Expressive completeness of separation logic with two variables and no separating conjunction

Stéphane Demri, Morgan Deters

► To cite this version:

Stéphane Demri, Morgan Deters. Expressive completeness of separation logic with two variables and no separating conjunction. CSL-LICS '14: JOINT MEETING OF the Twenty-Third EACSL Annual Conference on COMPUTER SCIENCE LOGIC, Jul 2014, Vienna, Austria. pp.37:1-37:10, 10.1145/2603088.2603142 . hal-03201035

HAL Id: hal-03201035

<https://hal.science/hal-03201035>

Submitted on 17 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.

Expressive Completeness of Separation Logic With Two Variables and No Separating Conjunction^{*}

Stéphane Demri

New York University, USA & CNRS, France
demri@cs.nyu.edu

Morgan Deters

New York University, USA
mdeters@cs.nyu.edu

Abstract

We show that first-order separation logic with one record field restricted to two variables and the separating implication (no separating conjunction) is as expressive as weak second-order logic, substantially sharpening a previous result. Capturing weak second-order logic with such a restricted form of separation logic requires substantial updates to known proof techniques. We develop these, and as a by-product identify the smallest fragment of separation logic known to be undecidable: first-order separation logic with one record field, two variables, and no separating conjunction.

Categories and Subject Descriptors F.4.1 [*Mathematical Logic and Formal Languages*]: Mathematical Logic

Keywords separation logic, expressive completeness

1. Introduction

Expressive completeness The literature is rich with results comparing the expressive power of non-classical logics with first-order or second-order logic. For instance, the celebrated Kamp’s Theorem amounts to stating that linear-time temporal logic (LTL) is equal in expressive power to first-order logic. More generally, we know the expressive completeness of Stavi connectives for general linear time, see e.g. [16]. This has been refined to the restriction to two variables, leading to the equivalence between unary LTL and FO2, see e.g. [14, 30]. In addition, there is a wealth of results relating first-order logic with two variables and non-classical logics, providing a neat characterization of the expressive power of many formalisms since first-order logic and second-order logic are queen logics. For instance, Boolean modal logic with converse and identity is as expressive as FO2 [26]. In the realm of interval temporal logics, we also know expressive completeness of metric propositional neighborhood logic with respect to the two-variable fragment of first-order logic for linear orders with successor function, interpreted over natural numbers [5].

^{*}Work partially supported by the EU Seventh Framework Programme (under grant PIOF-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).

In this paper, we compare separation logic restricted to two variables with (weak) second-order logic over concrete heaps.

Expressive power of separation logic Separation logic is used as an assertion language for Hoare-style proof systems about programs with pointers [28], and there is an ongoing quest for understanding its complexity and expressive power. Alternatively, there are a lot of activities to develop verification methods with decision procedures for fragments of practical use, see e.g. [10]. Many decision procedures have been designed for fragments of separation logics or abstract variants, from analytic methods [17, 20] to translation to theories handled by SMT solvers [27], passing via graph-based algorithms [19].

Theoretical issues for separation logic stem from the design of expressive fragments with relatively low complexity (see e.g. [10]) to the extension of known decidability results, see e.g. [2, 4, 22]. Indeed, it is known since [9] that first-order separation logic with two record fields (called herein 2SL) is undecidable (with a proof that does not require separating connectives and uses Trakhtenbrot’s Theorem). This is sharpened in [7] by showing that first-order separation logic with a unique record field (called herein 1SL) is also undecidable, as a consequence of the expressive equivalence between 1SL and weak second-order logic. More recently, 1SL restricted to two variables (1SL2) is shown undecidable too [12] but without touching the central question of expressive completeness—the purpose of the current paper. From the very beginning, the relationships between separation logic and second-order logic have been quite puzzling (see e.g. an interesting answer with infinite arbitrary structures in [23]). In this paper, we go one step further by showing that two variables suffice to get expressive completeness. As a consequence, 1SL2($\neg$) (that is, 1SL2 without separating conjunction) is undecidable too, which should not be confused with undecidability results from [8, 24] obtained in an alternative setting with propositional variables and no first-order quantification. It is fair to recall that separating implication has been less well-studied than separating conjunction in the literature, but its use for program verification is far more recognized nowadays; see e.g. [25, Section 1] for a recent, insightful analysis (see also [20, Section 8]).

Our contribution In this paper, we sharpen the main result in [7], namely we show that first-order separation logic with one record field, two quantified variables, and no separating conjunction is as expressive as weak second-order logic on heaps; in short, 1SL2($\neg$) $\equiv$ WSOL. Even though conjectured in [6, 7], it is surprising that two variables suffice, and that further we are able to drop the separating conjunction, thus obtaining expressive completeness and undecidability with only two variables and the magic wand operator. In doing so, we improve previous undecidability results about separation logic [7, 9, 12]. Because we forbid ourselves the use of many syntactic resources, this underlines even further the power of the magic wand. By way of comparison with [18, 21], we show

undecidability of a two-variable logic with second-order features, and our main undecidability result cannot be derived from [18, 21] since in 1SL models, we deal with a single functional binary relation. We believe that we have identified the core of separation logic as far as undecidability and expressive completeness are concerned, since, for instance, first-order separation logic with one record field, one quantified variable and an unbounded number of program variables, has recently been shown decidable and PSPACE-complete [13]. Below, we illustrate how the main result of the paper (Theorem 23) compares with known results from the literature (Section 2 contains definitions of the different logics).

In our proof, most of the difficulties are concentrated on the use of only two variables: we recycle variables as done for modal logics [15], but this is insufficient, especially when separating conjunction is also banished. So, as far as the proof for expressive completeness goes, we borrow some of the first principles from [7], but very quickly we are faced with serious problems when we need to identify in some heap at least $k > 0$ heap patterns (a typical example is to specify that at least $k > 0$ locations point to a given location).

Indeed, the standard way to identify such patterns is to use an unbounded number of variables or the separating conjunction. So, in the presence of only two variables and by using only the magic wand operator, instead of chopping the heap in k disjoint subheaps, we add $\mathcal{O}(k)$ new patterns so that the newly combined heap satisfies structural properties that witness the presence of the k patterns in the original heap. This high-level description has to be instantiated as many times as we have to identify different types of patterns, but this new point of view allows us to go far beyond what is known today (see e.g., the proof of Lemma 5). At times, it is not strictly necessary to introduce a radically new method, but instead we can be more thrifty in the way formulae are defined to express desirable properties; of course, this may come with more complex proofs and, above all, more ingenuity to design such formulae. It should be noted that the paper is structured in such a way that we provide more and more complex building blocks to establish our main results. Another contribution rests on the fact that we considerably simplify some of the technical insights borrowed from [7] and therefore the current paper proposes a self-contained proof of the equivalence between 1SL2(-*) and weak second-order logic that in many ways is much simpler than what has been done so far, even though our results are stronger. Extensions with program variables or with heaps having $k > 1$ record fields are presented in Section 5.5.

2. Preliminaries

2.1 First-order separation logic with one selector (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*, if their domains are disjoint; when this holds, we write $h_1 \uplus h_2$ to denote the disjoint union. *Locations* are elements of $\mathbb{N}$ and are denoted by l , possibly decorated with exponents or subscripts. We write $l_1 \rightarrow l_2 \rightarrow \dots \rightarrow l_m$ 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 $\{l' \in \mathbb{N} : h(l') = l\}$ made of *predecessors* of l (heap h is implicit in the expression $\#l$) and $\#l^*$ to denote the cardinal of $(\{l' \in \mathbb{N} : h(l') = l\} \setminus \{l\})$. So, $\#l = \#l^*$ iff (either $l \notin \text{dom}(h)$ or $h(l) \neq l$).

Let $\text{FVAR} = \{u_1, u_2, \dots\}$ be a countably infinite set of variables. Formulae of 1SL are defined by the grammar $\phi ::= u_i = u_j \mid u_i \hookrightarrow u_j \mid \phi \wedge \phi \mid \neg \phi \mid \phi * \phi \mid \phi \multimap \phi \mid \exists u_i \phi$. The connective $*$ is called the *separating conjunction* and $\multimap$ is called the *separating implication* (also known as the *magic wand*). We make use of standard notations for the derived connectives.

A *valuation* is a map $f : \text{FVAR} \rightarrow \mathbb{N}$. The satisfaction relation $\models$ is parameterized by valuations and is defined as follows (Boolean clauses are omitted):

- $h \models u_i = u_j$ iff $f(u_i) = f(u_j)$.
- $h \models u_i \hookrightarrow u_j$ iff $f(u_i) \in \text{dom}(h)$ and $h(f(u_i)) = f(u_j)$.
- $h \models \phi_1 * \phi_2$ iff there exist h_1, h_2 such that h_1 and h_2 are disjoint, $h = h_1 \uplus h_2$, $h_1 \models \phi_1$ and $h_2 \models \phi_2$.
- $h \models \phi_1 \multimap \phi_2$ iff for all h' , if h and h' are disjoint, and $h' \models \phi_1$ then $h \uplus h' \models \phi_2$.
- $h \models \exists u_i \phi$ iff there is $l \in \mathbb{N}$ such that $h \models_{f[u_i \mapsto l]} \phi$.

We also introduce so-called *septraction* operator $\multimap^*$: $\phi \multimap^* \psi$ is defined as $\neg(\phi * \neg\psi)$. So, $h \models \phi \multimap^* \psi$ iff there is h' disjoint from h such that $h' \models \phi$ and $h \uplus h' \not\models \psi$. Septraction states the existence of a disjoint heap satisfying a formula and for which its addition to the original heap satisfies another formula.

For $i \geq 1$, 1SL $_i$ denotes the fragment of 1SL restricted to i variables and 1SL $_i(-*)$ to denote its restriction when separating conjunction is disallowed. Let $\mathcal{L}$ be a logic among 1SL, 1SL $_i$, 1SL $_i(-*)$. The *satisfiability problem* for $\mathcal{L}$ takes as input a sentence ϕ from $\mathcal{L}$ and asks whether there is h such that $h \models \phi$ (regardless of valuation, as ϕ has no free variables).

Theorem 1. [7, 12] The satisfiability problem for 1SL is undecidable, even if restricted to 1SL $_2$.

