text{card}(\text{ran}(\mathfrak{h}_2) \setminus \text{dom}(\mathfrak{h}_1))$ and m is the number of predecessors of $f(\mathfrak{u})$ in $\mathfrak{h}_1$. We have the following properties:*

- (a) $\mathfrak{h} \models_f \neg(\text{KS} * \neg \text{forky}(\mathfrak{u}))$ iff $n \geq m$.
- (b) $\mathfrak{h} \models_f \neg(\text{KS1F} * \neg \text{forky}(\mathfrak{u}))$ iff $n \geq m - 1$.

In Figure 7, we present three heaps obtained by combining a segmented heap $\mathfrak{h}_2$ with collections of knives (corresponding to $\mathfrak{h}_3$ in the proof of Lemma 4.5). Edges labelled by ‘1’ are part of a fishbone heap $\mathfrak{h}_1$ (partially represented) whereas edges labelled by ‘2’ are part of a segmented heap $\mathfrak{h}_2$ so that no edge points to $f(\mathfrak{u})$ or to $f(\bar{u})$. The heap on the left (corresponding to $\mathfrak{h}_1 * \mathfrak{h}_2$ in Lemma 4.5) is obtained by adding a segmented heap $\mathfrak{h}_2$ whereas the heap in the middle (say $\mathfrak{h}_1 * \mathfrak{h}_2 * \mathfrak{h}_3$) is obtained then by adding a

Fig. 7. A segmented heap and collections of knives.

collection of knives h_3 so that every predecessor of $f(u)$ is the endpoint of a fork. Note that not all edges of the segmented heap are used to build forks. Similarly, the heap on the right (say $h_1 * h_2 * h'_3$) is obtained then by adding a collection of knives h'_3 to the heap $h_1 * h_2$ on the very left so that every predecessor of $f(\bar{u})$ is the endpoint of a fork.

PROOF. Let us provide the proof for (a). (The proof for (b), being analogous, is omitted.) So, let h be a heap with $h = h_1 \uplus h_2$ such that $h_1 \models_f \phi^*$ and $h_2 \models_f \text{seg} \wedge \sharp u = 0$. Moreover, $f(u)$ is on the main path, which entails that $h(f(u)) \neq f(u)$ (if $h(f(u))$ is defined at all) and $f(u)$ has at least one predecessor.

One can make the following (obvious) observations.

(O1). The heap h_1 has no knives and, h_1 and $h_1 \uplus h_2$ have no forks. (see Lemma 4.2 and Lemma 4.3).

(O2). $h_1 \uplus h_2$ may not satisfy ϕ^* but this is fine since we only need to focus on the number of predecessors of $f(u)$ (i.e., on the value m). Indeed, $h_1 \uplus h_2$ may contain knives (see the left heap in Figure 7). A knife $l_1 \rightarrow l_2 \rightarrow l_3$ in $h_1 \uplus h_2$ is made of $l_1 \in \text{dom}(h_2)$ and of $l_2 \in \text{dom}(h_1)$. This observation is not really used below but, hopefully, it could be helpful to better grasp how the heaps h_1 and h_2 are combined.

(O3). $f(u)$ has the same number of predecessors in h_1 and in $h_1 \uplus h_2$. This is due to the fact that $h_2 \models_f \sharp u = 0$.

(O4). For every $n \geq 0$, there is a disjoint heap h'_2 such that $h' = h_1 \uplus h'_2$, $h'_2 \models_f \text{seg} \wedge \sharp u = 0$ and $\text{card}(\text{ran}(h'_2) \setminus \text{dom}(h_1)) = n$. See the left heap in Figure 7 with $\text{card}(\text{dom}(h_2)) = 5$ and $\text{card}(\text{ran}(h_2) \setminus \text{dom}(h_1)) = 4$ (look at edges labelled by ‘2’). Once more, this observation is not used below but it will be in the proof of Proposition 4.6.

First, let us suppose that $h \models_f \neg(\text{KS} * \neg \text{forky}(u))$, i.e., ($\dagger$) there is a heap h_3 , disjoint from $h_1 \uplus h_2$, such that $(h_1 \uplus h_2) \uplus h_3 \models_f \text{forky}(u)$ and $h_3 \models_f \text{KS}$. Let us make additional observations.

- The only forks in $h_1 \uplus h_2 \uplus h_3$ whose endpoints are predecessors of $f(u)$ are those obtained with $l_1 \rightarrow l_2$ such that $l_1 \in \text{dom}(h_2)$ (so $h_2(l_1) = l_2$), $l_2 \notin \text{dom}(h_1)$, and $l'_1 \rightarrow l_2 \rightarrow l'_3$ is a knife from h_3 . This is due to (O1) and to the fact that all the predecessors of $f(u)$ in h have no predecessors since $f(u)$ is on the main path of h .

- The number of forks in $h_1 \uplus h_2 \uplus h_3$ whose endpoints are predecessors of $f(u)$ is therefore less or equal to $\text{card}(\text{ran}(h_2) \setminus \text{dom}(h_1))$.
- The number of predecessors of $f(u)$ in $h_1 \uplus h_2 \uplus h_3$ is greater or equal to the number of its predecessors in h_1 (by using (O3)). So, if $h_1 \uplus h_2 \uplus h_3 \models_f \text{forky}(u)$, then the number of predecessors of $f(u)$ in h_1 is smaller or equal to $\text{card}(\text{ran}(h_2) \setminus \text{dom}(h_1)) = n$, i.e. $n \geq m$.

