

Coming soon from Cambridge University Press

Program Logics for Certified Compilers



TABLE OF CONTENTS
and SAMPLE CHAPTER
of prepublication
manuscript
May 31, 2013

Andrew W. Appel

*with Robert Dockins, Aquinas Hobor,
Lennart Beringer, Josiah Dodds, Gordon Stewart,
Sandrine Blazy, and Xavier Leroy*

This is the preliminary manuscript of a book that will be published in late 2013, and will be citable as

Andrew W. Appel *et al.*,
Program Logics for Certified Compilers, Cambridge University Press, 2014.

Copyright © 2013 Andrew W. Appel

Contents

Road map	vii
Acknowledgments	viii
1 Introduction	9
I Generic Separation Logic	17
2 Hoare logic	18
3 Separation logic	24
4 Soundness of Hoare logic	33
5 Mechanized Semantic Library	41
6 Separation algebras	43
7 Operators on separation algebras	52
8 First-order separation logic	57
9 A little case study	63
10 Covariant recursive predicates	71
11 Share accounting	77
II Higher Order Separation Logic	83
12 Separation Logic as a logic	84
13 From separation algebras to separation logic	92
14 Simplification by rewriting	97
15 Introduction to step-indexing	102
16 Predicate implication and subtyping	107
17 General recursive predicates	112
18 Case Study: Separation logic with first-class functions	119

19 Data structures in indirection theory	131
20 Applying higher-order separation logic	138
21 Lifted Separation Logics	142
III Separation Logic for CompCert	149
22 Verifiable C	150
23 Expressions, values, and assertions	156
24 The VST Separation Logic for C light	161
25 Typechecking for Verifiable C	181
26 Derived rules and proof automation for C light	190
27 Proof of a program	200
28 More C programs	213
29 Dependently typed C programs	222
30 Concurrent separation logic	227
IV Operational Semantics of CompCert	237
31 CompCert	238
32 The CompCert memory model	242
33 How to specify a compiler	277
34 C light operational semantics	289
V Indirection Theory	295
35 Higher-order Hoare logic	296
36 Higher-order separation logic	304
37 Case study: Lambda-calculus with references	308
38 Semantic models of predicates-in-the-heap	332
VI Semantic model and soundness of Verifiable C	337
39 Separation algebra for CompCert	338
40 Share models	349
41 Juicy memories	360
42 Modeling the Hoare judgment	368
43 Modular structure of the development	376

VII Applications	392
45 Foundational static analysis	393
46 Heap theorem prover	408
Bibliography	424
Index	434

Road map

Readers interested in **the theory of separation logic** (with some example applications) should read Chapters 1–21. Readers interested in **the use of separation logic to verify C programs** should read Chapters 1–6 and 8–30. Those interested in **the theory of step-indexing** and **indirection theory** should read Chapters 35–39. Those interested in building models of **program logics** proved sound for **certified compilers** should read Chapters 40–46, though it would be helpful to read Chapters 1–39 as a warm-up.

Chapter 1

Introduction

An exciting development of the 21st century is that the 20th-century vision of *mechanized program verification* is finally becoming practical, thanks to 30 years of advances in logic, programming-language theory, proof-assistant software, decision procedures for theorem proving, and even Moore’s law which gives us everyday computers powerful enough to run all this software.

We can write functional programs in ML-like languages and prove them correct in expressive higher-order logics; and we can write imperative programs in C-like languages and prove them correct in appropriately chosen program logics. We can even prove the correctness of the verification toolchain itself: the compiler, the program logic, automatic static analyzers, concurrency primitives (and their interaction with the compiler). There will be few places for bugs (or security vulnerabilities) to hide.