2.2 Weak second-order logic (WSOL)

We consider a family $\text{SVAR} = (\text{SVAR}_i)_{i \geq 1}$ of second-order variables, denoted by $P, Q, R, \dots$ and interpreted as finite relations over $\mathbb{N}$. A second-order valuation f is an interpretation of the second-order variables such that for every $P \in \text{SVAR}_i$, $f(P)$ is a finite subset of $\mathbb{N}^i$. Formulae ϕ of WSOL are defined by $\phi ::= u_i = u_j \mid u_i \hookrightarrow u_j \mid \phi \wedge \phi \mid \neg \phi \mid \exists u_i \phi \mid \exists P \phi \mid P(u_1, \dots, u_n)$, where the P are second-order variables with $P \in \text{SVAR}_n$ for some $n \geq 1$. We write DSOL (dyadic second-order logic) to denote the restriction of WSOL to second-order variables in SVAR_2 . Like 1SL, models for WSOL are finite heaps and quantifications are done over all the locations. Satisfaction relation $\models$ is defined as follows (f is a hybrid valuation providing interpretation for any variable):

- $h \models \exists P \phi$ iff there is a *finite* relation $R \subseteq \mathbb{N}^n$ such that $h \models_{f[P \mapsto R]} \phi$ where $P \in \text{SVAR}_n$.
- $h \models P(u_1, \dots, u_n)$ iff $(f(u_1), \dots, f(u_n)) \in f(P)$.

The satisfiability problem for WSOL takes as input a sentence ϕ in WSOL and asks whether there is h such that $h \models \phi$. By Trakhtenbrot's Theorem (see e.g. [3, 29]), the satisfiability problem for DSOL (and therefore for WSOL) is undecidable since finite satisfiability for first-order logic with a unique binary relation symbol

is undecidable. Note that a monadic second-order variable can be simulated by a binary second-order variable from SVAR_2 , and this can be used to relativize a formula from DSOL in order to check finite satisfiability.

Theorem 2. [7] ISL, WSOL and DSOL have the same expressive power.

Consequently, in order to show that $\text{ISL2}(\ast)$ is as expressive as WSOL (our main result), it is sufficient to prove that every sentence from DSOL has an equivalent sentence in $\text{ISL2}(\ast)$. It is worth noting that Theorem 2 extends to $k > 1$ record fields [7], which actually requires a simpler proof. A similar adaptation is possible from our main result (see Section 5.5).

3. Expressing Properties in $\text{ISL2}(\ast)$

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 permuting the two variables.

In $\text{ISL2}(\ast)$ it is not difficult to state that the domain of the heap is empty or that it is a singleton—for example, we can write $\exists u ((\exists \bar{u} u \leftrightarrow \bar{u}) \wedge \forall \bar{u} (\bar{u} \neq u \Rightarrow \neg(\exists \bar{u} \bar{u} \leftrightarrow u)))$. However, to express that the domain contains exactly two locations, we run out of variables if we attempt to use a similar technique. Below, we propose a new and natural method to compare a number of predecessors against a constant or to express a reachability property. First, though, we define simple, standard formulae.

- u has a successor: $\text{alloc}(u) \stackrel{\text{def}}{=} \exists \bar{u} u \leftrightarrow \bar{u}$.
- u is an *isolated location* (it is not in $\text{dom}(h) \cup \text{ran}(h)$): $\text{isoloc}(u) \stackrel{\text{def}}{=} \neg \text{alloc}(u) \wedge \neg \exists \bar{u} \bar{u} \leftrightarrow u$.
- $\text{dom}(h)$ has exactly one location: $(\text{size} = 1) \stackrel{\text{def}}{=} \exists u (\text{alloc}(u) \wedge \forall \bar{u} (u \neq \bar{u} \Rightarrow \neg \text{alloc}(\bar{u})))$.
- u has at least one predecessor: $\sharp u > 0 \stackrel{\text{def}}{=} \exists \bar{u} \bar{u} \leftrightarrow u$. Naturally, we can also write: $\sharp u = 0 \stackrel{\text{def}}{=} \neg(\sharp u > 0)$.

Let us now proceed with more complicated constructions.

3.1 Counting z-predecessors

In this paper, we will make heavy use of counting predecessors of locations. Without an adequate number of variables to refer to locations of interest, we instead “remember” locations by installing a large (and unique) number of predecessors to a location in the heap. At another point in the formula, we can identify this location and operate on it, having not used any of the logic’s limited syntactic resources.

As one may imagine, counting and comparing numbers of predecessors is difficult to do in this logic. We first build up a method of counting a certain type of predecessor, then use it to formulate more involved constructions, and finally introduce a way to count the full number of predecessors of a location—that is, to compare its number of predecessors with some $k \geq 0$. After this is achieved, nearly the entirety of Section 4 extends this to comparing numbers of predecessors between two locations (rather than just a comparison to some k).

We first introduce the notion of *z-predecessors*, which are predecessors of a location that themselves have no predecessors (‘z’ is for ‘zero’). We can then write:

- u ’s predecessors are all z-predecessors: $\text{allzpred}(u) \stackrel{\text{def}}{=} \forall \bar{u} \bar{u} \leftrightarrow u \Rightarrow \sharp \bar{u} = 0$.

- u has zero z-predecessors:

$$\sharp_z u = 0 \stackrel{\text{def}}{=} \forall \bar{u} \bar{u} \leftrightarrow u \Rightarrow \sharp \bar{u} > 0.$$

- u has at most $k > 0$ z-predecessors:

$$\sharp_z u \leq k \stackrel{\text{def}}{=} (\text{size} = 1) \multimap \sharp_z u \leq k - 1 \text{ with } \sharp_z u \leq 0 \text{ defined as } \sharp_z u = 0.$$

Lemma 3. Let h be a heap, f be a valuation, and $k \in \mathbb{N}$.

$h \models_f \sharp_z u \leq k$ iff $\text{card}(\{l \in \mathbb{N} : \sharp l = 0, h(l) = f(u)\}) \leq k$.

Variants $\sharp_z u \bowtie k$ with $\bowtie \in \{=, <, >, \geq\}$ are easily defined.

3.2 A matter of forks and knives

To address the problem of referring to memory locations despite not having the variables to do so, we introduce the notion of *forks*. Forks are simple, recognizable shapes that we add to the heap with the magic wand. Forks are also a critical building block for comparing predecessors in Section 4. A *fork* in h is a sequence of distinct locations l, l_0, l_1, l_2 such that $h(l_0) = l$, $\sharp l_0 = 2$, $h(l_1) = h(l_2) = l_0$ and $\sharp l_1 = \sharp l_2 = 0$. The *endpoint* of the fork is l . The fork is *isolated* iff $\sharp l = 1$ and $l \notin \text{dom}(h)$. Similarly, a *knife* in h is a sequence of distinct locations l, l_0, l_1 such that $h(l_0) = l$, $\sharp l_0 = 1$, $h(l_1) = l_0$ and $\sharp l_1 = 0$. The *endpoint* of the knife is l . The knife is *isolated* iff $\sharp l = 1$ and $l \notin \text{dom}(h)$.

To identify fork and knife endpoints in a heap, we define:

$$\begin{aligned} \text{forkendpt}(u) &\stackrel{\text{def}}{=} \exists \bar{u} (\bar{u} \leftrightarrow u \wedge \sharp_z \bar{u} = 2 \wedge \text{allzpred}(\bar{u})) \\ \text{knifeendpt}(u) &\stackrel{\text{def}}{=} \exists \bar{u} (\bar{u} \leftrightarrow u \wedge \sharp_z \bar{u} = 1 \wedge \text{allzpred}(\bar{u})) \end{aligned}$$

Now, let $\text{forky}(u)$ be a formula stating that *all* predecessors of $f(u)$, possibly except $f(u)$, are endpoints of forks: $\text{forky}(u) \stackrel{\text{def}}{=} \forall \bar{u} ((\bar{u} \leftrightarrow u \wedge \bar{u} \neq u) \Rightarrow \text{forkendpt}(\bar{u}))$. Three forks, two endpoints, and a forky location are depicted at left. Next, let

$\text{antiforky}(u)$ be a formula stating that *no* predecessor of $f(u)$ is the endpoint of a fork, and let $\text{antiknify}(u)$ be a formula stating that *no* predecessor of $f(u)$ is the endpoint of a knife. We define these as

$\text{antiforky}(u) \stackrel{\text{def}}{=} \forall \bar{u} (\bar{u} \leftrightarrow u \Rightarrow \neg \text{forkendpt}(\bar{u}))$ and $\text{antiknify}(u) \stackrel{\text{def}}{=} \forall \bar{u} (\bar{u} \leftrightarrow u \Rightarrow \neg \text{knifeendpt}(\bar{u}))$. Note the asymmetry between $\text{forky}(u)$ and $\text{antiforky}(u)$ since $f(u)$ does not have to be the endpoint of a fork in $\text{forky}(u)$ (which would be then impossible to realize for all the predecessors of $f(u)$ if $f(u)$ were a self-loop). It is also easy to enforce that the heap is made of a single fork, which will be useful afterwards.

Lemma 4. There is a formula 1fork in $\text{ISL2}(\ast)$ such that for all h , we have $h \models \text{1fork}$ iff h is only made of a single, isolated fork.

Now let us build formulae that constrain the number of predecessors (not just z-predecessors): with $k \in \mathbb{Z}$, we define $\sharp u \leq k \stackrel{\text{def}}{=} \perp$ if $k < 0$, $\sharp u \leq 0 \stackrel{\text{def}}{=} \neg \exists \bar{u} \bar{u} \leftrightarrow u$, and for $k > 0$,

$$\sharp u \leq k \stackrel{\text{def}}{=} (u \leftrightarrow u \wedge \sharp u \leq k - 1) \vee (\neg(u \leftrightarrow u) \wedge \sharp u \leq k)$$

where $\sharp u \leq 0 \stackrel{\text{def}}{=} \neg \exists \bar{u} (\bar{u} \leftrightarrow u \wedge \bar{u} \neq u)$ and $\sharp u \leq k \stackrel{\text{def}}{=} (\sharp u = 0) \multimap (\text{antiforky}(u) \wedge \underbrace{\text{1fork} \multimap \dots \multimap \text{1fork}}_{k \text{ times}} \multimap \text{forky}(u))$

for $k > 0$. In a nutshell, $f(u)$ has at most $k > 0$ predecessors if one can make $f(u)$ antiforky without changing its predecessor count (always possible) and then adding k forks to make $f(u)$ forky (we distinguish the case when $f(u)$ is a self-loop).