Now, let us establish the other direction and let us suppose that $n \geq m$ and the predecessors of $f(u)$ are $p_1, \dots, p_m$. Let $l_1^1, l_1^2, \dots, l_n^1, l_n^2$ be locations such that $\{l_1^1, \dots, l_n^1\} = \text{ran}(h_2) \setminus \text{dom}(h_1)$ and for every $i \in [1, n]$, we have $h_2(l_i^2) = l_i^1$. Let us build h_3 so that it satisfies $(\dagger)$, which is quite easy to realize. Let $l_1^{\text{new}}, \dots, l_m^{\text{new}}$ be (new) locations that are not in $\text{dom}(h_1 \uplus h_2) \cup \text{ran}(h_1 \uplus h_2)$. We define h_3 so that it contains exactly m knives whose endpoints are exactly all the predecessors of $f(u)$. For every $i \in [1, m]$, we define $h_3(l_i^{\text{new}}) \stackrel{\text{def}}{=} l_i^1$ and $h_3(l_i^1) \stackrel{\text{def}}{=} p_i$ (well, that is possible because $l_i^1 \notin \text{dom}(h_1 \uplus h_2)$). It is easy to check that h_3 satisfies $(\dagger)$.

Consequently, $h \models_f \neg(\text{KS} * \neg \text{forky}(u))$ iff $n \geq m$. $\square$

Now, we are able to state the main proposition of this section that allows us to compare the numbers of predecessors for two locations on the main path of a fishbone heap. Let us introduce the following abbreviations: $\varphi_1(u, \bar{u}) \stackrel{\text{def}}{=} \text{seg} \wedge \#u = 0 \wedge \#\bar{u} = 0$, $\varphi_2(u) \stackrel{\text{def}}{=} \neg(\text{KS} * \neg \text{forky}(u))$ and $\varphi_3(u) \stackrel{\text{def}}{=} \neg(\text{KS1F} * \neg \text{forky}(u))$.

PROPOSITION 4.6. *Suppose $h_1 \models_f \phi^*$ and, $f(u)$ and $f(\bar{u})$ are on the main path of h_1 . We have the following equivalences:*

- $h_1 \models_f \varphi_1(u, \bar{u}) * (\varphi_2(\bar{u}) \Rightarrow \varphi_2(u))$ iff $\#u \leq \#\bar{u}$.
- $h_1 \models_f \varphi_1(u, \bar{u}) * (\varphi_2(\bar{u}) \Rightarrow \varphi_3(u))$ iff $\#u \leq \#\bar{u} + 1$.
- $h_1 \models_f \varphi_1(u, \bar{u}) * (\varphi_3(\bar{u}) \Rightarrow \varphi_2(u))$ iff $\#u \leq \#\bar{u} - 1$.

PROOF. By way of example, let us show the second property. The other cases are proved in a similar fashion. Let h_1 be a heap satisfying ϕ^* . The statements below are equivalent.

- (1) $h_1 \models_f (\varphi_1(u, \bar{u}) * (\varphi_2(\bar{u}) \Rightarrow \varphi_3(u)))$.
- (2) For every disjoint heap h_2 such that $h_2 \models_f \varphi_1(u, \bar{u})$, if $h_1 \uplus h_2 \models_f \varphi_2(\bar{u})$, then $h_1 \uplus h_2 \models_f \varphi_3(u)$. (by definition of $\models_f$)
- (3) For every $n \geq 0$, there is a disjoint heap h_2 with $\text{card}(\text{ran}(h_2) \setminus \text{dom}(h_1)) = n$ such that $h_2 \models_f \varphi_1(u, \bar{u})$ and if $h_1 \uplus h_2 \models_f \varphi_2(\bar{u})$, then $h_1 \uplus h_2 \models_f \varphi_3(u)$ (see (O4) in the proof of Lemma 4.5). This is possible by using the fact that one can add a segmented heap so that the resulting heap has n isolated memory cells. Indeed, given the heap h_1 , let us build a disjoint heap h_2 such that $h_2 \models_f \varphi_1(u, \bar{u})$ and $\text{dom}(h_2) = n$ for any fixed $n \geq 0$. Since $X = \text{dom}(h_1) \cup \text{ran}(h_2) \cup \{f(u), f(\bar{u})\}$ is a finite subset of $\mathbb{N}$, there are $2n$ distinct locations $l_1^1, l_1^2, \dots, l_n^1, l_n^2$ in $\mathbb{N} \setminus X$. We simply need to define h_2 such that $\text{dom}(h_2) \stackrel{\text{def}}{=} \{l_1^1, \dots, l_n^1\}$, $\text{ran}(h_2) \stackrel{\text{def}}{=} \{l_1^2, \dots, l_n^2\}$ and for all $i \in [1, n]$, we set $h_2(l_i^1) \stackrel{\text{def}}{=} l_i^2$.
- (4) for every $n \geq 0$, we have $n \geq \#\bar{u}$ in h implies $n \geq \#u - 1$ in h . (by Lemma 4.5)
- (5) $\#u \leq \#\bar{u} + 1$. $\square$

4.4. Constraints between locations at distance three

The goal of this section is the following: given a formula $\varphi(u, \bar{u})$ equal to either $\#u = \#\bar{u}$ or $\#u = \#\bar{u} + 1$ (in particular, this means that $\varphi(u, \bar{u})$ only deals with numbers of

predecessors and Section 4.3 explains how to define these formulae in 1SL2), we show how to define a formula in 1SL2, say $\varphi^{+3}(u)$, such that

$$h \models_f \varphi^{+3}(u) \text{ iff } h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u}),$$

assuming that $h^3(f(u))$ is defined and $h \models_f \text{dw}(\alpha, 2) \wedge (\#u \geq \alpha + 3)$. When $\varphi(u, \bar{u})$ is equal to $\#u = \#\bar{u}$ [resp. $\#u = \#\bar{u} + 1$], we write $\#u = \#u^{+3}$ [resp. $\#u = \#u^{+3} + 1$] instead of $\varphi^{+3}(u)$. Note that if we had three quantified variables, defining $\varphi^{+3}(u)$ would not require much work since the formula below does the job:

$$\exists u' (u \hookrightarrow u' \wedge \exists \bar{u} (u' \hookrightarrow \bar{u} \wedge \exists u' (\bar{u} \hookrightarrow u' \wedge \varphi(u, u'))))$$

Let us start our construction. To do so, let h be a heap and f be an assignment such that $h \models_f \text{dw}(\alpha, 2) \wedge (\#u^{+3} \geq 0) \wedge (\#u \geq \alpha + 3)$. In the statements below, this property is always satisfied.

