

A Machine-Checked Framework for Relational Separation Logic*

Juan Manuel Crespo¹ and César Kunz^{1,2}

¹ IMDEA Software Institute, Madrid, Spain

² Universidad Politécnica de Madrid

Abstract. Relational methods are gaining growing acceptance for specifying and verifying properties defined in terms of the execution of two programs— notions such as simulation, observational equivalence, non-interference, and continuity can be elegantly cast in this setting. In previous work, we have proposed *program product construction* as a technique to reduce relational verification to standard verification. This method hinges on the ability to interpret relational assertions as traditional predicates, which becomes problematic when considering assertions from relational separation logic. We report in this article an alternative method that overcomes this difficulty, defined as a relational weakest precondition calculus based on separation logic and formalized in the Coq proof assistant. The formalization includes an application to the formal verification of the Schorr-Waite graph marking algorithm. We discuss additional variants of relational separation logic inspired by the standard notions of partial and total correctness, and extensions of the logic to handle non-structurally equivalent programs.

1 Introduction

Separation logic [15, 23, 24] is a formalism devised to verify pointer programs using local reasoning; its extensions and variants have been used successfully in a variety of large scale programs [30] and smaller but challenging examples [17], including lock-free algorithms [13].

Relational reasoning, on the other hand, provides an effective means to understand program behavior: in particular, it allows one to establish that the same program behaves similarly on two different runs, or that two programs execute in a related fashion. Relational judgments are often formalized by quadruples $\{\varphi\} c_1 \sim c_2 \{\psi\}$, denoting that every pair of executions of c_1 and c_2 with initial states related by φ returns with final states related by ψ . Prime examples of relational properties include notions of simulation and observational equivalence, and 2-properties, such as non-interference and continuity.

Syntactic methods [7] have been developed to support relational reasoning. In particular, relational separation logic [29] is a variant of separation logic that supports

* Partially funded by European Projects FP7-231620 HATS and FP7-256980 NESSoS, Spanish project TIN2009-14599 DESAFIOS 10, Madrid Regional project S2009TIC-1465 PROMETIDOS. César Kunz is funded by a Juan de la Cierva Fellowship, MICINN, Spain.

reasoning about two pointer programs; it embodies the conventional wisdom that casting program correctness as an equivalence between two programs is often more beneficial than functional verification. More concretely, relational separation logic is intended to prove program correctness by showing the equivalence between the program to be verified and a reference implementation: e.g. Yang [29] provides an elegant proof in relational separation logic that the Schorr-Waite graph marking algorithm is equivalent to depth-first search.

However, these syntactic methods suffer from two important caveats: on the one hand, these logics confine reasoning to structurally equivalent programs with equivalent guards; on the other hand, tool support is negligible, with the exception of recent work by Aleks Nanevski *et al* [22]—which focuses mainly on the specification and proof of a rich set of security policies and its static enforcement. Although the relational postconditions used to describe such policies can be arbitrary relations between pairs of initial, final heaps and results, this tool seems to be specially tailored to reason about two runs of the same program, rather than about two different programs. To some extent it is possible to circumvent such restriction by casting two different programs P and P' as a single program with a guard deciding which program to execute, i.e. $\text{if } x \text{ then } P \text{ else } P'$. However, this approach seems a bit awkward and it is not at all clear whether doing this can enable reasoning in terms of relational invariants—which is essential to keep invariants simple.

In recent work [4] we propose a technique—*product program construction*—that reduces relational program reasoning to traditional program reasoning—even for non-structurally equivalent programs. Perhaps more importantly, it enables the use of traditional verification tools, circumventing two of the main issues of techniques supporting relational reasoning. However, this method relies on the ability to interpret relational assertions (predicates on two states) as traditional assertions (predicates on one state), but this is not straightforward when using assertions from relational separation logic. More precisely, an issue arises when trying to interpret the relational assertion

$$R = \begin{pmatrix} p \\ q \end{pmatrix}$$

(two heaps h_1 and h_2 are related by R if p holds in h_1 and q holds in h_2) as a predicate $p \star q$. Note that this interpretation induces a loss of information: predicate R holds for a fixed partition of the heap while the latter holds for any partition of the heap. This loss of information renders our method unsound. Indeed, $\{P \star \text{emp}\} \text{skip}; \text{skip}\{\text{emp} \star P\}$ is a valid separation logic judgment for all P , whereas the following relational judgment is not:

$$\begin{pmatrix} P \\ \text{emp} \end{pmatrix} \text{skip} \sim \text{skip} \begin{pmatrix} \text{emp} \\ P \end{pmatrix}$$

We present in this article an alternative approach that overcomes the difficulties of relational verification by product construction, based on a *weakest precondition calculus* for relational separation logic. The calculus is complete and formalized in the Coq proof assistant, and can be regarded as a first step towards providing tool support for relational methods that enables reasoning about heap manipulating programs.

The formalization provides a framework to reason about a small imperative language—using a deep embedding—with heap manipulating instructions very similar to the one

described in Yang’s article. We have formalized its semantics and provided a soundness proof of the relational weakest precondition. Local reasoning is supported by proving that the calculus is compatible with the frame rule. Also, we have defined an alternative calculus ensuring total relational correctness relying on *variants* (or ranking functions) defining a well-founded order on states.

The Coq formalization has been used to provide a formal proof of the equivalence of Depth First Search and the Schorr-Waite graph marking algorithm, reproducing the proof of the Schorr-Waite graph marking algorithm performed by Yang. We have extended Yang’s proof to the total relational correctness case, hence ensuring that both programs terminate. The formal proof of the Schorr-Waite algorithm required a slightly stronger loop invariant, indicating perhaps a small weakness in the original specification provided in Yang’s article.

We also introduce an extension of the relational calculus beyond structurally equivalent program, preserving relational reasoning over loop invariants, and thus retaining the aforementioned advantages. We illustrate the application of the calculus with the validation of a complex program optimization.

Contents. The rest of the paper is structured as follows: Section 2 describes the formalization of the relational weakest precondition calculus, instantiated with a simple programming setting. In this section, we briefly review relational separation logic and present the main properties of the calculus: soundness and framing. Also, we present a variant of the calculus that ensures termination of both programs. Section 3 presents our main case study, the proof of equivalence between the Schorr-Waite graph marking algorithm w.r.t. depth-first search. Section 4 describes an extension to non-structurally equivalent code.

2 Formalization of Relational Separation Logic

We start this section by introducing a simple program setting and then we provide an overview of relational separation logic. Afterwards, we develop our relational calculus based on weakest precondition computation.

