

Relational program logics in decomposed style

Lennart Beringer*

Department of Computer Science, Princeton University,
35 Olden Street, Princeton NJ 08540
eberinge@cs.princeton.edu

Abstract

Suppose you have developed a Hoare logic for your favorite programming language. You have justified the logic by defining an operational model of the language and constructing a soundness proof that interprets triples as partial correctness assertions. Now you want to interpret program analyses, validate security properties such as noninterference, or justify program transformations. You observe that all these properties are relational: they are naturally phrased over *pairs* of executions, for programs that are either identical or closely related. Is your program logic up to the task?

This article shows how to formally decompose termination-insensitive relational program logics into judgements from unary logics. We develop relational predicate transformers, present laws that govern their decomposition along the phrase structure, and relate them to their unary counterparts. We apply our findings to justify variants of Benton’s Relational Hoare Logic (RHL) for a language with objects, extend the logics to auxiliary state, derive a noninterference analysis in the style of Banerjee-Naumann, and develop relational interpretations of separation logic. As related executions do not have to refer to the same program syntax or employ the same notion of state, decomposition can in principle be applied to cross-language verification problems, as long as suitable one-execution logics exist.

Categories and Subject Descriptors D.3.1 [PROGRAMMING LANGUAGES]: Formal Definitions and Theory — Semantics; F.3.1 [LOGICS AND MEANINGS OF PROGRAMS]: Specifying and Verifying and Reasoning about Programs — Logics of programs; F.4.1 [MATHEMATICAL LOGIC AND FORMAL LANGUAGES]: Mathematical Logic — Mechanical theorem proving

General Terms Languages, Theory, Verification

Keywords Relational program logics, Decomposition, Noninterference, Self-composition, Separation logic, Relational predicate transformers

1. Introduction

Verification formalisms and tools based on Hoare logics are typically designed with one-execution properties in mind: their par-

* Supported in part by NSF grant CNS-0910448 and AFOSR award FA9550-09-1-0138.



tial or total correctness interpretation involves a single operational judgement.

However, many program properties are relational: they are naturally phrased as statements over pairs of executions, stipulating that the terminal states of two runs stand in relation S whenever the initial states are in relation R . Examples include not only “obviously relational” properties such as program transformations or noninterference [29], but also extensional interpretations of types, and program analyses that may be found in almost any compiler [9]. This article shows that unary verification technology nevertheless suffices in many cases for certifying relational properties in the foundational sense [3].

We introduce a novel technique called *relational decomposition* which reduces the task of proving a simulation diagram to the task of proving two separate unary assertions, one for $s \xrightarrow{P} t$ and one for $s' \xrightarrow{P'} t'$.

These assertions arise from witnesses ϕ that relate terminal states of executions to the left with initial states of executions to the right. Witnesses of this form were first introduced by Beringer and Hofmann [12] in order to derive interpretations of type systems for noninterference according to the approach of self-composition [8, 15, 20, 32]. The present paper establishes the general applicability of the technique, and some of its core properties:

1. We show that relational decomposition applies to arbitrary pairs of transition systems, rather than being limited to relational properties of single programs, or relationships between programs from the same language. Indeed, primed states in the above diagram are not required to be of the same type as non-primed states, and P and P' may be from different languages. Any witness that makes the two unary properties valid guarantees the simulation property (soundness).
2. We establish the completeness of decomposition in the general case: witnesses exist whenever the simulation property holds. The space of witnesses ϕ is characterized by an inclusion property between relational predicate transformers. We present laws that relate these transformers to their unary counterparts.



3. We extend relational decomposition by a generic parameterization mechanism which can be instantiated to communicate arbitrary additional information between pre- and postrelations.
4. We transfer relational decomposition to a setting where the operational model contains explicit error states, based on unary verification technology that exposes such states in assertions.
5. We instantiate relational decomposition to a number of relational and nonrelational formalisms, illustrating the previous items in concrete settings, for a language of simple objects. Specifically, we outline (i) variants of Benton’s relational Hoare logic (RHL, [9]) (ii) an interpretation of a type system for noninterference, for a security model in the style of Banerjee-Naumann [5], and (iii) relational interpretations of unary and relational separation logic [26, 27, 36].