The *u-3cut* of h is the minimal subheap $h_{3\text{cut}}$ of h (with respect to set inclusion of the domain) such that all the ancestors of $l' = h^3(f(u))$ in $\text{dom}(h)$ are also ancestors of l' in $h_{3\text{cut}}$. It looks like a clean cut from Section 3 but operated on $h^3(f(u))$ and on an $(\alpha, 2)$ -fishbone heap. As a consequence, $f(u)$ and l' have the same amount of predecessors in h and in the *u-3cut* heap. Moreover, if $f(u)$ has more than $\alpha + 3$ predecessors, then the *u-3cut* of h is also a $(\alpha, 2)$ -fishbone heap.

In Figure 8, the bottom left heap is the *u-3cut* of the heap at the top. When $h \models_f \#u^{+4} \geq 0$, the *almost u-3cut* of h is the minimal subheap of h containing the *u-3cut* heap and such that $\#u^{+4} = 1$ holds true. The *almost u-3cut* of h contains the edge from l' which is the only predecessor of the interpretation of u^{+4} . In Figure 8, the middle left heap is the *almost u-3cut* of the heap at the top. Below, we explain how to obtain the *u-3cut* of some heap, possibly via the construction of the *almost u-3cut*, if it exists.

Lemma 4.7 below states that all we need to define $\varphi^{+3}(u)$ is to be able to express a property in its *u-3cut*. In particular, the only location that is unallocated and on the main path is $h^3(f(u))$.

LEMMA 4.7. *Let $h \models_f \text{dw}(\alpha, 2) \wedge (\#u^{+3} \geq 0) \wedge (\#u \geq \alpha + 3)$ and h' be its *u-3cut* heap. Then, $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$ iff $h' \models_f (\exists \bar{u} \neg \text{alloc}(\bar{u}) \wedge \#\bar{u} \geq 1 \wedge \varphi(u, \bar{u}))$.*

PROOF. Let $l' = h^3(f(u))$ and $l = f(u)$. Let h' be the *u-3cut* heap of h . We have $(\dagger) \#l$ in h is equal to $\#l$ in h' and $\#l'$ in h is equal to $\#l'$ in h' . Indeed, the *u-3cut* heap h' is a subheap of h such that all the ancestors of l' in h are also ancestors of l' in h' and l is an ancestor of l' in h . Note also that l' is the unique location such that $h' \models_{[\bar{u} \mapsto l']} \neg \text{alloc}(\bar{u}) \wedge \#\bar{u} \geq 1$. So, $h' \models_f (\exists \bar{u} \neg \text{alloc}(\bar{u}) \wedge \#\bar{u} \geq 1 \wedge \varphi(u, \bar{u}))$ iff $h' \models_{f[\bar{u} \mapsto l']} \varphi(u, \bar{u})$ iff $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$ by $(\dagger)$. Note that we use the fact that $\varphi(u, \bar{u})$ specifies a property about the numbers of predecessors. $\square$

When h is equal to its *u-3cut*, i.e. when $(\#u^{+4} \geq 0)$ does not hold, we have $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$ iff $h \models_f \phi_{\text{UC}}(u)$ with

$$\phi_{\text{UC}}(u) \stackrel{\text{def}}{=} (\exists \bar{u} \neg \text{alloc}(\bar{u}) \wedge \#\bar{u} \geq 1 \wedge \varphi(u, \bar{u}))$$

Now, let us consider the case when h is not equal to its *u-3cut* (probably, the most common situation) and let us show how to separate the current heap so that we can isolate the *u-3cut* heap.

LEMMA 4.8. *Let $h \models_f \text{dw}(\alpha, 2) \wedge (\#u^{+4} \geq 0) \wedge (\#u \geq \alpha + 3)$ and $\phi(u)$ be an arbitrary formula. Then, $h \models_f 1\text{comp} * (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$ iff the *almost u-3cut* of h , say h' , satisfies: $h' \models_f \phi(u)$.*

The formula 1comp was introduced in Section 2.2, and it states that the heap is made of a unique connected component. The way h has to be divided to satisfy the formula is illustrated by the two heaps in the middle of Figure 8.

PROOF. Let h be heap such that $h \models_f \text{dw}(\alpha, 2) \wedge (\#u^{+4} \geq 0) \wedge (\#u \geq \alpha + 3)$. Let h' be the almost u -3cut heap of h and h'' be the heap such that $h = h' \uplus h''$. By construction of h' , it is easy to check that $h' \models_f 1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0)$. Similarly, $h'' \models_f 1\text{comp}$. This implies that $h \models_f 1\text{comp} * (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0))$.

So, suppose that the almost u -3cut heap of h satisfies: $h' \models_f \phi(u)$. This means that $h'' \models_f 1\text{comp}$ and $h' \models_f (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$. Hence, $h \models_f 1\text{comp} * (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$.

Now, suppose that $h \models_f 1\text{comp} * (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$. There are heaps h_1 and h_2 such that $h_2 \models 1\text{comp}$ and $h_1 \models_f (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$. In particular, this means that $h_1 \models_f 1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0)$.

Let us show that there is a unique pair (h_1, h_2) of heaps satisfying that property and $h_1 = h'$, which will entail that $h' \models_f \phi(u)$. First note that

$$\{f(u), h(f(u)), h^2(f(u)), h^3(f(u)), h^4(f(u))\} \subseteq \text{dom}(h_1) \quad h^5(f(u)) \notin \text{dom}(h_1)$$

Since $h_1 \models_f (\#u^{+4} = 1)$, all the predecessors of $h^4(f(u))$, apart from $h^3(f(u))$, are in $\text{dom}(h_2)$ and there are more than two such predecessors since $h^4(f(u))$ is on the main path of h and therefore has at least three predecessors in h .

Hence, h_1 contains also all the ancestors of $h^3(f(u))$, otherwise h_2 would have at least two distinct connected components. So, the u -3cut of h is also a subheap of h_1 .