The programming language presented in Figure 1 is a mild extension of the typical setting used in standard separation logic [23] to include list expressions. List values are rather uncommon in similar formalizations of imperative languages but are included here to ease the description of the Depth First Search (DFS) algorithm, which uses a stack as auxiliary data structure. In the figure, α stands for a list variable. We let BExp denote the set of boolean expressions and Stmt the set of statements.

State model. We let \mathcal{S} denote the set of states. A state comprises two components: the store and the heap. The store itself comprises two components to accommodate two types of expressions: natural numbers and lists. Each of the store components is modeled the usual way, as a finite mapping from scalar variables in Var_N to natural numbers and as a finite mapping from list variables in Var_L to lists of natural numbers. We assume that the sets of variables Var_N and Var_L are disjoint. We let $\text{upd}(x, n, s)$ stand for the result of updating the variable x to value n in the store s .

(integer expressions)	$E ::= 0 \mid 1 \mid E + E \mid E \times E \mid E - E \mid \text{hd}(L)$
(boolean expressions)	$B ::= \text{false} \mid B \Rightarrow B \mid E = E \mid E < E \mid L = L$
(list expressions)	$L ::= \alpha \mid \epsilon \mid E::L \mid \text{tl}(L)$
(instructions)	$I ::= x := \text{alloc}(E) \mid x := [E] \mid [E] := E \mid \text{free}(E) \mid$ $x := E \mid \alpha := L \mid \text{assert}(B)$
(statements)	$C ::= I \mid C_1; C_2 \mid \text{if } B \text{ then } C_1 \text{ else } C_2 \mid \text{while } b \text{ do } c \mid \text{skip}$

Fig. 1. Syntax of Programs

The heap is modeled as a finite mapping from locations (natural numbers) to values. The special location 0 is denoted null and cannot belong to the domain of a heap. Heaps are equipped with several operations such as look-up, free, fresh, disjoint union and interact in the expected way:

Expression	Meaning
$\text{freshn}(h, n)$	base location for a sequence of n consecutive free cells in h ;
$\text{look}(h, n)$	value of the cell n in the heap h ;
$\text{mut}(n, m, h)$	result of setting the contents of cell n of heap h to m ;
$\text{dealloc}(h, n)$	result of freeing cell n from heap h ;

Moreover, we let $\text{dom}(h)$ stand for the set of allocated locations of heap h , and $h_1 \uplus h_2$ denote the disjoint union of heaps h_1 and h_2 . In the actual Coq development, failure is captured in an error monad, but for simplicity we omit these details here. Much of the formalization is adapted from Nanevski *et al* [21].

Semantics of basic instructions. The semantics of an instruction i is modeled as a relation $\llbracket i \rrbracket$ on states; the rules are given in Fig. 2. The denotation of an instruction is a relation between states. States are noted as tuples (h, s_i, s_l) where h represents the heap and s_i and s_l denote the stores for integer and list variables, respectively. The instruction $x := \text{alloc}(E)$ evaluates the expression E to a natural number n and then allocates n free contiguous heap cells, initializes them with value 0 and sets the value of x to the first allocated cell. The look up instruction $x := [E]$ evaluates expression E to a location n and if it is allocated it updates the value of variable x to the contents of the heap cell n . The mutation instruction $[E_1] := E_2$ evaluates E_1 to a location n and if n is a valid location in the current heap, this is modified so that it maps n to the result of evaluating E_2 . A field access $x := E.f$ is used as a syntactic sugar of $x := [E+f]$, when the field identifier f represents a known offset. Similarly, we use $E_1.f := E_2$ as a syntax sugar of $[E_1+f] := E_2$. The instruction $\text{free}(E)$ releases the heap cell allocated at the location represented by E . The assert instruction has blocking semantics. The remaining assignments for integer and list variables are completely standard.

Semantics of commands. The semantics $\llbracket c \rrbracket$ of a command is defined as a relation on states (big step style), using as auxiliary definition the semantics of boolean expressions, modeled as a function from states to booleans. The definitions are standard and omitted. Also, we denote $\langle c, \mu \rangle \rightsquigarrow \langle c', \mu' \rangle$ the small-step command semantics and we use \rightsquigarrow^* for its reflexive transitive closure. Obviously these two semantic styles are sound and

$$\begin{array}{l}
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket x := \text{alloc}(E) \rrbracket \doteq s'_i = \text{upd}(x, m, s_i) \\
\quad \wedge s'_l = s_l \wedge h' = h \uplus \left(\biguplus_{i=0}^{n-1} (m+i) \mapsto 0 \right) \\
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket x := [E] \rrbracket \doteq s'_l = s_l \wedge h' = h \wedge s'_i = \text{upd}(x, \text{look}(h, n), s_i) \\
\quad \wedge n \in \text{dom}(h) \wedge n \in (\llbracket E \rrbracket (h, s_i, s_l)) \\
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket [E_1] := E_2 \rrbracket \doteq s'_i = s_i \wedge s'_l = s_l \wedge h' = \text{mut}(n, m, h) \\
\quad \wedge n \in \text{dom}(h) \wedge n \in (\llbracket E_1 \rrbracket (h, s_i, s_l)) \\
\quad \wedge m \in (\llbracket E_2 \rrbracket (h, s_i, s_l)) \\
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket \text{free}(E) \rrbracket \doteq s'_i = s_i \wedge s'_l = s_l \wedge h' = \text{dealloc}(h, n) \\
\quad \wedge n \in \text{dom}(h) \wedge n \in (\llbracket E \rrbracket (h, s_i, s_l)) \\
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket x := E \rrbracket \doteq h' = h \wedge s'_l = s_l \wedge s'_i = \text{upd}(x, n, s_i) \\
\quad \wedge n \in (\llbracket E \rrbracket (h, s_i, s_l)) \\
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket \alpha := L \rrbracket \doteq h' = h \wedge s'_i = s_i \wedge s'_l = \text{upd}(\alpha, xs, s_l) \\
\quad \wedge xs \in (\llbracket L \rrbracket (h, s_i, s_l)) \\
((h, s_i, s_l), (h', s'_i, s'_l)) \in \llbracket \text{assert}(B) \rrbracket \doteq \llbracket b \rrbracket (h, s_i, s_l) \wedge h = h' \wedge s_i = s'_i \wedge s_l = s'_l
\end{array}$$

Fig. 2. Semantics of basic instructions

complete w.r.t. each other, i.e. $\llbracket c \rrbracket \mu \mu'$ if and only if $\langle c, \mu \rangle \rightsquigarrow^* \langle \text{skip}, \mu' \rangle$. Also, we say that a command c is φ -safe if for any μ such that $\varphi \mu$ there exists μ' and c' such that $\langle c, \mu \rangle \rightsquigarrow \langle c', \mu' \rangle$, i.e., c is not stuck in φ -states.