The concrete instantiations are derived from suitable unary logics, and combine witnesses associated with subphrases to witnesses for compound phrases exactly as in [12], highlighting that the decomposition is in fact syntax- rather than type-driven.

The formal setup and principle of decomposition are as follows.

1.1 The principle of decomposition

For the purpose of this paper, a transition system \mathcal{T} over state space S and labels (thought of as program phrases) \mathcal{P} is a ternary relation $\mathcal{T} \subseteq S \times \mathcal{P} \times S$.

A transition system \mathcal{T} gives rise to a one-execution specification system with assertions $A \in \mathcal{A} \equiv S \Rightarrow S \Rightarrow \mathbb{T}$ (where \mathbb{T} is the set of truth values), interpreted by the partial-correctness condition

$$\models^{\mathcal{T}} P : A \equiv \forall s, t. (s, P, t) \in \mathcal{T} \rightarrow A s t.$$

The use of binary relations for one-execution properties follows the VDM discipline [19] and simplifies the treatment of auxiliary variables and quantifiers when compared to separate pre- and post-conditions.

The formal notion of simulation is as follows.

DEFINITION 1. For $\mathcal{T} \subseteq S \times \mathcal{P} \times S$ and $\mathcal{T}' \subseteq S' \times \mathcal{P}' \times S'$, let $R, S \subseteq S \times S'$. Programs $P \in \mathcal{P}$ and $P' \in \mathcal{P}'$ are $R \Rightarrow S$ -similar, notation $\models_{\mathcal{T}}^{\mathcal{T}'} P \sim P' : R \Rightarrow S$, if for all s, s', t , and t' with $(s, P, t) \in \mathcal{T}$ and $(s', P', t') \in \mathcal{T}'$, sRs' implies tSt' .

Unless explicitly remarked otherwise, we follow the convention that nonprimed entities (states, phrases, ...) refer to \mathcal{T} and primed entities to \mathcal{T}' . Relations from $S \times S'$ are often written infix.

Examples of properties of this shape for $S = S'$ include determinism (choose R and S to be equality), liveness of variables (choose R and S to be equality on live-in variables) and slicing, program transformations, and termination-insensitive versions of properties considered by [9]. Further instances arise when the condition $S = S'$ is dropped, including variations of compiler correctness, refinement, and abstract interpretation.

The core of relational decomposition consists of the unary assertions defined by the operators Dec_L and Dec_R

$$\begin{aligned} Dec_L R \phi s t &= \forall s'. sRs' \rightarrow t\phi s' \\ Dec_R S \phi s' t' &= \forall t. t\phi s' \rightarrow tSt' \end{aligned}$$

Given some $\phi \subseteq S \times S'$, operator Dec_L constructs a unary assertion for the upper left triangle, i.e. for the execution of P . In fact, the construction uniformly applies for all types S' subject to $R \subseteq S \times S'$. Similarly, Dec_R constructs a unary assertion for the lower right triangle (the execution of P'), uniformly for types S with $S \subseteq S \times S'$.

The operators are motivated by the following result.

LEMMA 1. (Soundness) Suppose $\models^{\mathcal{T}} P : Dec_L R \phi$ and $\models^{\mathcal{T}'} P' : Dec_R S \phi$. Then $\models_{\mathcal{T}}^{\mathcal{T}'} P \sim P' : R \Rightarrow S$.

Thus, the task of verifying $\models_{\mathcal{T}}^{\mathcal{T}'} P \sim P' : R \Rightarrow S$ is reduced to the task of exhibiting an arbitrary ϕ that satisfies the two one-execution properties.

