The Logic of Separation Logic: Models and Proofs

Frank S. de Boer[1,2], Hans-Dieter A. Hiep[1,2(✉)], and Stijn de Gouw[3]

[1] Centrum Wiskunde and Informatica (CWI), Amsterdam, The Netherlands
hdh@cwi.nl
[2] Leiden Institute of Advanced Computer Sciences (LIACS), Leiden, The Netherlands
[3] Open University (OU), Heerlen, The Netherlands

Abstract. The standard semantics of separation logic is restricted to finite heaps. This restriction already gives rise to a logic which does not satisfy compactness, hence it does not allow for an effective, sound and complete axiomatization. In this paper we therefore study both the general model theory and proof theory of the separation logic of finite and infinite heaps over arbitrary (first-order) models. We show that we can express in the resulting logic finiteness of the models and the existence of both countably infinite and uncountable models. We further show that a sound and complete sequent calculus still can be obtained by restricting the second-order quantification over heaps to first-order definable heaps.

1 Introduction

Separation logic [Rey02], in the sequel also referred to by SL, extends first-order logic with the separating connectives of conjunction and implication for reasoning about programs which feature the dynamic allocation of variables that are stored at locations of that part of the memory called the 'heap'. The *separating conjunction* allows to specify properties of a partition of the heap into two disjoint sub-heaps. The *separating implication* (also called 'the magic wand') allows to express properties of disjoint extensions of the heap. Both separating connectives involve a second-order quantification over heaps (which are represented by binary relations).

In this paper we study both the model theory and the proof theory of SL. The standard model of SL (as introduced in [Rey02]) extends the standard model of arithmetic with the so-called 'points-to' relation which provides a formalization of the heap in terms of the *graph* of a *finitely-based partial function*. This function assigns to each location of the heap its stored value, or is undefined if the location is not allocated. In the standard semantics of SL (here also called *weak* SL), the domains of heaps are finite, that is, only finitely many locations are allocated. Reasoning about finite heaps however requires an *infinitary* logic because the logic of finite heaps, and that of finite model theory in general, does not satisfy the compactness property: it is straightforward to express for each natural number that the domain of the heap contains at least that number of

R. Ramanayake and J. Urban (Eds.): TABLEAUX 2023, LNAI 14278, pp. 407–426, 2023.
https://doi.org/10.1007/978-3-031-43513-3_22

elements. It follows that every finite subset of this infinite set of sentences is satisfiable, but clearly no finite heap satisfies the entire set.

To study the general model and proof theory of *full* SL[1] we (1) extend its semantics to arbitrary first-order models and (2) generalize the notion of a heap to a partial function on the underlying domain of the given (first-order) model: no restrictions are imposed on the cardinality of the domain of heap, in contrast to weak SL which restricts to finite heaps. Our main model-theoretic results are that in this general setting we can express: (1) finiteness of models, (2) well-foundedness of the points-to relation, and (3) existence of countably infinite and uncountable models. As a consequence we have that full SL satisfies neither compactness nor the downward and upward Löwenheim-Skolem theorems (see [CK13]). Non-compactness implies that there does not exist an effective, sound and complete proof theory for SL. In fact, we will show that the well-foundedness of the points-to relation can already be expressed in full SL using only separating conjunction. Consequently, full SL without separating implication is already non-compact. For full SL without separating implication but in which separating conjunction only occurs positively, the fragment which we call separation logic light (SLL), we do have compactness, but its semantic consequence relation is not compact and therefore also does not allow for an effective, sound and complete proof theory.

The question thus arises whether there exists an *alternative* interpretation of SL that does allow for an effective, sound and complete proof theory. Clearly, the main complexity of SL stems from the (second-order) quantification over heaps (or sub-heaps, as in the case of the separating conjunction). For second-order logic a sound and complete axiomatization can be obtained by generalizing its semantics by means of so-called *general models*. Such models extend first-order models with a set of possible interpretations of the second-order variables. For example, instead of interpreting a monadic predicate over *all* possible subsets of the given first-order domain, a general model restricts its interpretation to a given set of such subsets. This generalization of the semantics of second-order logic allows for a sound and complete axiomatization by restricting to so-called Henkin models. A Henkin model is a general model for second-order logic which additionally satisfies the comprehension axiom

$$\exists R \forall x_1, \ldots, x_n (R(x_1, \ldots, x_n) \leftrightarrow \phi(x_1, \ldots, x_n))$$

for any second-order formula $\phi(x_1, \ldots, x_n)$ which does not contain the n-ary relation symbol R. In the *arithmetic* comprehension axiom $\phi(x_1, \ldots, x_n)$ is first-order.

Generalizing the semantics of SL accordingly in terms of a given set of possible heaps, which does not necessarily contain *all* heaps, we can formulate in SL the following version of the arithmetic comprehension axiom

$$\blacklozenge(\forall x, y((x \hookrightarrow y) \leftrightarrow \phi(x, y)))$$

[1] Here we adopt the terminology for second-order logic [Vää01], where the semantics of *full* second-order logic does not impose any restrictions on the *cardinality* of the interpretation of the predicates/relations, in contrast to *weak* second-order logic which restricts to *finite* interpretations (of the predicates/relations).

which expresses the existence of a heap such that its *graph*, as denoted by the points-to relation $\hookrightarrow$, satisfies the 'pure' first-order formula $\phi(x, y)$ (i.e., ϕ does not involve the separation connectives and the points-to relation). The $\blacklozenge$-modality (formally defined in Sect. 3) expresses the existence of a heap which satisfies the associated formula. Such an instance of the arithmetic comprehension axiom holds if there exists a heap which is characterized by the formula $\phi(x, y)$. We cannot generalize this axiom to arbitrary SL formulas because it is not obvious how to avoid contradictions like $\blacklozenge(\forall x, y((x \hookrightarrow y) \leftrightarrow \neg(x \hookrightarrow y)))$. Simply requiring that the points-to relation does not occur in $\phi(x, y)$ does not work because the separating connectives implicitly refer to it. Therefore, we introduce a new interpretation of SL that restricts the (second-order) quantification to *first-order definable* heaps. For this new interpretation we introduce a *sequent calculus* which is sound and complete. The completeness proof is based on the construction of a model for a *consistent* theory (a theory from which false is not derivable), following [Hen49]. From the completeness proof we further derive that this new interpretation satisfies both compactness and the downward Löwenheim-Skolem theorem. By the seminal theorem of Lindström we then infer that this new interpretation is as expressive as first-order logic.

Related Work. The model theory of SL has been focused mainly on finite heaps. For example, the computability and complexity results in [CYO01] depend on this assumption. Surprisingly, in [BDL12] the authors show that *weak* SL is as expressive as *weak* second-order logic [Man96], which is a semantics of second-order logic where quantification is restricted to finite relations. In [DD16] this result is further refined by the restriction to two variables and the separating implication (no separating conjunction) which still is as expressive as weak second-order logic. In [EIP20] the satisfiability problem for SL with k record fields has been studied for finite heaps, but over arbitrary first-order models. A tableaux method for a propositional fragment of SL has been developed in [GM10] which has been proven sound and complete. Extensions to first-order SL are discussed assuming finite heaps. In fact, the tableaux method introduced is based on a labelling mechanism for encoding finite heap structures.

In contrast, when investigating complete proof systems for SL the assumption of the finiteness of heaps has to be dropped, thus allowing for infinite heaps, because, as already observed above, finiteness leads to non-compactness. Our general model theory shows that this generalization of SL, *full* SL, is also non-compact, and therefore does not allow for a finitary sound and complete logic either. Consequently, to obtain such a logic one either has to syntactically restrict SL or further abstract or generalize its semantics. In [DLM21], for example, a sound and complete sequent calculus is described for a quantifier-free subset of SL. On the other hand, examples of further abstractions and generalizations are [HT16] and [Pym02], and both describe a finitary logic which is sound and complete. In [Pym02], models are based on very general preordered commutative monoids and there is no points-to relation. In [HT16], special commutative monoids called *separation algebras* are used to give semantics to the separating connectives. The elements of such separation algebras represent heaps as relations on the underlying (first-order) domain. This allows for a standard set-theoretic interpretation of the points-to relation. However, the semantics of separating conjunction is defined in terms of the abstract monoid, and as such is decoupled from the set-theoretic interpretation of the points-to relation. For example, a first-order specification (using plain conjunction) of an enumeration of the elements of the domain of a (finite) heap *as a set* does not in general correspond with an enumeration using separation conjunction.

