

Size Does Matter: Two Certified Abstractions to Disprove Entailment in Intuitionistic and Classical Separation Logic

Clément Hurlin
INRIA Sophia Antipolis –
Méditerranée and
University of Twente
clement.hurlin@inria.fr

François Bobot
LRI, Univ Paris-Sud and
INRIA Saclay – Île-de-France
bobot@lri.fr

Alexander J. Summers
Imperial College London
alexander.j.summers
@imperial.ac.uk

ABSTRACT

We describe an algorithm to disprove entailment between separation logic formulas. We abstract models of formulas by their size and check whether two formulas have models whose sizes are compatible. Given two formulas A and B that do not have compatible models, we can conclude that $A \not\vdash B$. We provide two different abstractions (of different precision) of models. Our algorithm is of interest wherever entailment checking is performed (such as in program verifiers).

1. INTRODUCTION

Separation logic [19] has recently made a dent in the verification of pointer manipulating programs. Its successes include simple pointer programs with parallelism [3], C programs [16], and object-oriented programs [9, 11]. Formulas in separation logic typically describe parts of the heap: in the context of program verifiers, method *preconditions* describe the heap space required by methods to execute correctly, and method *postconditions* describe the heap space passed back to the caller when the method returns. Separation logic comes in two flavours: *intuitionistic* separation logic is used to verify garbage-collected programs [18, 11] while *classical* separation logic is used to verify programs in which memory deallocation is performed manually [3, 16]. Contrary to intuitionistic separation logic, classical separation logic does not admit weakening, allowing one to reason about memory leaks.

We propose a novel algorithm to disprove entailment between formulas. We abstract formulas by the sizes of their possible models and use comparisons of these sizes to disprove entailment between formulas. Intuitively, to disprove an entailment $A \vdash B$, it suffices to show that there exists a model (i.e., a heap) of A whose size is smaller than the size of all models of B . We give two different ways of calculating sizes of models, which have differing complexities and precisions.

Our algorithm is of interest whenever entailment checking is performed, for example in program verifiers. In this particular context, since our algorithm's complexity is low, it can be used to quickly show that a method is not provable, perhaps because programmer-supplied specifications are incorrect.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IWACO '09, July 6 2009, Genova, Italy

Copyright 2009 ACM 978-1-60558-546-8/09/07 ...\$10.00.

2. BACKGROUND

We work with permission accounting separation logic based on that of Bornat et al. [5]. In particular, we reuse their structure of models to abstract over the permission models. A permission model \mathcal{M} consists of:

1. A commutative semigroup with total binary operator $+_{\mathcal{M}}$.
2. A minimal permission, m_0 , that satisfies the following property: $(\forall m)(m_0 +_{\mathcal{M}} m = m)$.
3. A maximal permission, m_W .
4. A total order on permissions, $\leq_{\mathcal{M}}$.
5. A subset of permissions which are defined to be *splittable*.
6. A pair of functions on permissions, *split1* and *split2*, which define the two parts into which a splittable permission may be split. These functions must satisfy the following property: $(\forall m)(\text{split1}(m) +_{\mathcal{M}} \text{split2}(m) = m)$.

The maximal permission m_W represents full (and exclusive) access to a location, whereas the minimal permission represents no access. Permissions in between these two extremes will allow read access but not write access to a location. Some (in some models, all) permissions are *splittable*: they can be divided into two smaller (in the sense of permission ordering) permissions. In particular, the maximal permission (which denotes exclusive access) may be split into smaller permissions, none of which allow full access to the location. However, permissions may also be *combined* (by addition): in this way it is possible for a full permission to be regained, if all partial permissions are combined together.

Later, we define heaps that only contain *valid* permissions. Valid permissions are permissions that can be obtained by splitting the maximal permission. As a corollary, any valid permission m satisfies the following inequality: $m_0 <_{\mathcal{M}} m \leq_{\mathcal{M}} m_W$. Our model, however, includes permissions that are bigger than m_W , because we need to sum permissions above m_W . That is why $+_{\mathcal{M}}$ is a total operator.