How can one find suitable witnesses ϕ ? In [12], we considered the special case of (termination-insensitive) *noninterference* [29], where R and S are state indistinguishability relations \sim_X with respect to sets X of (“low” security) program variables. We showed that suitable witnesses ϕ can be obtained mechanically from type systems for noninterference from the literature [18, 33], in a rule-by-rule fashion. We also showed the completeness of the characterization, by exhibiting a witness whenever noninterference holds. This constructed witness was phrased in terms of an operational judgement, similar to the standard argument in a proof of (relative) completeness proofs of program logics [21, 25]. However, we did not study whether the witnesses ϕ are unique, and our analysis was restricted to noninterference. Specifically, the development was hampered by the identification of \mathcal{T} with \mathcal{T}' , and the fact that $Dec_L R \phi$ and $Dec_R S \phi$ were combined to a single assertion.

1.2 Synopsis

Section 2 presents a more detailed analysis of decomposition witnesses, for general simulation diagrams rather than noninterference. We then derive a decomposed version of relational Hoare logic, for an extension of Benton’s language with heap objects.

Section 3 introduces some constructions on transition systems; these are subsequently used to derive a parameterization mechanism. An application to noninterference for the language with objects ensues, where partial bijections hide the nondeterminism of object allocation and are communicated from pre- to postrelations.

Section 4 considers simulations for operational semantics with explicit error states. We develop a unary program logic that exposes error states in assertions while at the same time retaining the applicability of decomposition. Having adapted the parameterized logic from Section 3 to this setting, we outline relational interpretations of unary and relational separation logic.

We conclude with a discussion of future and related work. The entire development presented has been formalized in Isabelle/HOL. A snapshot of the development is available at [35].

Before continuing, we emphasize that in contrast to Benton’s and Yang’s logics, all our interpretations are termination-insensitive: simulations are vacuously fulfilled if either execution fails to terminate. This design decision is motivated by the fact that even in the unary setting, proof techniques for termination (appropriate measures, i.e. variants) are rather different from those for partial-correctness properties (invariants). A second reason is that applications such as compiler verification actually relax termination-sensitivity to at least an asymmetric form [23]. Thus, termination appears sufficiently orthogonal to the functional aspects of relational behaviour to be treated separately. Nevertheless, we acknowledge that (and point out where) our design decision has repercussions on the proof rules we are able to derive.

2. Decomposition of relational Hoare logics

2.1 Relational decomposition operators

We first observe that Dec_L is covariant in ϕ and contravariant in R , i.e. for $\phi \subseteq \psi$ and $Q \subseteq R$, $Dec_L R \phi s t$ implies $Dec_L Q \psi s t$, while Dec_R is covariant in S and contravariant in ϕ , i.e. for $\psi \subseteq \phi$ and $S \subseteq T$, $Dec_R S \phi s' t'$ implies $Dec_R T \psi s' t'$. We also note the identity $Dec_R S \phi = Dec_L \phi^{-1} S^{-1}$.

Next, we characterize the witnesses suitable for establishing Lemma 1. To this end, consider the operators

$$\begin{aligned} \phi_L^{\mathcal{T}} P R &= \{(t, s'). \exists s. (s, P, t) \in \mathcal{T} \wedge sRs'\} \\ \phi_R^{\mathcal{T}'} P' S &= \{(t, s'). \forall t'. (s', P', t') \in \mathcal{T}' \rightarrow tSt'\} \end{aligned}$$

The former constructs a candidate for ϕ according to the upper left triangle in the diagram, given R and P . The latter constructs a (in general different) candidate ϕ according to the lower right triangle in the diagram, given S and P' .

By construction, these operators yield valid specifications for their defining triangles, i.e. they satisfy $\models^T P : Dec_L R (\phi_L^T P R)$ and $\models^{T'} P' : Dec_R S (\phi_R^{T'} P' S)$. Furthermore, they are covariant in their second argument, satisfy the laws in Figure 1, and relate to each other by $\phi_L^T P (\phi_R^{T'} P' T) \subseteq \phi_R^{T'} P' (\phi_L^T P T)$. In particular, $\phi_L^T P R$ is the *least* relation satisfying the left triangle, and $\phi_R^{T'} P' S$ is the *greatest* relation satisfying the right triangle:

LEMMA 2. *If $\models^T P : Dec_L R \phi$ then $\phi_L^T P R \subseteq \phi$. If $\models^{T'} P' : Dec_R S \phi$ then $\phi \subseteq \phi_R^{T'} P' S$.*

Thus, any witness ϕ from Lemma 1 is sandwiched between the two operators, i.e. satisfies $\phi_L^T P R \subseteq \phi \subseteq \phi_R^{T'} P' S$.