A sound and complete axiomatization of the points-to relation in the general context of first-order SL *respecting its standard set-theoretic interpretation* thus remains a main challenge.

Second-order logic allows for a straightforward translation of the (weak or full) semantics of SL, and one can use second-order logic to reason about validity in SL. This approach is followed for example by the IRIS project [JKJ+18] which formalizes the semantics of weak SL in the higher-order logic of Coq [HH14]. By restricting the semantics of the separating connectives to (first-order) definable heaps, our approach instead transforms a *compositional* second-order logical description of the semantics of SL into corresponding rules of a standard first-order sequent calculus. The resulting calculus allows us to reason, in a natural manner, in first-order logic about the (hierarchical) heap structures generated by the rules for the separating connectives. As such it does not involve the additional tree structures of the so-called *bunched contexts* of the sequent calculi of [HT16] and [Pym02]. Also [Kri08] avoids the use of bunched contexts in a modal sequent calculus for propositional SL, which is proven sound. However it is incomplete because it provides limited support for equational reasoning about the modal contexts (so-called 'worlds') associated with the SL formulas.

Plan of the Paper. In the next section we introduce the syntax and semantics of full SL. In Sect. 3 we investigate the expressiveness of full SL. Section 4 introduces a restriction of the semantics to definable heaps. In Sect. 5 we introduce the sequent calculus, and discuss soundness and completeness. Finally, in the conclusion section we wrap up, and discuss some future work.

2 Separation Logic

In this section we introduce the syntax of SL and define its classical semantics with respect to arbitrary first-order models. For an intuitive introduction to separation logic, see [Rey05]. Given a first-order signature of function and predicate symbols[2] and a countably infinite set of first-order variables $x, y, z, \ldots$, the first-order terms of this signature are denoted by $t, t', \ldots$.

We have the following inductive definition of formulas of separation logic.

Definition 1 (Syntax of SL). *We define*

$$p ::= (t_1 = t_2) \mid R(t_1, \ldots, t_n) \mid (\neg p) \mid (p \wedge q) \mid \exists x(p) \mid (p * q) \mid (p \mathrel{-\!\!*} q)$$

[2] We allow for a countably infinite set of such symbols.

where R is a n-ary relation symbol. As a special case we have the binary 'points-to' relation symbol $\hookrightarrow$ (also called the weak/loose points-to).

Let $M = (D, I)$ denote a first-order model, where D denotes the non-empty domain and I provides an interpretation of the function and predicate symbols as functions and relations over D. A valuation s assigns elements of the domain D of M to the first-order variables $x, y, z, \ldots$. We omit the standard inductive definition of the value $I_s(t)$ of a term t. Given a model $M = (D, I)$, we denote by $M, h, s \models p$ that p holds in the model M, under the interpretation $h \subseteq D \times D$ of the binary relation symbol $\hookrightarrow$, where h denotes a so-called *heap*, represented as the graph of a *partial function* with *finite domain.*

Definition 2 (Semantics of SL). *We have the following main cases.*

- $M, h, s \models (t \hookrightarrow t')$ *if and only if* $\langle I_s(t), I_s(t')\rangle \in h$.
- $M, h, s \models (p * q)$ *if and only if* $M, h_1, s \models p$ *and* $M, h_2, s \models q$, *for some heaps* $h_1, h_2 \subseteq D \times D$ *such that* $h = h_1 \cup h_2$ *and* $h_1 \perp h_2$.
- $M, h, s \models (p \mathrel{-\!\!*} q)$ *if and only if* $M, h', s \models p$ *implies* $M, h \cup h', s \models q$, *for all heaps* $h' \subseteq D \times D$ *such that* $h \perp h'$.

Other cases are the Tarksi-style semantics of classical logic [Yan01, Table 5.2].

In the above definition we use the set-theoretic operation of *union* of binary relations as sets of pairs. On the other hand, by $h_1 \perp h_2$ we denote that the *domains* of the relations h_1 and h_2 are *disjoint*[3]. As such, we can introduce the strict/tight points-to relation $\mapsto$ of SL, defined by $M, h, s \models t \mapsto t'$ if and only if $h = \{\langle I_s(t), I_s(t')\rangle\}$, as a derived concept: it can be expressed by $(t \hookrightarrow t') \wedge \forall x, y((x \hookrightarrow y) \rightarrow (x = t \wedge y = t'))$. The concept **emp** of the empty relation can also be expressed by $\forall x, y(x \not\hookrightarrow y)$. *Intuitionistic* SL only allows for the weak/loose points-to relation. The strict version cannot be expressed in intuitionistic SL because of its *monotonicity* property that the truth of a formula is preserved by extensions of the domain of the heap [Rey00]. In this article we focus on classical separation logic only.

Let $(x_i \hookrightarrow -)$ abbreviate $\exists y(x_i \hookrightarrow y)$. The sentences ϕ_n defined by

$$\exists x_1, \ldots, x_n((x_1 \hookrightarrow -) * \ldots * (x_n \hookrightarrow -))$$

then state that there exist at least n allocated elements of the underlying domain of the given first-order model. Note that the semantics of the separating conjunction implies that $x_i \neq x_j$ for $i \neq j$. It is also possible to formulate the same property using propositional conjunction instead of separating conjunction by explicitly stating this fact, that the variables are not aliases. Now collect all ϕ_n in a set. Clearly, every finite subset of this set of sentences is satisfied by a finite heap, but that there does not exist a finite heap satisfying all these sentences.

[3] The domain of an arbitrary relation $\mathcal{R} \subseteq D \times D$ is the set $d \in D$ for which there exists a $d' \in D$ such that $\langle d, d'\rangle \in \mathcal{R}$. Note that for heaps $h_1 \perp h_2$ is equivalent to $h_1 \cap h_2 = \emptyset$.

This simple counterexample to compactness provides the basic motivation to study the above semantics of SL extended to unbounded heaps, i.e. heaps which potentially have an infinite domain.

Further, for technical convenience only, we generalize the semantics to arbitrary *binary relations*. For an arbitrary (binary) relation $\mathcal{R} \subseteq D \times D$ on the underlying domain D of the given first-order model, we define $M, \mathcal{R}, s \models p$ as above, where the interpretation of the separating connectives ranges over arbitrary subsets of $D \times D$. In fact, in this generalized semantics, which we call *relational* SL, we can model the restriction to heaps simply by *syntactically* restricting the separating implication to assertions of the form $(p \wedge \mathit{fun}) \mathrel{-\!\!*} q$, where *fun* denotes the assertion $\forall x, y, z((x \hookrightarrow y \wedge x \hookrightarrow z) \rightarrow y = z)$. Let p' denote the result of restricting syntactically all occurrences of the separating implication in p to heaps (as described above). It follows that the evaluation of $p' \wedge \mathit{fun}$ is restricted to heaps.

It is worthwhile to observe here that there exists a straightforward formalization of relational SL in second-order logic. For any formula p as defined above we define inductively the second-order formula $p(R)$, where R is a binary relation.

Definition 3 (Logical formalization of relational SL). *We have the following main cases.*

- $(t \hookrightarrow t')(R) = R(t, t')$,
- $(p * q)(R) = \exists R_1, R_2(R = R_1 \uplus R_2 \wedge p(R_1) \wedge q(R_2))$,
- $(p \mathrel{-\!\!*} q)(R) = \forall R_1, R_2((R_2 = R_1 \uplus R \wedge p(R_1)) \rightarrow q(R_2))$.

Here we denote by $R = R_1 \uplus R_2$, for any binary relation symbols R, R_1, R_2, the conjunction of the formulas $\forall x, y(R(x, y) \leftrightarrow (R_1(x, y) \vee R_2(x, y)))$ and $\forall x, y, z(\neg R_1(x, y) \vee \neg R_2(x, z))$. We denote by $M, s \models \phi$ the standard truth definition of a second-order formula ϕ, where the evaluation s additionally interprets the second-order variables. Correctness of this translation, that is, $M, \mathcal{R}, s \models p$ if and only if $M, s[R := \mathcal{R}] \models p(R)$ (where $s[R := \mathcal{R}]$ denotes the update of s which assigns to the binary variable R the relation $\mathcal{R}$), can be established by a straightforward induction on p.