Our paper's results have been established for two concrete permissions models: Boyland's fractional permissions [6, 5] and the counting model of Bornat et al. [5].

In the fractional permission model, permissions are rational numbers. The binary operator and the ordering are the usual $+$ and \leq on rationals, the minimal permission is 0 and the maximal permission is 1. All permissions are splittable, and splitting simply divides permissions in half: $\text{split1}(m) = \text{split2}(m) = \frac{m}{2}$.

In the counting model, permissions were originally represented by integers, but we re-encode them in a different manner to allow the use of a standard order. We encode permissions as a pair

(n, i) , where n is a natural number (in fact, only 0 or 1), and i is an integer. Intuitively, the first of the pair indicates whether this permission can be *split*, to give out further (read) permissions, while the second indicates a number of read permissions (positive to represent read permissions which have been obtained, and negative to represent those given out). The minimal permission is $(0, 0)$, and the maximal permission is $(1, 0)$. Permissions are added by pairwise addition of their components, and ordered by the lexicographic ordering. A permission is splittable if its first component is 1. The split functions are defined by: $\text{split1}((1, i)) = (1, i - 1)$ and $\text{split2}((1, i)) = (0, 1)$.

We map Bornat's counting permissions onto our pairs model, by the following rules (in which n represents a positive integer):

$$0 \mapsto (1, 0) \quad n \mapsto (1, -n) \quad -n \mapsto (0, n)$$

We study both intuitionistic and classical separation logic [15]. Both flavours contains pure (heap independent) and spatial (heap dependent) formulas:

n	\in	\mathbb{N}	
a, v	\in	$n \mid n+n \mid n-n \mid \dots$	addresses and values
m			abstract permissions
Π	$::=$	$a = a \mid a \neq a \mid \text{true} \mid \dots$	pure formulas
Σ^I	$::=$	$a \overset{m}{\hookrightarrow} v$	intuitionistic spatial formulas
Σ^C	$::=$	$\text{emp} \mid a \overset{m}{\mapsto} v$	classical spatial formulas

In intuitionistic separation logic, the atomic spatial formula is the *harpoon* $a \overset{m}{\hookrightarrow} v$ while in classical separation logic, the atomic spatial formula is the *points-to* predicate $a \overset{m}{\mapsto} v$. The harpoon $a \overset{m}{\hookrightarrow} v$ has a dual meaning. Firstly, $a \overset{m}{\hookrightarrow} v$ asserts that the heap contains a cell at address a with content v . Secondly, $a \overset{m}{\hookrightarrow} v$ asserts permission m to the cell at address a . If $m = m_W$, $a \overset{m}{\hookrightarrow} v$ asserts write and read authorization to the cell at address a , otherwise it asserts *readonly* authorization. The points-to predicate $a \overset{m}{\mapsto} v$ has the same meaning as the harpoon but it enforces that the heap contains *only* the cell at address a . Intuitionistic (respectively classical) separation logic is obtained by taking Σ to be Σ^I (respectively Σ^C) below:

$$A, B ::= \Pi \mid \Sigma \mid A \star A \mid A \wedge A \mid A \vee A \quad \text{formulas}$$

In $A \star B$, \star is the *separating conjunction*. $A \star B$ represents a heap consisting of two *separate* subheaps A and B . $A \wedge B$ is the usual logical conjunction, which represents two different views of a heap. We do not include the standard implication \Rightarrow because this simplifies the situation by avoiding bunched contexts in the proof system [17]. The fact that most separation-logic-based program verifiers do not include a standard implication [3, 16, 11] supports our choice. In addition, we use a language without variables, but describe how variables interfere with our algorithm in Sec. 7.