Conversely, either operator is suitable as a witness:

LEMMA 3. (Completeness) *Suppose $\models_{\mathcal{T}}^{T'} P \sim P' : R \Rightarrow S$. Then*

- $\models^T P : Dec_L R (\phi_L^T P R)$ and $\models^{T'} P' : Dec_R S (\phi_R^{T'} P' S)$
- $\models^T P : Dec_L R (\phi_R^{T'} P' S)$ and $\models^{T'} P' : Dec_R S (\phi_L^T P R)$

Combining Lemmas 2 and 3, we obtain that $\models_{\mathcal{T}}^{T'} P \sim P' : R \Rightarrow S$ is equivalent to the inclusion $\phi_L^T P R \subseteq \phi_R^{T'} P' S$.

While the operators are defined from the relational perspective, they are intimately connected with the unary transformers

Strongest postcondition :

$$SP_P^T(X) = \{t. \exists s \in X. (s, P, t) \in \mathcal{T}\}$$

Weakest liberal precondition :

$$WLP_{P'}^{T'}(Y') = \{s'. \forall t'. (s', P', t') \in \mathcal{T}' \rightarrow t' \in Y'\}$$

where X and Y' are state sets from \mathcal{T} and \mathcal{T}' , respectively. Each operator arises from the appropriate unary transformer if one keeps the state in the opposite execution invariant:

$$\begin{aligned} \phi_L^T P R &= \{(t, s'). t \in SP_P^T(\{s. sRs'\})\} \\ \phi_R^{T'} P' S &= \{(t, s'). s' \in WLP_{P'}^{T'}(\{t'. tSt'\})\} \end{aligned}$$

Thus, $R \Rightarrow S$ -similarity is soundly and completely characterized by the inclusion¹

$$\begin{aligned} \{(t, s'). t \in SP_P^T(\{s. sRs'\})\} \\ \subseteq \{(t, s'). s' \in WLP_{P'}^{T'}(\{t'. tSt'\})\}. \end{aligned} \quad (1)$$

Finally, we note that the relational transformers

Strongest postrelation :

$$\begin{aligned} SR_{P, P'}^{T, T'}(R) = \\ \{(t, t'). \exists s s'. (s, P, t) \in \mathcal{T} \wedge (s', P', t') \in \mathcal{T}' \wedge sRs'\} \end{aligned}$$

Weakest liberal preration :

$$\begin{aligned} WLR_{P, P'}^{T, T'}(S) = \\ \{(s, s'). \forall t t'. (s, P, t) \in \mathcal{T} \rightarrow (s', P', t') \in \mathcal{T}' \rightarrow tSt'\} \end{aligned}$$

satisfy the following properties.

¹Replacing $WLP_{P'}^{T'}(Y)$ by the weakest precondition operator

$$WP_{P'}^{T'}(Y') = \{s'. \exists t' \in Y'. (s', P', t') \in \mathcal{T}'\}$$

yields a similar characterization for simulation diagrams where each \mathcal{T} -execution implies the existence of a corresponding \mathcal{T}' -execution, i.e. the property that $(s, P, t) \in \mathcal{T}$ and sRs' imply the existence of some t' with $(s', P', t') : \mathcal{T}'$ and tSt' . We do not study this variant any further in the present paper but observe that diagrams of this shape are considered in Leroy's CompCert project [23] and in Schmidt's work on relational abstract interpretation [30].