3 Model Theory: Compactness and Countability

To explore the general model theory of SL we introduce the modalities $\blacksquare p$ and $\Box p$ as abbreviations of $\mathbf{true} * (\mathbf{emp} \wedge (\mathbf{true} \mathrel{-\!\!*} p))$ and $\neg(\mathbf{true} * \neg p)$, respectively[4]. For $M = (D, I)$ we have $M, \mathcal{R}, s \models \blacksquare p$ if and only if $M, \mathcal{R}', s \models p$, for *every* $\mathcal{R}' \subseteq D \times D$. Further, we have $M, \mathcal{R}, s \models \Box p$ if and only if $M, \mathcal{R}', s \models p$, for *every* sub-relation $\mathcal{R}'$ of $\mathcal{R}$ (that is, $\mathcal{R}' \subseteq \mathcal{R}$). By $\blacklozenge p$ we denote the formula $\neg \blacksquare \neg p$. It follows that $M, \mathcal{R}, s \models \blacklozenge p$ if and only if $M, \mathcal{R}', s \models p$, for *some* $\mathcal{R}' \subseteq D \times D$.

Characterizing Finite Models. The above $\blacksquare$-modality allows to express that the domain D of a model $M = (D, I)$ is finite, by asserting that every injective

[4] We note that $\blacksquare$ and $\blacklozenge$ are, respectively, $\Box$ and $\Diamond$ in [HT16]. However in [HT16] they are introduced not as abbreviations but as *primitive* concepts.

function $f : D \to D$ is a surjection: Let *inj* be the conjunction of the formulas *fun* (as defined above), $\forall x, y, z((x \hookrightarrow z \wedge y \hookrightarrow z) \to x = y)$, and $\forall x \exists y(x \hookrightarrow y)$. We have that $M, \mathcal{R}, s \models \mathit{inj}$ if and only if $\mathcal{R} : D \to D$ is injective (note that the domain of $\mathcal{R}$ is D because $M, \mathcal{R}, s \models \forall x \exists y(x \hookrightarrow y)$). And so $M, \mathcal{R}, s \models \blacksquare(\mathit{inj} \to \forall x \exists y(y \hookrightarrow x))$ if and only if D is finite. Note that the occurrences of $\hookrightarrow$ in the scope of the $\blacksquare$-modality are universally bounded, and the interpretation of $\hookrightarrow$ thus ranges over *all* $\mathcal{R} \subseteq D \times D$.

Characterizing Countable Infinity. We next show that countability of the underlying domain of a model can be expressed, using the above two modalities. We will be working with chains related by $\hookrightarrow$, and in that sense we speak of a *predecessor* of x, being any y such that $(y \hookrightarrow x)$, and *successor* of x, being any y such that $(x \hookrightarrow y)$. Let *enum* be the conjunction of the following formulas:

- the above formula *inj*,
- the formula $\exists! x \forall y(y \not\hookrightarrow x)$[5], which states the existence of a unique *minimal* element (that is, an element that has no predecessor),
- the formula $\square(\mathbf{emp} \vee \exists x((x \hookrightarrow -) \wedge \forall y((y \hookrightarrow -) \to (y \not\hookrightarrow x)))$, which expresses that the points-to relation $\hookrightarrow$ is *well-founded.*

Note that a relation $\mathcal{R}$ is well-founded iff *every* (non-empty) sub-relation of $\mathcal{R}$ has a minimal element (with respect to that sub-relation). This fact can be expressed by the use of the formula *enum*. Let $M, \mathcal{R}, s \models \mathit{enum}$. We show that $\mathcal{R}$ encodes an enumeration $\langle d_n \rangle_n$ of D (still we have $M = (D, I)$). We define the sequence $\langle d_n \rangle_n$ by induction on n: for d_0 we take the (unique) minimal element, and for d_{n+1} we take the unique element $d \in D$ such that $\langle d_n, d \rangle \in \mathcal{R}$. Note that *inj* implies that every element of D has a unique 'successor' and that $d_{n+1} \notin \{d_0, \ldots, d_n\}$. Well-foundedness ensures that every element of D appears in the enumeration $\langle d_n \rangle_n$. Because otherwise we can construct an infinite descending chain of elements not appearing in the enumeration $\langle d_n \rangle_n$ (since d_0 denotes the unique minimal element with respect to the functional interpretation $\mathcal{R}$ of $\hookrightarrow$, it follows that for any $d \in D$ which does not appear in the enumeration $\langle d_n \rangle_n$ there exists a $d' \in D$ which also does not appear in the enumeration $\langle d_n \rangle_n$ and $\langle d', d \rangle \in \mathcal{R}$).

We thus have that $M, \mathcal{R}, s \models \mathit{enum}$ implies that the domain of M is countably infinite. The formula $\blacklozenge \mathit{enum}$ further abstracts from the current interpretation of the points-to relation $\hookrightarrow$, so that if the domain of M is countably infinite then $M, \mathcal{R}, s \models \blacklozenge \mathit{enum}$, for arbitrary $\mathcal{R}$ (and s).

The class of uncountable models is characterized by $\neg(\blacklozenge \mathit{enum} \vee \mathit{fin})$, where *fin* denotes the above formula which characterizes the class of finite models.

Summarizing, the logic of full SL is neither compact nor does it satisfy the Löwenheim-Skolem theorem because it can distinguish between countable and uncountable models. Further, we observe that the above expressiveness results do not depend on the interpretation of the points-to relation as an arbitrary relation. That is, these results also hold for the semantics restricted to (infinite) heaps.

[5] $\exists! xp$ is an abbreviation of $\exists x(p \wedge \forall y(p[y/x] \to y = x))$, where $p[y/x]$ denotes the substitution of x by y.

Interestingly, since we can express that the points-to relation $\hookrightarrow$ is well-founded (see above), even restricting to the separating conjunction gives rise to non-compactness: given a countably infinite set of individual constants c_n, $n \geq 0$, let Γ consist of the above formula $\Box(\mathbf{emp} \vee \exists x((x \hookrightarrow -) \wedge \forall y((y \hookrightarrow -) \rightarrow (y \not\hookrightarrow x)))$ and the formulas $c_{n+1} \hookrightarrow c_n$, $n \geq 0$. Clearly, every finite subset of Γ is satisfiable but Γ itself is not. Note that we do not need to require that all the $c_i \neq c_j$, for every $i \neq j$, because in case the formulas $c_{n+1} \hookrightarrow c_n$, $n \geq 0$, are satisfied and additionally $c_i = c_j$ holds, for some $i \neq j$, we have a loop in the interpretation of $\hookrightarrow$. Further, restricting SL to separating conjunction also does not satisfy the *upward* Löwenheim-Skolem theorem, because, as argued above, $M, \mathcal{R}, s \models enum$ implies (infinite) countability of the domain of M.

Separation Logic Light. What about further restricting to *positive* occurrences of the separating conjunction? Since we then can push negation inside, this restriction can be formally defined by the following syntax describing SLL ('separation logic light'):

$$p ::= (\neg)R(t_1, \ldots, t_n) \mid (p \vee q) \mid (p \wedge q) \mid \exists x(p) \mid \forall x(p) \mid (p * q)$$

Here R denotes either a n-ary relation symbol or the points-to relation $\hookrightarrow$. Thus, in this version of SL, negation can only be applied to atomic formulas. To show that the notion of satisfiability of SLL is compact, we introduce the following first-order translation $p@R$, where R is a binary predicate different from $\hookrightarrow$, $\circ$ denotes conjunction/disjunction, and Q denotes the existential/universal quantifier.

$$\begin{array}{ll}
(\neg)R(t_1, \ldots, t_n)@R' & = (\neg)R(t_1, \ldots, t_n) \\
(t \hookrightarrow t')@R & = R(t, t') \\
(p \circ q)@R & = p@R \circ q@R \\
Qx(p)@R & = Qx(p@R) \\
(p * q)@R & = R = R_1 \uplus R_2 \wedge p@R_1 \wedge q@R_2
\end{array}$$

The binary relation symbols R_1 and R_2 are 'fresh'. It follows that p is satisfiable if and only if $p@R$ is satisfiable. More precisely, $M, \mathcal{R}, s \models p$ if and only if there exists a (first-order) model M' such that $M', s \models p@R$. Consequently, compactness of first-order logic implies compactness of SLL: Let Γ be an infinite set of formulas of SLL and $\Gamma' = \{p@R \mid p \in \Gamma\}$[6], for some binary relation symbol R. If every finite subset of Γ is satisfiable, so is every finite subset of Γ'. By the compactness of first-order logic Γ' is satisfiable, and so is Γ. Along the same lines it follows that if Γ is satisfiable then there exists a model $M = (D, I)$ such that D is *countable* and $M, \mathcal{R}, s \models p$, for every $p \in \Gamma$.