Lemma 5. Let $k \in \mathbb{N}$, h be a heap and f be a valuation. (I) $h \models_f \#u \leq k$ iff $\#f(u)^* \leq k$; (II) $h \models_f \#u \leq k$ iff $\#f(u) \leq k$.

Thus we can express in $\text{ISL2}(\ast)$ the following properties:

- $f(u)$ has at least k predecessors: $\#u \geq k \stackrel{\text{def}}{=} \neg(\#u \leq k - 1)$.
- $f(u)$ has exactly k predecessors: $\#u = k \stackrel{\text{def}}{=} \#u \leq k \wedge \#u \geq k$.

Lonely memory cells Now, we can express a useful property on the heap itself. The formula below states that $f(u)$ is an isolated cell: $\text{isocell}(u) \stackrel{\text{def}}{=} \text{alloc}(u) \wedge \#u = 0 \wedge (\exists \bar{u} \ u \hookrightarrow \bar{u} \wedge \#\bar{u} = 1 \wedge \neg \text{alloc}(\bar{u}))$. A heap h is *segmented* whenever $\text{dom}(h) \cap \text{ran}(h) = \emptyset$ and no location has strictly more than one predecessor. Otherwise said, all the memory cells are isolated. Being segmented can be naturally expressed in $\text{ISL2}(\ast)$ too:

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

3.3 Expressing reachability

In ISL, reachability can be expressed, and [12] gives a technique for doing so with the two-variable restriction, itself a variant of material from [7, 11]. Without the separating conjunction, we can still specify reachability from $f(u)$ to $f(\bar{u})$, but we need a new technique: we put a fork on $f(u)$, propagate this fork forward in the heap, and finally check whether $f(\bar{u})$ is the endpoint of some fork. This is analogous to the way reachability is handled with a monadic second-order predicate, but here, a finite set of locations is identified by propagation of forks. We define $\text{propagate}(u)$ as the formula characterizing the property that $f(u)$ is the endpoint of some fork, and that the property of being the endpoint of some fork is propagated along memory cells. The reachability predicate is written $\text{reach}(u, \bar{u})$:

$$\begin{aligned} \text{propagate}(u) &\stackrel{\text{def}}{=} \text{forkendpt}(u) \wedge \\ &\quad \forall u \forall \bar{u} ((u \hookrightarrow \bar{u} \wedge \text{forkendpt}(u)) \Rightarrow \text{forkendpt}(\bar{u})) \\ \text{reach}(u, \bar{u}) &\stackrel{\text{def}}{=} \top \ast (\text{propagate}(u) \Rightarrow \text{forkendpt}(\bar{u})) \end{aligned}$$

The purpose of “ $\top \ast$ ” is to destroy forks in the original heap (if any) and to augment it with new forks such that $\text{propagate}(u)$ is satisfied. Of course, a disjoint subheap could change the very reachability property we are testing. But some such subheaps do not, and the universal quantification of the magic wand ensures that we are on the safe side: if for *all* combined heaps satisfying $\text{propagate}(u)$, $f(\bar{u})$ is the endpoint of a fork, then we can conclude that $f(u)$ reaches $f(\bar{u})$ in the original heap.

Lemma 6. Given a heap h and valuation f , we have $h \models_f \text{reach}(u, \bar{u})$ iff $h^k(f(u)) = f(\bar{u})$ for some $k \geq 0$.

It is notable that $\text{ISL2}(\ast)$ has no need for built-in reachability predicates, in contrast to formalisms from e.g. [21]. In the rest of the paper, we generalize, in a sense, what was done here in an *ad hoc* manner for reachability, so that any second-order property can be represented in $\text{ISL2}(\ast)$.

4. Comparing Numbers of Predecessors

The main goal of this section is to define in $\text{ISL2}(\ast)$ a formula expressing that $\#f(u) + k \leq \#f(\bar{u}) + k'$, where $k, k' \in \mathbb{N}$, for any heap h and valuation f . Without the restriction on the number of variables, we know that such properties can be expressed in $\text{ISL}(\ast)$ [7]. Note that arithmetical constraints on list lengths can be found in [4] but this is primitive in the logical formalism. By contrast, we show that $\#f(u) + k \leq \#f(\bar{u}) + k'$ can be expressed in $\text{ISL2}(\ast)$ itself. Since a property $\#f(u) + k \leq \#f(\bar{u}) + k'$ requires a formula with $\mathcal{O}(k + k')$ variables in $\text{ISL}(\ast)$ according to [7], herein we need to circumvent this issue by proposing an alternative

way to express the key properties that are helpful to state that $\#f(u) + k \leq \#f(\bar{u}) + k'$. Below, we still use first principles from [7] to construct a formula in $\text{ISL2}(\ast)$ —mainly, how to build a fork from a knife and an isolated memory cell—but we will need to bypass the serious problem of having only two variables at our disposal, without permitting ourselves any use of the separating conjunction.

4.1 Principles and difficulties with $\text{ISL2}(\ast)$

Let h be a heap and f be a valuation for which we wish to check if $\#f(u)^* + k \leq \#f(\bar{u})^* + k'$ holds (it is easy to conclude if $\#f(u) + k \leq \#f(\bar{u}) + k'$ holds, see the proof of Theorem 10). Below, we explain which extensions of h must be performed to achieve this. We mainly describe first principles from [7], but the reader should be warned that in several places, we propose a simplified alternative, apart from the fact that all the formulae need to be part of the restricted fragment $\text{ISL2}(\ast)$.

4.1.1 Preparing the heap

The first step consists in preparing the heap by destroying any forks and knives at $f(u)$ and $f(\bar{u})$, and ensuring there are no isolated memory cells—these properties will be necessary in later steps—while maintaining the number of predecessors at $f(u)$ and $f(\bar{u})$. To do this, we augment h with a heap h_p so that (a) h_p is disjoint from h ; (b) $f(u)$ has the same number of predecessors in h and in $h \uplus h_p$, say $\#f(u)^* = m \geq 0$; (c) $f(\bar{u})$ has the same number of predecessors in h and in $h \uplus h_p$, say $\#f(\bar{u})^* = \bar{m} \geq 0$; (d) In $h \uplus h_p$, $f(u)$ has no predecessor that is an endpoint of a fork or knife; (e) In $h \uplus h_p$, $f(\bar{u})$ has no predecessor that is an endpoint of a fork or knife; (f) $h \uplus h_p$ has no isolated memory cell. This step is always possible, since to destroy the structure of an isolated memory cell, a fork or a knife, it is sufficient to add a memory cell at the adequate position. We take advantage of the formulae $\text{antiforky}(\cdot)$ and $\text{antiknify}(\cdot)$ to establish properties (d) and (e); the others are straightforward.

4.1.2 Addition of a segmented heap

This step consists in checking whether, for each segmented heap h_s satisfying certain properties, the condition $(\mathcal{P})$ defined below holds true. The heap h_s must be segmented and also must satisfy the following: (a) h_s is disjoint from $h \uplus h_p$; (b) $f(u)$ and $f(\bar{u})$ have no predecessors in h_s ; (c) In $h \uplus h_p \uplus h_s$, neither $f(u)$ nor $f(\bar{u})$ has a predecessor that is an endpoint of a fork or knife (they are *antiforky* and *antiknify*). Let n be the number of isolated memory cells in $h \uplus h_p \uplus h_s$. Note that for each $q \geq 0$, it is possible to build h_s^q so that $h \uplus h_p \uplus h_s^q$ has exactly q isolated memory cells and h_s^q satisfies the above conditions.

Figure 1: The construction of forks used for comparison.

In order to construct forks in $h \uplus h_p \uplus h_s$ whose endpoints are predecessors of $f(u)$ or $f(\bar{u})$, either we can augment the heap with a fork, or we can augment the heap with a knife so that its combination with an isolated memory cell in $h \uplus h_p \uplus h_s$ leads to a fork whose endpoint is a predecessor of $f(u)$ or $f(\bar{u})$. See Figure 1, which depicts two locations, each with three predecessors. Imagine

we are performing the comparison $\widehat{\#f(u)}^* \leq \widehat{\#f(\bar{u})}^*$ (presently equivalent to $\widehat{\#f(u)} \leq \widehat{\#f(\bar{u})}$). First, a segmented heap is added (allocated locations are marked with a white circle in the figure). The left half of the figure shows the addition of a collection of knives through these segments to make $f(\bar{u})$ fork; with the same segments, it must then be possible to add a collection of knives to make $f(u)$ fork. Indeed it is; if this is true for *all* such segmented heaps, then it must be that $\widehat{\#f(u)}^* \leq \widehat{\#f(\bar{u})}^*$. The segments from the segmented heap should really be seen as “potential forks.” When k or k' is nonzero, we consider additional forks on one or both sides to compensate for the offset.

This illustrates the principle behind the definition of the following property ($\mathcal{P}$): If there is a heap $h_{[k]}$ disjoint from $h \uplus h_p \uplus h_s$ made of isolated knives and k isolated forks so that $f(\bar{u})$ is fork in $h \uplus h_p \uplus h_s \uplus h_{[k]}$, then there is a heap $h_{[k']}$ disjoint from $h \uplus h_p \uplus h_s$ made of isolated knives and k' isolated forks so that $f(u)$ is fork in $h \uplus h_p \uplus h_s \uplus h_{[k']}$. Note that by the conditions satisfied by $h_{[k]}$ or by $h_{[k']}$, the number of predecessors of $f(\bar{u})$ [resp. $f(u)$] in $h_{[k]}$ [resp. $h_{[k']}$] is necessarily zero (there is no need to specify explicitly that we do not add predecessors). The number of forks in $h \uplus h_p \uplus h_s \uplus h_{[k]}$ whose endpoints are predecessors of $f(\bar{u})$ is bounded by $n + k$ and similarly, the number of forks in $h \uplus h_p \uplus h_s \uplus h_{[k']}$ whose endpoints are predecessors of $f(u)$ is bounded by $n + k'$. Note also that the number of predecessors of $f(u)$ is the same in h and in $h \uplus h_p \uplus h_s \uplus h_{[k']}$ and the number of predecessors of $f(\bar{u})$ is the same in h and in $h \uplus h_p \uplus h_s \uplus h_{[k]}$. So, $n + k \geq \bar{m}$ implies $n + k' \geq m$, i.e. $n \geq \bar{m} - k$ implies $n \geq m - k'$.