LEMMA 4. *We have*

1. $\phi_L^T P (WLR_{P, P'}^{T, T'}(S)) \subseteq \phi_R^{T'} P' S$
2. $\phi_L^T P R \subseteq \phi_R^{T'} P' (SR_{P, P'}^{T, T'}(S))$
3. $WLR_{P, P'}^{T, T'}(S) = \{(s, s'). s' \in WLP_{P'}^{T'}(\{t'. s \in WLP_P^T(\{t. tSt'\})\})\} = \{(s, s'). s \in WLP_P^T(\{t'. s' \in WLP_{P'}^{T'}(\{t'. tSt'\})\})\}$
4. $SR_{P, P'}^{T, T'}(R) = \{(t, t'). t' \in SP_{P'}^{T'}(\{s'. t \in SP_P^T(\{s. sRs'\})\})\} = \{(t, t'). t \in SP_P^T(\{s. t' \in SP_{P'}^{T'}(\{s'. sRs'\})\})\}$

2.2 Application: Relational Hoare Logic in decomposed form

We instantiate the above development to the situation where \mathcal{T} and \mathcal{T}' coincide and are equal to the operational judgement of an imperative language with objects.

We assume infinite and distinct categories of variables $x, y \in \mathcal{X}$, field identifiers $f \in \mathcal{F}$, class identifiers $c \in \mathcal{C}$, and locations $\ell \in \mathcal{L}$. The space of finite partial functions from A to B is denoted by $A \rightarrow B$, and the space of total functions by $A \Rightarrow B$. Operations and constructions such as update, domain, and range, are defined and denoted in the standard fashion. A value $v \in \mathcal{V}$ is either an integer value i , a location ℓ , or Null. Value expressions $e \in \mathcal{E}$ are value constants, variables, or binary operators, while boolean expressions b are binary predicates over values.

The syntax is given by the grammar

$$C \in \mathcal{P} ::= \mathbf{Skip} \mid x := e \mid x := \mathbf{new} \ c \ \iota \mid x := y.f \mid x.f := e \mid C; D \mid \mathbf{While} \ b \ \mathbf{do} \ C \mid \mathbf{If} \ b \ \mathbf{then} \ C \ \mathbf{else} \ D$$

where $\iota \in \mathcal{F} \rightarrow \mathcal{E}$ specifies the initialization of fields in the absence of a formalized class system.

The operational semantics is defined over the semantic domains of objects, heaps, stores, and states

$$\begin{aligned} o \in \mathcal{O} &\equiv \mathcal{C} \times (\mathcal{F} \rightarrow \mathcal{V}) & s \in \mathcal{R} &\equiv \mathcal{X} \Rightarrow \mathcal{V} \\ h \in \mathcal{H} &\equiv \mathcal{L} \rightarrow \mathcal{O} & \sigma \in \Sigma &\equiv \mathcal{R} \times \mathcal{H} \end{aligned}$$

²We write $\llbracket e \rrbracket_s$ and $\llbracket b \rrbracket_s$ for the (heap-independent) evaluation of value and boolean expressions, respectively, and map the former operation over initialization maps in the expected manner.

The transition system $\mathcal{T}_{\text{Obj}} \subseteq \Sigma \times \mathcal{P} \times \Sigma$, with pretty-printed judgements $\sigma \xrightarrow{C} \tau$, is defined essentially as in [34] for the heap-free instruction forms, and by rules such as

$$\text{OPNEW} \frac{\ell \notin \text{locs}(s, h)}{(s, h) \xrightarrow{x := \mathbf{new} \ c \ \iota} (s[x \mapsto \ell], h[\ell \mapsto (c, \llbracket \iota \rrbracket_s)])}$$

for the object-manipulating instructions, where $\text{locs} \ \sigma$ denotes the set of all locations ℓ occurring in σ . For now, attempts to access dangling pointers or undefined fields of allocated objects result in the absence of an execution.

In accordance with the setup of Section 1.1, we introduce a unary logic with judgements of the form $\triangleright C : A$ where $A : \Sigma \Rightarrow \Sigma \Rightarrow \mathbb{T}$. Selected proof rules are given in Figure 2. The logic can be proven sound and complete (the latter relative to the ambient logic HOL) using standard techniques [21, 25].