Note however that compactness of the satisfiability relation does not imply that the (semantic) consequence relation is compact. In fact, non-compactness of the consequence relation for SLL follows directly from the above argument

[6] Note that Γ' may require the introduction of an infinite number of fresh (binary) relation symbols. This is however no problem because first-order logic allows for a countably infinite set of function and relation symbols.

involving well-founded relations: Let $\mathit{\Gamma}$ denote the set formulas $c_{n+1} \hookrightarrow c_n$, $n \geq 0$. It follows that $\mathit{\Gamma} \models \mathbf{true} * (\neg\mathbf{emp} \wedge \forall x((x \hookrightarrow -) \rightarrow \exists y(y \hookrightarrow x)))$. But clearly, there does not exist a finite subset $\mathit{\Gamma}_0$ of $\mathit{\Gamma}$ such that $\mathit{\Gamma}_0 \models \mathbf{true} * (\neg\mathbf{emp} \wedge \forall x((x \hookrightarrow -) \rightarrow \exists y(y \hookrightarrow x)))$.

Some Open Problems. The question remains whether restricting to separating conjunction satisfies the *downward* Löwenheim-Skolem theorem. A counterexample to the downward Löwenheim-Skolem theorem would be the expressibility of uncountable models. This seems to require the $\blacksquare p$ modality (and thus the separating implication).

Another interesting question is whether we can express finiteness of the domain of the current interpretation of the points-to relation, that is, does there exist a formula p in SL such that $M, \mathcal{R}, s \models p$ if and only if the domain of the relation $\mathcal{R}$ is finite?

A main open problem is a formalization of the relation between full SL and second-order logic. Intuitively, one of the main differences is the *local perspective* of SL, which is determined by the current heap. Remarkably, as already mentioned in the introduction, [BDL12] presents a rather intricate encoding of (dyadic) weak second-order logic into weak SL. Apparently this restriction to finite heaps allows to break the local perspective. Our conjecture however is that full SL is strictly less expressive than (dyadic) second-order logic. To illustrate how subtle this difference may be, consider the following extension of separation logic with a *binding* operator $\downarrow R(p)$ which binds the binary variable R in the evaluation of p to the current interpretation of the points-to relation. In other words, it corresponds to a bounded (second-order) quantification $\exists R((R = \hookrightarrow) \wedge p)$, where, $R = \hookrightarrow$ abbreviates the first-order formula $\forall x, y(R(x, y) \leftrightarrow (x \hookrightarrow y))$. Alternatively, we can directly define $M, \mathcal{R}, s \models \downarrow R(p)$ if and only if $M, \mathcal{R}, s[R := \mathcal{R}] \models p$. This definition thus assumes an extension of the valuation s to (binary) second-order variables. The expressive power of this binding operator lies in that it allows to 'break the spell' of the local perspective since the bound binary variable allows in the local context of the current interpretation of the points-to relation to refer to those 'outer' ones that have generated it (by the separating connectives). This extension of SL allows for a simple, compositional translation of (dyadic) second-order logic. We have the following main case which translates $\exists R(\phi)$, where ϕ a dyadic second-order formula (which is assumed not to contain occurrences of the points-to relation of SL), into the SL formula $\blacklozenge(\downarrow R(p))$.

4 Separation Logic of Definable Binary Relations

In this section we restrict the interpretation of the separating connectives to first-order definable binary relations. By ϕ we now denote a first-order formula which does not contain occurrences of the points-to relation $\hookrightarrow$ of SL. We omit the standard inductive truth definition $M, s \models \phi$ of a first-order formula ϕ.

By $\phi(x_1, \ldots, x_n)$ we denote that the free (first-order) variables of ϕ are among the distinct variables $x_1, \ldots, x_n$. A formula $\phi(x, y)$ is called a *binary* formula.

A binary formula is also simply denoted by ϕ, omitting its free variables x and y. Given a model $M = (D, I)$, and a first-order formula $\phi(x, y)$, we denote by $Rel_M(\phi)$ the relation $\{\langle s(x), s(y)\rangle \mid M, s \models \phi\} \subseteq D \times D$. Note that the evaluation of $\phi(x, y)$ only depends on the values of its free variables x and y, that is, $M, s \models \phi$ if and only if $M, s' \models \phi$, where $s(x) = s'(x)$ and $s(y) = s'(y)$. By $\phi(t, t')$ we denote the result of replacing in $\phi(x, y)$ the variables x and y by t and t', respectively (if necessary renaming bound variables to ensure that the variables of t and t' do not become bound).

Definition 4 (First-order definability). *Given a model $M = (D, I)$, a relation $\mathcal{R} \subseteq D \times D$ is* first-order definable *if $\mathcal{R} = Rel_M(\phi)$, for some binary formula $\phi(x, y)$.*

Note that, given a model $M = (D, I)$, $I(R) = Rel_M(R)$, that is, for any binary relation symbol R its interpretation $I(R)$ is trivially a first-order definable relation. We generalize the definition of $R = R_1 \uplus R_2$ to arbitrary binary formulas: we denote by $\phi = \phi_1 \uplus \phi_2$ that the binary formulas $\phi_1(x, y)$ and $\phi_2(x, y)$ represent a partition of the binary formula $\phi(x, y)$ which is expressed by the conjunction of $\forall x, y(\phi(x, y) \leftrightarrow (\phi_1(x, y) \lor \phi_2(x, y)))$ and $\forall x, y, z(\neg\phi_1(x, y) \lor \neg\phi_2(x, z))$. The latter formula, which states that the domains of the binary relations represented by $\phi_1(x, y)$ and $\phi_1(x, y)$ are disjoint, we abbreviate by $\phi_1 \perp \phi_2$.

In the sequel we denote by $M, \mathcal{R}, s \models p$ the *restriction* of the relational semantics of full SL (Definition 2 extended to binary relations) such that instead of quantifying over arbitrary binary relations, the separating connectives involve quantification over first-order definable binary relations. It is worthwhile to observe here that, as for Henkin models of second-order logic [Hen50], the implicit second-order quantification depends on the underlying signature of function and relation symbols. Extending or restricting the signature affects the semantics of formulas of the 'old' signature.

5 Sequent Calculus

To reason about the implicit quantification over definable (binary) relations, we introduce *rooted* assertions of the form $p@\phi$, where ϕ denotes a binary formula and p is a formula of SL (see Definition 1). We define $M, s \models p@\phi$ if and only if $M, \mathcal{R}, s \models p$, where $\mathcal{R} = Rel_M(\phi)$. The variables x and y of the binary formula $\phi(x, y)$ are thus implicitly bound by the @-operator, that is, $M, s \models p@\phi$ if and only if $M, s' \models p@\phi$, for any s and s' such that $s(z) = s'(z)$, for any free variable occurring in p.

Note that the separating connectives are interpreted in terms of relations which are definable by first-order formulas which do not involve the points-to relation $\hookrightarrow$. This allows for the following alternative *predicative* definition[7] of the semantics of the separating connectives in rooted assertions (used in both the soundness and completeness proofs). Here $\psi \perp \phi$, for the binary formulas $\psi(x, y)$ and $\phi(x, y)$, denotes the formula $\forall x, y, z(\neg\psi(x, y) \lor \neg\phi(x, z))$.

[7] For a foundational discussion concerning predicativity, see [Cro17].

