

Advances in Programming Languages

Lecture 14: Separation Logic

Ian Stark

School of Informatics
The University of Edinburgh

Monday 5 November 2018
Semester 1 Week 8



Outline

- 1 Opening
- 2 Propositions in Separation Logic
- 3 Hoare Triples in Separation Logic
- 4 Closing

Outline

- 1 Opening
- 2 Propositions in Separation Logic
- 3 Hoare Triples in Separation Logic
- 4 Closing

Topic: Augmented Languages for Correctness and Certification

The next block of lectures cover some language techniques and tools for improving program correctness:

- Specification and Verification
- Practical Java tools for Correctness
- Separation Logic
- Augmented Programming and Certifying Correctness

The focus here is not necessarily on changing what a program does, or making it do that thing faster, or using less memory, or less power. Instead we want to make sure that what it does is the right thing to do.

This lecture introduces an enrichment of Hoare Logic that supports local reasoning about code that works with linked-pointer data structures in memory.

Homework from Monday

JML is just one of many frameworks for specifying and verifying code in various programming languages.

For example: Java annotations, FindBugs™, QuickCheck, the C specification language ACSL and Frama-C platform, Spec#, ...

Your homework is to use a small part of one such system.

1. Read This

Find out about *Type Annotations and Pluggable Type Systems* by reading the Oracle Java Tutorial lesson with that name.

2. Do this

Try out the `@NonNull` and `@Nullable` annotations on the *Checker Framework* live demo:

- <http://eisop.uwaterloo.ca/live>

Outline

1 Opening

2 Propositions in Separation Logic

3 Hoare Triples in Separation Logic

4 Closing

Plan of Work

- The store/heap model for linked datastructures
- Decomposing a linked heap into heaplets
- Assertions about the store and heap
- New operations: star * and magic wand \rightarrow^*
- Representing datastructures with this logic

Store

A Hoare triple $\{P\} C \{Q\}$ involves assertions P and Q which talk about the *state* of an imperative program.

$$\{a > 3\} \text{ b := a+a } \{b > 6\}$$

$$\{(d > z) \wedge (d' > z) \wedge (z \geq 0)\} \text{ c := d*d' } \{c > z^2\}$$

$$\{\text{true}\} \text{ while } i > 0 \text{ do } i := i-1 \{i \leq 0\}$$

In these examples the state is the values of the variables mentioned — a , b , c , d , d' and i . This is usually described as an *environment* or *store*, formalised as a mapping s from variable names to stored values.

$$s \in \text{Store} \stackrel{\text{def}}{=} \text{VarName} \rightarrow \text{Int}$$

Heap

Most imperative languages also provide larger and more complex datastructures than a simple store of variables. Things like objects, linked lists, trees, or graphs are typically implemented using a *heap*: an area of memory over which either the programmer or the language runtime has very flexible control. The following lists some distinctive characteristics of the heap.

- The heap contains cells that can hold values.
- Each cell has an index or address that uniquely identifies its location in the heap.
- Cells in consecutive locations may be grouped into structured records.
- Cells may contain index values (*pointers*) that identify other cells in the heap.
- Variables in the store can also contain pointers.
- Records and pointers can be used to build complex linked datastructures.
- Portions of the heap can be arbitrarily allocated and deallocated as a program executes.

Modelling the Heap

To model the behaviour of a heap, we can combine the store s with another partial map h recording which heap cells are allocated and what values they hold. For simplicity, we'll assume cells have integer addresses and contain integer values.

$$s \in \text{Store} \stackrel{\text{def}}{=} \text{VarName} \rightarrow \text{Int}$$

$$h \in \text{Heap} \stackrel{\text{def}}{=} \text{Int} \rightarrow \text{Int}$$

$$(s, h) \in \text{State} \stackrel{\text{def}}{=} \text{Store} \times \text{Heap}$$

The *domain* of the map h , written $\text{dom}(h)$, indicates which heap cells have been allocated. It's possible to have *dangling pointers* in both the heap and the store, where an address points to a heap cell that's not currently in use.

Operations on Heaps and Heaplets

A heap is partial, as only some cells will be allocated, and may include dangling pointers.

$$h \in \text{Heap} \stackrel{\text{def}}{=} \text{Int} \rightarrow \text{Int}$$

Two heaps h_1 and h_2 are *disjoint* $h_1 \perp h_2$ if they allocate non-overlapping sets of cells. In this case we can combine them to get their *union* $h_1 \cdot h_2$.

$$h_1 \perp h_2 \stackrel{\text{def}}{\iff} \text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$$

$$h_1 \cdot h_2 \stackrel{\text{def}}{=} h_1 \cup h_2 \text{ provided that } h_1 \perp h_2$$

Going in the other direction, it's possible to separate a larger heap h into smaller *heaplets* $h = h_1 \cdot h_2$, usually in many different ways. Note that the heaplets h_1 and h_2 may have dangling pointers crossing between them.

Assertions about Heaps

Empty	emp	The heap is empty, no cells are allocated.
Singleton	$e \mapsto e'$	The heap has a single allocated cell, at address e , containing value e' .
Separating Conjunction “Star”	$P * P'$	The heap can be split into two disjoint heaplets, with one satisfying P and the other satisfying P' .
Separating Implication “Magic Wand”	$P \multimap Q$	If the heap is combined with any disjoint heap h' that satisfies P , then the combined heap will satisfy Q .