Now, it is easy to check that if any location in $\text{dom}(h'')$ that is not a predecessor of $h^4(f(u))$ were in $\text{dom}(h_1)$, then h_1 would have more than two connected components. Hence, h_1 is the almost u -3cut heap of h and therefore $h' \models_f \phi(u)$. $\square$

Let us build on Lemma 4.8 so as to be able to specify properties on the u -3cut heap.

LEMMA 4.9. *Let $h \models_f \text{dw}(\alpha, 2) \wedge (\#u^{+4} \geq 0) \wedge (\#u \geq \alpha + 3)$ and $\phi(u)$ be the formula $(\text{size} = 1) * (\neg(\#u^{+4} \geq 0) \wedge \phi_{UC}(u))$. Then, $h \models_f 1\text{comp} * (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$ iff the u -3cut of h , say h' , satisfies: $h' \models_f \phi_{UC}(u)$.*

We write $\phi_{AUC}(u)$ to denote the formula $1\text{comp} * (1\text{comp} \wedge (\#u^{+4} = 1) \wedge \neg(\#u^{+5} \geq 0) \wedge \phi(u))$ with $\phi(u)$ equal to $(\text{size} = 1) * (\neg(\#u^{+4} \geq 0) \wedge \phi_{UC}(u))$.

The proof for Lemma 4.9 is also by an easy verification by observing that an almost u -3cut heap is equal to the u -3cut plus one memory cell (see Figure 8).

By combining Lemma 4.7–4.9, we get the following proposition by performing a case analysis depending whether $\#u^{+4} \geq 0$ holds true or not on the heap h .

PROPOSITION 4.10. *Let h be a heap and f be an assignment such that $h \models_f \text{dw}(\alpha, 2) \wedge (\#u^{+3} \geq 0) \wedge (\#u \geq \alpha + 3)$. We have $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$ iff $h \models_f \varphi^{+3}(u)$ with the formula $\varphi^{+3}(u)$ defined below:*

$$\varphi^{+3}(u) \stackrel{\text{def}}{=} (\neg(\#u^{+4} \geq 0) \wedge \phi_{UC}(u)) \vee ((\#u^{+4} \geq 0) \wedge \phi_{AUC}(u))$$

PROOF. We distinguish two cases depending whether h is itself its u -3cut or not.

Case 1: $\neg(\#u^{+4} \geq 0)$, i.e. h is its own u -3cut heap. By Lemma 4.7, if $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$, then $h \models_f \phi_{UC}(u)$ and therefore $h \models_f \varphi^{+3}(u)$. Conversely, if $h \models_f \varphi^{+3}(u)$, then $h \models_f \phi_{UC}(u)$ since $(\#u^{+4} \geq 0)$ does not hold on h . Again, by Lemma 4.7, we get that $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$.

Case 2: $(\#u^{+4} \geq 0)$. By Lemma 4.7, if $h \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$, then $h' \models_f \phi_{UC}(u)$ where h' is the u -3cut of h . By Lemma 4.9, this implies that $h \models_f \phi_{AUC}(u)$ and therefore $h \models_f$

Fig. 8. How to get a u-3cut – Decomposition in two stages.

$\varphi^{+3}(u)$ (thanks to its second disjunction). Conversely, if $\mathfrak{h} \models_f \varphi^{+3}(u)$, then $\mathfrak{h} \models_f \phi_{AUC}(u)$ since $(\#u^{+4} \geq 0)$ holds on $\mathfrak{h}$. Again, by Lemma 4.9, we get that $\mathfrak{h}' \models_{f[\bar{u} \mapsto h^3(f(u))]} \phi_{UC}(u)$ and by Lemma 4.7, we conclude $\mathfrak{h} \models_{f[\bar{u} \mapsto h^3(f(u))]} \varphi(u, \bar{u})$. $\square$

Note that the reasoning performed in this section cannot be extended to an arbitrary formula $\varphi(u, \bar{u})$ since taking a u-3cut or an almost u-3cut preserves the number of predecessors of $f(u)$ and $h^3(f(u))$ but may not preserve more general properties. Nevertheless, this is sufficient for our needs in Section 4.5.

4.5. Master reduction from halting problem for Minsky machines

We shall use formulae of the form $\varphi^{+3}(u)$ when $\varphi(u, \bar{u})$ expresses one of the following arithmetical constraints: $\#u = \#\bar{u}$, $\#u = \#\bar{u} + 1$ and $\#\bar{u} = \#u + 1$ (see Section 4.4). For each instruction $I \in [1, \alpha - 1]$, we build a formula ϕ_I so that the Minsky machine M halts iff the formula $\phi^* \wedge \bigwedge_{I \in [1, \alpha - 1]} \varphi_I$ is satisfiable in 1SL2. It remains to define φ_I for each instruction I .

If instruction I is of the form “ $I: C_j := C_j + 1; \text{goto } J$ ” then we need to check the following properties:

- (1) If a location l encodes the instruction I on the main path (i.e. $\#l = I + 2$) and $h^3(l)$ is defined, then the location $h^3(l)$ encodes the instruction J .
- (2) If a location l encodes the value for the counter C_j in a configuration with instruction I (i.e., $\#l \geq \alpha + 3$ and the j th ancestor of l has $I + 2$ predecessors) and $h^3(l)$ is defined, then $\#l + 1 = \#h^3(l)$.

- (3) Similarly, if a location l encodes the value for the counter C_{3-j} (i.e., the counter C_{3-j} is not updated after instruction I) in a configuration with instruction I (i.e., $\#l \geq \alpha + 3$ and the j th ancestor of l has $I + 2$ predecessors) and $h^3(l)$ is defined, then $\#l = \#h^3(l)$.