THEOREM 1. $\triangleright C : A \iff \models^{\mathcal{T}_{\text{Obj}}} C : A$.

²The use of s, t, \dots for stores as well as for abstract states of transition systems should not lead to confusion, as instantiations to the concrete language are always discussed separately from the abstract treatment.

$$\begin{array}{l} \phi_L^T P(R \cap T) \subseteq \phi_L^T P R \cap \phi_L^T P T \quad \phi_L^T P(U \cup T) = \phi_L^T P U \cup \phi_L^T P T \quad \phi_L^T P \emptyset = \emptyset \\ \phi_R^T P'(S \cap T) \supseteq \phi_R^T P' S \cap \phi_R^T P' T \quad \phi_R^T P'(U \cup T) \supseteq \phi_R^T P' U \cup \phi_R^T P' T \quad \phi_R^T P' \emptyset = \{(t, s'). \forall t'. (s', P', t') \notin T'\} \end{array}$$

Figure 1. Interaction between $\phi_L^T P R$ and $\phi_R^{T'} P' S$ and set operations

$$\begin{array}{c} \frac{}{\triangleright x := e : \lambda(s, h) \tau. \tau = (s[x \mapsto \llbracket e \rrbracket_s], h)} \quad \frac{}{\triangleright x := \mathbf{new} \ c \ \iota : \lambda(s, h) \tau. \exists \ell \notin \mathit{locs}(s, h). \tau = (s[x \mapsto \ell], h[\ell \mapsto (c, \llbracket \ell \rrbracket_s)])} \\ \frac{}{\triangleright x := y.f : \lambda(s, h) \tau. \exists \ell \ c \ F \ v. \ s \ y = \ell \wedge h \ \ell = (c, F) \wedge Ff = v \wedge \tau = (s[x \mapsto v], h)} \quad \frac{\triangleright C : A \quad \triangleright D : B}{\triangleright \mathbf{If} \ b \ \mathbf{then} \ C \ \mathbf{else} \ D : \lambda(s, h) \tau. (\llbracket b \rrbracket_s \rightarrow A(s, h) \tau) \wedge (\neg \llbracket b \rrbracket_s \rightarrow B(s, h) \tau)} \\ \frac{\triangleright C : A \quad \triangleright D : B}{\triangleright C; D : \lambda \sigma \tau. \exists \rho. A \sigma \rho \wedge B \rho \tau} \quad \frac{\triangleright C : B \quad \forall s \ h. \neg \llbracket b \rrbracket_s \rightarrow A(s, h) \quad \forall s \ h \ \rho \ \tau. \llbracket b \rrbracket_s \rightarrow B(s, h) \rho \rightarrow A \rho \tau \rightarrow A(s, h) \tau}{\triangleright \mathbf{While} \ b \ \mathbf{do} \ C : \lambda \sigma \ t \ k. A \sigma(t, k) \wedge \neg \llbracket b \rrbracket_t} \end{array}$$

Figure 2. Selected proof rules for logic $\triangleright C : A$.

Instantiating the relational decomposition operators to $\mathcal{T} = \mathcal{T}_{\text{Obj}}$ and/or $\mathcal{T}' = \mathcal{T}_{\text{Obj}}$ yields decomposition laws such as

$$\begin{array}{l} \phi_L^{\mathcal{T}_{\text{Obj}}} C; D \ R = \phi_L^{\mathcal{T}_{\text{Obj}}} D (\phi_L^{\mathcal{T}_{\text{Obj}}} C \ R) \\ \phi_R^{\mathcal{T}'_{\text{Obj}}} C'; D' \ S = \phi_R^{\mathcal{T}'_{\text{Obj}}} C' (\phi_R^{\mathcal{T}'_{\text{Obj}}} D' \ S) \\ WLR_{C;D, C';D'}^{\mathcal{T}_{\text{Obj}}, \mathcal{T}'_{\text{Obj}}}(S) = WLR_{C, C'}^{\mathcal{T}_{\text{Obj}}, \mathcal{T}'_{\text{Obj}}}((WLR_{D, D'}^{\mathcal{T}_{\text{Obj}}, \mathcal{T}'_{\text{Obj}}}(S))) \end{array}$$


where in the first two cases, the type of the opposite transition system is only constrained by the type of the relations R and S .