Because our algorithms are defined in Coq [10], we represent heaps as lists (which are well supported by Coq's standard library). A list entry is a triplet of an address, a valid permission, and a value:

$$\begin{aligned} c &::= (a, m, v) && \text{heap cells} \\ h &::= [] \mid c :: h && \text{heaps} \end{aligned}$$

We define a projection operator to extract a value from a cell: $\text{val}(a, m, v) = v$ and we write $h[a]$ to denote h 's set of heap cells whose address are a :

$$\begin{aligned} [][a] &= \emptyset \\ ((a, m, v) :: h)[a'] &= \begin{cases} \{(a, m, v)\} \cup h[a] & \text{if } a = a' \\ h[a] & \text{otherwise} \end{cases} \end{aligned}$$

We write $h(a)$ to denote the sum of permissions to a occurring in h :

$$\begin{aligned} [](a) &= m_0 \\ ((a, m, v) :: h)(a') &= \begin{cases} m +_{\mathcal{M}} h(a) & \text{if } a = a' \\ h(a) & \text{otherwise} \end{cases} \end{aligned}$$

Conjunction of two heaps is simply list concatenation (written $h @ h'$) and compatibility of heaps is standard:

$$h \# h' \triangleq (\forall a) \left(\begin{array}{l} h(a) +_{\mathcal{M}} h'(a) \leq m_W \text{ and} \\ (\forall c \in h[a], \forall c' \in h'[a])(\text{val}(c) = \text{val}(c')) \end{array} \right)$$

The semantics of formulas is standard. In particular, the semantics of pure formulas is left unaxiomatized:

$$\begin{aligned} h &\models \text{iff } \Pi && \text{oracle}(\Pi) \\ h &\models \text{iff } \text{emp} && h = [] \\ h &\models \text{iff } a \overset{m}{\hookrightarrow} v && (\exists h')(h = (a, m, v) @ h') \\ h &\models \text{iff } a \overset{m}{\mapsto} v && h = (a, m, v) \\ h &\models \text{iff } A \star B && (\exists h_1, h_2)(h_1 \# h_2, h = h_1 @ h_2, \\ &&& h_1 \models \text{iff } A, \text{ and } h_2 \models B) \\ h &\models \text{iff } A \wedge B && h \models A \text{ and } h \models B \\ h &\models \text{iff } A \vee B && h \models A \text{ or } h \models B \end{aligned}$$

We write $A \vdash B$ to denote that A entails B . Appendix A shows \vdash 's definition, i.e., our proof system. It is driven by natural deduction rules that are common to the logic of bunched implications [17] and linear logic [20]. The only non-standard rules are (Splitting) and (Merging) [13] that lift permission splitting to formulas. The proof system is sound w.r.t. to the semantics above:

THEOREM 1 (SOUNDNESS OF THE PROOF SYSTEM). *If $A \vdash B$, then for all h , $h \models A$ implies $h \models B$.*

Because we use it later, we spell out Thm. 1's contraposition:

THEOREM 2 (CONTRAPOSITION OF THM. 1). *If there exists h such that $h \models A$ and $h \not\models B$, then $A \not\vdash B$.*

3. PRELIMINARIES

We abstract models of separation logic formulas (i.e., heaps) by their size. We use two abstractions of different precisions: the first abstraction is a *whole heap* abstraction while the second abstraction is a *per cell* abstraction. For the whole heap abstraction, the size of a heap is simply the sum of all permissions occurring in this heap:

$$\text{size}_{\mathcal{M}}([]) = m_0 \quad \text{size}_{\mathcal{M}}((a, m, v) :: h) = m +_{\mathcal{M}} \text{size}_{\mathcal{M}}(h)$$

For the per cell abstraction, the size of a heap is a *permission table* i.e., the corresponding heap without values:

$$\begin{aligned} pt_c &::= (a, m) && \text{permission table cells} \\ pt &::= [] \mid pt_c :: pt && \text{permission tables} \end{aligned}$$

$$\text{size}_{\mathcal{P}}([]) = [] \quad \text{size}_{\mathcal{P}}((a, m, v) :: h) = (a, m) :: \text{size}_{\mathcal{P}}(h)$$