The properties can be expressed by the formula φ_I below:

$$\begin{aligned} \forall u (\#u^{+3} \geq 0) \Rightarrow & \overbrace{[(\#u = I + 2) \Rightarrow (\#u^{+3} = J + 2)]}^{(1)} \wedge \\ & \overbrace{((\#u \geq \alpha + 3) \wedge (\#u^{-j} = I + 2) \Rightarrow (\#u = \#u^{+3} - 1))}^{(2)} \wedge \\ & \overbrace{((\#u \geq \alpha + 3) \wedge (\#u^{j-3} = I + 2) \Rightarrow (\#u = \#u^{+3}))}^{(3)} \end{aligned}$$

Each subformula decorated by a curly bracket with (i) expresses exactly the property (i) above. Note that $\#u^{+3} = J + 2$ states that the number of predecessors of $h^3(f(u))$ is $J + 2$, which is quite easy to express in 1SL2 (see Section 2.2). By contrast, the formula $\#u = \#u^{+3} - 1$ states that the number of predecessors of $h^3(f(u))$ is equal to the number of predecessors of $f(u)$ plus one, which requires the more sophisticated formulae introduced in Section 4.3 and in Section 4.4.

Similarly, let I be the instruction “ I : if $C_j = 0$ then goto J_1 else ($C_j := C_j - 1$; goto J_2)” then φ_I is defined as follows:

$$\begin{aligned} \forall u (\#u^{+3} \geq 0) \Rightarrow & \overbrace{[(\#u = I + 2) \wedge (\#u^{+j} = \alpha + 3) \Rightarrow (\#u^{+3} = J_1 + 2)]}^{(4)} \wedge \\ & \overbrace{((\#u = I + 2) \wedge (\#u^{+j} > \alpha + 3) \Rightarrow (\#u^{+3} = J_2 + 2))}^{(5)} \wedge \\ & \overbrace{((\#u > \alpha + 3) \wedge (\#u^{-j} = I + 2) \Rightarrow (\#u^{+3} = \#u - 1))}^{(6)} \wedge \\ & \overbrace{((\#u = \alpha + 3) \wedge (\#u^{-j} = I + 2) \Rightarrow (\#u^{+3} = \alpha + 3))}^{(7)} \wedge \\ & \overbrace{((\#u \geq \alpha + 3) \wedge (\#u^{j-3} = I + 2) \Rightarrow (\#u = \#u^{+3}))}^{(8)} \end{aligned}$$

The subformula decorated by a curly bracket with (4) states that if a location l encodes the instruction I and $h^j(l)$ has $\alpha + 3$ predecessors (i.e., counter C_j has value zero), then the location $h^3(l)$ has $J_1 + 2$ predecessors (i.e., the next instruction is J_1). Similarly, the subformula decorated by a curly bracket with (5) states that if a location l encodes the instruction I and $h^j(l)$ has strictly more than $\alpha + 3$ predecessors (i.e., counter C_j has non-zero value), then the location $h^3(l)$ has $J_2 + 2$ predecessors (i.e., the next instruction is J_2). Moreover, the subformula decorated by a curly bracket with (6) states that if a location l has at least $\alpha + 3$ predecessors and its j th ancestor has $I + 2$ predecessors

(i.e., counter C_j has non-zero value and we are really dealing with instruction I), then the number of predecessors of $h^3(l)$ is equal to the number of predecessors of l minus one, which corresponds to encode a decrement on counter C_j . Subformulae (7) and (8) admit a similar reading.

It is now easy to show the following lemma since we have seen that all the constraints between consecutive configurations can be encoded in 1SL2, assuming that the heap encodes a data word in $([1, \alpha] \times \mathbb{N}^2)^+$.

LEMMA 4.11. *M has a halting run iff*

$$\text{dw}(\alpha, 2) \wedge \phi_{first} \wedge \phi_{last} \wedge \bigwedge_{I \in [1, \alpha]} \varphi_I$$

is satisfiable in 1SL2.

Below, we conclude by the main result of the paper.

THEOREM 4.12. *1SL2 satisfiability problem is undecidable.*

We know that if the number of quantified variables is not restricted, 1SL($\ast$) is undecidable too [Brochenin et al. 2012] and recently the satisfiability problem for 1SL2($\ast$) has been shown undecidable as well [Demri and Deters 2014], but this requires a far more complex proof passing via an equivalence to weak second-order logic.

4.6. Note on a variant proof using $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha,1}(<, +1, =, \sim, \prec)$

As mentioned earlier, there exist many formalisms to specify properties about data words; among them can be found first-order languages. Below, we recall a few standard definitions as well as the main results. Finally, we sketch the proof of a reduction from an undecidable variant of first-order logic on data words into 1SL2. These results show interesting relationships between first-order logics on data words and separation logics.

Let us present the first-order language $\text{FO2}_{\alpha,\beta}(<, +1, =, \sim, \prec)$ to interpret data words in $([1, \alpha] \times \mathbb{N}^\beta)^+$ following developments from [Bojańczyk et al. 2011]. Most of the time, a fragment of the full language is needed, but it is helpful to provide the most general definition once and uniformly.

Let $\text{FO2}_{\alpha,\beta}(<, +1, =, \sim, \prec)$ be the set of formulae defined below:

$$\phi ::= a(v) \mid v \sim_j v \mid v \prec_j v \mid v < v \mid v = 1 + (v) \mid v = v \mid \neg \phi \mid \phi \wedge \phi \mid \exists v \phi$$

with $v ::= u_1 \mid u_2$, $j \in [1, \beta]$ and $a \in [1, \alpha]$. When $\beta = 0$, this implies that there is no atomic formula using $\sim_j$ or $\prec_j$. We write $\text{FO2}_{\alpha,\beta}(<, +1, =, \sim)$ to denote the restriction of $\text{FO2}_{\alpha,\beta}(<, +1, =, \sim, \prec)$ without $\prec$. Formulae in $\text{FO2}_{\alpha,\beta}(<, +1, =, \sim, \prec)$ are interpreted over data words

$$\text{dw} = (a^1, \mathfrak{d}_1^1, \dots, \mathfrak{d}_\beta^1) \cdots (a^L, \mathfrak{d}_1^L, \dots, \mathfrak{d}_\beta^L)$$

in $([1, \alpha] \times \mathbb{N}^\beta)^+$ via the satisfaction relation $\models_f$ parameterized by $f : \{u_1, u_2\} \rightarrow [1, L]$ (Boolean clauses are omitted, and $i, i' \in \{1, 2\}$):

