

Separation Logic with One Quantified Variable

D. Larchey-Wendling²

Joint work with S. Demri¹, D. Galmiche² and D. Méry²

¹ NYU – CNRS and ² LORIA – CNRS – UL

CSR, June 2014

Overview

- 1 Separation Logic in a Nutshell
- 2 Separation Logic 1SL1
- 3 Expressive Completeness
- 4 Some remarks on MC and SAT

Separation logic

- Introduced by Ishtiaq, Reynolds, O'Hearn, Pym.
- Extension of Hoare Logic by J.C. Reynolds with separating connectives.
- Reasoning about the heap with a strong form of locality built-in.
- $\phi * \psi$ is true whenever the heap can be divided into two disjoint parts, one satisfies ϕ , the other one ψ .
- $\phi \text{-} * \psi$ is true whenever ϕ is true for a (fresh) disjoint heap, ψ is true for the combined heap.

Hoare triples

- Hoare triple: $\{\phi\} \text{ PROG } \{\psi\}$ (total correctness).
- Rule of constancy:

$$\frac{\{\phi\} \text{ PROG } \{\psi\}}{\{\phi \wedge \psi'\} \text{ PROG } \{\psi \wedge \psi'\}}$$

where no variable free in ψ' is modified by PROG .

- Unsoundness of the rule of constancy with pointers:

$$\frac{\{(\exists z. x \mapsto z)\} [x] := 4 \{x \mapsto 4\}}{\{(\exists z. x \mapsto z) \wedge y \mapsto 3\} [x] := 4 \{x \mapsto 4 \wedge y \mapsto 3\}}$$

(when $x = y$)

$x \mapsto z$: “memory has a unique memory cell $x \mapsto z$ ”

When separation logic enters into the play

- Reparation with frame rule:

$$\frac{\{\phi\} \text{ PROG } \{\psi\}}{\{\phi * \psi'\} \text{ PROG } \{\psi * \psi'\}}$$

where no variable free in ψ' is modified by PROG.

- Strengthening precedent (SP)

$$\frac{\phi \Rightarrow \psi' \quad \{\psi'\} \text{ PROG } \{\psi\}}{\{\phi\} \text{ PROG } \{\psi\}}$$

- Checking entailment/validity/satisfiability in separation logic is a building block of the verification process.

Memory states for n SL (n record fields)

- Program variables $PVAR = \{x_1, x_2, x_3, \dots\}$.
- Memory state:
 - Store $\mathfrak{s} : PVAR \rightarrow Val$.
 - Heap $\mathfrak{h} : Loc \rightarrow Val^n$ with finite domain.
($Loc = \{l, l', \dots\}$, $Val = \mathbb{N} \uplus Loc \uplus \{nil\}$)
- Simplification: $Loc = Val = \mathbb{N}$ (like low level memory).
- Disjoint heaps: $dom(\mathfrak{h}_1) \cap dom(\mathfrak{h}_2) = \emptyset$ (noted $\mathfrak{h}_1 \perp \mathfrak{h}_2$).
- When $\mathfrak{h}_1 \perp \mathfrak{h}_2$, $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \stackrel{\text{def}}{=} \mathfrak{h}_1 \uplus \mathfrak{h}_2$.

Syntax and semantics for $n\text{SL}$

- Quantified variables $F\text{VAR} = \{u_1, u_2, u_3, \dots\}$.
- Expressions: $e ::= x_i \mid u_j$
- Atomic formulae: $\pi ::= e = e' \mid e \hookrightarrow e_1, \dots, e_n \mid \text{emp}$
- Formulae in $n\text{SL}$

$$\phi ::= \perp \mid \pi \mid \phi \wedge \psi \mid \neg\phi \mid \phi * \psi \mid \phi \text{-*} \psi \mid \exists u_j \phi$$

- $(s, h) \models_f \text{emp} \stackrel{\text{def}}{\iff} \text{dom}(h) = \emptyset$.
- $(s, h) \models_f e = e' \stackrel{\text{def}}{\iff} [e] = [e']$, with $[x_i] \stackrel{\text{def}}{=} s(x_i)$ and $[u_j] \stackrel{\text{def}}{=} f(u_j)$.
- $(s, h) \models_f e \hookrightarrow e_1, \dots, e_n \stackrel{\text{def}}{\iff} h([e]) = ([e_1], \dots, [e_n])$.