As for heaps, we write $pt(a)$ to denote the sum of permissions to a occurring in pt . Sizes are equipped with an order. For the whole heap abstraction, we use the order on permissions $\leq_{\mathcal{M}}$ while for the per cell abstraction, we use an order on permission tables $\leq_{\mathcal{P}}$:

$$pt \leq_{\mathcal{P}} pt' \triangleq (\forall a)(pt(a) \leq pt'(a))$$

Sizes are equipped with minimum and maximum functions. For the whole heap abstraction, we use standard definitions according to the corresponding order on permissions $\leq_{\mathcal{M}}$:

$$\text{smin}_{\mathcal{M}}(m, m') = \begin{cases} m & \text{iff } m \leq_{\mathcal{M}} m' \\ m' & \text{otherwise} \end{cases}$$

$$\text{smax}_{\mathcal{M}}(m, m') = \begin{cases} m' & \text{iff } m \leq_{\mathcal{M}} m' \\ m & \text{otherwise} \end{cases}$$

For the per cell abstraction, we use definitions that merge permission tables point-wise:

$$\begin{aligned} \text{smin}_{\mathcal{P}}(pt, pt') &\triangleq \{pt'' \mid (\forall a)(pt''(a) = \text{smin}_{\mathcal{M}}(pt(a), pt'(a)))\} \\ \text{smax}_{\mathcal{P}}(pt, pt') &\triangleq \{pt'' \mid (\forall a)(pt''(a) = \text{smax}_{\mathcal{M}}(pt(a), pt'(a)))\} \end{aligned}$$

>From now on, we subscript functions when we speak about a concrete abstraction while we use functions without subscripts when we describe properties of both abstractions.

4. THE DISPROVING ALGORITHM FOR INTUITIONISTIC SEPARATION LOGIC

To disprove entailment between formulas in intuitionistic separation logic, we search for models with bounded sizes. Formally, we use min and max functions that satisfy the following property:

THEOREM 3 (PROPERTY OF min AND max). *If $h \models A$, then there exist h_s, h_r such that $h = h_s @ h_r$ and $h_s \models A$ and $\min(A) \leq \text{size}(h_s) \leq \max(A)$.*

Intuitively, Thm. 3 states that, given a satisfiable formula A (i.e., there exists h such that $h \models A$), there exists a “small” model h_s of A whose size is in between $\min(A)$ and $\max(A)$ (the possibility of a superfluous heap portion h_r reflects the semantics of intuitionistic separation logic). Then, we use min and max to disprove entailment as follows:

THEOREM 4 (min/max IS A DISPROVING ALGORITHM). *If A is satisfiable and $\min(B) \leq \max(A)$ does not hold, then $A \not\vdash B$.*

PROOF. Because A is satisfiable, there exists h such that $h \models A$. Then, by Thm. 3, it follows that there exists h_s such that $h_s \models A$ and $\text{size}(h_s) \leq \max(A)$.

Now, suppose $h_s \models B$. Then, by Thm. 3, it follows that $\min(B) \leq \text{size}(h_s)$. From $\text{size}(h_s) \leq \max(A)$, by transitivity, it follows that $\min(B) \leq \max(A)$. This contradicts, however, the hypothesis that $\min(B) \leq \max(A)$ does not hold. Hence, $h_s \not\models B$.

From $h_s \models A$ and $h_s \not\models B$, by Thm. 2, it follows that $A \not\vdash B$. \square

The reader might wonder why we formulate Thm. 4 with a negation on the \leq operator instead of using a $<$ operator. The reason is that, for permission tables, $\neg(pt \leq_{\mathcal{P}} pt')$ is *not* equivalent to $pt' <_{\mathcal{P}} pt$ (where $<_{\mathcal{P}}$ would be defined in a way similar to $\leq_{\mathcal{P}}$).

In the remainder of this section, we show definitions of min and max for the whole heap abstraction and the per cell abstraction.