$$\begin{aligned} & \neg \text{dw} \models_f a(u_i) \stackrel{\text{def}}{\Leftrightarrow} a^{f(u_i)} = a, \\ & \neg \text{dw} \models_f u_i \sim_j u_{i'} \stackrel{\text{def}}{\Leftrightarrow} \mathfrak{d}_j^{f(u_i)} = \mathfrak{d}_j^{f(u_{i'})}, \\ & \neg \text{dw} \models_f u_i \prec_j u_{i'} \stackrel{\text{def}}{\Leftrightarrow} \mathfrak{d}_j^{f(u_i)} < \mathfrak{d}_j^{f(u_{i'})}, \\ & \neg \text{dw} \models_f u_i = u_{i'} \stackrel{\text{def}}{\Leftrightarrow} f(u_i) = f(u_{i'}), \\ & \neg \text{dw} \models_f u_i = 1 + (u_{i'}) \stackrel{\text{def}}{\Leftrightarrow} f(u_i) = f(u_{i'}) + 1, \end{aligned}$$

- $\mathfrak{dw} \models_f u_i < u_{i'} \stackrel{\text{def}}{\Leftrightarrow} f(u_i) < f(u_{i'}),$
- $\mathfrak{dw} \models_f \exists u_i \phi \stackrel{\text{def}}{\Leftrightarrow} \text{there is } p \in [1, L] \text{ such that } \mathfrak{dw} \models_{f[u_i \mapsto p]} \phi.$

A sentence ϕ in $\text{FO2}_{\alpha, \beta}(<, +1, =, \sim, \prec)$ is satisfiable $\stackrel{\text{def}}{\Leftrightarrow}$ there is a data word $\mathfrak{dw}$ in $([1, \alpha] \times \mathbb{N}^\beta)^+$ such that $\mathfrak{dw} \models \phi$ (no need to specify a variable assignment since ϕ is closed).

Let us recall major results about FO2 on data words; $\text{FO2}_{\alpha, 0}(<, +1, =)$ was introduced in Section 2.3 and the others just above.

THEOREM 4.13.

- (I) *The satisfiability problem for $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha, 0}(<, +1, =)$ is NEXPTIME-complete [Etesami et al. 1997] (see also [Weis 2011, Corollary 2.2.4]).*
- (II) *The satisfiability problem for $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha, 1}(<, +1, =, \sim)$ is decidable and closely related to the reachability problem for Petri nets [Bojańczyk et al. 2011; David 2009, Theorem 3].*
- (III) *The satisfiability problem for $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha, 2}(<, +1, =, \sim)$ is undecidable [Bojańczyk et al. 2011; David 2009, Proposition 27].*
- (IV) *The satisfiability problem for $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha, 1}(<, +1, =, \sim, \prec)$ is undecidable [Bojańczyk et al. 2011; David 2009].*

Theorem 4.13(IV) shall be used in this section but decidability can be regained, as shown in [Schwentick and Zeume 2012], where finite satisfiability of FO2 over data words with a linear order on the positions and a linear order and a corresponding successor relation on the data values shown in EXPSPACE [Schwentick and Zeume 2012].

A slightly simpler undecidability proof can be also obtained from the undecidability of the satisfiability problem for $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha, 1}(<, +1, =, \sim, \prec)$ on data words [Bojańczyk et al. 2011] (see Theorem 4.13(IV)). In order to obtain a reduction from the halting problem for Minsky machines, we had to deal with encoding of instructions, which is a bit tedious in some places but the gain has been to obtain a *master reduction*. Master reductions are understood as reductions from decision problems involving, for instance, Turing machines, Minsky machines or any other standard class of computational models. These reductions are always preferred because no intermediate decision problems are involved and therefore this limits the sources of flaws for example. Nevertheless, the reduction from the satisfiability problem $\bigcup_{\alpha \geq 1} \text{FO2}_{\alpha, 1}(<, +1, =, \sim, \prec)$ is a bit simpler if undecidability has to be explained in the most concise way.

Let us briefly provide the main ingredients for such a proof. We define a logarithmic-space translation t as follows. A position u in the data word corresponds to a location on the main path of the fishbone encoding the same position but for the (unique) part related to the (unique) datum. In the translation process, we freely use macros defined earlier ($i, j \in \{1, 2\}$).

- t is homomorphic for Boolean connectives,
- $t(u_i = u_j) \stackrel{\text{def}}{=} u_i = u_j,$
- $t(u_i < u_j) \stackrel{\text{def}}{=} \text{ls}(u_i, u_j) \wedge u_i \neq u_j,$
- $t(u_j = 1 + (u_i)) \stackrel{\text{def}}{=} \top * (\text{ls}'(u_i, u_j) \wedge (\#u_i^{+3} = 1) \wedge \neg(\#u_i^{+4} \geq 0)),$
- $t(a(u_i)) \stackrel{\text{def}}{=} \exists u_{3-i} (u_{3-i} \hookrightarrow u_i) \wedge (\#u_{3-i} = a + 2),$
- $t(u_i \sim_1 u_j) \stackrel{\text{def}}{=} \#u_i = \#u_j,$
- $t(u_i \prec_1 u_j) \stackrel{\text{def}}{=} \#u_i + 1 \leq \#u_j,$

— $t(\exists u_i \phi) \stackrel{\text{def}}{=} \exists u_i (\sharp u_i \geq \alpha + 3) \wedge t(\phi)$.