Separating conjunction

$$\mathbf{L}_{*} \qquad \frac{\varGamma, \phi = R_1 \uplus R_2, p@R_1, q@R_2 \Rightarrow \varDelta}{\varGamma, (p * q)@\phi \Rightarrow \varDelta}$$

$$\mathbf{R}_{*} \qquad \frac{\varGamma \Rightarrow \varDelta, \phi = \phi_1 \uplus \phi_2 \quad \varGamma \Rightarrow \varDelta, p@\phi_1 \quad \varGamma \Rightarrow \varDelta, q@\phi_2}{\varGamma \Rightarrow \varDelta, (p * q)@\phi}$$

Separating implication

$$\mathbf{L}_{-\!\!*} \qquad \frac{\varGamma \Rightarrow \varDelta, \phi \perp \psi \quad \varGamma \Rightarrow \varDelta, p@\psi \quad \varGamma, q@(\phi \vee \psi) \Rightarrow \varDelta}{\varGamma, (p -\!\!* q)@\phi \Rightarrow \varDelta}$$

$$\mathbf{R}_{-\!\!*} \qquad \frac{\varGamma, R \perp \phi, p@R \Rightarrow \varDelta, q@(\phi \vee R)}{\varGamma \Rightarrow \varDelta, (p -\!\!* q)@\phi}$$

Points-to rules

$$\frac{\varGamma, p[\phi/\hookrightarrow] \Rightarrow \varDelta}{\varGamma, p@\phi \Rightarrow \varDelta} \qquad \frac{\varGamma \Rightarrow p[\phi/\hookrightarrow], \varDelta}{\varGamma \Rightarrow p@\phi, \varDelta}$$

Fig. 1. Sequent calculus. The binary relation symbols R_1, R_2 and R introduced in the rules $\mathbf{L}_{*}$ and $\mathbf{R}_{-\!\!*}$ are 'fresh'. In the points-to rules p denotes a basic formula (which does not contain occurrences of the separating connectives).

Lemma 1. *We have*

- *$M, s \models (p * q)@\phi$ if and only if there exist binary formulas ϕ_1 and ϕ_2 such that $M, s \models \phi = \phi_1 \uplus \phi_2$, $M, s \models p@\phi_1$, and $M, s \models q@\phi_2$.*
- *$M, s \models (p -\!\!* q)@\phi$ if and only if $M, s \models \psi \perp \phi$ and $M, s \models p@\psi$ implies $M, s \models q@(\phi \vee \psi)$, for all binary formulas ψ.*

We now develop a calculus for sequents $A_1, \ldots, A_n \Rightarrow B_1, \ldots, B_m$, where each A_i, $i = 1, \ldots, n$, and B_j, $j = 1, \ldots, m$, is constructed from first-order formulas and rooted assertions, which can be further composed using propositional connectives and quantification of first-order variables. This calculus is an extension of standard first-order sequent calculus (including cut), where the standard rules are applicable with respect to top-level propositional connectives and quantifiers. Figure 1 shows the left and right rules for separating conjunction and implication. These rules closely follow the translation in Definition 3 of SL into second-order logic, eliminating the explicit second-order quantification by applying the standard proof rules for second-order quantification (which themselves are straightforward generalizations of the rules for first-order quantification, instantiating the second-order variables by formulas). The binary relation symbols R_1, R_2 and R introduced in the rules $\mathbf{L}_{*}$ and $\mathbf{R}_{-\!\!*}$ are 'fresh' binary relation symbols, that is, they must not appear in the formulas of the conclusion of the rules.

We also have rules which allow classical reasoning under rooted assertions: $(p \circ q)@\phi \leftrightarrow (p@\phi) \circ (q@\phi)$, where $\circ$ denotes binary propositional connectives, e.g., conjunction, disjunction, and implication, $(\neg p)@\phi \leftrightarrow \neg(p@\phi)$, and $(\exists x p)@\phi \leftrightarrow \exists x(p@\phi)$ (and similarly $(\forall x p)@\phi \leftrightarrow \forall x(p@\phi)$). Further, we have $\forall x, y(\phi \leftrightarrow \psi) \rightarrow (p@\phi \leftrightarrow p@\psi)$. It is straightforward to validate these rules, but we omit the details of the semantics $M, s \models A$, which follows the standard Tarski-style classical semantics, given the semantics of rooted assertions which may appear in the place of atomic formulas.

In the so-called 'points-to' rules of Fig. 1 the formula p does not involve occurrences of the separating connectives. Such a formula of SL we call *basic*. Note that it differs from pure first-order formulas in that basic formulas additionally may involve the points-to relation. For such formulas we denote by $p[\phi/\hookrightarrow]$, for any binary formula $\phi(x, y)$, the result of replacing every atomic assertion $(t \hookrightarrow t')$ in p by $\phi(t, t')$, which is a pure first-order formula. It follows that $M, s \models p[\phi/\hookrightarrow]$ if and only if $M, Rel_M(\phi), s \models p$, for any basic formula p.

Example Proofs

$$\dfrac{\dfrac{\dfrac{\dfrac{\Gamma \Rightarrow q@R, R_1 \perp R_2 \quad \Gamma \Rightarrow q@R, p@R_1 \quad \Gamma, q@(R_1 \vee R_2) \Rightarrow q@R}{R = R_1 \uplus R_2, p@R_1, (p \mathrel{-\!\!*} q)@R_2 \Rightarrow q@R}\,\mathbf{L}_{-\!*}}{(p * (p \mathrel{-\!\!*} q))@R \Rightarrow q@R}\,\mathbf{L}_*}{\Rightarrow (p * (p \mathrel{-\!\!*} q))@R \rightarrow q@R}}{\Rightarrow ((p * (p \mathrel{-\!\!*} q)) \rightarrow q)@R}$$

As a first example of the use of the sequent calculus, above we have a derivation of the sequent $\Rightarrow ((p * (p \mathrel{-\!\!*} q)) \rightarrow q)@R$ which represents the validity of $(p * (p \mathrel{-\!\!*} q)) \rightarrow q$. This derivation essentially consists of an application of the rule $\mathbf{L}_*$ followed by an application of the rule $\mathbf{L}_{-\!*}$. In this derivation Γ denotes the formulas $R = R_1 \uplus R_2, p@R_1$ generated by the application of rule $\mathbf{L}_*$. The second premise of the application of the rule $\mathbf{L}_{-\!*}$ is derivable from an instance of the axiom $\Gamma, A \Rightarrow A, \Delta$. Note that ψ (in the $\mathbf{L}_{-\!*}$ rule) is instantiated with R_1. The first and third premise follows from the fact that $R = R_1 \uplus R_2$ reduces to $R_1 \perp R_2$ and $R = R_1 \cup R_2$ (that part of the proof is not shown above).

Next we show how to use the calculus in reasoning about the equivalence of weakest preconditions that arise in the practice of verifying the correctness of heap manipulating programs. Let p denote the weakest precondition $(u \hookrightarrow -) \wedge (z = 0 \triangleleft u = v \triangleright v \hookrightarrow z)$ of the heap update $[u] := 0$ which ensures the postcondition $v \hookrightarrow z$ after assigning the value 0 to the location denoted by the variable u (here $\phi \triangleleft b \triangleright \psi$ abbreviates $(b \wedge \phi) \vee (\neg b \wedge \psi)$) (in [dBHdG23] a dynamic logic extension of SL is introduced which generates this weakest precondition). The standard rule for backwards reasoning in [Rey02] gives the weakest precondition $(u \mapsto -) * (u \mapsto 0 \mathrel{-\!\!*} v \hookrightarrow z)$, which we denote by p'. These preconditions are equivalent because both are the weakest.

Surprisingly, a proof of the implication $p' \rightarrow p$ however exceeds the capability of all the automatic SL provers in the benchmark competition for SL [SNPR+19].

In particular, of the automatic provers, only the CVC4-SL tool [RISK16] supports the fragment of SL that includes the separating implication connective. However, from our own experiments with that tool, we found that it produces an incorrect counter-example and reported this as a bug to one of the maintainers of the project (Andrew Reynolds). In fact, the latest version, CVC5-SL, reports the same input as 'unknown', indicating that the tool is incomplete. In the case of (semi) interactive SL provers (such as Iris [JKJ+18], and VerCors [AH21,MRH22] that uses Viper [MSS16] as a back-end) we sought out expertise and collaborated in our search for a tool-supported proof of the above equivalence. Even after personally visiting the Iris team in Nijmegen (lead by Robbert Krebbers) and the VerCors team in Twente (lead by Marieke Huisman), we were unable to guide the tools to produce a proof of $p' \rightarrow p$. The problem here seems similar to that of [HT16], in that their semantics of separating connectives, which are formalized in terms of abstract monoids, are not compatible with the set-theoretic interpretation of the points-to relation.

In fact, the equivalence between the above two formulas can be expressed in quantifier-free separation logic, for which a complete axiomatization of all valid formulas has been given in [DLM21]. In the sequent calculus we can express the equivalence of p and p' in terms of the sequent $\mathit{fun}(R) \Rightarrow (p \leftrightarrow p')@R$. Here R is an arbitrary binary relation symbol used to represent the current interpretation of the points-to relation. We abbreviate $\forall x, y, z((R(x, y) \wedge R(x, z)) \rightarrow y = z)$ by $\mathit{fun}(R)$. A proof of the above sequent amounts to proving the sequents $\mathit{fun}(R), p'@R \Rightarrow p@R$ and $\mathit{fun}(R), p@R \Rightarrow p'@R$. Below we present a high-level proof of the first sequent, abstracting from some basic first-order reasoning in the calculus.