Semantics for nSL

- $(s, h) \models_f \phi_1 * \phi_2 \stackrel{\text{def}}{\Leftrightarrow} h = h_1 \boxplus h_2, (s, h_1) \models_f \phi_1, (s, h_2) \models_f \phi_2$
for some h_1, h_2 .
- $(s, h) \models_f \phi_1 \multimap \phi_2 \stackrel{\text{def}}{\Leftrightarrow}$ for all h' , if $h \perp h'$ and $(s, h') \models_f \phi_1$
then $(s, h \boxplus h') \models_f \phi_2$.
- $(s, h) \models_f \exists u_j \phi \stackrel{\text{def}}{\Leftrightarrow}$ there is $l \in \mathbb{N}$ such that $(s, h) \models_{f'} \phi$
where $f' = f[u_j \mapsto l]$ is the assignment equal to f except
that u_j takes the value l .
- Satisfiability problem:
 - input:** formula ϕ in nSL
 - question:** are there (s, h) and f such that $(s, h) \models_f \phi$?

Satisfiability in fragments of $n\text{SL}$

- $n\text{SL}$: n record fields, unrestricted quantification
- $n\text{SL}_i$: n record fields, at most i quantified variables
- $n\text{SL}_0$ decidable and PSPACE-complete [Calcagno et al., 01]
- $n\text{SL}$ undecidable for $n \geq 2$, by encoding finitary SAT of classical logic with a single binary relation [Calcagno et al., 01]
- 1SL and $1\text{SL}(-*)$ undecidable [Brochenin, Demri & Lozes 08] by reduction to WSOL
- 1SL_2 undecidable [Demri & Deters, submitted] by reduction to Minsky machines
- Our focus is on 1SL_1 : decidability and complexity


Summary of our contributions on 1SL1

- 1SL1 = one record, one quantified var., q program vars.
- decomposition of heaps: core, loops, predecessors...
- given a bound α , a finite set of test formulae Test_α
 - test the structure of the core + cardinality constraints
 - SAT of Boolean comb. of Test_α is NP-complete
- if two heaps cannot be distinguished by Test_α , they cannot be distinguished by any ϕ s.t. $\text{th}(q, \phi) \leq \alpha$
- ϕ (with $\text{th}(q, \phi) \leq \alpha$) equiv. to Bool. comb. of Test_α
- model check w.r.t. equiv. classes of heaps (w.r.t. Test_α)
- give an abstract representation for these classes
- PSPACE algorithm for abstract MC and SAT

Separation Logic 1SL1

Memory states (one field)

- Memory state (s, h) :
 - Store $s : \text{PVAR} \rightarrow \mathbb{N}$.
 - Heap $h : \mathbb{N} \rightarrow \mathbb{N}$ with finite domain.
Graph of a unary function with finite domain.



Specialization for 1SL1

(one field, one quantified variable)

- Expressions: $e ::= x_i \mid \boxed{u}$
- Atomic formulae: $\pi ::= e = e' \mid \boxed{e \hookrightarrow e'} \mid \text{emp}$
- Formulae in 1SL1

$$\phi ::= \perp \mid \pi \mid \phi \wedge \psi \mid \neg \phi \mid \phi * \psi \mid \phi \multimap \psi \mid \boxed{\exists u \phi}$$


- $(s, h) \models_{\perp} \text{emp} \stackrel{\text{def}}{\iff} \text{dom}(h) = \emptyset.$
- $(s, h) \models_{\perp} e = e' \stackrel{\text{def}}{\iff} [e] = [e'], \text{ with } [x_i] \stackrel{\text{def}}{=} s(x_i) \text{ and } [u] \stackrel{\text{def}}{=} \perp.$
- $(s, h) \models_{\perp} e \hookrightarrow e' \stackrel{\text{def}}{\iff} [e] \in \text{dom}(h) \text{ and } h([e]) = [e'].$