2.1 Relational Calculus

We introduce in this section the relational calculus establishing the validity of relational specifications. Relational judgments are formalized as quadruples of the form $\{\varphi\} c_1 \sim c_2 \{\psi\}$, where φ and ψ are relations on states and c_1 and c_2 are programs, establishing a relation over every pair of executions of c_1 and c_2 , as formalized in the following definition:

Definition 1 (valid relational judgment). *Two commands c_1 and c_2 satisfy the pre and post-relation φ and ψ , denoted by the judgment $\models \{\varphi\} c_1 \sim c_2 \{\psi\}$ if for all states μ_1, μ_2 s.t. $\llbracket \varphi \rrbracket \mu_1 \mu_2$ one of the following holds:*

- c_1 diverges with initial state μ_1 iff c_2 diverges with initial state μ_2 ; or
- for all states μ'_1 and μ'_2 s.t. $\llbracket c_1 \rrbracket \mu_1 \mu'_1$ and $\llbracket c_2 \rrbracket \mu_2 \mu'_2$ we have $\llbracket \psi \rrbracket \mu'_1 \mu'_2$.

Assertions. Rather than representing assertions as syntactic objects, we have modeled them as relations between states. All of the assertions presented in Yang's work have a straightforward interpretation as state relations. The definition of some of them is shown in Figure 3. We let P, Q stand for relational assertions and p, q for standard separation logic assertions.

Adopting a shallow embedding of assertions provides extra flexibility by not committing beforehand to a particular logical language, and allows inheriting all the features of Coq's rich higher-order language. This proved to be convenient when defining

$$\begin{aligned}
\text{Same } st_1 \ st_2 &\doteq st_1.h = st_2.h \\
\text{emp2 } st_1 \ st_2 &\doteq st_1.h = \text{empty} \wedge st_2.h = \text{empty} \\
(P \star Q) \ st_1 \ st_2 &\doteq \exists h_{11} \ h_{12} \ h_{21} \ h_{22}. \\
&\quad st_1.h = h_{11} \uplus h_{12} \wedge st_2.h = h_{21} \uplus h_{22} \\
&\quad \wedge P(h_{11}, st_1.s_i, st_1.s_l)(h_{21}, st_2.s_i, st_2.s_l) \\
&\quad \wedge Q(h_{12}, st_1.s_i, st_1.s_l)(h_{22}, st_2.s_i, st_2.s_l) \\
\binom{p}{q} \ st_1 \ st_2 &\doteq p \ st_1 \wedge q \ st_2
\end{aligned}$$

Fig. 3. Definition of some Relational Assertions

$$\begin{aligned}
\text{wp}(x := \text{alloc}(E)) \ \varphi(h, s_i, s_l) &\doteq \forall n \ m. \ n \in (\llbracket E \rrbracket(h, s_i, s_l)) \wedge m = \text{freshn}(h, n) \Rightarrow \\
&\quad \varphi(h \uplus \biguplus_{i=0}^{n-1} (m+i) \mapsto 0, s_i, s_l) \\
\text{wp}(x := [E]) \ \varphi(h, s_i, s_l) &\doteq \forall n. \ n \in (\llbracket E \rrbracket(h, s_i, s_l)) \Rightarrow \\
&\quad \varphi(h, \text{upd}(x, \text{look}(h, n), s_i), s_l) \\
\text{wp}([E_1] := E_2) \ \varphi(h, s_i, s_l) &\doteq \forall n \ m. \ n \in (\llbracket E_1 \rrbracket(h, s_i, s_l)) \wedge m \in (\llbracket E_2 \rrbracket(h, s_i, s_l)) \Rightarrow \\
&\quad n \in \text{dom}(h) \wedge \varphi(\text{mut}(n, m, h), s_i, s_l) \\
\text{wp}(\text{free}(E)) \ \varphi(h, s_i, s_l) &\doteq \forall n. \ n \in (\llbracket E \rrbracket(h, s_i, s_l)) \Rightarrow \\
&\quad n \in \text{dom}(h) \wedge \varphi(\text{dealloc}(h, n), s_i, s_l) \\
\text{wp}(x := E) \ \varphi(h, s_i, s_l) &\doteq \forall n. \ n \in (\llbracket E \rrbracket(h, s_i, s_l)) \Rightarrow \varphi(h, \text{upd}(x, n, s_i), s_l) \\
\text{wp}(\alpha := L) \ \varphi(h, s_i, s_l) &\doteq \forall xs. \ xs \in (\llbracket L \rrbracket(h, s_i, s_l)) \Rightarrow \varphi(h, s_i, \text{upd}(\alpha, xs, s_l)) \\
\text{wp}(\text{assert}(B)) \ \varphi(h, s_i, s_l) &\doteq \llbracket B \rrbracket(h, s_i, s_l) \wedge \varphi(h, s_i, s_l)
\end{aligned}$$

Fig. 4. Weakest Precondition of basic instructions

a weakest precondition calculus ensuring termination, in which a well-founded relation must be provably decreasing throughout loop iterations—see Subsection 2.2.

Weakest precondition calculus for basic instructions. Most program verification tools rely on weakest precondition calculi rather than program logics: concretely, the prevailing means to verify programs against a pre-condition and a post-condition is to generate a set of proof obligations using a weakest precondition calculus, and finally to discharge the proof obligations using automatic or interactive provers. Our formalization supports a similar methodology for relational judgments, and provides a weakest precondition calculus that computes a set of proof obligations from relational judgments. The weakest precondition of a basic instruction i w.r.t. to a state predicate ϕ is again, a state predicate (a function taking states and returning propositions). Here instead of using λ -abstractions we write the state on the left side as arguments to the wp function. Moreover, by abuse of notation we use pattern matching, i.e. a state is noted as a tuple. The definition of the weakest precondition of the basic instructions is provided in figure 4. Its definition is straightforward and obviously sound w.r.t. the semantics.