LEMMA 4.14. *Let ϕ be a formula in $\text{FO2}_{\alpha,1}(<, +1, =, \sim, <)$.*

- (I) *For every data word $\mathfrak{dw}$ in $([1, \alpha] \times \mathbb{N})^+$, $\mathfrak{dw} \models \phi$ iff $\mathfrak{h}_{\mathfrak{dw}} \models \mathbf{dw}(\alpha, 1) \wedge t(\phi)$.*
- (II) *Let $\mathfrak{h}$ be a heap such that $\mathfrak{h} \models \mathbf{dw}(\alpha, 1) \wedge t(\phi)$, then there is a data word $\mathfrak{dw}$ in $([1, \alpha] \times \mathbb{N})^+$ such that $\mathfrak{h}$ and $\mathfrak{h}_{\mathfrak{dw}}$ are isomorphic and $\mathfrak{dw} \models \phi$.*

As a corollary, 1SL2 is undecidable, since for any ϕ in $\text{FO2}_{\alpha,1}(<, +1, =, \sim, <)$, ϕ is satisfiable iff $\mathbf{dw}(\alpha, 1) \wedge t(\phi)$ is satisfiable in 1SL2. Note that the gain with the proof sketched above is not that significant; however, we do not need the material from Section 4.4. (Still, though, we need to express arithmetical constraints between the numbers of predecessors.) Furthermore, we need to formally prove Lemma 4.14, which is of a complexity comparable to the material in Section 4.5. Observe also that Section 4.4 is interesting in its own right because it establishes a result about the expressive power of 1SL2 by performing surgical cuts.

5. CONCLUSION

In the paper, we have shown that two-variable first-order separation logic (1SL2) is undecidable by designing a simple master reduction from the halting problem for Minsky machines and, if we drop the magic wand operator (a fragment called 1SL2(*)), we get decidability but with non-elementary complexity. Our contribution is related to the hardness results when only two variables are used (no program variables, only two quantified variables, no unary predicate symbols). Heavy recycling of variables is done, following the classical result for modal logic [Gabbay 1981]. In order to get these results, we have shown a simple encoding of data words with multiple attributes that is used both for the undecidability and for the non-elementarity result. This nice relationship with data logics [Bojańczyk et al. 2011] is also complemented by the fact that we have shown how to encode propositional interval temporal logic (PITL) [Moszkowski 2004] into 1SL2(*). Hence, we believe we have established promising bridges between data logic(s), interval temporal logic(s), and separation logic(s), apart from providing additional evidence of the importance of intervals in logics for computer science. Moreover, we have introduced a separation modal logic MLH. While its decidability status remains open, we have proved that its restriction to separating conjunction is decidable with non-elementary complexity. This has been not only an opportunity to better understand the expressive power of separation logic fragments but also to significantly populate 1SL2's inner circle (see Figure 1 for an overview).

ACKNOWLEDGMENT

We would like to thank Emanuel Kieroński, Ben Moszkowski, and Guido Sciavicco. for insightful exchanges about interval temporal logics and two-variable first-order logic(s). Furthermore, we are indebted to two anonymous referees for many helpful and insightful suggestions that helped us to improve the quality of this manuscript.

References

- J. Allen. 1983. Maintaining knowledge about temporal intervals. *Commun. ACM* 26, 11 (1983), 832–843.
- K. Bansal, R. Brochenin, and E. Lozes. 2009. Beyond Shapes: Lists with Ordered Data. In *FOSSACS'09 (Lecture Notes in Computer Science)*, Vol. 5504. Springer, 425–439.
- S. Benaïm, M. Benedikt, W. Charatonik, E. Kieroński, R. Lenhardt, F. Mazowiecki, and J. Worrell. 2013. Complexity of two-variable logic over finite trees. In *ICALP'13 (Lecture Notes in Computer Science)*, Vol. 7966. Springer, 74–88.
- P. Blackburn, M. de Rijke, and Y. Venema. 2001. *Modal Logic*. Cambridge University Press.

- M. Bojańczyk, C. David, A. Muscholl, Th. Schwentick, and L. Segoufin. 2011. Two-variable logic on data words. *ACM Transactions on Computational Logic* 12, 4 (2011), 27.
- E. Börger, E. Grädel, and Y. Gurevich. 1997. *The Classical Decision Problem*. Springer.
- P. Bouyer. 2002. A Logical Characterization of Data Languages. *Inform. Process. Lett.* 84, 2 (2002), 75–85.
- R. Brochenin, S. Demri, and E. Lozes. 2012. On the Almighty Wand. *Information and Computation* 211 (2012), 106–137.
- J. Brotherston, C. Fuhs, N. Gorogiannis, and J. Navarro Perez. 2014. A Decision Procedure for Satisfiability in Separation Logic with Inductive Predicates. In *CSL-LICS'14*. ACM.
- J. Brotherston and M. Kanovich. 2014. Undecidability of Propositional Separation Logic and Its Neighbours. *Journal of the Association for Computing Machinery* 61, 2 (2014).
- C. Calcagno, P. O'Hearn, and H. Yang. 2001. Computability and Complexity Results for a Spatial Assertion Language for Data Structures. In *FSTTCS'01 (Lecture Notes in Computer Science)*, Vol. 2245. Springer, 108–119.
- W. Charatonik, E. Kieroński, and F. Mazowiecki. 2014. Decidability of Weak Logics with Deterministic Transitive Closure. In *CSL-LICS'14*. ACM.
- B. Cook, Ch. Haase, J. Ouaknine, M. Parkinson, and J. Worrell. 2011. Tractable Reasoning in a Fragment of Separation Logic. In *CONCUR'11 (Lecture Notes in Computer Science)*. Springer, 235–249.
- J.-R. Courtault and D. Galmiche. 2013. A Modal BI Logic for Dynamic Resource Properties. In *LFCS'13 (Lecture Notes in Computer Science)*, Vol. 7734. Springer, 134–148.
- C. David. 2009. *Analyse de XML avec données non-bornées*. Ph.D. Dissertation. LIAFA, Université Paris VII.
- A. Dawar, Ph. Gardner, and G. Ghelli. 2007. Expressiveness and complexity of graph logic. *Information and Computation* 205, 3 (2007), 263–310.
- M. de Rijke. 1992. The modal logic of inequality. *The Journal of Symbolic Logic* 57, 2 (1992), 566–584.
- N. Decker, P. Habermehl, M. Leucker, and D. Thoma. 2014. Ordered Navigation on Multi-attributed Data Words. In *CONCUR'14 (Lecture Notes in Computer Science)*, Vol. 8704. 497–511.
- A. Degtyarev, M. Fisher, and A. Lisitsa. 2002. Equality and monodic first-order temporal logic. *Studia Logica* 72, 2 (2002), 147–156.
- S. Demri and M. Deters. 2014. Expressive completeness of separation logic with two variables and no separating conjunction. In *CSL-LICS'14*. ACM, 37.
- S. Demri, D. Galmiche, D. Larchey-Wendling, and D. Mery. 2014. Separation Logic with One Quantified Variable. In *CSR'14 (Lecture Notes in Computer Science)*, Vol. 8476. Springer, 125–138.
- S. Demri and Ph. Schnoebelen. 2002. The complexity of Propositional Linear Temporal Logics in Simple Cases. *Information and Computation* 174, 1 (2002), 84–103.
- K. Etessami, M. Vardi, and Th. Wilke. 1997. First-Order Logic with Two variables and Unary Temporal logics. In *LICS'97*. IEEE, 228–235.
- D. Figueira. 2010. *Reasoning on words and trees with data*. Ph.D. Dissertation. ENS de Cachan.
- D. Gabbay. 1981. Expressive Functional Completeness in Tense Logic. In *Aspects of Philosophical Logic*. Reidel, 91–117.
- D. Gabbay and V. Shehtman. 1993. Undecidability of modal and intermediate first-order logics with two individual variables. *The Journal of Symbolic Logic* 58, 3 (1993), 800–823.
- E. Grädel, Ph. Kolaitis, and M. Vardi. 1997a. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic* 3, 1 (1997), 53–69.
- E. Grädel, M. Otto, and E. Rosen. 1997b. Two-Variable logic with counting is Decidable. In *LICS'97*. IEEE, 306–317.
- E. Grädel, M. Otto, and E. Rosen. 1999. Undecidability Results on Two-Variable Logics. *Archive for Mathematical Logic* 38, 4–5 (1999), 313–354.
- J. Halpern. 1995. The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence* 75, 2 (1995), 361–372.
- D. Harel, D. Kozen, and J. Tiuryn. 2000. *Dynamic Logic*. MIT Press.
- M. Hennessy and R. Milner. 1980. On observing nondeterminism and concurrency.. In *ICALP'80 (Lecture Notes in Computer Science)*, Vol. 85. Springer, 299–309.
- R. Iosif, A. Rogalewicz, and J. Simacek. 2013. The Tree Width of Separation Logic with Recursive Definitions. In *CADE'13 (Lecture Notes in Computer Science)*, Vol. 7898. Springer, 21–38.
- A. S. Kahr, E. F. Moore, and H. Wang. 1962. Entscheidungsproblem reduced to the $\forall \exists \forall$ case. *Proc. Nat. Acad. Sci. U. S. A.* 48, 3 (1962), 365–377.