This book explains how to construct powerful and expressive program logics based on Separation Logic and Indirection Theory. It is accompanied by an open-source machine-checked formal model and soundness proof, the *Verified Software Toolchain*¹ (VST), formalized in the Coq proof assistant. The VST components include the theory of *separation logic* for reasoning about pointer-manipulating programs; *indirection theory* for reasoning with “step-indexing” about first-class function pointers, recursive types,

¹<http://vst.cs.princeton.edu>

recursive functions, dynamic mutual-exclusion locks, and other higher-order programming; a *Hoare logic* (separation logic) with full reasoning about control-flow and data-flow of the C programming language; theories of *concurrency* for reasoning about programming models such as Pthreads; theories of *compiler correctness* for connecting to the CompCert verified C compiler; theories of *symbolic execution* for implementing foundationally verified static analyses. VST is built in a modular way, so that major components apply very generally to many kinds of separation logics, Hoare logics, and step-indexing semantics.

One of the major demonstration applications comprises certified program logics and certified static analyses for the *C light* programming language. *C light* is compiled into assembly language by the CompCert² certified optimizing compiler. [59] Thus, the VST is useful for verified formal reasoning about programs that will be compiled by a verified compiler. But Parts I, II, and V of this book show principles and Coq developments that are quite independent of CompCert and have already been useful in other applications of separation logics.

PROGRAM LOGICS FOR CERTIFIED COMPILERS. Software is complex and prone to bugs. We would like to reason about the correctness of programs, and even to prove that the behavior of a program adheres to a formal specification. For this we use program logics: rules for reasoning about the behavior of programs. But programs are large and the reasoning rules are complex; what if there is a bug in our proof (in our application of the rules of the program logic)? And how do we know that the program logic itself is sound—that when we conclude something using these rules, the program will really behave as we concluded? And once we have reasoned about a program, we compile it to machine code; what if there is a bug in the compiler?

We achieve soundness by formally verifying our program logics, static analyzers, and compilers. We prove soundness theorems based on foundational specifications of the underlying hardware. We check all proofs by machine, and connect the proofs together end-to-end so there are no gaps.

²<http://compcert.inria.fr>

DEFINITIONS. A *program* consists of instructions written in a *programming language* that direct a computer to perform a task. The *behavior* of a program, *i. e.* what happens when it executes, is specified by the *operational semantics* of the programming language. Some programming languages are *machine languages* that can directly execute on a computer; others are *source languages* that require translation by a *compiler* before they can execute.

A *program logic* is a set of formal rules for *static* reasoning about the behavior of a program; the word *static* implies that we do not actually execute the program in such reasoning. *Hoare Logic* is an early and still very important program logic. *Separation Logic* is a 21st-century variant of Hoare Logic that better accounts for pointer and array data structures.

A compiler is *correct* with respect to the specification of the operational semantics of its source and its target languages if, whenever a source program has a particular defined behavior, and when the compiler translates that program, then the target program has a *corresponding* behavior. [36] The correspondence is part of the correctness specification of the compiler, along with the two operational semantics. A compiler is *proved correct* if there is a formal proof that it meets this specification. Since the compiler is itself a program, this formal proof will typically be using the rules of a program logic for the implementation language of the compiler.

Proofs in a logic (or program logic) can be written as derivation trees in which each node is the application of a rule of the system. The validity of a proof can be checked using a computer program. A *machine-checked proof* is one that has been checked in this way. Proof-checking programs can be quite small and simple, [12] so one can reasonably hope to implement a proof-checker free of bugs.

It is inconvenient to construct derivation trees “by hand.” A *proof assistant* is a tool that combines a proof checker with a user interface that assists the human in building proofs. The proof assistant may also contain algorithms for proof automation, such as *tactics* and *decision procedures*.

A *certified compiler* is one proved correct with a machine-checked proof. A *certified program logic* is one proved sound with a machine-checked proof. A *certified program* is one proved correct (using a program logic) with a machine-checked proof.