Weakest precondition calculus for 2-statements. Our weakest precondition calculus wp_2 operates on 2-statements, which combine two structurally equivalent statements into a single construction. Formally, the set Stm_2 of 2-statements is defined

$$\begin{aligned}
\text{wp}_2 \langle i_1, i_2 \rangle \phi &= \text{wp } i_1 (\lambda m_1. \text{wp } i_2 (\lambda m_2. \phi m_1 m_2)) \\
\text{wp}_2 (c_1; c_2) \phi &= \text{wp}_2 c_1 (\text{wp}_2 c_2 \phi) \\
\text{wp}_2 (\text{if } \langle b, b' \rangle \text{ then } c_1 \text{ else } c_2) \phi &= \Psi_{b,b'} \wedge (b_{\langle 1 \rangle} \Rightarrow \text{wp}_2 c_1 \phi) \wedge (\neg b_{\langle 1 \rangle} \Rightarrow \text{wp}_2 c_2 \phi) \\
\text{wp}_2 (\text{while } \langle b, b' \rangle \text{ do } c) \phi &= \exists \varphi. \varphi \wedge \forall m_1, m_2. \Psi_{b,b'} m_1 m_2 \wedge \Psi_\varphi m_1 m_2 \wedge \Psi_\phi m_1 m_2
\end{aligned}$$

where

$$\begin{aligned}
\Psi_{b,b'} &\doteq b_{\langle 1 \rangle} \Leftrightarrow b'_{\langle 2 \rangle} && \text{guard equivalence} \\
\Psi_\varphi &\doteq \varphi \wedge b_{\langle 1 \rangle} \Rightarrow \text{wp}_2 c \varphi && \text{invariant preservation} \\
\Psi_\phi &\doteq \varphi \wedge \neg b_{\langle 1 \rangle} \Rightarrow \phi && \text{valid postcondition}
\end{aligned}$$

Fig. 5. Relational weakest precondition calculus

inductively by the clauses: i) if i_1 and i_2 are instructions, then $[i_1, i_2]$ is a 2-statement; ii) if c, c_1, c_2 are 2-statements and b, b' are boolean expressions, then $c_1; c_2$, and if $\langle b, b' \rangle$ then c_1 else c_2 , and while $\langle b, b' \rangle$ do c are 2-statements. Each 2-statement yields two structurally equivalent statements; we write $c \triangleright (c_1, c_2)$ to denote that c is a 2-statement whose left and right components are the statements c_1 and c_2 respectively. Conversely, any two structurally equivalent statements yield a 2-statement. Intuitively, a 2-statement encodes the simultaneous execution of its components, restricting the calculus to structurally similar programs. In Section 4 we explain how to remove this restriction by the application of a preliminary program transformation.

The weakest precondition calculus wp_2 is defined inductively on the structure of 2-statements; the rules are given in Figure 5, where $b_{\langle 1 \rangle}$ and $b_{\langle 2 \rangle}$ respectively denote the interpretation of the expression b in the first and second memories, and the extension of connectives to relations is defined in the usual way.

Frame rule. The frame rule lies at the very heart of any separation logic based verification framework, being the cornerstone of so called “local reasoning”. In order to support modular verification, we have shown that it holds on the framework presented in this paper. Let P, Q , and R be relational assertions and c a 2-statement. Then, if

1. the proposition $\forall st_1, st_2. P \ st_1 \ st_2 \Rightarrow \text{wp}_2 c \ Q \ st_1 \ st_2$ holds; and
2. R is independent of the variables modified by c

then the following proposition holds:

$$\forall st_1, st_2. (P \star R) \ st_1 \ st_2 \Rightarrow \text{wp}_2 c \ (Q \star R) \ st_1 \ st_2$$

Intuitively, the hypothesis 1 and 2 implies that the only part of state that program c is allowed to inspect or operate on is described by P and any other part R will remain unchanged after its execution. This simplifying result has been systematically used to ease the verification of the Schorr-Waite algorithm.

One of the most challenging aspects of characterizing the frame rule in our setting is the fact that there is no syntax for assertions, so the customary side condition on R is formulated semantically by defining the **modset** of c , i.e. the set of modified variables, and requiring the validity of R to be independent of it.

Soundness. The calculus is sound, i.e. for all statements c_1 and c_2 , and 2-statement c s.t. $c \triangleright (c_1, c_2)$, and assertions φ and ψ ,

$$\llbracket \varphi \rrbracket \subseteq \llbracket \text{wp}_2 c \psi \rrbracket \implies \vDash \{\varphi\} c_1 \sim c_2 \{\psi\}$$

Moreover, the weakest precondition calculus is sound and complete w.r.t. relational separation logic, i.e. for all statements c_1 and c_2 , and 2-statement c s.t. $c \triangleright (c_1, c_2)$, and assertions φ and ψ ,

$$\llbracket \varphi \rrbracket \subseteq \llbracket \text{wp}_2 c \psi \rrbracket \iff \vdash \{\varphi\} c_1 \sim c_2 \{\psi\}$$

where $\vdash \{\varphi\} c_1 \sim c_2 \{\psi\}$ is used to denote that the judgment is derivable in relational separation logic.

2.2 Total Correctness

One can modify the weakest precondition calculus wp_2 to enforce total correctness. To this end, one must provide for each while statement a variant relation between pairs of initial and final states, and prove that it is a well-founded order (i.e. no infinite descending chains) and that it decreases with each iteration. The clause for the while statement is modified accordingly:

$$\begin{aligned} \text{wp}_2^{\text{tc}} (\text{while } \langle b, b' \rangle \text{ do } c) \phi \mu_1 \mu_2 \doteq \\ \exists \varphi, \exists \mu, \varphi \wedge \forall m_1, m_2. (\Psi_{b, b'} m_1 m_2 \wedge \Psi_\varphi m_1 m_2 \wedge \Psi_\phi m_1 m_2) \wedge \\ \text{wellfounded}(\mu) \wedge \forall m_1, m_2. (\varphi m_1 m_2 \wedge \llbracket b \rrbracket m_1 m_2 \Rightarrow \\ \text{wp}_2 c (\lambda s_1, s_2. \mu (s_1, s_2) (m_1, m_2))) m_1 m_2) \end{aligned}$$

where $\Psi_{b, b'}$, Ψ_φ , and Ψ_ϕ are defined as in Figure 5 (replacing wp_2 by wp_2^{tc}), and μ stands for the variant relation. The predicate $\text{wellfounded}(\mu)$ requires μ to be well-founded to establish the termination of the loop. Notice that we use wp_2 instead wp_2^{tc} in the last line of the formulae above, to avoid redundancy on the verification of termination of c , which is already established by Ψ_φ . Then, assuming termination of instructions, we can prove total correctness, i.e. for all statements c_1 and c_2 , and 2-statement c s.t. $c \triangleright (c_1, c_2)$, and assertions φ and ψ , and memories μ_1 and μ_2 s.t. $\varphi \mu_1 \mu_2$,

$$\llbracket \varphi \rrbracket \subseteq \llbracket \text{wp}_2^{\text{tc}} c \psi \rrbracket \implies \exists \mu'_1, \mu'_2. \llbracket c_1 \rrbracket \mu_1 \mu'_1 \wedge \llbracket c_2 \rrbracket \mu_2 \mu'_2 \wedge \psi \mu'_1 \mu'_2$$

Note that the shallow embedding of assertions plays a crucial role here, a partial application of the variant is used as argument for the wp . This would not be possible if we had established a syntax for the formulae through a deep embedding.