- E. Kieroński, J. Michaliszyn, I. Pratt-Hartmann, and L. Tendera. 2012. Two-Variable First-Order Logic with Equivalence Closure. In *LICS'12*. IEEE, 431–440.
- D. Larchey-Wendling and D. Galmiche. 2013. Nondeterministic Phase Semantics and the Undecidability of Boolean BI. *ACM Transactions on Computational Logic* 14, 1 (2013).
- H. Lewis. 1980. Complexity Results for classes of quantificational formulas. *J. Comput. System Sci.* 21 (1980), 317–353.
- C. Lutz and U. Sattler. 2002. The complexity of reasoning with Boolean Modal Logics. In *Advances in Modal Logics 2000, Volume 3*. World Scientific, 329–348.
- M. L. Minsky. 1967. *Computation: finite and infinite machines*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Ch. Morgan. 1976. Methods for automated theorem proving in non classical logics. *IEEE Trans. Comput.* 25, 8 (1976), 852–862.
- M. Mortimer. 1975. On language with two variables. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* 21 (1975), 135–140.
- B. Moszkowski. 1983. *Reasoning about digital circuits*. Technical Report STAN-CS-83-970. Dept. of Computer Science, Stanford University, Stanford, CA.
- B. Moszkowski. 2004. A Hierarchical Completeness Proof for Propositional Interval Temporal Logic with Finite Time. *Journal of Applied Non-Classical Logics* 14, 1–2 (2004), 55–104.
- M. Otto. 2001. Two Variable First-Order Logic over Ordered Domains. *The Journal of Symbolic Logic* 66, 2 (2001), 685–702.
- L. Pacholski, W. Szwast, and L. Tendera. 1997. Complexity of Two-Variable logic with counting. In *LICS'97*. IEEE, 318–327.
- A. Prior. 1967. *Past, Present and Future*. Oxford University Press.
- M. Rabin. 1969. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.* 41 (1969), 1–35.
- J. C. Reynolds. 2002. Separation logic: a logic for shared mutable data structures. In *LICS'02*. IEEE, 55–74.
- Th. Schwentick and Th. Zeume. 2012. Two-Variable Logic and Two Order Relations. *Logical Methods in Computer Science* 8 (2012), 1–27.
- D. Scott. 1962. A decision method for validity of sentences in two variables. *The Journal of Symbolic Logic* 27 (1962), 377.
- L. Stockmeyer. 1974. *The complexity of decision problems in automata theory and logic*. Ph.D. Dissertation. Department of Electrical Engineering, MIT.
- W. Szwast and L. Tendera. 2013. FO² with one transitive relation is decidable. In *STACS'13 (LIPIcs)*, Vol. 20. 317–328.
- J. van Benthem. 1976. *Correspondence Theory*. Ph.D. Dissertation. Mathematical Institute, University of Amsterdam.
- Ph. Weis. 2011. *Expressiveness and Succinctness of First-Order Logic on Finite Words*. Ph.D. Dissertation. University of Massachusetts.
- H. Yang. 2001. *Local Reasoning for Stateful Programs*. Ph.D. Dissertation. University of Illinois, Urbana-Champaign.
- Zhou Chaochen. 2008. In *Logics of Specification Languages*, D. Björner and M. Henson (Eds.). Springer, Chapter Reviews on “Duration Calculus”, 609–611. Monographs in Theoretical Computer Science. An EATCS Series.