Semantics for 1SL1


- $(s, h) \models_l \phi_1 * \phi_2 \stackrel{\text{def}}{\iff} h = h_1 \uplus h_2, (s, h_1) \models_l \phi_1, (s, h_2) \models_l \phi_2$
for some h_1, h_2 .
- $(s, h) \models_l \phi_1 \multimap \phi_2 \stackrel{\text{def}}{\iff}$ for all h' , if $h \perp h'$ and $(s, h') \models_l \phi_1$ then $(s, h \uplus h') \models_l \phi_2$.
- $(s, h) \models_l \exists u \phi \stackrel{\text{def}}{\iff}$ there is $l' \in \mathbb{N}$ such that $(s, h) \models_{l'} \phi$.
- Satisfiability problem:
 - input:** formula ϕ in 1SL1
 - question:** are there (s, h) and l such that $(s, h) \models_l \phi$?
- Between 1SL0 (PSPACE) and 1SL2 (undecidable)

Simple properties stated in 1SL1

- The domain of the heap has at least k elements:
 $\neg \text{emp} * \dots * \neg \text{emp}$ (k times).
- The variable x_i is allocated in the heap:
 $\text{alloc}(x_i) \stackrel{\text{def}}{=} (x_i \leftrightarrow x_i) * \perp$.
- The variable x_i points to a location that is a loop:
 $\text{toLoop}(x_i) \stackrel{\text{def}}{=} \exists u (x_i \leftrightarrow u \wedge u \leftrightarrow u)$.




- The variable x_i points to a location that is allocated:
 $\text{toAlloc}(x_i) \stackrel{\text{def}}{=} \exists u (x_i \leftrightarrow u \wedge \text{alloc}(u))$.



More properties


- Variables x_i and x_j point to a shared location:

$$\text{conv}(x_i, x_j) \stackrel{\text{def}}{=} \exists u (x_i \hookrightarrow u \wedge x_j \hookrightarrow u).$$



- there is a location between x_i and x_j :

$$\text{inbetween}(x_i, x_j) \stackrel{\text{def}}{=} \exists u (x_i \hookrightarrow u \wedge u \hookrightarrow x_j).$$




What Else?

Partition one: loops, predecessors, etc.

- $\text{pred}(\mathfrak{s}, \mathfrak{h}) \stackrel{\text{def}}{=} \bigcup_i \text{pred}(\mathfrak{s}, \mathfrak{h}, i)$ with
 $\text{pred}(\mathfrak{s}, \mathfrak{h}, i) \stackrel{\text{def}}{=} \{\ell' : \mathfrak{h}(\ell') = \mathfrak{s}(x_i)\}$ for every $i \in [1, q]$.
- $\text{loop}(\mathfrak{s}, \mathfrak{h}) \stackrel{\text{def}}{=} \{\ell \in \text{dom}(\mathfrak{h}) : \mathfrak{h}(\ell) = \ell\}$.
- $\text{rem}(\mathfrak{s}, \mathfrak{h}) \stackrel{\text{def}}{=} \text{dom}(\mathfrak{h}) \setminus (\text{pred}(\mathfrak{s}, \mathfrak{h}) \cup \text{loop}(\mathfrak{s}, \mathfrak{h}))$.
- $\boxed{\text{dom}(\mathfrak{h}) = \text{rem}(\mathfrak{s}, \mathfrak{h}) \uplus (\text{pred}(\mathfrak{s}, \mathfrak{h}) \cup \text{loop}(\mathfrak{s}, \mathfrak{h}))}$.

Partition two: introducing the core

- $\text{ref}(s, h) \stackrel{\text{def}}{=} \text{dom}(h) \cap s(\mathcal{V})$; $\text{acc}(s, h) \stackrel{\text{def}}{=} \text{dom}(h) \cap h(s(\mathcal{V}))$.
- $\heartsuit(s, h) \stackrel{\text{def}}{=} \text{ref}(s, h) \cup \text{acc}(s, h)$; $\overline{\heartsuit}(s, h) \stackrel{\text{def}}{=} \text{dom}(h) \setminus \heartsuit(s, h)$.