Instantiating both transition systems with \mathcal{T}_{Obj} , we now derive a relational logic with judgements $C \sim C' : R \Rightarrow S$, which we interpret (in contrast to Benton [9], but in accordance with Definition 1) in the termination-insensitive style. By virtue of the previous section, several formal interpretations of these judgements are compatible with this interpretation. The derivability from the unary program logic is most explicit if we define $C \sim C' : R \Rightarrow S$ to be a shorthand for

$$\exists \phi. \triangleright C : \mathit{Dec}_L R \phi \wedge \triangleright C' : \mathit{Dec}_R S \phi \quad (2)$$

and then establish the proof rules as derived lemmas. Figure 3 shows some proof rules for pairs of structurally identical phrases. These are similar to the rules given (for the heap-free fragment of the language) by Benton [9]. As in loc. cit., the loop rule is restricted to the case where both iterations proceed in lock-step. Our rule for conditionals (not shown) allows the executions to proceed along different control paths and consequently has hypotheses for all four possible combinations. The simplified rule that enforces identical evaluation of the branches arises as a special case.

The derivation of the rules exhibits witnesses ϕ as mandated by equation (2). By the results of the previous section, witnesses for the axioms may be chosen as $\phi_L^{\mathcal{T}_{\text{Obj}}} C \ R$ or $\phi_R^{\mathcal{T}'_{\text{Obj}}} C' \ S$, or any relation between the two. Witnesses for compound phrases are constructed from the constituents' witnesses using the same constructions (modulo the addition of the heap) as in [12]. For example, the witness for pairs of compositions arises from the witnesses of the constituent squares as $\phi = \{(\tau, \sigma'). \exists \rho. \rho \zeta \sigma' \wedge (\forall \rho'. \rho T \rho' \rightarrow \tau \pi \rho')\}$. The wit-



ness for the while rule, $\Phi_{(b', R, \phi)}^{\text{While}}$, is constructed as indicated in Figure 3, where ϕ is the witness of $C \sim C' : U \Rightarrow R$. All proofs for compound phrases respect decomposition in that no proof of a conjunct $\mathit{Dec}_L \dots$ uses a hypothesis for $\mathit{Dec}_R \dots$, and vice versa.

In addition to these rules, we have derived rules where the two phrases may be of different shape, including Benton's rules of falsity, consequence, common branch elimination, and dead code elimination – see Figure 4. Carrying a unary judgement in the hypothesis, the dead-code rule applies to arbitrary phrases C whereas Benton only considers the specializations for assignment and while. Conclusions of DEADL may be promoted to phrase compositions using COMPSKIP. We omit the similar rules for handling dead code and common branches in phrases to the right of \sim . Rule UNARY injects a pair of unary judgements into the relational world. This rule is unsound in the termination-insensitive setting of [9]. On the other hand, Benton's rule of transitivity

$$\frac{C \sim C' : R \Rightarrow S \quad C' \sim C'' : R \Rightarrow S}{C \sim C'' : R \Rightarrow S} \quad \mathit{PER}(R \Rightarrow S)$$

where $\mathit{PER}(R \Rightarrow S)$ indicates that the function space $R \Rightarrow S$ is a partial equivalence relation, is unsound in the termination-insensitive setting, as the hypotheses are vacuously satisfied if C' diverges but C and C'' converge. As decomposition witnesses orientate the simulation relation, Benton's rule of symmetry $\frac{C \sim C' : R \Rightarrow S \quad \mathit{PER}(R \Rightarrow S)}{C' \sim C : R^{-1} \Rightarrow S^{-1}}$ can only be derived if we exploit the semantic symmetry of the simulation relation and use the formal completeness of the program logic. An alternative is to modify the interpretation, by conjoining property (2) with

$$\exists \psi. \triangleright C' : \mathit{Dec}_L R^{-1} \psi \wedge \triangleright C : \mathit{Dec}_R S^{-1} \psi. \quad (3)$$

The resulting interpretation is immediately symmetric and also allows the derivation of the rules above, except for transitivity.