4.1 Whole Heap Abstraction

min and max are defined as follows:

$$\begin{aligned} \min(\Pi) &= m_0 \\ \min(a \xrightarrow{m} v) &= m \\ \min(A \star B) &= \min(A) +_{\mathcal{M}} \min(B) \\ \min(A \wedge B) &= \text{smax}_{\mathcal{M}}(\min(A), \min(B)) \\ \min(A \vee B) &= \text{smin}_{\mathcal{M}}(\min(A), \min(B)) \end{aligned}$$

$$\begin{aligned} \max(\Pi) &= m_0 \\ \max(a \xrightarrow{m} v) &= m \\ \max(A \star B) &= \max(A) +_{\mathcal{M}} \max(B) \\ \max(A \wedge B) &= \max(A) +_{\mathcal{M}} \max(B) \\ \max(A \vee B) &= \text{smax}_{\mathcal{M}}(\max(A), \max(B)) \end{aligned}$$

4.2 Per Cell Abstraction

min and max are defined as follows:

$$\begin{aligned} \min(\Pi) &= \square \\ \min(a \xrightarrow{m} v) &= (a, m) :: \square \\ \min(A \star B) &= \min(A) @ \min(B) \\ \min(A \wedge B) &= \text{smax}_{\mathcal{P}}(\min(A), \min(B)) \\ \min(A \vee B) &= \text{smin}_{\mathcal{P}}(\min(A), \min(B)) \end{aligned}$$

$$\begin{aligned} \max(\Pi) &= \square \\ \max(a \xrightarrow{m} v) &= (a, m) :: \square \\ \max(A \star B) &= \max(A) @ \max(B) \\ \max(A \wedge B) &= \text{smax}_{\mathcal{P}}(\max(A), \max(B)) \\ \max(A \vee B) &= \text{smax}_{\mathcal{P}}(\max(A), \max(B)) \end{aligned}$$

Although max’s definitions are very similar for the two abstractions, there is one case where they differ: in case \wedge of max. In analogy with the per cell abstraction, one could expect $\max(A \wedge B)$ to be $\text{smax}_{\mathcal{M}}(\max(A), \max(B))$ in the whole heap abstraction. To see why this is unsound, however, one can consider a formula where the right and left hand sides of a \wedge represent separate parts of the heap such as $42 \xrightarrow{1} _ \wedge 47 \xrightarrow{1} _$ (where $_$ denotes irrelevant values).

5. THE DISPROVING ALGORITHM FOR CLASSICAL SEPARATION LOGIC

To disprove entailment between formulas in classical separation logic, we pursue a slightly different goal than for intuitionistic separation logic: instead of exhibiting *one* model whose size is bounded, we search for bounds on the size of *all* models.

Because we allow pure formulas, which do not constrain the size of heaps modelling them, we cannot always find an upper-bound on the size of heaps. Hence we add a distinguished “maximal” size that we write ∞ (both for the whole heap and for the per cell abstraction). The meaning of ∞ is axiomatized as follows:

$$(\forall m)(m \leq_{\mathcal{M}} \infty) \quad (\forall pt)(pt \leq_{\mathcal{P}} \infty)$$

THEOREM 5 (PROPERTY OF min AND max). *If $h \models A$, then $\min(A) \leq \text{size}(h) \leq \max(A)$.*

Then, we use min and max to disprove entailment as follows:

THEOREM 6 (min/max IS A DISPROVING ALGORITHM). *If A is satisfiable and $\min(B) \leq \max(A)$ does not hold, then $A \not\vdash B$.*

PROOF. Similar to the proof of Thm. 4. \square

In the remainder of this section, we give definitions of min and max for the whole heap abstraction and the per cell abstraction. For min, we only show its definition in case of a formula emp ; for other cases min’s definition is similar to the intuitionistic case (see Sec. 4.1 and 4.2).

5.1 Whole Heap Abstraction

min and max are defined as follows:

$$\min(\text{emp}) = m_0$$

$$\begin{aligned} \max(\Pi) &= \infty \\ \max(\text{emp}) &= m_0 \\ \max(a \xrightarrow{m} v) &= m \\ \max(A \star B) &= \max(A) +_{\mathcal{M}} \max(B) \\ \max(A \wedge B) &= \text{smin}_{\mathcal{M}}(\max(A), \max(B)) \\ \max(A \vee B) &= \text{smax}_{\mathcal{M}}(\max(A), \max(B)) \end{aligned}$$

For these definitions to be sound, we need to assume that formulas appearing on the left-hand side of the magic wand are satisfiable. While this is harmful from a theoretical point of view, the literature so far on program verification that uses the magic wand in examples [13, 14] suggests this is a plausible assumption for program verifiers.