3 Verification of the Schorr-Waite Algorithm

The Schorr-Waite graph marking algorithm is a widely used case study, see Section 5. Yang [29] uses relational separation logic to prove the equivalence between the Schorr-Waite algorithm and depth-first search, and convincingly argues that the proof in relational separation logic is more elegant and more concise than an earlier functional

verification [28] of the SW algorithm in separation logic. In this section we report on a machine-checked proof of the Schorr-Waite algorithm using the weakest precondition calculus described in the previous section. The structure of the proof is similar to Yang's pen-and-paper proof [29]; one difference is that we prove total correctness rather than co-termination.

Algorithm and relational specification. DFS traverses a binary tree marking every node in a depth-first basis. In order to backtrack the tree traversal, it uses a stack as an auxiliary storage to keep track of the parent nodes that need to be revisited. The Schorr-Waite algorithm optimizes the space needed by DFS by removing the stack. The set of nodes to be revisited are encoded as a transformation on the heap structure: pushing a node in the stack is implemented as an inversion of the left edge that is traversed, removing a node from the stack is defined as restoring the original edge. Figure 6 shows a 2-statement merging the Schorr-Waite algorithm (marked with a gray shadow) with DFS.

Verification. We have used the Coq framework to verify the 2-statement in Figure 6 against the specification:

$$\begin{aligned} Pre &\doteq \text{Same} \wedge c = c' \wedge \left(\begin{array}{l} \text{noDang } G \wedge c \in G \cup \{\text{nil}\} \\ \text{noDang } G \wedge c' \in G \cup \{\text{nil}\} \end{array} \right) \\ Post &\doteq \text{Same} \end{aligned}$$

where c and c' represent the corresponding tree roots and G denotes the set of tree nodes. The predicate $\text{noDang } G$ states that G is a set of non-dangling pointers closed under heap reachability:

$$\text{noDang } G \doteq \forall_* x \in G. \exists lr. (x \mapsto l, r, -, -) \wedge l \in G \cup \{\text{nil}\} \wedge r \in G \cup \{\text{nil}\}$$

The additional condition $c \in G$ implies that the set of tree nodes reachable from the root c is a subset of G . The specification states that under initial heaps with the same tree structure with root c , SW and DFS terminate with the same final states.

The application of the wp_2 function to the 2-statement and the postcondition above returns a verification condition that contains an existential quantification for the loop invariant. We have used a slightly modified version of the invariant proposed by Yang [29]:

$$\text{Same} \star \text{uniq } \alpha \wedge \text{Stack } p c \alpha \wedge p = p' \wedge \left(\begin{array}{l} \text{noDang } G \wedge p \in G \wedge c \in G \\ \text{noDang } G \wedge p' \in G \wedge \alpha \subseteq G \cup \{\text{nil}\} \end{array} \right)$$

Basically, the invariant establishes that no dangling pointers can be introduced during the algorithms execution, and provides a relation between the auxiliary stack storage used by DFS and its representation in the Schorr-Waite algorithm. This relation is formalized by the predicate Stack :

$$\begin{aligned} \text{Stack } p c \epsilon &\doteq c = \text{nil} \\ \text{Stack } p c a :: \alpha &\doteq \exists n_0, x. \text{Stack } c n_0 \alpha \star c = a \wedge \\ &\left[\left(\begin{array}{l} c \mapsto n_0, x, \text{Marked}, \text{Left} \\ c \mapsto p, x, \text{Marked}, \text{Left} \end{array} \right) \vee \left(\begin{array}{l} c \mapsto x, n_0, \text{Marked}, \text{Right} \\ c \mapsto x, p, \text{Marked}, \text{Right} \end{array} \right) \right] \end{aligned}$$

```

if  $\langle c \neq \text{nil}, c' \neq \text{nil} \rangle$  then
  [
     $p := c.Left;$            $p' := c'.Left;$ 
     $c.Mark := \text{Marked}$     $c'.Mark := \text{Marked};$ 
     $c.Current := isLeft;$   $c'.Current := isLeft;$ 
     $c.Left := \text{nil}$        $\alpha := c'::\epsilon$ 
  ]
else
  [
     $p := \text{nil}, p' := \text{nil};$ 
     $\alpha := \epsilon$ 
  ]
fi
while  $\langle c \neq \text{nil}, \alpha \neq \epsilon \rangle$  do
  if  $\langle p \neq \text{nil}, p' \neq \text{nil} \rangle$  then
    [  $m := p.Mark, m' := p'.Mark$  ]
  else
    [  $m := \text{Marked}, m' := \text{Marked}$  ]
  fi
  if  $\langle p \neq \text{nil} \wedge m \neq \text{Marked}, p' \neq \text{nil} \wedge m' \neq \text{Marked} \rangle$  then
    [
       $t := p.Left;$ 
       $p.Left := c;$            $\alpha := p'::\alpha;$ 
       $c := p;$                $p'.Mark := \text{Marked};$ 
       $p := t;$                $p'.Current := isLeft;$ 
       $c.Mark := \text{Marked};$   $p' := p'.Left$ 
       $c.Current := isLeft$ 
    ]
  else
    [  $d := c.Current, d' := (hd \alpha).Current$  ]
    if  $\langle d = isLeft, d' = isLeft \rangle$  then
      [
         $t := c.Left;$ 
         $c.Left := p;$ 
         $p := c.Right;$ 
         $c.Right := t;$ 
         $c.Current := isRight$ 
         $(hd \alpha).Current := isRight;$ 
         $p' := (hd \alpha).Right$ 
      ]
    else
      [
         $t := p;$ 
         $p := c;$ 
         $c := p.Right;$ 
         $p.Right := t$ 
         $p' := hd \alpha;$ 
         $\alpha := tl \alpha$ 
      ]
    fi
  fi
fi
done

```

Fig. 6. Schorr-Waite and DFS 2-statement

In particular, when $c \neq \text{nil}$, c is the top element in the stack α and p its left or right child. The remaining stack elements are related inductively. The difference with respect to Yang's invariant consists on the predicate $\text{uniq } \alpha$, that states that the list α does not contain repeated elements. The need for this extra condition became evident when discharging the verification conditions in the Coq proof assistant.