Locations outside of the core

- Locations in the core are easy to identify thanks to program variables.
- $\text{pred}_{\heartsuit}(\mathfrak{s}, \mathfrak{h}, i) \stackrel{\text{def}}{=} \text{pred}(\mathfrak{s}, \mathfrak{h}, i) \setminus \heartsuit(\mathfrak{s}, \mathfrak{h})$.
- $\text{loop}_{\heartsuit}(\mathfrak{s}, \mathfrak{h}) \stackrel{\text{def}}{=} \text{loop}(\mathfrak{s}, \mathfrak{h}) \setminus \heartsuit(\mathfrak{s}, \mathfrak{h})$.
- $\text{rem}_{\heartsuit}(\mathfrak{s}, \mathfrak{h}) \stackrel{\text{def}}{=} \text{rem}(\mathfrak{s}, \mathfrak{h}) \setminus \heartsuit(\mathfrak{s}, \mathfrak{h})$.
- $\text{dom}(\mathfrak{h}) = \heartsuit(\mathfrak{s}, \mathfrak{h}) \uplus \text{pred}_{\heartsuit}(\mathfrak{s}, \mathfrak{h}) \uplus \text{loop}_{\heartsuit}(\mathfrak{s}, \mathfrak{h}) \uplus \text{rem}_{\heartsuit}(\mathfrak{s}, \mathfrak{h})$.

Test formulae

- Equality $\stackrel{\text{def}}{=} \{x_i = x_j \mid i, j \in [1, q]\}$.
- Pattern $\stackrel{\text{def}}{=} \{x_i \hookrightarrow x_j, \text{conv}(x_i, x_j), \text{inbetween}(x_i, x_j) \mid i, j \in [1, q]\} \cup \{\text{toalloc}(x_i), \text{toloop}(x_i), \text{alloc}(x_i) \mid i \in [1, q]\}$.
- Extra^u $\stackrel{\text{def}}{=} \{u \hookrightarrow u, \text{alloc}(u)\} \cup \{x_i = u, x_i \hookrightarrow u, u \hookrightarrow x_i \mid i \in [1, q]\}$.
- Size _{α} $\stackrel{\text{def}}{=} \{\#\text{pred}_{\heartsuit}^i \geq k \mid i \in [1, q], k \in [1, \alpha]\} \cup \{\#\text{loop}_{\heartsuit} \geq k, \#\text{rem}_{\heartsuit} \geq k \mid k \in [1, \alpha]\}$.
- Test _{α} ^u $\stackrel{\text{def}}{=} \text{Equality} \cup \text{Pattern} \cup \text{Size}_{\alpha} \cup \text{Extra}^u \cup \{\perp\}$.

Counting loops outside of the core

- Needed for expressing test formulae in 1SL1 !

- $T \stackrel{\text{def}}{=} \{\text{alloc}(x_1), \dots, \text{alloc}(x_q)\} \cup \{\text{toalloc}(x_1), \dots, \text{toalloc}(x_q)\}.$

- $f: T \rightarrow \{0, 1\}.$

$$\phi_f \stackrel{\text{def}}{=} \bigwedge \{\psi \mid \psi \in T \text{ and } f(\psi) = 1\} \wedge \bigwedge \{\neg\psi \mid \psi \in T \text{ and } f(\psi) = 0\}$$

- $\# \text{loop}_{\heartsuit} \geq k \stackrel{\text{def}}{=} \bigvee_f \phi_f \wedge \left(\phi_f^{\text{pos}} * (\# \text{loop} \geq k) \right)$ with
 - ϕ_f^{pos} = the positive part of ϕ_f .
 - $\# \text{loop} \geq k \stackrel{\text{def}}{=} (\exists u \ u \leftrightarrow u) * \dots * (\exists u \ u \leftrightarrow u)$ (k times).