By an application of $\mathbf{L}_*$ to derive the sequent $\mathit{fun}(R), p'@R \Rightarrow p@R$ it suffices to derive

$$\mathit{fun}(R), R = R_1 \uplus R_2, (u \mapsto -)@R_1, (u \mapsto 0 \mathrel{-\!\!*} v \hookrightarrow z)@R_2 \Rightarrow p@R$$

for some fresh R_1 and R_2. Let $\psi(x, y)$ denote the binary formula $x = u \wedge y = 0$. Further, let Γ denote the set of formulas $\mathit{fun}(R), R = R_1 \uplus R_2, (u \mapsto -)@R_1$. By an application of the rule $\mathbf{L}_{-\!*}$ it then suffices to prove the following sequents (from $\Gamma \Rightarrow \Delta$ we can derive $\Gamma \Rightarrow A, \Delta$ by right-weakening). First we prove $\Gamma \Rightarrow R_2 \cap \psi = \emptyset$: By the points-to rules the rooted assertion $(u \mapsto -)@R_1$ (appearing in Γ) reduces to $\exists z(R_1(u, z) \wedge \forall x, y(R_1(x, y) \rightarrow x = u \wedge y = z))$ (the forall-part of the formula is due to the 'strict' points-to which states that the domain contains u as its only location). Further, $R_2 \cap \psi = \emptyset$ logically boils down to $\neg\exists x, y(R_2(x, y) \wedge (x = u \wedge y = 0))$, that is, $\neg R_2(u, 0)$, which in basic first-order logic follows from $\exists z R_1(u, z)$ and the assumptions $R = R_1 \uplus R_2$ and $\mathit{fun}(R)$.

Second, we prove $\Gamma \Rightarrow (u \mapsto 0)@\psi$: By the points-to rules $(u \mapsto 0)@\psi$ (using the expanded definition ϕ of $u \mapsto 0$ and the definition of the substitution $\phi[\psi/\hookrightarrow]$) reduces to $(u = u) \wedge (0 = 0) \wedge \forall x, y((x = u \wedge y = 0) \rightarrow (x = u \wedge y = 0))$ which is equivalent to **true**.

And, finally, we prove $\Gamma, (v \hookrightarrow z)@(R_2 \vee \psi) \Rightarrow p@R$: First note that (again, by the points-to rules)

$$((u \hookrightarrow -) \wedge (z = 0 \triangleleft u = v \triangleright v \hookrightarrow z))@R$$

reduces to

$$(\exists z R(u,z)) \wedge (z = 0 \triangleleft u = v \triangleright R(v,z)))$$

The assertion $\exists z R(u,z)$ clearly follows from the assumptions $R = R_1 \uplus R_2$ and $(u \mapsto -)@R_1$ in Γ. To prove $z = 0 \triangleleft u = v \triangleright R(v,z)$, we first reduce the assumption $(v \hookrightarrow z)@(R_2 \vee \psi)$ to $R_2(v,z) \vee (v = u \wedge z = 0)$. Now, if $v = u$ then $\neg R_2(v,z)$, because of the assumptions $fun(R)$, $R = R_1 \uplus R_2$ and $(u \mapsto -)@R_1$. So we have that $z = 0$. Otherwise, we have $R_2(v,z)$, and thus $R(v,z)$, because $R = R_1 \uplus R_2$.

Soundness and Completeness. We denote by $\vdash \Gamma \Rightarrow \Delta$ that there exists a proof of the sequent $\Gamma \Rightarrow \Delta$. To define $\models \Gamma \Rightarrow \Delta$, let σ denote a substitution which assigns to every binary relation symbol R of the sequent $\Gamma \Rightarrow \Delta$ a binary formula ϕ. Such a substitution σ simply replaces occurrences of $R(t,t')$ by $\phi(t,t')$, where $\sigma(R) = \phi(x,y)$. By $\models \Gamma \Rightarrow \Delta$ we then denote that $M, s \models \bigwedge \Gamma\sigma$ (that is, $M, s \models A\sigma$, for every $A \in \Gamma$) implies $M, s \models \bigvee \Delta\sigma$ (that is, $M, s \models B\sigma$, for some $B \in \Delta$), for every M, s and every substitution σ.

In the soundness proof below we use these substitutions to instantiate the fresh binary relation symbols introduced in the rules $\mathbf{L}_*$ and $\mathbf{R}_{-*}$. Note that updating the interpretation of these symbols (as provided by M) would affect the semantics of the separating connectives if binary formulas would refer to these fresh binary relation symbols (note that they are only supposed not to appear in formulas of the conclusion of the rules $\mathbf{L}_*$ and $\mathbf{R}_{-*}$).

We generalize the above notions of derivability and validity to possibly infinite Γ: $\Gamma \vdash \Delta$ indicates that $\vdash \Gamma' \Rightarrow \Delta$, for some finite $\Gamma' \subseteq \Gamma$, and $\Gamma \models \Delta$ indicates that for every substitution σ we have that $M, s \models \Gamma\sigma$ (that is, $M, s \models A\sigma$, for every $A \in \Gamma$) implies $M, s \models B\sigma$, for some $B \in \Delta$.

Theorem 1 (Soundness). *We have that* $\vdash \Gamma \Rightarrow \Delta$ *implies* $\models \Gamma \Rightarrow \Delta$.

Proof. We prove that the rules for the separating connectives preserve validity. The points-to rules are sound because $M, Rel_M(\phi), s \models p$ if and only if $M, s \models p[\phi/\hookrightarrow]$, for any basic formula p (note that $p[\phi/\hookrightarrow]$ is a pure first-order formula which does not depend on the heap).

$\mathbf{L}_*$: Let $M, s \models \Gamma\sigma$ and $M, s \models (p\sigma * q\sigma)@\phi\sigma$. We have to show that $M, s \models \bigvee \Delta\sigma$. By Lemma 1, there exist ϕ_1 and ϕ_2 such that $M, s \models (\phi\sigma) = \phi_1 \uplus \phi_2$, $M, s \models p\sigma@\phi_1$, and $M, s \models q\sigma@\phi_2$. Let $\sigma' = \sigma[R_1, R_2 := \phi_1, \phi_2]$. Since R_1 and R_2 are fresh and as such do not appear in $\Gamma, (p * q)@\phi$, it follows that $M, s \models \Gamma'\sigma'$, where $\Gamma' = \Gamma, \phi = R_1 \uplus R_2, p@R_1, q@R_2$. By the validity of the premise we thus obtain that $M, s \models \bigvee \Delta\sigma'$. Since R_1 and R_2 also do not appear in Δ, we conclude that $M, s \models \bigvee \Delta\sigma$.

$\mathbf{R}_*$: Let $M, s \models \Gamma\sigma$ and suppose that $M, s \not\models \bigvee \Delta\sigma$. From the validity of the premises it then follows that $M, s \models \phi\sigma = (\phi_1 \uplus \phi_2)\sigma$, $M, s \models p\sigma@\phi_1\sigma$, and $M, s \models q\sigma@\phi_2\sigma$, By Lemma 1 we conclude $M, s \models (p\sigma * q\sigma)@\phi\sigma$.

$\mathbf{L}_{-\!*}$: Let $M, s \models \Gamma\sigma$ and $M, s \models (p\sigma \mathrel{-\!*} q\sigma)@\phi\sigma$, and suppose that $M, s \not\models \bigvee \Delta\sigma$. From the validity of the first two premises it then follows that $M, s \models \phi\sigma \perp \psi\sigma$ and $M, s \models p\sigma@\psi\sigma$. By Lemma 1 again, it follows that $M, s \models q\sigma@(\phi\sigma \vee \psi\sigma)$. By the validity of the third premise we thus derive that $M, s \not\models \bigvee \Delta\sigma$, which a contradicts our assumption.