We believe this definition could be adapted for the intuitionistic semantics, but it is a non-trivial task. In a nutshell, the universal quantification in the magic wand’s semantics and the way Thm. 3 is formulated (i.e., with an existential quantifier) prevent a simple proof by induction, and we leave this for future work.

7.2 Quantification

We distinguish between quantification on values and permissions. We add the following formulas to Sec. 2’s language:

$$A, B ::= \dots \mid \exists v.A \mid \exists m.A \mid \forall v.A \mid \forall m.A$$

The semantics is standard and we omit the corresponding proof rules which are also standard. To define \min and \max , we introduce a “very small permission” $\varepsilon \in \mathcal{M}$. The meaning of ε is axiomatized by: $(\forall m)(m <_{\mathcal{M}} m +_{\mathcal{M}} \varepsilon)$

For both abstractions \min and \max are defined as follows:

$$\begin{aligned} \min(\exists v.A) &= \min(A) & \min(\exists m.A) &= \min(A[\varepsilon/m]) \\ \max(\exists v.A) &= \max(A) & \max(\exists m.A) &= \max(A[m_W/m]) \\ \min(\forall v.A) &= \min(A) & \min(\forall m.A) &= \min(A[m_W/m]) \\ \max(\forall v.A) &= \max(A) & \max(\forall m.A) &= \max(A[m_W/m]) \end{aligned}$$

7.3 Variables

In real-world flavours of separation logic, variables are used. For the whole heap abstraction, variables (including quantification over variables) do not create any extra difficulties for our work. For the per cell abstraction, it is unclear at this stage how variables can be handled elegantly, and we leave this for future work.

8. COMPLEXITY

The complexity of our techniques is linear for the whole heap abstraction, and $O(n(\log n))$ for the per cell abstraction (where n is the input formula’s size). To our knowledge, there is no complexity result for entailment checking in separation logic. We believe the complexity of our disproving algorithm is low enough to be useful when entailment must be checked quickly.

This is supported by the fact that integrating both the magic wand and quantifiers do not increase our algorithm’s complexity, while they typically increase by orders of magnitude the complexity of model-checking or validity-checking [8, 7].

9. SOUNDNESS

All theorems from Sec. 2, 4, and 5 have been mechanically checked with Coq [10]. Addition of the magic wand (Sec. 7.1) for the classical semantics has also been mechanically checked. Proofs about the proof system (Thms. 1 and 2); properties of permissions, heaps, permissions tables etc. are 3090 lines long. The proofs of Thms 3, 4, 5 and 6 are 620 lines long. The proofs have been engineered so that certified implementations of the two abstractions could be extracted. Proof scripts are available online [1].

Proofs for quantifiers (Sec. 7.2) have been checked only on pen and paper.

10. RELATED WORK AND CONCLUSION

Related Work.

In the context of program verifiers, checking entailment between separation logic formulas has been studied for while languages [4] and object-oriented languages [11, 9]. In these fragments, the magic wand \rightarrow^* is omitted and special predicates for describing data structures are present.

Properties (decidability, undecidability, and complexity) both for model checking and for validity have been studied [8, 2, 7]. To our knowledge, disproving entailment has only been studied in a fully-fledged tableaux procedure [12]. The cited work’s algorithm for checking entailment can generate either proofs or counter-models, whereas our algorithm focuses on disproving entailment and does not generate counter-models.

Conclusion.

We have presented new techniques for disproving entailment between separation logic formulas, which we believe to be particularly relevant in the context of automated provers for separation logic. Because our algorithm’s complexity is low, it is of interest wherever entailment checking needs to be checked quickly. We have provided two different techniques (of different precision), each of which can be applied to both the intuitionistic and classical flavours of separation logic. We provide mechanical proofs of the soundness of our techniques [1].

Acknowledgments.