Total correctness. We have also developed a total correctness argument for the Schorr-Waite algorithm using the total correctness version of the weakest precondition calculus presented earlier. Then, we extended the proof with the addition of a variant relation, a lexicographic order similar to the one used by Giorgino et al [12]: let (st_1, st_2) and (st'_1, st'_2) be pairs of states, then $\text{var}(st_1, st_2)(st'_1, st'_2)$ iff one of the following holds:

- the number of unmarked nodes in (st_1, st_2) is smaller than the number of unmarked nodes in (st'_1, st'_2) ,
- the number of unmarked nodes in (st_1, st_2) and (st'_1, st'_2) is the same but the number of nodes in (st_1, st_2) with *Current* field set to *isLeft* is smaller than in (st'_1, st'_2) , or
- the number of unmarked nodes and the number of nodes of (st_1, st_2) and (st'_1, st'_2) is the same but the size of the stack α in (st_1, st_2) is smaller than in (st'_1, st'_2) .

We showed that this is a well-founded order and proved that it holds for the pre and post states of the loop body using the wp_2 calculus. In particular note that of the three ways to construct the order, the first one corresponds to a push, the second one to a swing and the third one to a pop operation.

4 Beyond Structurally Equivalent Programs

A common caveat of syntactic relational methods is the limited support for non structurally equivalent programs. Although this restriction can be circumvented in the setting of relational separation logic by using the embedding rule, the ability to reason in terms of relational loop invariants is still not supported.

In this section, we present a different strategy that cleanly extends the weakest precondition based calculus presented in Section 2 to cope with non structurally equivalent code. We enhance the previous formalism through a preliminary transformation that is performed on the programs to be verified. This syntactic transformation can yield structurally equivalent programs while retaining some semantic properties that ensure that the relational validity on the transformed programs also holds on the original programs.

Let us first make precise the notion of refinement we adopt. We say that c is a refinement of c' , noted $c \succcurlyeq c'$ if the following conditions hold for all μ, μ', μ'' and σ :

- if $\llbracket c' \rrbracket \mu \mu'$ then $\llbracket c \rrbracket \mu \mu'$;
- if $\llbracket c' \rrbracket \mu \mu'$ and $\llbracket c \rrbracket \mu \mu''$ then $\mu' = \mu''$ and
- if c is σ -safe then c' is σ -safe.

We know, under this rather weak definition of refinement, that in order to establish a relational property on two programs c_1 and c_2 , it is sufficient to establish such property for any two programs c'_1 and c'_2 s.t. $c_1 \succcurlyeq c'_1$ and $c_2 \succcurlyeq c'_2$:

$$\begin{array}{c}
\text{(RO)} \frac{}{\vdash c; d \succcurlyeq d; c} \text{ if } \text{fv}(c) \cap \text{fv}(d) = \emptyset \quad \text{(SK)} \frac{}{\vdash c; \text{skip} \succcurlyeq c} \\
\text{(IF1)} \frac{}{\vdash \text{if } b \text{ then } c_1 \text{ else } c_2 \succcurlyeq \text{assert}(b); c_1} \quad \text{(IF2)} \frac{}{\vdash \text{if } b \text{ then } c_1 \text{ else } c_2 \succcurlyeq \text{assert}(-b); c_2} \\
\text{(WHU)} \frac{}{\vdash \text{while } b \text{ do } c \succcurlyeq \text{assert}(b); c; \text{while } b \text{ do } c} \\
\text{(WHS)} \frac{}{\vdash \text{while } b \text{ do } c \succcurlyeq \text{assert}(b); c; \text{assert}(-b)} \\
\text{(IFM)} \frac{}{\vdash \text{if } b \text{ then } c; \text{if } b' \text{ then } c' \succcurlyeq \text{assert}(b \Leftrightarrow b'); \text{if } b \text{ then } c; \text{assert}(b'); c} \\
\text{(LRS)} \frac{}{\vdash \text{for } i=m \text{ to } n \text{ by } k \text{ do } c \succcurlyeq \text{assert}(m \leq n' \leq n); \text{for } i=m \text{ to } n' \text{ by } k \text{ do } c; \\ \text{for } j=i \text{ to } n \text{ by } k \text{ do } c[j/i]} \\
\text{(LT)} \frac{}{\vdash \text{for } i=0 \text{ to } n \text{ by } 1 \text{ do } c \succcurlyeq \text{assert}(n \bmod k = 0); \\ \text{for } i=0 \text{ to } n \text{ by } k \text{ do } (\text{for } j=0 \text{ to } k \text{ by } 1 \text{ do } c[i+j/i])} \\
\text{(R-RI)} \frac{}{\vdash \text{for } i=m \text{ to } n \text{ by } k \text{ do } c \succcurlyeq \text{assert}\left(\left\lceil \frac{n-m}{k} \right\rceil = \left\lceil \frac{n'-m'}{k} \right\rceil\right); \\ \text{for } i=m' \text{ to } n' \text{ by } k \text{ do } c[(i-m'+m)/i]}
\end{array}$$

Fig. 7. Syntactic refinement rules (excerpt)

Lemma 1. For all programs c_1 and c_2 , and c'_1, c'_2 such that $c_1 \succcurlyeq c'_1$ and $c_2 \succcurlyeq c'_2$, if $\models \{\varphi\} c'_1 \sim c'_2 \{\psi\}$ then $\models \{\varphi\} c_1 \sim c_2 \{\psi\}$, provided c'_1 and c'_2 are φ -safe.

Figure 7 provides a set of syntactic rules deriving a refinement relation. For clarity, we introduce the statement for $i = m$ to n by k do c as a syntax sugar for statement $i := m$; while $i < n$ do c ; $i := i + k$. As can be seen in the figure, the rules consist of basic structure transformations. The most complex rules are perhaps (LRS) and (LT), which perform loop range splitting and loop tiling, respectively. The set of refinement rules in Figure 7 is sound, i.e., it induces a refinement relation:

Lemma 2. For all statements c and c' , if $\vdash c \succcurlyeq c'$ then $c \succcurlyeq c'$.