$\mathbf{R}_{-\!*}$: Let $M, s \models \Gamma\sigma$ and suppose that $M, s \not\models \bigvee \Delta\sigma$. We have to show that $M, s \models (p\sigma \mathrel{-\!*} q\sigma)@\phi\sigma$. Let ψ be such that $M, s \models \psi \perp (\phi\sigma)$ and $M, s \models p\sigma@\psi$. Further, let R be a fresh variable and $\sigma' = s[R := \psi]$. It follows that $M, s \models \Gamma'\sigma'$, where $\Gamma' = \Gamma, R \perp \phi, p@R$ and $M, s \not\models \bigvee \Delta\sigma'$. And so we derive from the validity of the premise of the rule that $M, s \models q\sigma@(\phi\sigma \cup \psi)$. Since ψ was arbitrarily chosen, by Lemma 1 again we conclude that $M, s \models (p\sigma \mathrel{-\!*} q\sigma)@\phi\sigma$. □

As a corollary we obtain that $\Gamma \vdash \Delta$ implies $\Gamma \models \Delta$.

Following the completeness proof of first-order logic as described in [Hen49], it suffices to show that every consistent set of formulas is satisfiable (the so-called 'model existence theorem'). A set of formulas Γ is consistent if $\Gamma \not\vdash \emptyset$. We first show that every consistent set of formulas can be extended to a maximal consistent set. To this end we assume an infinite set of 'fresh' binary relation symbols R that do not appear in Γ. We construct for any consistent set Γ a maximal consistent extension Γ^∞, assuming an enumeration of all formulas A (which also covers all first-order formulas). We define $\Gamma_0 = \Gamma$ and Γ_{n+1} satisfies the general rule: if $\Gamma_n, A_n \not\vdash \emptyset$ then $\Gamma_n \cup \{A_n\} \subseteq \Gamma_{n+1}$, otherwise $\Gamma_{n+1} = \Gamma_n$. Additionally, in case A_n is added and A_n is of the form $\exists x A$ or a rooted assertion $(p * q)@\phi$ or $\neg(p \mathrel{-\!*} q)@\phi$, we also include corresponding *witnesses* in Γ_{n+1}:

- If A_n is of the form $\exists x A$ we additionally add $A(y)$, where $A(y)$ results from replacing all free occurrences of x in A by the fresh variable y which does not appear in Γ_n.
  Note that $A(y)$ can indeed be added consistently because from $\Gamma_n, A(y) \vdash \emptyset$ we would derive $\Gamma_n, \exists x A \vdash \emptyset$, which contradicts the assumption that $\Gamma_n, \exists x A \not\vdash \emptyset$.
- If A_n is of the form $(p * q)@\phi$ we additionally add the formulas $\phi = R_1 \uplus R_2, R_1 \perp R_2, p@R_1$, and $q@R_2$, where R_1 and R_2 are fresh (e.g., not appearing in Γ_n).
  Note that these formulas can indeed be added consistently because from $\Gamma_n, \phi = R_1 \uplus R_2, R_1 \perp R_2, p@R_1, q@R_2 \vdash \emptyset$ we would derive $\Gamma_n, (p * q)@\phi \vdash \emptyset$ (by rule $\mathbf{L}_*$).
- If A_n is of the form $\neg(p \mathrel{-\!*} q)@\phi$ (which is equivalent to $\neg((p \mathrel{-\!*} q)@\phi)$) we additionally add the formulas $R \perp \phi, p@R(x, y)$, and $\neg q@(\phi \vee R)$, where R is fresh (e.g., not appearing in Γ_n).
  Note that these formulas can indeed be added consistently because from $\Gamma_n, R \perp \phi, p@R(x, y), \neg q@(\phi \vee R) \vdash \emptyset$ we would derive $\Gamma_n \vdash (p \mathrel{-\!*} q)@\phi$ (by rule $\mathbf{R}_{-\!*}$), which contradicts the assumption that $\Gamma_n, \neg(p \mathrel{-\!*} q)@\phi \not\vdash \emptyset$.

We define $\Gamma^\infty = \bigcup_n \Gamma_n$. By construction Γ^∞ is maximal consistent. Given a maximal consistent set of formulas Γ, let $M_\Gamma = (D, I)$, where D is the set of equivalences classes $[t] = \{t' \mid t = t' \in \Gamma\}$. For any function symbol f and relation symbol R (excluding the points-to relation $\hookrightarrow$) we define

- $I(f)([t_1], \ldots, [t_n]) = [f(t_1, \ldots, t_n)]$,
- $I(R)([t_1], \ldots, [t_n]) = \textbf{true}$ if and only if $R(t_1, \ldots, t_n) \in \Gamma$.

The above interpretation of the function and relational symbols is well-defined because its definition does not depend on the choice of the representatives (this follows from the equality axioms).

Given a maximal consistent set of formulas Γ and the model $M_\Gamma = (D, I)$, a corresponding valuation s assigns to every variable x an equivalence class $[t]$. However, in the sequel we will represent such a valuation by a *substitution* s which simply assigns to each variable a term. The value $I_s(x)$ of a variable x then is given by the equivalence class $[s(x)]$ of the term $s(x)$.

Given a substitution s, for any term t and formula A (of the sequent calculus) we denote by ts and As the result of replacing every free occurrence of a (first-order) variable x in t and A by $s(x)$. Note that $(p@\phi)s = ps@\phi$, because the meaning of $p@\phi$ does not depend on the free variables x and y of the binary formula $\phi(x, y)$.

Given a maximal consistent set of formulas Γ and the model $M_\Gamma = (D, I)$, it follows that $I_s(t) = [ts]$, for every term t and substitution s.

Lemma 2. *Given a maximal consistent set of formulas Γ and the model $M_\Gamma = (D, I)$, we have $M, s \models A$ if and only if $As \in \Gamma$, for every formula A and substitution s.*

Proof. The proof proceeds by induction on the following well-founded ordering $A < B$ on formulas of the sequent calculus: Let $\#A = (n, m)$, where n denotes the number of occurrences of the separating connectives and the @-binding operator of A and m denotes the number of occurrences of the (standard) first-order logical operations of A. Then $A < B$ if $\#A < \#B$, where the latter denotes the lexicographical ordering on $\mathbb{N} \times \mathbb{N}$ (w.r.t. the standard 'smaller than' ordering on the natural numbers). We treat the following main cases (for notational convenience M denotes the model M_Γ).

- Let $M, s \models A$, where A denotes the formula $(p * q)@\phi$. By Lemma 1 there exist ϕ_1 and ϕ_2 such that $M, s \models \phi = \phi_1 \uplus \phi_2$, $M, s \models p@\phi_1$ and $M, s \models q@\phi_2$. From the induction hypothesis it follows that $ps@\phi_1, qs@\phi_2, \phi = \phi_1 \uplus \phi_2 \in \Gamma$ (note that the first-order formula $\phi = \phi_1 \uplus \phi_2$ does not contain free variables, and thus is not affected by the substitution s). So we derive by rule $\mathbf{R}_*$ that $\Gamma \vdash (ps * qs)@\phi$. By maximal consistency of Γ, we then conclude that $(ps * qs)@\phi \in \Gamma$, that is, $As \in \Gamma$.
  On the other hand, let $As \in \Gamma$. That is, $(ps * qs)@\phi \in \Gamma$. By construction $\phi = R_1 \uplus R_2, ps@R_1, qs@R_2 \in \Gamma$, for some witnesses R_1 and R_2. By the induction hypothesis it then follows that $M, s \models p@R_1$ and $M, s \models p@R_2$. Further, the induction hypothesis gives $M, s \models \phi = R_1 \uplus R_2$ (again, note that the formula $\phi = R_1 \uplus R_2$ has no free variables, and thus is not affected by the substitution s). We conclude by Lemma 1 that $M, s \models (p * q)@\phi$.
- Let $M, s \models A$, where A denotes the formula $(p \mathrel{-\!\!*} q)@\phi$. Suppose $As \notin \Gamma$. By the maximal consistency of Γ, we then have $\neg(ps \mathrel{-\!\!*} qs)@\phi \in \Gamma$. By

construction $R \perp \phi, ps@R, \neg qs@(\phi \vee R) \in \Gamma$, for some witness R, which contradicts $M, s \models (p \twoheadrightarrow q)@\phi$ (after application of the induction hypothesis and using Lemma 1 again).