We thank Didier Galmiche and Marieke Huisman for their fruitful comments about this paper. Hurlin was supported in part by IST-FET-2005-015905 Mobius project and ANR-06-SETIN-010 ParSec project.

11. REFERENCES

- [1] Coq proof scripts: <http://tinyurl.com/aez9sj>.
- [2] J. Berdine, C. Calcagno, P. W. O’Hearn. A decidable fragment of separation logic. In K. Lodaya, M. Mahajan, eds., *Foundations of Software Technology and Theoretical Computer Science*, vol. 3821 of *Lecture Notes in Computer Science*. Springer-Verlag, 2004.
- [3] J. Berdine, C. Calcagno, P. W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In F. S. de Boer, M. M. Bonsangue, S. Graf, W.-P. de Roever, eds., *Formal Methods for Components and Objects*, vol. 4111 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
- [4] J. Berdine, C. Calcagno, P. W. O’Hearn. Symbolic execution with separation logic. In K. Yi, ed., *Asian Programming Languages and Systems Symposium*, vol. 3780 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
- [5] R. Bornat, P. W. O’Hearn, C. Calcagno, M. Parkinson. Permission accounting in separation logic. In J. Palsberg, M. Abadi, eds., *Principles of Programming Languages*. ACM Press, 2005.
- [6] J. Boyland. Checking interference with fractional permissions. In R. Cousot, ed., *Static Analysis Symposium*, vol. 2694 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
- [7] R. Brochenin, S. Demri, E. Lozes. On the almighty wand. In M. Kaminski, S. Martini, eds., *Computer Science Logic*. Springer-Verlag, 2008.
- [8] C. Calcagno, H. Yang, P. W. O’Hearn. Computability and complexity results for a spatial assertion language for data

$$\begin{array}{c}
\frac{}{A \vdash A} \text{ (Id)} \quad \frac{A \vdash B}{A, C \vdash B} \text{ (Weak)} \quad \frac{\text{oracle}(\Pi)}{\text{true} \vdash \Pi} \text{ (Oracle)} \\
\frac{A_1 \vdash B_1 \quad A_2 \vdash B_2}{A_1, A_2 \vdash B_1 \star B_2} \text{ (\star Intro)} \quad \frac{A \vdash B_1 \star B_2 \quad C, B_1, B_2 \vdash D}{A, C \vdash D} \text{ (\star Elim)} \\
\frac{A, B_1 \vdash B_2}{A \vdash B_1 \multimap B_2} \text{ (\multimap Intro)} \quad \frac{A \vdash C_1 \multimap C_2 \quad B \vdash C_1}{A, B \vdash C_2} \text{ (\multimap Elim)} \\
\frac{A \vdash B_1 \quad A \vdash B_2}{A \vdash B_1 \wedge B_2} \text{ (\wedge Intro)} \quad \frac{A \vdash B_1 \wedge B_2}{A \vdash B_1} \text{ (\wedge Elim 1)} \quad \frac{A \vdash B_1 \wedge B_2}{A \vdash B_2} \text{ (\wedge Elim 2)} \\
\frac{A \vdash B_1}{A \vdash B_1 \vee B_2} \text{ (\vee Intro 1)} \quad \frac{A \vdash B_2}{A \vdash B_1 \vee B_2} \text{ (\vee Intro 2)} \\
\frac{\vee \text{ does not occur in } A}{A \vdash \text{split1}(A) \star \text{split2}(A)} \text{ (Splitting)} \quad \frac{\vee \text{ does not occur in } A}{\text{split1}(A) \star \text{split2}(A) \vdash A} \text{ (Merging)}
\end{array}$$

Figure 1: Proof System

structures. In R. Hariharan, M., V. Vinay, eds., *Foundations of Software Technology and Theoretical Computer Science*, vol. 2245 of *Lecture Notes in Computer Science*. Springer-Verlag, 2001.