A *static analysis* algorithm calculates properties of the behavior of a program without actually running it. A static analysis is *sound* if, whenever it claims some property of a program, that property holds on all possible behaviors (in the operational semantics). The proof of soundness can be done using a (sound) program logic, or it can be done directly with respect to the operational semantics of the programming language. A *certified static analysis* is one that is proved sound with a machine-checked proof—either the static analysis program is proved correct, or each run of the static analysis generates a machine-checkable proof about a particular instance.

In Part I we will review Hoare logics, operational semantics, and separation logics. For a more comprehensive introduction to Hoare logic, the reader can consult Huth and Ryan [52] or many other books; For operational semantics, see Harper [45, Parts I & II] or Pierce [73]. For an introduction to theorem-proving in Coq, see Pierce’s *Software Foundations*[74] which also covers applications to operational semantics and Hoare logic.

THE VST SEPARATION LOGIC FOR C LIGHT is a higher-order impredicative concurrent separation logic certified with respect to CompCert. *Separation Logic* means that its assertions specify heap-domain footprints: the assertion $(p \mapsto x) * (q \mapsto y)$ describes a memory with exactly two disjoint parts; one part has only the cell at address p with contents x , and the other has only address q with contents y , with $p \neq q$. *Concurrent Separation Logic* is an extension that can describe shared-memory concurrent programs with Dijkstra-Hoare synchronization (e.g., Pthreads). *Higher-order* means that assertions can use existential and universal quantifiers, the logic can describe pointers to functions and mutex locks, and recursive assertions can describe recursive data types such as lists and trees. *Impredicative* means that the \exists and \forall quantifiers can even range over assertions containing quantifiers. *Certified* means that there is a machine-checked proof of soundness with respect to the operational semantics of a source language of the CompCert C compiler.

A separation logic has assertions $p \mapsto x$ where p ranges over a particular address type A , x ranges over a specific type V of values, and the assertion as a whole can be thought of as a predicate over some specific type of

“heaps” or “computer memories” M . Then the logic will have theorems such as $(p \mapsto x) * (q \mapsto y) \vdash (q \mapsto y) * (p \mapsto x)$.

We will write down *generic separation logic* as a theory parameterizable by types such as A, V, M , and containing generic axioms such as $P * Q \vdash Q * P$. For a particular instantiation such as CompCert C light, we will instantiate the generic logic with the types of C values and C expressions.

[Chapter 3](#) will give an example of an informal program verification in “pencil-and-paper” Separation Logic. Then [Part V](#) shows the VST tools applied to build a foundationally sound toolchain for a toy language, with a machine-verified separation-logic proof of a similar program. [Part III](#) demonstrates the VST tools applied to the C language, connected to the CompCert compiler, and shows machine-checked verification C programs.