4.1.3 Checking the resulting heap

By checking step 2 for all $n \geq 0$, we get that for all $n \geq 0$, we have $n \geq \bar{m} - k$ implies $n \geq m - k'$, which entails that $\bar{m} - k \geq m - k'$, i.e. $m + k \leq \bar{m} + k'$, whenever $m - k' \geq 0$ and $\bar{m} - k \geq 0$. Universal quantification over n is simulated in a formula by using separating implication. When $m < k'$ or $\bar{m} < k$, we make a dedicated case analysis (see the proof of Theorem 10). Below, we present the technical developments.

4.2 Building blocks

Apart from forks, we introduce the notions of *collections* and *large forks* (instrumental in the proof of Lemma 7 below). A *large fork* is a sequence of distinct locations $l_1, \dots, l_5$ such that l_1, l_2 , and l_3 have no predecessors, $h(l_1) = h(l_2) = h(l_3) = l_4$, $\#l_4 = 3$ and $h(l_4) = l_5$. Location l_5 is called the *endpoint* of the large fork, and as with forks and knives, the large fork is called *isolated* iff $\#l_5 = 1$ and $l_5 \notin \text{dom}(h)$. A heap h is a *collection of knives and forks* $\stackrel{\text{def}}{=}$ there is no location in $\text{dom}(h)$ that does not belong to an isolated knife or to an isolated fork. Similarly, a heap h is a *collection of knives and large forks* $\stackrel{\text{def}}{=}$ there is no location in $\text{dom}(h)$ that does not belong to an isolated knife or isolated large fork.

Lemma 7. There are formulae ksfs , kslfs and ksfs_k ($k \geq 0$) in $\text{ISL2}(\ast)$ such that for every heap h , (1) $h \models \text{ksfs}$ iff h is a collection of knives and forks, (2) $h \models \text{kslfs}$ iff h is a collection of knives and large forks, (3) $h \models \text{ksfs}_k$ iff h is a collection of knives and forks with exactly k forks.

Here is the way we can use the formulae from Lemma 7.

Lemma 8. Let $k \geq 0$, h be a heap and f be a valuation such that $h \models_f \text{antiforky}(\bar{u}) \wedge \text{antiknify}(u)$, h has n isolated memory cells and $m = \widehat{\#f(u)}^*$. We have $h \models_f (\text{ksfs}_k \multimap \text{forky}(u))$ iff $n \geq m - k$.

The proof of Lemma 8 uses principles similar to what has been done in [7, 12] except that all formulae belong to $\text{ISL2}(\ast)$ and we propose a simplified version of the lemma and proof (note also our use of $\widehat{\#f(u)}^*$). Let anti be $\text{antiforky}(u) \wedge \text{antiknify}(u) \wedge \text{antiforky}(\bar{u}) \wedge \text{antiknify}(\bar{u})$ and let $\text{comp}(u, \bar{u}, k, k')$ be

$$[(\text{seg} \wedge \#u = 0 \wedge \# \bar{u} = 0) \multimap$$

$$(\text{anti} \Rightarrow ([\text{ksfs}_k \multimap \text{forky}(\bar{u})] \Rightarrow [\text{ksfs}_{k'} \multimap \text{forky}(u)]))]$$

Proposition 9. Let $k, k' \geq 0$, f and h such that $h \models_f \text{anti} \wedge \neg \exists u \text{isocell}(u)$, $\widehat{\#f(u)}^* - k' \geq 0$ and $\widehat{\#f(\bar{u})}^* - k \geq 0$. We have $h \models_f \text{comp}(u, \bar{u}, k, k')$ iff $\widehat{\#f(u)}^* + k \leq \widehat{\#f(\bar{u})}^* + k'$.

Here is the main result in this section using $\text{comp}(u, \bar{u}, k, k')$.

Theorem 10. For $k, k' \geq 0$, there is ϕ in $\text{ISL2}(\ast)$ (of linear size in $k + k'$) such that for all h, f , $(h \models_f \phi \text{ iff } \widehat{\#f(u)}^* + k \leq \widehat{\#f(\bar{u})}^* + k')$.

We denote such a formula ϕ as $\#u + k \leq \# \bar{u} + k'$ and, as usual, easily extend to $<, \geq, >, =$. The structure of the proof of Theorem 10 is similar to the structure of the proof of [7, Theorem 5.5], except that now all formulae are in $\text{ISL2}(\ast)$ instead of being defined in the less-constrained $\text{ISL}(\ast)$. Moreover, our case analysis is quite different and we also provide several simplifications, making our new proof even more valuable.

Proof. Below, we show that for $k, k' \geq 0$, there is a formula $\phi_{k,k'}$ in $\text{ISL2}(\ast)$ (of linear size in $k + k'$) such that for every heap h and valuation f , we have $h \models_f \phi_{k,k'}$ iff $\widehat{\#f(u)}^* + k \leq \widehat{\#f(\bar{u})}^* + k'$. Once we have such a $\phi_{k,k'}$, we easily define ϕ as:

$$(u \hookrightarrow u \wedge \bar{u} \hookrightarrow \bar{u} \wedge \phi_{k,k'}) \vee (\neg(u \hookrightarrow u) \wedge \bar{u} \hookrightarrow \bar{u} \wedge \phi_{k,k'+1}) \vee (u \hookrightarrow u \wedge \neg(\bar{u} \hookrightarrow \bar{u}) \wedge \phi_{k+1,k'}) \vee (\neg(u \hookrightarrow u) \wedge \neg(\bar{u} \hookrightarrow \bar{u}) \wedge \phi_{k,k'})$$

So it only remains to define the formulae $\phi_{k,k'}$ with $k, k' \geq 0$. By Proposition 9, we have the following property for any h, f :

($\star$) When h satisfies $\text{anti} \wedge \neg \exists u \text{isocell}(u)$, $\widehat{\#f(u)}^* - k' \geq 0$ and $\widehat{\#f(\bar{u})}^* - k \geq 0$, we have $h \models_f \text{comp}(u, \bar{u}, k, k')$ iff $\widehat{\#f(u)}^* + k \leq \widehat{\#f(\bar{u})}^* + k'$.

Even though the original heap h does not satisfy the formula $\text{anti} \wedge \neg \exists u \text{isocell}(u)$, it can be safely extended to satisfy this property without modifying the number of predecessors of $f(u)$ and $f(\bar{u})$.

Whenever $\widehat{\#f(u)}^* - k' \geq 0$ and $\widehat{\#f(\bar{u})}^* - k \geq 0$, we have the following equivalences:

- (1) $h \models_f (\#u = 0 \wedge \# \bar{u} = 0) \multimap (\text{anti} \wedge (\neg \exists u \text{isocell}(u)) \wedge \text{comp}(u, \bar{u}, k, k'))$.
- (2) There is h' disjoint from h such that $h' \models_f (\#u = 0 \wedge \# \bar{u} = 0)$, $h \uplus h' \models_f \text{anti} \wedge \neg \exists u \text{isocell}(u)$ and $h \uplus h' \models_f \text{comp}(u, \bar{u}, k, k')$.
- (3) There is h' disjoint from h such that $h' \models_f (\#u = 0 \wedge \# \bar{u} = 0)$ and $h \uplus h' \models_f \text{anti} \wedge \neg \exists u \text{isocell}(u)$ and $\widehat{\#f(u)}^* + k \leq \widehat{\#f(\bar{u})}^* + k'$ (in $h \uplus h'$) by ($\star$).
- (4) $\widehat{\#f(u)}^* + k \leq \widehat{\#f(\bar{u})}^* + k'$ in h .

Observe that $\widehat{\#f(u)}^*$ and $\widehat{\#f(\bar{u})}^*$ in h are equal to their values in $h \uplus h'$ since $h' \models_f (\#u = 0 \wedge \# \bar{u} = 0)$. Moreover, (4) implies (3) since it is always possible to extend a model satisfying $\text{anti} \wedge \neg \exists u \text{isocell}(u)$ while preserving $\widehat{\#f(u)}^*$ and $\widehat{\#f(\bar{u})}^*$.

W.l.o.g., we can assume that $k \times k' = 0$.

Case $k = 0$ and $k' \geq 0$. So $\widehat{\#f(\bar{u})}^* - k \geq 0$ and we need to make a case analysis depending on the satisfaction of $\widehat{\#f(u)}^* \geq k'$. Note

that if $\#f(u)^* < k'$, then obviously $\#f(u)^* \leq \#f(\bar{u})^* + k'$. So, we write $\phi_{k,k'}$ to denote the formula below:

$$(\#u < k') \vee ((\#u \geq k') \wedge ((\#u = 0 \wedge \#u = 0) \rightarrow \neg (\text{anti} \wedge (\neg \exists u \text{ isocell}(u)) \wedge \text{comp}(u, \bar{u}, 0, k'))))$$

Formulae of the form $\#u \geq k'$ (and variants) can be defined thanks to Lemma 5.

Case $k' = 0$ and $k \geq 0$. So $\#f(u)^* - k' \geq 0$ and we need to make a case analysis depending on the satisfaction of $\#f(u)^* \geq k$. Note that if $\#f(u)^* < k$, then $\#f(u)^* + k \leq \#f(\bar{u})^*$ cannot hold. So, we write $\phi_{k,k'}$ to denote the formula below:

$$(\#u \geq k) \wedge ((\#u = 0 \wedge \#u = 0) \rightarrow \neg (\text{anti} \wedge (\neg \exists u \text{ isocell}(u)) \wedge \text{comp}(u, \bar{u}, k, 0))) \quad \square$$

We can now find locations in a heap with a maximal number of predecessors, and we conclude this section with a definition useful in later constructions. Let $\text{maxdeg}(u) \stackrel{\text{def}}{=} \neg \exists \bar{u} \#u > \#u$. For all h and f , we have $h \models \text{maxdeg}(u)$ iff $\#f(u) = \max(\{l : l \in \mathbb{N}\})$.

5. Expressive Completeness for 1SL2(−*)

In order to express sentences in DSOL by sentences in 1SL2(−*), a hybrid valuation is encoded in the heap by building a disjoint *valuation heap* that takes care of pairs of locations (for interpretation of second-order variables) and that takes care of locations (for interpretation of first-order variables). In principle, this makes sense since every heap has a finite domain and therefore there is always an infinite set of locations that is not in its domain. This leaves enough room to encode a finite amount of information such as the interpretation of second-order variables when they are interpreted by finite sets. We can easily add to the original heap with the magic wand; this permits us to create and update the valuation heap. However, we then must always be able to distinguish between the original heap and the valuation heap.