- [9] W. Chin, C. David, H. Nguyen, S. Qin. Enhancing modular OO verification with separation logic. In G. C. Necula, P. Wadler, eds., *Principles of Programming Languages*. ACM Press, 2008.
- [10] Coq development team. The Coq proof assistant reference manual V8.0. Technical Report 255, INRIA, France, 2004. <http://coq.inria.fr/doc/main.html>.
- [11] D. DiStefano, M. Parkinson. jStar: Towards practical verification for Java. In *ACM Conference on Object-Oriented Programming Systems, Languages, and Applications*, vol. 43. ACM Press, 2008.
- [12] D. Galmiche, D. Mery. Tableaux and resource graphs for separation logic. *Journal of Logic and Computation*, 2008.
- [13] C. Haack, C. Hurlin. Separation logic contracts for a Java-like language with fork/join. In J. Meseguer, G. Rosu, eds., *Algebraic Methodology and Software Technology*, vol. 5140 of *Lecture Notes in Computer Science*. Springer-Verlag, 2008.
- [14] C. Haack, C. Hurlin. Resource usage protocols for iterators. *Journal of Object Technology*, 2009. To appear.
- [15] S. Ishtiaq, P. W. O’Hearn. BI as an assertion language for mutable data structures. In *Principles of Programming Languages*, 2001.
- [16] B. Jacobs, F. Piessens. The verifast program verifier. Technical Report CW520, Katholieke Universiteit Leuven, 2008.
- [17] P. W. O’Hearn, D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2), 1999.
- [18] M. Parkinson. Local reasoning for Java. Technical Report UCAM-CL-TR-654, University of Cambridge, 2005.
- [19] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science*. IEEE Press, 2002.
- [20] P. Wadler. A taste of linear logic. In *Mathematical Foundations of Computer Science*, 1993.

APPENDIX

A. PROOF SYSTEM

Fig. 1 shows our proof system (where weakening is only included for the intuitionistic flavour).

In the proof rules (Splitting) and (Merging), we use split1 and split2 (defined below) to lift permission splitting to formula splitting. The conditions on the proof rules for splitting and merging, ensure that no disjunctions occur within the formulas concerned. This is because splitting and merging is unsound for disjunctions (and formulas with the magic wand). As we only split formulas without these connectives, we need not define those cases for the definition of splitting formulas:

$$\begin{array}{ll}
\text{split1}(\Pi) & = \Pi \\
\text{split1}(\text{emp}) & = \text{emp} \\
\text{split1}(a \overset{m}{\hookrightarrow} v) & = a \overset{\text{split1}(m)}{\hookrightarrow} v \\
\text{split1}(a \overset{m}{\mapsto} v) & = a \overset{\text{split1}(m)}{\mapsto} v \\
\text{split1}(A \star B) & = \text{split1}(A) \star \text{split1}(B) \\
\text{split1}(A \wedge B) & = \text{split1}(A) \wedge \text{split1}(B) \\
\\
\text{split2}(\Pi) & = \Pi \\
\text{split2}(\text{emp}) & = \text{emp} \\
\text{split2}(a \overset{m}{\hookrightarrow} v) & = a \overset{\text{split2}(m)}{\hookrightarrow} v \\
\text{split2}(a \overset{m}{\mapsto} v) & = a \overset{\text{split2}(m)}{\mapsto} v \\
\text{split2}(A \star B) & = \text{split2}(A) \star \text{split2}(B) \\
\text{split2}(A \wedge B) & = \text{split2}(A) \wedge \text{split2}(B)
\end{array}$$

B. OMITTED CASES

The definitions of min and max from Sec. 5.2 are completed as follows (where $-\mathcal{P}$ is defined by applying the unary $-\mathcal{M}$ operator point-wisely to permission tables):

$$\begin{array}{l}
\min(A \multimap B) = \begin{cases} \min(B) @ -\mathcal{P}(\max(A)) & \text{if } \min(B) \neq \infty \text{ and } \\ & \max(A) \neq \infty \\ \square & \text{otherwise} \end{cases} \\
\max(A \multimap B) = \begin{cases} \max(B) @ -\mathcal{P}(\min(A)) & \text{if } \max(B) \neq \infty \text{ and } \\ & \min(A) \neq \infty \\ \infty & \text{otherwise} \end{cases}
\end{array}$$