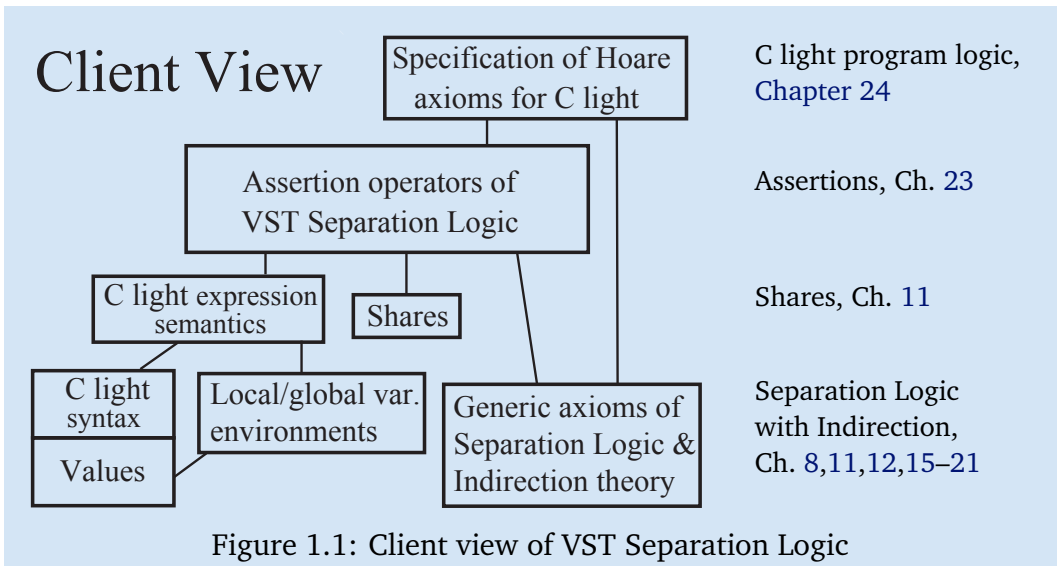


Figure 1.1: Client view of VST Separation Logic

FIGURE 1.1 SHOWS THE *client view* of the VST Separation Logic for C light—that is, the specification of the axiomatic semantics. Users of the program logic will reason directly about CompCert values (integers, floats, pointers) and C-light expression evaluation. Users do not see the operational semantics of C-light commands, or CompCert memories. Instead, they use

the axiomatic semantics—the Hoare judgment and its reasoning rules—to reason indirectly about memories via assertions such as $p \mapsto x$.

The modular structure of the *client view* starts (at bottom left of Fig. 1.1) with the specification of the *C light* language, a subset of C chosen for its compatibility with program-verification methods. We have C values (such as integers, floats, and pointers); the abstract syntax of C light, and the mechanism of evaluating C light expressions. The client view treats statements such as assignment and looping *abstractly* via an axiomatic semantics (Hoare logic), so it does not expose an operational semantics.

At bottom right of Figure 1.1 we have the operators and axioms of Separation Logic and of Indirection Theory. At center are the assertions of our program logic for C light, which (as the diagram shows) make use of C-light expressions and of our logical operators. At top, the Hoare axioms for C light complete the specification of the program logic.

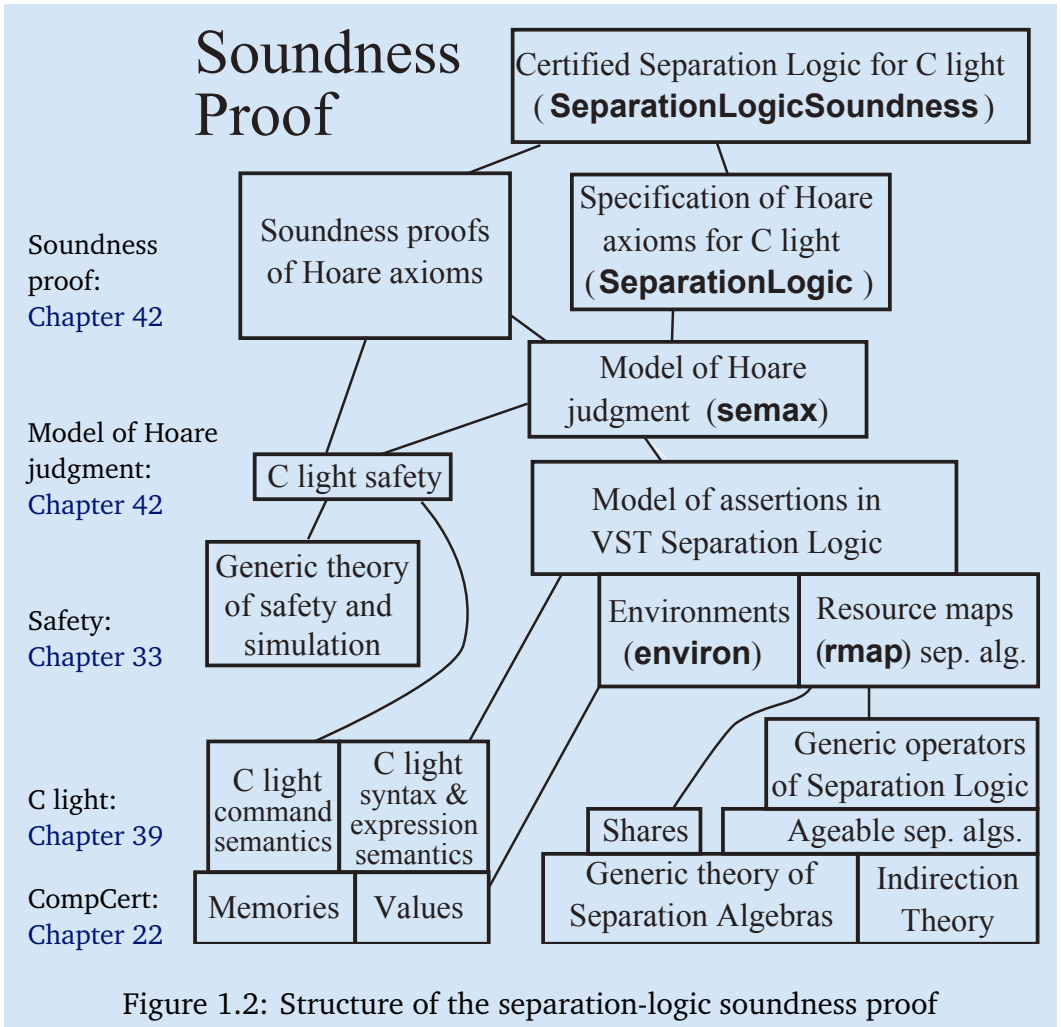
Readers primarily interested in *using* the VST tools may want to read Parts I through III, which explain the components of the client view.