The main idea to build such a valuation heap rests on the fact that a pair of locations (l, l') belongs to the interpretation of a second-order variable P_i whenever l and l' can be identified by special patterns involving l and l' that uniquely characterize the interpretation by P_i . Similarly, a location l is the interpretation of a first-order variable whenever l can be identified in the valuation heap thanks to some dedicated pattern around l .

Before explaining further the general principles, let us first provide more information about the above-mentioned patterns. An *entry of degree $d \geq 2$* is a sequence of distinct locations $l_1, \dots, l_d, l_{\text{ind}}, l$ such that $h(l_1) = \dots = h(l_d) = l_{\text{ind}}, \#l_{\text{ind}} = d, \#l_1 = \dots = \#l_d = 0$ and $h(l_{\text{ind}}) = l$. The location l is called the *element*, l_{ind} the *index* and the locations $l_1, \dots, l_d$ the *pins*. Entries generalize the notions of forks and large forks from Section 3 and are called *markers* in [7]. See an entry of degree 4 in the middle of Figure 2. So, the pair of locations (l, l') is identified as part of the interpretation of P_i when l and l' are elements of entries with very large degree. The above-mentioned special patterns are therefore entries, but we require that the degree of the respective entries for l and l' satisfy arithmetical constraints, which is possible thanks to Theorem 10, and which allows us to relate l with l' .

Then, the principle of the translation consists in building the valuation heap on demand (typically when a quantification appears) and to find special patterns involving entries with large degree whenever an atomic formula needs to be evaluated.

Apart from our essential restriction to 1SL2(−*) and therefore the need for encoding also first-order valuations, these principles have been introduced in [7] to translate DSOL formulae into

1SL(−*) formulae. However, because we are restricted to two first-order variables and because we also require that the separating conjunction is banished, we present below a different way to apply these principles so that we can show that 1SL2(−*) is expressively equivalent to DSOL.

This high-level description of the formula translation and of the encoding of some hybrid valuation in the heap hides many of the details, which can be found below. However, before explaining how we apply these principles within 1SL2(−*), let us emphasize the most obvious and difficult problems to be solved: (I) we must be able to distinguish the pairs of locations from distinct second-order variables, (II) we also need to encode first-order valuations, and (III) the main problem is certainly to access the original heap properly without interference from the valuation heap.

5.1 Left and right parentheses

We introduce variants of entries that are used as delimiters. A *left j -parenthesis of degree $d \geq 3$ with $j \geq 0$* is a sequence of distinct locations $l'_{j+1}, \dots, l'_1, l_1, \dots, l_d, l_{\text{ind}}$ such that (u) $h(l_1) = \dots = h(l_d) = l_{\text{ind}}; \#l_{\text{ind}} = d; \#l'_{j+1} = \dots = \#l_d = 0$, (v) $l_{\text{ind}} \notin \text{dom}(h); l'_{j+1} \rightarrow l'_j \rightarrow \dots \rightarrow l'_1 \rightarrow l_1; \#l'_j = \dots = \#l'_1 = \dots = \#l_1 = \#l_{\text{ind}} = 1$, and (w) $\#l_{\text{ind}} = 0$. The location l_{ind} is called the *index*. The heap at the left of Figure 2 presents a left j -parenthesis of degree 3. A *right j -parenthesis of degree $d \geq 3$ with $j \geq 0$* is a sequence of distinct locations $l'_{j+1}, \dots, l'_1, l'_{j+1}, \dots, l'_1, l_1, \dots, l_d, l_{\text{ind}}$ such that (u), (v), $\#l'_{j+1} = 0, \#l'_j = \dots = \#l'_1 = \dots = \#l_1 = \#l_{\text{ind}} = 1$, and $l'_{j+1} \rightarrow l'_j \rightarrow \dots \rightarrow l'_1 \rightarrow l_2$. The location l_{ind} is also called the *index*. The heap at the right of Figure 2 presents a right j -parenthesis of degree 5. A j -parenthesis can be understood as an entry, except that the index location is not allocated, and containing one or two paths of length $j + 1$, depending whether it is a left or a right parenthesis.

Lemma 11. For all $j \geq 0$, there is a formula $\text{lp}_j(u)$ [resp. $\text{rp}_j(u)$] in 1SL2(−*) such that for all heaps h and valuations f , we have $h \models \text{lp}_j(u)$ [resp. $h \models \text{rp}_j(u)$] iff $f(u)$ is the index of some left [resp. right] j -parenthesis in h .

In several places, we need to identify the indices from entries as well as their pins. Let $\text{eindex}(u)$ be $(\#u \geq 2) \wedge \text{allzpred}(u) \wedge \exists \bar{u} u \hookrightarrow \bar{u}$ that characterizes indices from entries. Let $\text{epin}(u)$ characterize pins from entries: $\text{epin}(u) \stackrel{\text{def}}{=} \exists \bar{u} u \hookrightarrow \bar{u} \wedge \text{eindex}(\bar{u})$. Similarly, we need to characterize the locations from parentheses. We already know how to identify their indices (Lemma 11). It remains to identify the other locations via the formula $\text{onp}_i(u)$ to characterize the locations on some i -parenthesis: roughly speaking, such locations are exactly those that can reach the index of some i -parenthesis in less than $i+2$ steps. Let $\text{onp}_i(u) \stackrel{\text{def}}{=} \bigvee_{j=0}^{i+2} \text{dist}_i(j, u)$ with $\text{dist}_i(0, u) \stackrel{\text{def}}{=} \text{lp}_i(u) \vee \text{rp}_i(u)$ and $\text{dist}_i(j+1, u) \stackrel{\text{def}}{=} \exists \bar{u} (u \hookrightarrow \bar{u}) \wedge \text{dist}_i(j, \bar{u})$ for all $j \geq 0$.

Lemma 12. Let h be a heap, f be a valuation and $i \geq 0$. Then, $h \models \text{onp}_i(u)$ iff $f(u)$ is on some i -parenthesis in h .

5.2 The role of parentheses

First, we introduce the interval $[1, K]$ ($K \in \mathbb{N} \setminus \{0\}$) assuming that for each $j \in [1, K]$, either P_j or u_j occurs in the DSOL formula to be translated (but not both of them). So, the developments below are relative to a finite set of first-order and second-order variables and this is concretized by $[1, K]$ (always possible since a formula has a finite number of variables).

Let us come back to parentheses and assume that X is a subset of $[0, K]$. In an X -well-formed heap h (see Definition 18 below), the parentheses play the following role. For each $j \in X$, we have

the index location lp_j from a distinguished left j -parenthesis and the index location rp_j from a distinguished right j -parenthesis. Moreover, let $d_j^l = \# \text{lp}_j$ and $d_j^r = \# \text{rp}_j$ (in $\mathfrak{h}$). When $j \in X$ is related to a first-order variable, we require that $d_j^r = d_j^l + 2$ and there is an entry of degree $d_j^l + 1$ such that its element is understood as the interpretation of the variable u_j (see Figure 2 with $d_j^r = 5$ and $d_j^l = 3$). That explains why the parentheses are viewed as delimiters. Similarly, let $\{(l_1, l'_1), \dots, (l_\beta, l'_\beta)\}$ be a finite set of pairs of locations, understood as the interpretation of a second-order variable P_j with $j \in X$. In $\mathfrak{h}$, there are 2β entries whose respective degrees are exactly $\{d_j^l + 3(i-1) + 1, d_j^l + 3(i-1) + 2 : i \in [1, \beta]\}$ with $d_j^r = d_j^l + 3\beta + 1$. A pair of entries of respective degrees $d_j^l + 3(i-1) + 1$ and $d_j^l + 3(i-1) + 2$ have exactly as elements l_i and l'_i respectively, which allows to encode the pair (l_i, l'_i) . All this underlying encoding makes sense only if the left and right parentheses as well as the entries whose degrees are related to their degrees are uniquely determined (see Condition (1) in Definition 14, below). For this reason, we introduce a left 0-parenthesis and a right 0-parenthesis with $d_0^r = d_0^l + 1$ (0 is not a variable index), the degree d_0^l is strictly greater than the degree of any location in the original heap, all degrees d_j^l with $j \neq 0$ are strictly greater than d_0^l and finally, the above-mentioned entries and parentheses are the only ones with their respective degrees. This guarantees that any entry from a pair of entries with successive degrees serving for the interpretation of a second-order variable, cannot serve twice for another pair or for another variable. Below, we provide the technical developments. A heap $\mathfrak{h}$ is *made of entries and parentheses only* $\stackrel{\text{def}}{=}$ every location in $\text{dom}(\mathfrak{h})$ belongs either to a left i -parenthesis for some $i \geq 0$, to a right i -parenthesis for some $i \geq 0$, or to an entry. Given a heap $\mathfrak{h}$ made of entries and parentheses only, we write $\text{ispect}(\mathfrak{h})$ to denote $\{\#l : l \text{ is the index of some entry or parenthesis in } \mathfrak{h}\}$ and call this the *index spectrum* of $\mathfrak{h}$.

Figure 2: Encoding $[u_j \mapsto l]$.

Let $\mathfrak{h}_B$ be a heap such that $\alpha = \max(\{\#l : l \in \mathbb{N}\})$. We have seen in Section 4 that it is possible to characterize the locations that witness this maximal value α thanks to $\text{maxdeg}(\mathfrak{u})$. A valuation heap $\mathfrak{h}_V$ for $\mathfrak{h}_B$ is made of entries and parentheses only whose degrees are greater than $\max(3, \alpha + 1)$. The heap $\mathfrak{h}_V$ satisfies the following simple conditions (more constraints will follow): $\min(\text{ispect}(\mathfrak{h}_V))$ is greater than $\max(3, \alpha + 1)$ and it is witnessed by the degree of some left 0-parenthesis; each degree in $\text{ispect}(\mathfrak{h}_V)$ is witnessed by exactly one entry or parenthesis. Formula $\text{imin}(\mathfrak{u})$ below is satisfied in $\mathfrak{h} = \mathfrak{h}_B \uplus \mathfrak{h}_V$ by a location l witnessing the minimal value in $\text{ispect}(\mathfrak{h}_V)$:

$$\text{imin}(\mathfrak{u}) \stackrel{\text{def}}{=} \text{lp}_0(\mathfrak{u}) \wedge (\forall \bar{u} ((\bar{u} \neq \mathfrak{u}) \wedge \text{lp}_0(\bar{u})) \Rightarrow \# \bar{u} < \# \mathfrak{u})$$

Note that thanks to Section 4, we know that it is possible to compare numbers of predecessors as expressed above. So, $\text{imin}(\mathfrak{u})$ holds when $\mathfrak{f}(\mathfrak{u})$ is the unique location that is the index of some left 0-parenthesis with greatest degree.