Deciding satisfiability for test formulae

- Satisfiability of conjunctions of $\text{Test}_\alpha^u / \neg \text{Test}_\alpha^u$ can be checked in polynomial time (with bounds in binary).
- Polynomial-time decision based on a saturation algorithm (see rules)

$$\frac{\phi \vdash x_j \leftrightarrow x \quad \phi \vdash x \leftrightarrow y \quad \phi \vdash x = y}{\phi \vdash \text{toloop}(x_i)}$$

$$\frac{\phi \vdash \text{conv}(x_i, x_j) \quad \phi \vdash \text{toloop}(x_i)}{\phi \vdash \text{toloop}(x_j)}$$

$$\frac{\phi \vdash \neg \text{alloc}(x_i)}{\phi \vdash \neg \text{toloop}(x_i)}$$

- Satisfiability problem for Boolean combinations of test formulae in the set $\bigcup_{\alpha \geq 1} \text{Test}_\alpha^u$ is NP-complete.

Expressive Completeness

Memory threshold

- for any formula of 1SL1 with at most q program variables
- $\text{th}(\mathbf{q}, \phi) \stackrel{\text{def}}{=} 1$ for every atomic formula ϕ .
- $\text{th}(\mathbf{q}, \phi_1 \wedge \phi_2) \stackrel{\text{def}}{=} \max(\text{th}(\mathbf{q}, \phi_1), \text{th}(\mathbf{q}, \phi_2))$.
- $\text{th}(\mathbf{q}, \neg\phi_1) \stackrel{\text{def}}{=} \text{th}(\mathbf{q}, \phi_1)$ and $\text{th}(\mathbf{q}, \exists u \phi_1) \stackrel{\text{def}}{=} \text{th}(\mathbf{q}, \phi_1)$.
- $\text{th}(\mathbf{q}, \phi_1 * \phi_2) \stackrel{\text{def}}{=} \text{th}(\mathbf{q}, \phi_1) + \text{th}(\mathbf{q}, \phi_2)$.
- $\text{th}(\mathbf{q}, \phi_1 \multimap \phi_2) \stackrel{\text{def}}{=} \mathbf{q} + \max(\text{th}(\mathbf{q}, \phi_1), \text{th}(\mathbf{q}, \phi_2))$.
- For all ϕ built over $\{x_1, \dots, x_q\}$, $1 \leq \text{th}(\mathbf{q}, \phi) \leq \mathbf{q} \times |\phi|$.

α -equivalence, correctness of abstraction

- α -equivalence: indistinguishability with respect to test formula $\psi \in \text{Test}_\alpha^u$:

$$(\mathfrak{s}, \mathfrak{h}, \mathfrak{l}) \simeq_\alpha (\mathfrak{s}', \mathfrak{h}', \mathfrak{l}') \text{ whenever } (\mathfrak{s}, \mathfrak{h}) \models_{\mathfrak{l}} \psi \text{ iff } (\mathfrak{s}', \mathfrak{h}') \models_{\mathfrak{l}'} \psi$$

- Cardinality constraints are precise up to α .

$$\text{if } \boxed{(\mathfrak{s}, \mathfrak{h}, \mathfrak{l}) \simeq_\alpha (\mathfrak{s}', \mathfrak{h}', \mathfrak{l}')}$$

then

$$\boxed{(\mathfrak{s}, \mathfrak{h}) \models_{\mathfrak{l}} \phi \text{ iff } (\mathfrak{s}', \mathfrak{h}') \models_{\mathfrak{l}'} \phi}$$

for any ϕ s.t. $\text{th}(\mathfrak{q}, \phi) \leq \alpha$

- Hence formulae of threshold below α do not distinguish more memory states than those formulae in Test_α^u