THE SOUNDNESS PROOF OF THE VST SEPARATION LOGIC is constructed by reasoning in the *model* of separation logic. Figure 1.2 shows the structure of the soundness proof. At bottom left is the specification of C-light operational semantics. We have a generic theory of safety and simulation for shared-memory programs, and we instantiate that into the “C light safety” theory.

At bottom right (Fig. 1.2) is the theory of *separation algebras*, which form models of separation logics. The assertions of our logic are predicates on the *resource maps* that, in turn, model CompCert memories. The word *predicate* is a technical feature of our Indirection Theory that implicitly accounts for “resource approximation,” thus allowing higher-order reasoning about circular structures of pointers and resource invariants.

We construct a semantic model of the Hoare judgment, and use this to prove sound all the judgment rules of the Separation Logic. All this is encapsulated in a Coq module called SeparationLogicSoundness.

Parts IV through VI explain the components of Figure 1.2, the semantic model and soundness proof of higher-order impredicative separation logic for CompCert C light.



The Coq development of the Verified Software Toolchain is available at vst.cs.princeton.edu and is structured in a root directory with several subdirectories:

compcert: A few files copied from the CompCert verified C compiler, that comprise the specification of the *C light* programming language.

sepcomp: Theory of how to specify shared-memory interactions of CompCert-compiled programs.

msl: Mechanized Software Library, the theory of separation algebras, share accounting, and generic separation logics.

veric: The program logic: a higher-order splittable-shares concurrent separation logic for C light.

floyd: A proof-automation system of lemmas and tactics for semiautomated application of the program logic to C programs.

progs: Applications of the program logic to sample programs.

veristar: A heap theorem prover using resolution and paramodulation.

A proof development, like any software, is a living thing: it is continually being evolved, edited, maintained, and extended. We will not tightly couple this book to the development; we will just explain the key mathematical and organizational principles, illustrated with snapshots from the Coq code.

Part I

Generic Separation Logic

SYNOPSIS: Separation logic is a formal system for static reasoning about pointer-manipulating programs. Like Hoare logic, it uses assertions that serve as preconditions and postconditions of commands and functions. Unlike Hoare logic, its assertions model anti-aliasing via the disjointness of memory heaplets. Separation algebras serve as models of separation logic. We can define a calculus of different kinds of separation algebras, and operators on separation algebras. Permission shares allow reasoning about shared ownership of memory and other resources. In a first-order separation logic we can have predicates to describe the contents of memory, anti-aliasing of pointers, and simple (covariant) forms of recursive predicates. A simple case study of straight-line programs serves to illustrate the application of separation logic.