Example: vectorization of sum. Figure 8 presents a simple algorithm that computes the sum of the values of the node elements in a singly linked list. A program vectorization consists on relying on special purpose *SIMD* (single instruction, multiple data) instructions, taking advantage of the associativity and commutativity of the arithmetic computation performed in a program loop. Intuitively, for this particular example the vectorization consists in grouping the loop iterations in chunks of 4 iterations, and performing 4 addition operations simultaneously with the `_mm_add_epi32` instruction. Figure 9 shows the vectorized algorithm. Let n denote the length of the linked list pointed by `head`. The first loop iterates $n \div 4$ times and computes the summation of the first $4 \times (n \div 4)$ elements of the linked list, storing it in the 128-bits vector `sum`. The second loop computes the summation of the remaining $n \bmod 4$ elements and stores it in variable `rest`. The final value is computed by adding to the variable `rest` the partial results stored in the bit vector `sum`.

$$\frac{\text{sum}(\text{list* head}, \text{int size})}{\begin{array}{l} \text{rest} := 0; \\ \text{for } i=0 \text{ to } \text{size} \text{ by } 1 \text{ do} \\ \quad \text{rest} := \text{rest} + \text{head.val}; \\ \quad \text{head} := \text{head.next}; \end{array}}$$
Fig. 8. Original version of `sum` algorithm

By applying a sequence of refinement steps over the original program one can obtain a pair of structurally similar programs. Then, providing a relational invariant becomes much simpler than verifying each of the programs functionally. Indeed, assume that the predicate $\text{EqList}(\text{head}, \text{head}', \text{size})$ holds as precondition, with inductive predicate EqList is defined by the following clauses:

$$\begin{array}{l} \text{EqList}(l_1, l_2, 0) \quad \doteq l_1 = l_2 = \text{null} \\ \text{EqList}(l_1, l_2, n+1) \doteq \text{EqList}(l'_1, l'_2, n) \wedge \exists v, l'_1, l'_2. l_1 \mapsto (v, l'_1) \wedge l_2 \mapsto (v, l'_2) \end{array}$$

Then, in order to verify that original and vectorized algorithms compute the same value, i.e., that $\text{rest} = \text{rest}'$ holds as a relational postcondition, it is sufficient to establish the validity of loop invariants of the form:

$$\text{sum}[0] + \text{sum}[1] + \text{sum}[2] + \text{sum}[3] = \text{rest}$$

and

$$\text{rest}' + \text{sum}[0] + \text{sum}[1] + \text{sum}[2] + \text{sum}[3] = \text{rest}$$

Notice that these relational loop invariants are much simpler than those required in a functional verification of the algorithm.

5 Related Work

Relational methods and program verification techniques have been intimately connected since their origins. In particular, methods based on program refinement, program equivalence, and logical relations have been used widely to reason about program correctness. In this respect, it is perhaps surprising that relational program logics have only been introduced recently. Benton [7] develops a relational Hoare logic for a small imperative language and shows how program optimizations can be validated using relational reasoning. Other relational logics include Yang's relational separation logic [29] and Barthe, Grégoire and Zanella's probabilistic relational Hoare logic [5]. More recently, Nanevski, Banerjee and Garg developed a relational separation logic for Hoare type theory [22]. It extends Yang's logic to a richer programming and specification language, and is tailored for reasoning information flow; the logic is formalized in the Coq proof assistant; in contrast to our formalization, it uses a shallow embedding of programs. Independently, Beringer [8] provided a reconstruction of relational separation logic based on a notion of decomposition that allows reducing relational program logics to standard program logics; the soundness of the logic is formalized in the Isabelle

```

ssesum(list* head', int size)
sum = _mm_set1_epi32(0);
for i=0 to size - 3 by 4 do
  curr := _mm_insert_epi32(curr, head'.val, 0);
  head' := head'.next;
  curr := _mm_insert_epi32(curr, head'.val, 1);
  head' := head'.next;
  curr := _mm_insert_epi32(curr, head'.val, 2);
  head' := head'.next;
  curr := _mm_insert_epi32(curr, head'.val, 3);
  head' := head'.next;
  sum := _mm_add_epi32(sum, curr);
rest' := 0;
for j=i to size by 1 do
  rest' := rest'+head'.val;
  head' := head'.next;
rest' := rest'+_mm_extract_epi32(sum, 0)+_mm_extract_epi32(sum, 1)+
  _mm_extract_epi32(sum, 2) + _mm_extract_epi32(sum, 3);

```

Fig. 9. SSE optimized version of `sum` algorithm

proof assistant. In addition to these general-purpose logics, specialized relational logics have been developed for specific properties, and especially information flow [1].

On the formalization side, there have been many machine-checked accounts of separation logic in proof assistants, e.g. [3, 26], including some frameworks designed to support automated reasoning in separation logic [2, 11, 18, 20]. Moreover, the Schorr-Waite algorithm is a classical example in program verification, and has been verified formally using a variety of tools and techniques. Suzuki [25] provides an early machine-supported proof of the Schorr-Waite algorithm using an automated verifier for pointer programs. More recently, Bornat [10] provides a machine-checked proof of the algorithm in the Jape proof assistant. Subsequently, Mehta and Nipkow [19], Hubert and Marché [14], Bubel [6], Jacobs and Piessens [16] formalize the algorithm in Isabelle, Caduceus, KeY and VeriFast respectively. More recently, Giorgino et al [12] prove the correctness and termination of the algorithm in Isabelle, using refinement. All these formalizations use standard program logics.

6 Conclusion

Relational separation logic is a powerful tool devised for reasoning about the relation between heap manipulating programs. To the best of our knowledge, we have formalized in the Coq proof assistant the first certified weakest precondition calculus for relational separation logic. We illustrated its usefulness and scalability by proving a challenging case study: the correctness of the Schorr-Waite graph marking algorithm.

The Coq development has been done using `ssreflect` library which greatly improves the conciseness of the proofs. For example, the relational weakest precondition,

soundness proofs, the definition and specification and proof of the Schorr-Waite graph marking algorithm and Depth First Search take 1586 lines of specification and 3538 lines of proofs. We believe that the formalization of the verification setting and the formal proof of the algorithms poses no significant overhead over hand-written proofs.