Lemma 13. Let $\mathfrak{f}$ be a valuation and $\mathfrak{h}$ be a heap. We have $\mathfrak{h} \models \text{imin}(\mathfrak{u})$ iff $\mathfrak{f}(\mathfrak{u})$ is an index of some left 0-parenthesis and there is no other location $l \neq \mathfrak{f}(\mathfrak{u})$ such that $\#l \geq \# \mathfrak{f}(\mathfrak{u})$ and l is the index of some left 0-parenthesis.

The proof is by an easy verification by using Lemma 11. Once a heap $\mathfrak{h}$ satisfies $\exists \mathfrak{u} \text{imin}(\mathfrak{u})$, the unique location l_0 such that $\mathfrak{h} \models_{[\mathfrak{u} \mapsto l_0]} \text{imin}(\mathfrak{u})$ (say with $\#l_0 = d_0$) plays the role of a delimiter between the original heap and the part of the heap that encodes the hybrid valuation. We have seen that an index spectrum is defined for heaps made of entries and parentheses only. This is fine, but below we adapt the definition to heaps $\mathfrak{h}$ satisfying $\exists \mathfrak{u} \text{imin}(\mathfrak{u})$. Let us define the set $\text{spect}(\mathfrak{h})$ as

$$\{\#l : l \text{ is an index of some entry or parenthesis in } \mathfrak{h}\} \cap [d_0, +\infty[$$

The set $\text{spect}(\mathfrak{h})$ is called the *spectrum* of $\mathfrak{h}$. This illustrates how l_0 and $\#l_0 = d_0$ play the role of separator between the original heap and the valuation heap.

The subheap encoding the valuation is made of parentheses and entries and we shall need to identify the indices of such patterns. The formula $\text{lindex}(\mathfrak{u})$ defined below does the job:

$$(\exists \bar{u} \text{imin}(\bar{u}) \wedge \# \bar{u} \leq \# \mathfrak{u}) \wedge$$

$$\left(\left(\bigvee_{i \in [0, K]} (\text{lp}_i(\mathfrak{u}) \vee \text{rp}_i(\mathfrak{u})) \right) \vee \text{eindex}(\mathfrak{u}) \right)$$

($\mathfrak{u}$ is interpreted as a *large index*). Entries and parentheses with large indices are also called large entries and parentheses, respectively. It is easy to define a large index that is also the index of a left [resp. right] parenthesis. Let $\text{llp}_i(\mathfrak{u}) \stackrel{\text{def}}{=} \text{lindex}(\mathfrak{u}) \wedge \text{lp}_i(\mathfrak{u})$ and $\text{lrp}_i(\mathfrak{u}) \stackrel{\text{def}}{=} \text{lindex}(\mathfrak{u}) \wedge \text{rp}_i(\mathfrak{u})$ (see Lemma 11). The large index with a maximal degree can be also characterized as follows:

$$\text{maxlindex}(\mathfrak{u}) \stackrel{\text{def}}{=} (\forall \bar{u} \text{lindex}(\bar{u}) \Rightarrow (\# \bar{u} \leq \# \mathfrak{u})) \wedge \text{lindex}(\mathfrak{u})$$

Below, we state how the parentheses are organized.

Definition 14. Let $X = \{i_0, \dots, i_s\} \subseteq [0, K]$ with $0 = i_0 < i_1 < \dots < i_s$. A heap $\mathfrak{h}$ is *X-almost-well-formed* $\stackrel{\text{def}}{=}$

1. For every $j \in [0, s]$, there is a unique location l_j^l [resp. l_j^r] such that $\mathfrak{h} \models_{[\mathfrak{u} \mapsto l_j^l]} \text{llp}_{i_j}(\mathfrak{u})$ [resp. $\mathfrak{h} \models_{[\mathfrak{u} \mapsto l_j^r]} \text{lrp}_{i_j}(\mathfrak{u})$].
2. For every $j \in [0, s]$, $\#l_j^l < \#l_j^r$, and $\#l_j^r = \#l_{j-1}^r + 1$.
3. For every $j \in [1, s]$, we have $\#l_j^l = \#l_{j-1}^l + 1$.
4. $\mathfrak{h} \models_{[\mathfrak{u} \mapsto l_j^r]} \text{maxlindex}(\mathfrak{u})$.
5. For every $j \in [1, s]$, if i_j is the index of a first-order variable, then $\#l_j^l = \#l_j^r - 2$ (see Figure 2).
6. For every $j \in ([1, K] \setminus X)$, there is no location l such that $\mathfrak{h} \models_{[\mathfrak{u} \mapsto l]} \text{llp}_j(\mathfrak{u}) \vee \text{lrp}_j(\mathfrak{u})$.

The degrees are organized as follows and they all belong to the spectrum of $\mathfrak{h}$ (below we let $d_j^l = \#l_j^l$ and $d_j^r = \#l_j^r$).

$$\begin{array}{ccccccc} \models \text{imin}(\mathfrak{u}) & & & & & & \models \text{maxlindex}(\mathfrak{u}) \\ d_0^l < d_0^r < d_1^l < d_1^r < d_2^l < d_2^r < \dots < d_s^l < d_s^r \\ \parallel & \parallel & & \parallel & & \parallel & \\ d_0^l + 1 & d_0^r + 1 & & d_1^r + 1 & & d_{s-1}^r + 1 & \end{array}$$

Lemma 15. There is awfh_X in $\text{ISL2}(\ast)$ such that $\mathfrak{h} \models \text{awfh}_X$ iff $\mathfrak{h}$ is *X-almost-well-formed*.

Let $\mathfrak{h}$ be an *X-almost-well-formed* heap for some $\{0\} \subseteq X \subseteq [0, K]$ and $i \in X$. We write $\text{vind}_i(\mathfrak{u})$ to denote

$$\text{lindex}(\mathfrak{u}) \wedge \text{eindex}(\mathfrak{u}) \wedge (\exists \bar{u} \text{llp}_i(\bar{u}) \wedge \# \bar{u} < \# \mathfrak{u}) \wedge (\exists \bar{u} \text{lrp}_i(\bar{u}) \wedge \# \bar{u} > \# \mathfrak{u})$$

It characterizes indices whose degree is strictly between the degree of some large left i -parenthesis and the degree of some large right i -parenthesis. We write $\text{degrees}(i, h)$ to denote the set: $\{\#l : h \models_{[u \rightarrow l]} \text{vind}_i(u), l \in \mathbb{N}\}$.

Lemma 16. Let h be such that $h \models \exists u \text{imin}(u)$ and $i \geq 0$ be such that there are unique locations lp and rp with $h \models_{[u \rightarrow lp]} \text{llp}_i(u)$ and $h \models_{[u \rightarrow rp]} \text{lrp}_i(u)$. For every $l \in \mathbb{N}$, we have $h \models_{[u \rightarrow l]} \text{vind}_i(u)$ iff l is the index of some entry and $\#lp < \#l < \#rp$.

The translation of $P_i(u_j, u_k)$ is therefore: $\exists u (\text{on}_j(u) \wedge \exists \bar{u} (\bar{u} \hookrightarrow u \wedge \text{vind}_i(\bar{u}) \wedge \exists u (\#u = \# \bar{u} + 1 \wedge \text{vind}_i(u) \wedge \exists \bar{u} (\bar{u} \hookrightarrow u \wedge \text{on}_k(\bar{u}))))$, where $\text{on}_j(u) \stackrel{\text{def}}{=} \exists \bar{u} (\bar{u} \hookrightarrow u) \wedge \text{vind}_j(\bar{u})$. Formula $\text{on}_j(u)$ holds true when u is interpreted as the element of the unique entry attached to u_j . These definitions take advantage of the fact that there are unique large left and right parentheses for each variable index. Figure 3 illustrates the constraints when $j < i < k$. From left to right, it represents explicitly a left j -parenthesis, then a right j -parenthesis, then a left i -parenthesis, a right i -parenthesis and a left k -parenthesis, followed finally by a right k -parenthesis. Other entries and parentheses are present in the figure, but they are represented by dots in order to focus on the memory cells relevant to evaluate the formula obtained by translation of $P_i(u_j, u_k)$. The degrees of parentheses and entries increase from left to right.

5.3 Taking care of valuations

Now that we have a way of identifying that part of the heap that encodes our valuation, we turn our attention to encoding the valuation itself. Below, we introduce a condition for a subheap to be “glued” to an existing valuation. We distinguish three cases. A *local 0-valuation* is a heap made of a left 0-parenthesis of degree d and a right 0-parenthesis of degree $d + 1$ only, for some $d \geq 3$. Let $i \in [1, K]$ be the index of some first-order variable. A *local i -valuation* is a heap made of a left i -parenthesis of degree d , an entry of degree $d + 1$ and a right i -parenthesis of degree $d + 2$ only, for some $d \geq 3$. Let $i \in [1, K]$ be the index of some second-order variable. A *local i -valuation* is a heap h such that

1. every location l in $\text{dom}(h)$ belongs either to a left i -parenthesis, to a right i -parenthesis, or to an entry,
2. h contains a unique left [resp. right] i -parenthesis,
3. $\min(\text{ispect}(h))$ is the degree of some left i -parenthesis,
4. $\max(\text{ispect}(h))$ is the degree of some right i -parenthesis,
5. $\text{ispect}(h)$ is of the form below for some $\alpha \geq 3, \beta \geq 0, \{\alpha\} \cup \{\alpha + 3(i - 1) + 1, \alpha + 3(i - 1) + 2 : i \in [1, \beta]\} \cup \{\alpha + 3\beta + 1\}$ (when $\beta = 0$, $\text{ispect}(h)$ is equal to $\{\alpha, \alpha + 1\}$),
6. there are no two distinct indices with the same degree.