Part II

Higher Order Separation Logic

SYNOPSIS: Instead of reasoning directly on the model (that is, separation algebras), we can treat Separation Logic as a syntactic formal system, that is, a logic. We can implement proof automation to assist in deriving separation-logic proofs.

Reasoning about recursive functions, recursive types, and recursive predicates can lead to paradox if not done carefully. Step-indexing avoids paradoxes by inducting over the number of remaining program-steps that we care about. Indirection theory is a kind of step-indexing that can serve as models of higher-order Hoare logics. Using indirection theory we can define general (not just covariant) recursive predicates.

Recursive data structures such as lists and trees are easily modeled in indirection theory, but the model is not the same one conventionally used, as it inducts over “age”—the approximation level, the amount of information left in the model—rather than list-length or tree-depth. A tiny pointer/continuation language serves as a case study for Separation Logic with first-class function-pointers, modeled in indirection theory. The proof of a little program in the case-study language illustrates the application of separation logic with function pointers.

Part III

Separation Logic for CompCert

SYNOPSIS: *Verifiable C* is a style of C programming suited to separation-logic verifications; it is similar to the C light intermediate language of the CompCert compiler. We show the assertion language of separation-logic predicates for specifying states of a C execution. The judgment form *semax* of the axiomatic semantics relates a C command to its precondition postconditions, and for each kind of command there is an inference rule for proving its *semax* judgments. We illustrate with the proof of a C program that manipulates linked lists, and we give examples of other programs and how they can be specified in the Verifiable C program logic. Shared-memory concurrent programs with Dijkstra-Hoare synchronization can be verified using the rules of concurrent separation logic.

Part IV

Operational Semantics of CompCert

SYNOPSIS: Specification of the interface between CompCert and its clients such as the VST Separation Logic for C light, or clients such as proved-sound static analyses and abstract interpretations. This specification takes the form of an operational semantics with a nontrivial memory model. The need to preserve the compiler's freedom to optimize the placement of data (in memory, out of memory) requires the ability to rename addresses and adjust block sizes. Thus the specification of shared-memory interaction between subprograms (separately compiled functions, or concurrent threads) requires particular care, to keep these renamings consistent.

Part V

Indirection Theory

SYNOPSIS: Indirection theory gives a clean interface to higher-order step indexing. Many different semantic features of programming languages can be modeled in indirection theory. The models of indirection theory use dependent types to stratify quasirecursive predicates, thus avoiding paradoxes of self-reference. Lambda calculus with mutable references serves as a case study to illustrate the use of indirection theory models.

When defining both Indirection and Separation one must take extra care to ensure that aging commutes over separation. We demonstrate how to build an axiomatic semantics with using higher-order separation logic, for the pointer/continuation language introduced in the case study of [Part II](#).

Part VI

Semantic model and soundness of Verifiable C

SYNOPSIS: To prove soundness of the Verifiable C separation logic, we first give a model of mpred as $\text{pred}(\text{rmap})$, that is, predicates on resource maps. We give a model for permission-shares using trees of booleans. We augment the C light operational semantics with juicy memories that keep track of resources as well as “dry” values. We give a semantic model of the Hoare judgment, using the continuation-passing notion of “guards.” We use this semantic model to prove all the Hoare rules. Our model and proofs have a modular structure, so that they can be ported to other programming languages (especially in the CompCert family).

Part VII

Applications

SYNOPSIS: In [Part III](#) we showed how to apply a program logic interactively to a program, using tactics. Here we will show a different use of program logics: we build automatic static analyses and decision procedures as efficient functional programs, and prove their soundness using the rules of the program logic.