On the other hand, let $As \in \Gamma$. To show that $M, s \models (p \twoheadrightarrow q)@\phi$, let $M, s \models \phi \perp \psi$ and $M, s \models p@\psi$, for some binary formula ψ. By the induction hypothesis we have that $\phi \perp \psi, ps@\psi \in \Gamma$. Suppose that $qs@(\phi \vee \psi) \notin \Gamma$, that is $\neg qs@(\phi \vee \psi) \in \Gamma$ (Γ is maximal consistent), and thus $\Gamma, qs@(\phi \vee \psi) \vdash \emptyset$. Applying rule $\mathbf{L}_{\twoheadrightarrow}$ we then derive $\Gamma, (ps \twoheadrightarrow qs)@\phi \vdash \emptyset$, which contradicts the consistency of Γ ($(ps \twoheadrightarrow qs)@\phi \in \Gamma$). So we have that $qs@(\phi \vee \psi) \in \Gamma$, that is, $M, s \models q@(\phi \vee \psi)$, by the induction hypothesis. Since ψ is chosen arbitrarily, it follows by Lemma 1 that $M, s \models (p \twoheadrightarrow q)@\phi$.

- Let A be a formula $p@\phi$, where p denotes a basic formula. Let $\mathcal{R} = Rel_M(\phi)$. We then have $M, s \models p@\phi$ iff (by definition)
  $M, \mathcal{R}, s \models p$ iff (straightforward induction on p)
  $M, s \models p[\phi/\hookrightarrow]$ iff (induction hypothesis for $p[\phi/\hookrightarrow]$)
  $ps[\phi/\hookrightarrow] \in \Gamma$ iff (by the points-to rules)
  $ps@\phi \in \Gamma$. Note that applying the substitution s to $p@\phi$ and $p[\phi/\hookrightarrow]$ results in $ps@\phi$ and $ps[\phi/\hookrightarrow]$. □

The downward Löwenheim-Skolem property follows. It should be noted that we cannot remove from the constructed model the binary relation symbols which are introduced as witnesses, as these determine the notion of first-order definability.

Theorem 2 (Completeness). *We have that $\Gamma \models \Delta$ implies $\Gamma \vdash \Delta$.*

Compactness follows. We thus derive (by Lindström's theorem [Vää10]) that this version of SL is as expressive as first-order logic.

6 Conclusion

We investigated the expressiveness of full SL over arbitrary first-order models. We have shown that restricting the quantification over first-order definable heaps gives rise to a semantic consequence relation that can be captured by a sound and complete extension of the standard sequent calculus for first-order logic.

The main question remains what is the exact relationship between full SL which allows for infinite heaps and second-order logic. In [KR04] a translation is given of general second-order logic in a first-order logic with *spatial conjunction*. Spatial conjunction (as defined in [KR04]) allows to split a global set of *arbitrary* relations. As such it goes beyond the *local* scope of separating conjunction which is restricted to the points-to relation. We conjecture that second-order logic is strictly more expressive than full SL.

Acknowledgements. The authors thank the anonymous referees for providing many constructive and useful suggestions for improvement.

References

[AH21] Armborst, L., Huisman, M.: Permission-based verification of red-black trees and their merging. In: 2021 IEEE/ACM 9th International Conference on Formal Methods in Software Engineering (FormaliSE), pp. 111–123. IEEE (2021)

[BDL12] Brochenin, R., Demri, S., Lozes, E.: On the almighty wand. Inf. Comput. **211**, 106–137 (2012)

[CK13] Chang, C.C., Keisler, H.J.: Model Theory: Third Edition. Dover Books on Mathematics. Dover Publications (2013)

[Cro17] Crosilla, L.: Predicativity and Feferman. In: Jäger, G., Sieg, W. (eds.) Feferman on Foundations: Logic, Mathematics, Philosophy. OCL, vol. 13, pp. 423–447. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63334-3_15

[CYO01] Calcagno, C., Yang, H., O'Hearn, P.W.: Computability and complexity results for a spatial assertion language for data structures. In: Hariharan, R., Vinay, V., Mukund, M. (eds.) FSTTCS 2001. LNCS, vol. 2245, pp. 108–119. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45294-X_10

[dBHdG23] de Boer, F., Hiep, H.-D., de Gouw, S.: Dynamic separation logic. In: Mathematical Foundations of Programming Semantics (MFPS) (2023, to appear)

[DD16] Demri, S., Deters, M.: Expressive completeness of separation logic with two variables and no separating conjunction. ACM Trans. Comput. Log. **17**(2), 12 (2016)

[DLM21] Demri, S., Lozes, É., Mansutti, A.: A complete axiomatisation for quantifier-free separation logic. Log. Methods Comput. Sci. **17**(3) (2021)

[EIP20] Echenim, M., Iosif, R., Peltier, N.: The Bernays-Schönfinkel-Ramsey class of separation logic with uninterpreted predicates. ACM Trans. Comput. Log. **21**(3), 19:1–19:46 (2020)

[GM10] Galmiche, D., Méry, D.: Tableaux and resource graphs for separation logic. J. Log. Comput. **20**(1), 189–231 (2010)

[Hen49] Henkin, L.: The completeness of the first-order functional calculus. J. Symb. Log. **14**(3), 159–166 (1949)

[Hen50] Henkin, L.: Completeness in the theory of types. J. Symb. Logic **15**(2), 81–91 (1950)

[HH14] Huet, G.P., Herbelin, H.: 30 years of research and development around Coq. In: Jagannathan, S., Sewell, P. (eds.) The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2014, San Diego, CA, USA, 20–21 January 2014, pp. 249–250. ACM (2014)

[HT16] Hóu, Z., Tiu, A.: Completeness for a first-order abstract separation logic. In: Igarashi, A. (ed.) APLAS 2016. LNCS, vol. 10017, pp. 444–463. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47958-3_23

[JKJ+18] Jung, R., Krebbers, R., Jourdan, J.-H., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: a modular foundation for higher-order concurrent separation logic. J. Funct. Program. **28** (2018)

[KR04] Kuncak, V., Rinard, M.C.: On spatial conjunction as second-order logic. CoRR, cs.LO/0410073 (2004)

[Kri08] Krishnaswami, N.R.: A modal sequent calculus for propositional separation logic (2008)

[Man96] Manzano, M.: Extensions of First-Order Logic, vol. 19. Cambridge University Press, Cambridge (1996)

[MRH22] Monti, R.E., Rubbens, R., Huisman, M.: On deductive verification of an industrial concurrent software component with VerCors. In: Margaria, T., Steffen, B. (eds.) ISoLA 2022. LNCS, vol. 13701, pp. 517–534. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19849-6_29

[MSS16] Müller, P., Schwerhoff, M., Summers, A.J.: Viper: a verification infrastructure for permission-based reasoning. In: Jobstmann, B., Leino, K.R.M. (eds.) VMCAI 2016. LNCS, vol. 9583, pp. 41–62. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49122-5_2

[Pym02] Pym, D.J.: The semantics and proof theory of the logic of bunched implications. In: Applied Logic Series (2002)

[Rey00] Reynolds, J.C.: Intuitionistic reasoning about shared mutable data structure. In: Davies, J., Roscoe, B., Woodcock, J. (eds.) Millennial Perspectives in Computer Science, Cornerstones of Computing, pp. 303–321. Macmillan Education (2000)

[Rey02] Reynolds, J.C.: Separation logic: a logic for shared mutable data structures. In: Proceedings of the 17th IEEE Symposium on Logic in Computer Science (LICS 2002), Copenhagen, Denmark, 22–25 July 2002, pp. 55–74. IEEE Computer Society (2002)

[Rey05] Reynolds, J.C.: An overview of separation logic. In: Meyer, B., Woodcock, J. (eds.) VSTTE 2005. LNCS, vol. 4171, pp. 460–469. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-69149-5_49

[RISK16] Reynolds, A., Iosif, R., Serban, C., King, T.: A decision procedure for separation logic in SMT. In: Artho, C., Legay, A., Peled, D. (eds.) ATVA 2016. LNCS, vol. 9938, pp. 244–261. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46520-3_16

[SNPR+19] Sighireanu, M., et al.: SL-COMP: competition of solvers for separation logic. In: Beyer, D., Huisman, M., Kordon, F., Steffen, B. (eds.) TACAS 2019. LNCS, vol. 11429, pp. 116–132. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17502-3_8

[Vää01] Väänänen, J.: Second-order logic and foundations of mathematics. Bull. Symb. Logic **7**(4), 504–520 (2001)

[Vää10] Väänänen, J.: Lindström's theorem. Universal Logic: An Anthology, pp. 231–236 (2010)

[Yan01] Yang, H.: Local reasoning for stateful programs. Ph.D. thesis, University of Illinois at Urbana-Champaign. (Technical Report UIUCDCS-R-2001-2227) (2001)