Since local i -valuations are typically heaps that are added to the current heap to encode the interpretation of a variable, it is essential to be able to characterize them by $\text{ISL2}(\ast)$ formulae. This is the purpose of the result below.

Lemma 17. Let $i \in [0, K]$. There is $\text{localval}_i(u)$ in $\text{ISL2}(\ast)$ such that $h \models \text{localval}_i(u)$ iff h is a local i -valuation and $f(u)$ is the index of its left i -parenthesis.

The definition for X -almost-well-formed heaps mainly takes care of parentheses. In Definition 18, constraints on the degrees of large indices are specified.

Definition 18. Let $X = \{i_0, \dots, i_s\} \subseteq [0, K]$ with $0 = i_0 < i_1 < \dots < i_s$. A heap h is *X -well-formed* $\stackrel{\text{def}}{=}$

1. h is X -almost-well-formed,
2. for every $j \in [1, s]$, if i_j is the index of a first-order variable, then $\text{degrees}(i_j, h)$ is a singleton,

3. for every $j \in [1, s]$, if i_j is the index of a second-order variable, $\text{degrees}(i_j, h)$ is the set below for some $\alpha_j \geq 3, \beta_j \geq 0$: $\{\alpha_j + 3(i - 1) + 1, \alpha_j + 3(i - 1) + 2 : i \in [1, \beta_j]\}$,
4. for every location l such that $h \models_{[u \rightarrow l]} \text{lindex}(u)$, there is no $l' \neq l$ such that $h \models_{[u \rightarrow l']} \text{lindex}(u)$ and $\#l = \#l'$.

When h is X -well-formed, we write $h = h_B \uplus h_V$ such that $\text{dom}(h_V)$ is made of entries and parentheses of degree $d \geq \#l_0$ for some $l_0 \in \mathbb{N}$ such that $h \models_{[u \rightarrow l_0]} \text{imin}(u)$. By Definition 18, we have $\text{spect}(h) = \text{ispect}(h_V)$ and clearly the decomposition is unique since l_0 is unique.

Again, well-formed heaps can be characterized by formulae in $\text{ISL2}(\ast)$ whose size is linear in K .

Lemma 19. Given $\{0\} \subseteq X \subseteq [0, K]$, there is a formula wfh_X in $\text{ISL2}(\ast)$ such that $h \models \text{wfh}_X$ iff h is X -well-formed.

Let us define a valuation from a valuation heap.

Definition 20. Let h be an X -well-formed heap for some $\{0\} \subseteq X \subseteq [0, K]$. For every second-order $i \in X$, we define $\mathfrak{V}_h(P_i) \stackrel{\text{def}}{=} \{(h_V(l), h_V(l')) : \#l = \#l' + 1, \#l, \#l' \in \text{degrees}(i, h), l, l' \text{ are index locations}\}$ and for every first-order $i \in X$, $\mathfrak{V}_h(u_i) \stackrel{\text{def}}{=} h_V(l)$ where l is the unique index location such that $\#l \in \text{degrees}(i, h)$. We say that $\mathfrak{V}_h$ is the valuation *extracted* from h .

Below, we present an essential result stating how heaps can be composed when a new variable needs to be interpreted. The formulae involved to compose the X -well-formed heap h and the local i -valuation heap h' are directly used in the translation of quantified formulae (see Section 5.4). Lemma 21 is only used in the proof of Lemma 22 but it is stated here in order to emphasize how the heaps can be safely composed.

Lemma 21 (Composition). Let f be a valuation, h be an X -well-formed heap with $\{0\} \subseteq X \subseteq [0, K]$, $i \in [1, K] \setminus X$ with $i > \max(X)$, and h' be a disjoint heap such that:

1. $h \models_f \text{imin}(u) \wedge \text{isoloc}(\bar{u})$,
2. $h' \models_f \text{localval}_i(\bar{u})$,
3. $h \uplus h' \models_f \text{wfh}_{X \cup \{i\}} \wedge \text{imin}(u) \wedge \text{llp}_i(\bar{u})$.

Then, $\text{spect}(h \uplus h') = \text{spect}(h) \uplus \text{ispect}(h')$.

The proof of Lemma 21 is quite combinatorial and this is the place where we check that the original heap cannot be confused with the valuation heap (and the other way around). It is important to guarantee, as the proof does, that adding a new part of the valuation does not destroy what has been built so far.

5.4 A reduction from DSOL into $\text{ISL2}(\ast)$

Below, we define a translation from a sentence ϕ in DSOL into a sentence in $\text{ISL2}(\ast)$ that uses only logspace. Without any loss of generality, we assume that (1) two occurrences of quantified variables in ϕ have distinct variable indices (e.g., P_4 and u_4 cannot both occur in ϕ and “ $\forall u_4$ ” cannot occur more than once) and (2) if $\exists u_i \psi_1$ is a subformula of $\exists u_j \psi_2$, then $i > j$ and this holds for any combination of first-order/second-order variables. We may rename variables so that these conditions are satisfied. We assume that the variable indices are among $[1, K]$.

The translation of the formula ϕ , written $T(\phi)$, first applies a top-level translation $t_{\text{top}}(\cdot)$ which takes care of initializing the valuation heap; then, a recursive map $t(\cdot)$ is applied. So, $T(\phi) \stackrel{\text{def}}{=} t_{\text{top}}(\phi)$ where $t_{\text{top}}(\phi)$ is defined as follows:

$$t_{\text{top}}(\phi) \stackrel{\text{def}}{=} \exists u \text{isoloc}(u) \wedge (\text{localval}_0(u) \multimap (\text{wfh}_{\{0\}} \wedge$$

$$\text{imin}(u) \wedge (\forall \bar{u} ((u \neq \bar{u}) \wedge \neg \text{lrp}_0(\bar{u})) \Rightarrow (\# \bar{u} < \#u) \wedge t(\{0\}, \phi)))$$

Figure 3: How the translation of $P_i(u_j, u_k)$ works ($j < i < k$): $(l, l') \in \mathfrak{V}_h(P_i)$.

The first step of the translation consists in adding 0-parentheses so that the heap that evaluates $t(\{0\}, \phi)$ is $\{0\}$ -well-formed. The recursive map $t(\cdot)$ is inductively defined as follows ($X \subseteq [0, K]$, ψ subformula of ϕ):

- $t(X, \cdot)$ is homomorphic for Boolean connectives,
- $t(X, u_i = u_j) \stackrel{\text{def}}{=} \exists u \text{ on}_i(u) \wedge \text{on}_j(u)$,
- $t(X, u_i \hookrightarrow u_j) \stackrel{\text{def}}{=} \exists u \exists \bar{u} (\text{on}_i(u) \wedge \text{on}_j(\bar{u}) \wedge u \hookrightarrow \bar{u})$,
- $t(X, P_i(u_j, u_k)) \stackrel{\text{def}}{=} \exists u (\text{on}_j(u) \wedge \exists \bar{u} (\bar{u} \hookrightarrow u \wedge \text{vind}_i(\bar{u}) \wedge \exists u' (\#u' + 1 \wedge \text{vind}_i(u') \wedge \exists \bar{u}' (u' \hookrightarrow \bar{u}' \wedge \text{on}_k(\bar{u}')))))$,
- For the quantifier $\exists u_i$, we choose l and l' such that l is the index of the left 0-parenthesis and l' is an isolated location. We construct a new heap that is a local i -valuation while enforcing that the index of the left 0-parenthesis is preserved and l' becomes the index of the unique large left i -parenthesis (see Lemma 21).
 $t(X, \exists u_i \psi) \stackrel{\text{def}}{=} \exists u \exists \bar{u} ((\text{imin}(u) \wedge \text{isoloc}(\bar{u})) \wedge (\text{localval}_i(\bar{u}) \xrightarrow{*} (\text{wfh}_{X \cup \{i\}} \wedge \text{imin}(u) \wedge \text{llp}_i(\bar{u}) \wedge t(X \cup \{i\}, \psi))))$

- The translation with second-order variables is analogous (the formula $\text{localval}_i(\bar{u})$ is actually defined differently, see the proof of Lemma 17): $t(X, \exists P_i \psi) \stackrel{\text{def}}{=} \exists u \exists \bar{u} ((\text{imin}(u) \wedge \text{isoloc}(\bar{u})) \wedge (\text{localval}_i(\bar{u}) \xrightarrow{*} (\text{wfh}_{X \cup \{i\}} \wedge \text{imin}(u) \wedge \text{llp}_i(\bar{u}) \wedge t(X \cup \{i\}, \psi))))$.

Every subformula $t(X, \psi)$ has no free variable from $\text{fr}(\psi) \subseteq X$ where $\text{fr}(\psi)$ denotes the set of variable indices in ψ from either first-order or second-order *free* variables. As noted by one anonymous referee, a standard trick is to convert first-order variables into second-order ones so that the proof has only to deal with one type of variable. Herein, we do not quite eliminate first-order variables but we provide a uniform treatment for first-order quantifications and second-order quantifications, which essentially amounts to dealing with a single type of encoding.

Below, we state the correctness lemma that allows us to get Theorem 23.

Lemma 22. Let ϕ be a DSOL sentence of the above form, ψ be one of its subformulae and $(\text{fr}(\psi) \cup \{0\}) \subseteq X \subseteq [0, K]$. Let $h = h_B \uplus h_V$ be a X -well-formed heap and $\mathfrak{V}_h$ be the valuation extracted from h . Then, $h_B \models_{\mathfrak{V}_h} \psi$ iff $h \models t(X, \psi)$.

Here is the main result of the paper.

Theorem 23. For every sentence ϕ in DSOL, for every heap h , we have $h \models \phi$ iff $h \models T(\phi)$, so WSOL and $\text{ISL2}(\ast)$ have the same expressive power.

Observe that $T(\phi)$ can be computed in logspace (to do this, one must check the size of all the formulae built in the proofs). So, the restriction to two variables does not reduce the expressive power, unlike restrictions in [14], for instance.

We get the ultimate undecidability result below (no separating conjunction, two quantified variables, one record field).

Corollary 24. Satisfiability problem for $\text{ISL2}(\ast)$ is undecidable.

The absence of program variables in $\text{ISL2}(\ast)$ makes the proof of Corollary 24 even more difficult to design, which is perfect to obtain the sharpest undecidability result. An expressivity result with program variables is briefly presented in Section 5.5.

It is also possible to establish the following consequences.