Derived Assertions about Heaps

$$e \hookrightarrow e' \stackrel{\text{def}}{\iff} e \mapsto e' * \mathbf{true}$$

$$e \mapsto - \stackrel{\text{def}}{\iff} \exists x. e \mapsto x$$

$$e \hookrightarrow - \stackrel{\text{def}}{\iff} (e \mapsto -) * \mathbf{true}$$

$$e \mapsto e_1, \dots, e_n \stackrel{\text{def}}{\iff} (e \mapsto e_1) * ((e + 1) \mapsto e_2) \cdots * ((e + (n - 1)) \mapsto e_n)$$

$$e \hookrightarrow e_1, \dots, e_n \stackrel{\text{def}}{\iff} (e \mapsto e_1, \dots, e_n) * \mathbf{true}$$

Datastructures in Separation Logic

The language of separation logic makes it possible to describe properties of datastructures built using pointers in the heap: tuples, records, lists, trees, graphs, objects, ...

Lists

Assertion $\text{list } xs \ e$ describes a heap made of a linked list that starts at the cell pointed to by e and contains exactly the values in xs .

$$\begin{aligned} \text{list } [] \ i &\stackrel{\text{def}}{=} \mathbf{emp} \wedge (i = 0) \\ \text{list } (x::xs) \ i &\stackrel{\text{def}}{=} \exists j . (i \mapsto x, j) * (\text{list } xs \ j) \end{aligned}$$

The use of separating conjunction “ $*$ ” here means that every node in the list uses a distinct part of the heap. The end of the list is signalled by a null pointer (value 0).

Datastructures in Separation Logic

The language of separation logic makes it possible to describe properties of datastructures built using pointers in the heap: tuples, records, lists, trees, graphs, objects, ...

Binary Trees

Assertion $\text{tree } t \ e$ describes a heap containing a binary tree with structure matching t and root at the cell pointed to by e .

$$\begin{aligned} \text{tree } (\text{Leaf } x) \ i &\stackrel{\text{def}}{=} i \mapsto 0, x \\ \text{tree } (\text{Branch } (s, t)) \ i &\stackrel{\text{def}}{=} \exists j, k . (i \mapsto 1, j, k) * (\text{tree } s \ j) * (\text{tree } t \ k) \end{aligned}$$

Tags 0 and 1 distinguish between leaves and internal nodes. Again, the separating conjunction “*” means that the tree really is a tree: none of the subtrees overlap or share nodes.

Outline

- 1 Opening
- 2 Propositions in Separation Logic
- 3 Hoare Triples in Separation Logic**
- 4 Closing

Triples

In separation logic, a triple $\{P\} \text{ C } \{Q\}$ now has assertions P and Q that make statements about the whole state (s, h) not just the store s .

Assertions that don't depend on the heap are *pure* assertions; those that do depend on the contents of the heap are *spatial* assertions.

The same rules of inference apply as before, together with additional rules for the new heap operations. For example, the star and magic wand operators satisfy the following:

$$\frac{s, h \vdash P * (P \multimap Q)}{s, h \vdash Q}$$

Heap Inference Rules

Any language operation relating to the heap gets an appropriate inference rule.

$$\frac{}{\{\mathbf{emp}\} \ v := \mathbf{alloc}(e_1, \dots, e_n) \ \{v \mapsto e_1, \dots, e_n\}}$$

$$\frac{}{\{(e \mapsto -)\} \ [e] := e' \ \{e \mapsto e'\}}$$

$$\frac{}{\{e \mapsto -\} \ \mathbf{dispose} \ e \ \{\mathbf{emp}\}}$$

These can also be elaborated to account for the remaining heap.

$$\frac{}{\{(e \mapsto -) * r\} \ [e] := e' \ \{e \mapsto e' * r\}}$$

$$\frac{}{\{(e \mapsto -) * ((e \mapsto e') \multimap p)\} \ [e] := e' \ \{p\}}$$

Frame Rule

Crucially, separation logic makes it possible to reason *compositionally*, where *local reasoning* about use of the heap can be extended to reasoning about a complete program. Central to this is the *frame rule*:

$$\frac{\{P\} \text{ C } \{Q\}}{\{P * R\} \text{ C } \{Q * R\}} \quad \text{Where C does not modify any of } \text{dom}(R).$$

This *modularity* is one thing that makes it feasible to scale and automate separation logic to checking and verifying large programs.

Outline

1 Opening

2 Propositions in Separation Logic

3 Hoare Triples in Separation Logic

4 Closing

Applications

Program Analysis

Separation logic has been extensively applied in academic research and industrial practice. As an enrichment of Hoare logic, it's suitable for annotating programs or libraries, specifying desired behaviour, statically checking properties, or proving program correctness.

It's been used to prove correctness of numerous pointer-manipulation algorithms for efficient processing of complex data structures.

The *Space Invader* analyser checks pointer manipulation in C code; Microsoft's *SLayer* verifies correctness of device drivers; Facebook's *Infer* tool statically checks C, C++, Java and Objective C.

Other Directions

Separation logic has applications to reasoning about concurrent code, where $P * Q$ describes properties P and Q of distinct threads. There are higher-order versions of separation logic; *Hoare Type Theory* incorporating these ideas into a type system; and an mathematical *Logic of Bunched Implications* that provided the groundwork for separation logic.

1. Watch this

Peter O'Hearn: "Continuous Reasoning: Scaling the impact of formal methods"
Plenary talk at the Federated Logic Conference 2018.

Video: <https://is.gd/continuousreasoning>

More information: <https://www.floc2018.org/speaker/peter-ohearn>

2. Do this

Complete and submit your written assignment. If you have questions then send me mail or post on Piazza.