In the future, it would be interesting to formalize the modular proof of the algorithm reported in [9] and to prove the equivalence between different implementations of ADTs; for the latter, we believe that the extensions to non-structurally equivalent code will prove crucial. Another line of work is to extend our formalization to reason about concurrent separation logic [27] and verify the correctness of lock-free algorithms [13].

Acknowledgement. The authors would like to thank Aleksander Nanevski for introducing `ssreflect` to us and patiently explaining some of its main features.

References

1. Amtoft, T., Bandhakavi, S., Banerjee, A.: A logic for information flow in object-oriented programs. In: Morrisett, G., Peyton Jones, S. (eds.) *Principles of Programming Languages*, pp. 91–102. ACM, New York (2006)
2. Appel, A.: Tactics for separation logic (January 2006) (unpublished manuscript), <http://www.cs.princeton.edu/~appel/papers/septacs.pdf>
3. Appel, A.W., Blazy, S.: Separation logic for small-step cminor. In: Schneider, K., Brandt, J. (eds.) *TPHOLs 2007*. LNCS, vol. 4732, pp. 5–21. Springer, Heidelberg (2007)
4. Barthe, G., Crespo, J.M., Kunz, C.: Relational verification using product programs. In: Butler, M., Schulte, W. (eds.) *FM 2011*. LNCS, vol. 6664, pp. 200–214. Springer, Heidelberg (2011)
5. Barthe, G., Grégoire, B., Zanella Béguelin, S.: Formal certification of code-based cryptographic proofs. In: Shao, Z., Pierce, B.C. (eds.) *Principles of Programming Languages*, pp. 90–101. ACM Press, New York (2009)
6. Beckert, B., Hähnle, R., Schmitt, P.H. (eds.): *Verification of Object-Oriented Software. The KeY Approach*. LNCS (LNAI), vol. 4334. Springer, Heidelberg (2007)
7. Benton, N.: Simple relational correctness proofs for static analyses and program transformations. In: Jones, N.D., Leroy, X. (eds.) *Principles of Programming Languages*, pp. 14–25. ACM Press, New York (2004)
8. Beringer, L.: Relational program logics in decomposed style (2010) (submitted)
9. Bodík, R., Chandra, S., Galenson, J., Kimelman, D., Tung, N., Barman, S., Rodarmor, C.: Programming with angelic nondeterminism. In: *Principles of Programming Languages*, pp. 339–352 (2010)
10. Bornat, R.: Proving pointer programs in hoare logic. In: Backhouse, R.C., Oliveira, J.N. (eds.) *MPC 2000*. LNCS, vol. 1837, pp. 102–126. Springer, Heidelberg (2000)
11. Gast, H.: Lightweight separation. In: Mohamed, O., Muñoz, C., Tahar, S. (eds.) *TPHOLs 2008*. LNCS, vol. 5170, pp. 199–214. Springer, Heidelberg (2008)
12. Giorgino, M., Strecker, M., Matthes, R., Pantel, M.: Verification of the Schorr-Waite algorithm - From trees to graphs (January 2010)
13. Gotsman, A., Cook, B., Parkinson, M.J., Vafeiadis, V.: Proving that non-blocking algorithms don't block. In: Shao, Z., Pierce, B.C. (eds.) *Principles of Programming Languages*, pp. 16–28. ACM, New York (2009)
14. Hubert, T., Marché, C.: A case study of c source code verification: the schorr-waite algorithm. In: Aichernig, B., Beckert, B. (eds.) *Software Engineering and Formal Methods*, pp. 190–199. IEEE Computer Society, Los Alamitos (2005)

15. Ishtiaq, S., O'Hearn, P.: BI as an assertion language for mutable data structures. In: *Principles of Programming Languages*, pp. 14–26 (2001)
16. Jacobs, B., Piessens, F.: The VeriFast program verifier. Technical Report CW-520, Katholieke Universiteit Leuven (2008)
17. Jacobs, B., Smans, J., Piessens, F.: Verifying the composite pattern using separation logic. In: *Workshop on Specification and Verification of Component-Based Systems, Challenge Problem Track* (November 2008)
18. McCreight, A.: Practical tactics for separation logic. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) *TPHOLS 2009*. LNCS, vol. 5674, pp. 343–358. Springer, Heidelberg (2009)
19. Mehta, F., Nipkow, T.: Proving pointer programs in higher-order logic. *Inf. Comput.* 199(1-2), 200–227 (2005)
20. Myreen, M.O.: Separation logic adapted for proofs by rewriting. In: Kaufmann, M., Paulson, L.C. (eds.) *ITP 2010*. LNCS, vol. 6172, pp. 485–489. Springer, Heidelberg (2010)
21. Nanevski, A., Vafeiadis, V., Berdine, J.: Structuring the verification of heap-manipulating programs. In: Hermenegildo, M., Palsberg, J. (eds.) *Principles of Programming Languages*, pp. 261–274. ACM, New York (2010)
22. Nanevski, A., Banerjee, A., Garg, D.: Verification of information flow and access control policies with dependent types. In: *2011 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Los Alamitos (2011)
23. O'Hearn, P.W., Reynolds, J., Yang, H.: Local reasoning about programs that alter data structures. In: Fribourg, L. (ed.) *CSL 2001 and EACSL 2001*. LNCS, vol. 2142, pp. 1–19. Springer, Heidelberg (2001)
24. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: *Logic in Computer Science*, pp. 55–74. IEEE Computer Society, Los Alamitos (2002)
25. Suzuki, N.: *Automatic Verification of Programs with Complex Data Structures*. PhD thesis, Stanford University (1976)
26. Tuch, H., Klein, G., Norrish, M.: Types, bytes, and separation logic. In: Hofmann, M., Felleisen, M. (eds.) *Principles of Programming Languages*, pp. 97–108. ACM, New York (2007)
27. Vafeiadis, V., Parkinson, M.J.: A marriage of rely/guarantee and separation logic. In: Caires, L., Vasconcelos, V.T. (eds.) *CONCUR 2007*. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007)
28. Yang, H.: *Local reasoning for stateful programs*. PhD thesis, University of Illinois, Urbana, IL, USA (2001)
29. Yang, H.: Relational separation logic. *Theoretical Computer Science* 375(1-3), 308–334 (2007)
30. Yang, H., Lee, O., Berdine, J., Calcagno, C., Cook, B., Distefano, D., O'Hearn, P.: Scalable shape analysis for systems code. In: Gupta, A., Malik, S. (eds.) *CAV 2008*. LNCS, vol. 5123, pp. 385–398. Springer, Heidelberg (2008)