Corollary 25. (I) Let ϕ be a sentence in ISL . There is an equivalent sentence in $\text{ISL2}(\ast)$ of polynomial size in the size of ϕ . (II) $\text{ISL2}(\ast)$ is strictly more expressive than $\text{ISL}(\ast)$.

Corollary 25(II) follows from Theorem 23 and $\text{ISL}(\ast)$ is strictly less expressive than MSO [1, Corollary 5.3].

Our main results are Theorem 23 and Corollary 24, significantly improving previously known results (see the figure in Section 1). As far as the translation into $\text{ISL2}(\ast)$ is concerned (see the current section but it uses in essential ways the formulae of Section 4), the lack of variables is partially compensated by the introduction of the parentheses in order to constrain sufficiently the valuation heap. More importantly, we have shown that this is a viable solution in $\text{ISL2}(\ast)$ despite only having two variables (see the proof of Lemma 21). This was not at all clear at the outset, and of course, in view of the complexity of the final proof, this led to lengthy arguments to show correctness of the whole enterprise.

5.5 Adding an unbounded number of program variables

We consider ISL with program variables, which is a strict extension of ISL and therefore undecidability for $\text{ISL2}(\ast)$ is still valid in the presence of program variables. Let $\text{PVAR} = \{x_1, x_2, \dots\}$ be a countably infinite set of *program variables*. A *memory state* is a pair (s, h) such that $s : \text{PVAR} \rightarrow \mathbb{N}$ and h is a heap. Formulae of *ISL with program variables* are built from expressions of the form $e ::= x \mid u$ where $x \in \text{PVAR}$ and $u \in \text{FVAR}$, and atomic formulae of the form $\pi ::= e = e' \mid e \hookrightarrow e'$. Formulae are defined by $\phi ::= \pi \mid \phi \wedge \psi \mid \neg \phi \mid \phi \ast \psi \mid \phi \dashv \ast \psi \mid \exists u \phi$, where $u \in \text{FVAR}$. A *valuation* is a map $f : \text{FVAR} \rightarrow \mathbb{N}$. The satisfaction relation $\models$ is extended as follows:

- $(s, h) \models_f e = e'$ iff $[e] = [e']$, with $[x] \stackrel{\text{def}}{=} s(x)$, $[u] \stackrel{\text{def}}{=} f(u)$. Obviously, program variables can be understood as free quantified variables interpreted rigidly.
- $(s, h) \models_f e \hookrightarrow e'$ iff $[e] \in \text{dom}(h)$ and $h([e]) = [e']$.

The satisfiability problem takes as input a sentence from 1SL with program variables, in which the only free variables are program variables. The version of DSOL with program variables is defined similarly when models are memory states.

Theorem 26. There is a translation T such that for every sentence ϕ in DSOL with program variables, the sentence $T(\phi)$ in 1SL2($\ast$) with program variables (of polynomial-size in the size of ϕ) is such that for all (s, h) , we have $(s, h) \models \phi$ iff $(s, h) \models T(\phi)$.

Actually, the translation T is defined as a variant of the one in Section 5.4 that takes into account program variables. The variant is pretty natural since program variables do not require any encoding. For instance, the map $t(\cdot)$ is extended as follows (more details in the proof of Theorem 26): $t(X, u_i = x) \stackrel{\text{def}}{=} \exists u (\text{on}_i(u) \wedge u = x)$ and $t(X, u_i \hookrightarrow x) \stackrel{\text{def}}{=} \exists u (\text{on}_i(u) \wedge u \hookrightarrow x)$.

Extension with k record fields So far, memory cells have a unique record field but it is possible to extend our results to $k > 1$ record fields, along the lines of [7, Section 7]. Let $k\text{SL2}(\ast)$ be the logic in which heaps are partial functions $h : \mathbb{N} \rightarrow \mathbb{N}^k$ with finite domain and atomic formulae include $u_j \xrightarrow{i} \bar{u}_{j'}$ ($i \in [1, k]$); $k\text{WSOL}$ and $k\text{DSOL}$ are defined similarly. One can show that every sentence in $k\text{DSOL}$ has an equivalent sentence in $k\text{SL2}(\ast)$. To do so, we need to adapt the definitions from Sections 3 and 4 so that memory cells involved in the valuation heap are relevant only with respect to $\xrightarrow{1}$ (and comparing numbers of predecessors is performed only with respect to $\xrightarrow{1}$). Details are tedious because many notions need to be redefined relatively to $\xrightarrow{1}$ but the encoding of valuations is based on the same general principles as for 1SL2($\ast$).

6. Conclusion

We have shown that 1SL2($\ast$) is as expressive as weak second-order logic on concrete heaps (Theorem 23). As a consequence, the satisfiability problem for 1SL2($\ast$) is undecidable (Corollary 24) and we have identified the undecidable core of separation logic, apart from illustrating further the power of separating implication when interpreted on concrete heaps. We only use two variables, and our results also exclude separating conjunction, which is quite remarkable in view of the restricted number of variables. As far as the proofs are concerned, we used first principles from [7] but we had to provide non-trivial adaptations to fit the restricted fragment 1SL2($\ast$). However, this illustrates the robustness of those principles since they could be applied by using the proof techniques developed in the present paper. Other semantical variants are possible and will be the subject of future work.

Acknowledgments

We would like to thank the anonymous referees for their suggestions and remarks to improve the quality of the paper.

References

- [1] T. Antonopoulos and A. Dawar. Separating graph logic from MSO. In *FOSSACS'09*, volume 5504 of *LNCS*, pages 63–77. Springer, 2009.
- [2] T. Antonopoulos, N. Gorogiannis, C. Haase, M. Kanovich, and J. Ouaknine. Foundations for decision problems in separation logic with general inductive predicates. In *FOSSACS'14*, volume 8412 of *LNCS*, pages 411–425. Springer, 2014.
- [3] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.
- [4] M. Bozga, R. Iosif, and S. Perarnau. Quantitative separation logic and programs with lists. *Journal of Automated Reasoning*, 45(2):131–156, 2010.
- [5] D. Bresolin, D. Della Monica, V. Goranko, A. Montanari, and G. Sciavicco. Metric propositional neighborhood logics: Expressiveness, decidability, and undecidability. In *ECAI'10*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 695–700. IOS Press, 2010.
- [6] R. Brochenin. *Separation Logic: Expressiveness, Complexity, Temporal Extension*. PhD thesis, LSV, ENS Cachan, September 2013.
- [7] R. Brochenin, S. Demri, and E. Lozes. On the Almighty Wand. *Information and Computation*, 211:106–137, 2012.
- [8] J. Brotherston and M. Kanovich. Undecidability of propositional separation logic and its neighbours. In *LICS'10*, pages 130–139. IEEE, 2010.
- [9] C. Calcagno, P. O'Hearn, and H. Yang. Computability and complexity results for a spatial assertion language for data structures. In *FSTTCS'01*, volume 2245 of *LNCS*, pages 108–119. Springer, 2001.
- [10] B. Cook, C. Haase, J. Ouaknine, M. Parkinson, and J. Worrell. Tractable reasoning in a fragment of separation logic. In *CONCUR'11*, volume 6901 of *LNCS*, pages 235–249. Springer, 2011.
- [11] A. Dawar, P. Gardner, and G. Ghelli. Expressiveness and complexity of graph logic. *Information and Computation*, 205(3):263–310, 2007.
- [12] S. Demri and M. Deters. Two-variable separation logic and its inner circle, September 2013. Under submission.
- [13] S. Demri, D. Galmiche, D. Larchey-Wendling, and D. Mery. Separation logic with one quantified variable. In *CSR'14*, volume 8476 of *LNCS*, pages 125–138. Springer, 2014.
- [14] K. Etessami, M. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logics. In *LICS'97*, pages 228–235. IEEE, 1997.
- [15] D. Gabbay. Expressive functional completeness in tense logic. In *Aspects of Philosophical Logic*, pages 91–117. Reidel, 1981.
- [16] D. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic - Mathematical Foundations and Computational Aspects, Vol. 1*. OUP, 1994.
- [17] D. Galmiche and D. Méry. Tableaux and resource graphs for separation logic. *Journal of Logic and Computation*, 20(1):189–231, 2010.
- [18] E. Grädel, M. Otto, and E. Rosen. Undecidability results on two-variable logics. *Arch. Mathematical Logic*, 38(4–5):313–354, 1999.
- [19] C. Haase, S. Ishtiaq, J. Ouaknine, and M. Parkinson. SeLogger: A tool for graph-based reasoning in separation logic. In *CAV'13*, volume 8044 of *LNCS*, pages 790–795. Springer, 2013.
- [20] Z. Hou, R. Clouston, R. Goré, and A. Tiu. Proof search for propositional abstract separation logics via labelled sequents. In *POPL'14*, pages 465–476. ACM, 2014.
- [21] N. Immerman, A. Rabinovich, T. Reps, M. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *CSL'04*, volume 3210 of *LNCS*, pages 160–174. Springer, 2004.
- [22] R. Iosif, A. Rogalewicz, and J. Simacek. The tree width of separation logic with recursive definitions. In *CADE'13*, volume 7898 of *LNCS*, pages 21–38. Springer, 2013.
- [23] V. Kuncak and M. Rinard. On spatial conjunction as second-order logic. Technical Report MIT-CSAIL-TR-2004-067, MIT CSAIL, October 2004.
- [24] D. Larchey-Wendling and D. Galmiche. The undecidability of Boolean BI through phase semantics. In *LICS'10*, pages 140–149. IEEE, 2010.
- [25] W. Lee and S. Park. A proof system for separation logic with magic wand. In *POPL'14*, pages 477–490. ACM, 2014.
- [26] C. Lutz, U. Sattler, and F. Wolter. Modal logic and the two-variable fragment. In *CSL'01*, volume 2142 of *LNCS*, pages 247–261. Springer, 2001.
- [27] R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using SMT. In *CAV'13*, volume 8044 of *LNCS*, pages 773–789. Springer, 2013.
- [28] J. Reynolds. Separation logic: a logic for shared mutable data structures. In *LICS'02*, pages 55–74. IEEE, 2002.
- [29] B. Trakhtenbrot. Impossibility of an algorithm for the decision problem in finite classes. *AMS Translations, Series 2*, 23:1–5, 1963.
- [30] P. Weis. *Expressiveness and Succinctness of First-Order Logic on Finite Words*. PhD thesis, University of Massachusetts, 2011.