Quantifier elimination

- Any ϕ in 1SL1 (with q program variables) is equivalent to a Boolean combination ϕ' of test formulae in $\text{Test}_{\text{th}(\mathbf{q}, \phi)}^u$.
- $\alpha = \text{th}(\mathbf{q}, \phi)$.
- $\mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l}) \stackrel{\text{def}}{=} \left[\begin{array}{l} \{\psi \mid \psi \in \text{Test}_{\alpha}^u \text{ and } (\mathfrak{s}, \mathfrak{h}) \models_{\mathfrak{l}} \psi\} \\ \cup \{\neg\psi \mid \psi \in \text{Test}_{\alpha}^u \text{ and } (\mathfrak{s}, \mathfrak{h}) \not\models_{\mathfrak{l}} \psi\} \end{array} \right]$
- Finiteness of Test_{α}^u entails $\mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l})$ is finite and $\bigwedge \mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l})$ is a well-defined atom.
- $(\mathfrak{s}', \mathfrak{h}') \models_{\mathfrak{l}'} \bigwedge \mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l})$ iff $(\mathfrak{s}, \mathfrak{h}, \mathfrak{l}) \simeq_{\alpha} (\mathfrak{s}', \mathfrak{h}', \mathfrak{l}')$.
 $\mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l})$ characterizes $(\mathfrak{s}, \mathfrak{h}, \mathfrak{l})$ up to α .
- $\bigwedge \mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l})$ spans a finite domain.
- $\phi' \stackrel{\text{def}}{=} \bigvee \{\bigwedge \mathcal{S}(\mathfrak{s}, \mathfrak{h}, \mathfrak{l}) \mid (\mathfrak{s}, \mathfrak{h}) \models_{\mathfrak{l}} \phi\}$ equivalent to ϕ .

non-constructive proof !

Corollaries

- Any satisfiable ϕ in 1SL1 has a polynomial-size model.
- 1SL2 is strictly more expressive than 1SL1.
- Test_α^u formulae cannot distinguish the two models below

$$x_1 \rightarrow \bullet \rightarrow \bullet \rightarrow x_2 \quad | \quad x_1 \rightarrow \bullet \rightarrow \bullet \quad \circ \rightarrow x_2$$

- hence neither can 1SL1.
- but 1SL2 can: $\exists u \exists v (x_1 \leftrightarrow u \wedge u \leftrightarrow v \wedge v \leftrightarrow x_2)$

Some remarks on MC and SAT


MC and SAT in 1SL1

- to check $(s, h) \models_I \phi_1 * \phi_2$ we need to verify:

$$(s, h') \not\models_I \phi_1 \text{ or } (s, h \boxplus h') \models_I \phi_2 \quad \boxed{\text{for any}} \quad h' \perp h$$

- $(s, \emptyset) \models_I \neg(\top * \neg\phi)$ iff there exists h s.t. $(s, h) \models_I \phi$.
- $(\exists h, (s, h) \models_I \top * (\text{emp} \wedge \phi))$ iff $(s, \emptyset) \models_I \phi$
- hence (MC) \iff (SAT) in SL.
- for MC: transform the $\boxed{\text{for any}}$ into finite quantification
- indeed, given α , the test formula Test_α^u
 - are finitely many, as well as their Boolean combinations
 - hence only finitely many classes for $(s, h, l) \simeq_\alpha (s', h', l')$
- any formula s.t. $\text{th}(q, \phi) \leq \alpha$, the value of $(s, h) \models_I \phi$ only depends of the class of (s, h, l)
- transform (infinite) "for any" into (finite) "for any class"

Abstract memory states \approx atoms of Test_α^u



$$l = 2, \tau = 2, p_1 = 1, p_2 = p_3 = p_4 = 0.$$

Abstract memory state: $\alpha = ((V, E), l, \tau, p_1, \dots, p_q)$.

$V_{\text{par}} \subseteq V$ partition of $\{x_1, \dots, x_q\}$.

Abstract Model Checking in 1SL1

- we then prove that abstraction "commutes" with MC
- we describe abstract composition/decomposition of heaps
- we present a MC algorithm on abstract memory states
- this MC algorithm runs in PSPACE
- PSPACE-hardness already holds for 1SL0
- hence MC in 1SL1 is PSPACE-complete
- the same complexity holds for SAT

Concluding remarks

- Quantifier elimination property for 1SL1 formulae.
- Conjunction of test formulae decidable in polynomial time.
- Satisfiability and model-checking problems for 1SL1 are PSPACE-complete.
- Also, restriction to q program variables in polynomial time.
- Possible extension with $k > 1$ record fields.