Completeness is also used when deriving rules that contain conclusions with phrases that are subphrases of phrases in hypotheses, thus reversing the standard subphrase orientation that is obeyed by our unary logic. An instance of this phenomenon are **Skip**-elimination rules like $\frac{\mathbf{Skip}; C \sim C' : R \Rightarrow S}{C \sim C' : R \Rightarrow S}$. An alternative to the use of the formal completeness result would be to work entirely at the level of semantic validity.

THEOREM 2. *The rules in Figures 3 and 4 are derivable as discussed and thus sound with respect to Definition 1.*

$$\begin{array}{c}
R = \{((s, h), (s', h')). \forall \ell \notin \text{locs}(s, h). \forall \ell' \notin \text{locs}(s', h'). (s[x \mapsto \ell], h[\ell \mapsto (c, \llbracket \ell \rrbracket_s)]) S(s'[x' \mapsto \ell'], h'[\ell' \mapsto (c', \llbracket \ell' \rrbracket_{s'})])\} \\
x := \mathbf{new} \ c \ \iota \sim x' := \mathbf{new} \ c' \ \iota' : R \Rightarrow S \\
R = \{((s, h), (s', h')). \forall \ell \ell' \ c \ c' \ F \ F' \ v \ v'. (s \ y = \ell \wedge s' \ y' = \ell' \wedge h \ \ell = (c, F) \wedge h' \ \ell' = (c', F') \wedge F \ f = v \wedge F' \ f' = v') \\
\longrightarrow (s[x \mapsto v], h) S(s'[x' \mapsto v'], h')\} \\
x := y.f \sim x' := y'.f' : R \Rightarrow S \\
\hline
C \sim C' : U \Rightarrow R \quad R = T \cap \{((s, h), (s', h')). \llbracket b \rrbracket_s = \llbracket b' \rrbracket_{s'}\} \quad U = R \cap \{((s, h), (s', h')). \llbracket b \rrbracket_s\} \quad S = R \cap \{((s, h), (s', h')). \neg \llbracket b \rrbracket_s\} \\
\mathbf{While} \ b \ \mathbf{do} \ C \sim \mathbf{While} \ b' \ \mathbf{do} \ C' : R \Rightarrow S
\end{array}$$

The operator $\Phi_{(b', R, \phi)}^{\mathbf{While}}$ is defined as the least fixed point of the (monotone) functional

$$\psi \mapsto \{(\tau, (t', k')). (\llbracket b' \rrbracket_{t'} \rightarrow (\exists \sigma. \sigma \phi(t', k') \wedge (\forall \sigma'. \sigma R \sigma' \rightarrow \tau \psi \sigma')) \wedge (\neg \llbracket b' \rrbracket_{t'} \rightarrow \tau R(t', k')) \}$$

Figure 3. RHL rules for identically shaped phrases (excerpt)

$$\begin{array}{c}
\text{COMBRL} \frac{C \sim C' : U \Rightarrow S \quad U = R \cap \{((s, h), (s', h')). \llbracket b \rrbracket_s\} \quad D \sim C' : T \Rightarrow S \quad T = R \cap \{((s, h), (s', h')). \neg \llbracket b \rrbracket_s\}}{\mathbf{If} \ b \ \mathbf{then} \ C \ \mathbf{else} \ D \sim C' : R \Rightarrow S} \\
\text{DEADL} \frac{\triangleright C : A \quad \forall \sigma \sigma' \tau. A \sigma \tau \rightarrow \sigma R \sigma' \rightarrow \tau S \sigma'}{C \sim \mathbf{Skip} : R \Rightarrow S} \\
\text{FALSE} \frac{}{C \sim C' : \emptyset \Rightarrow S} \quad \text{COMPSKIP} \frac{C \sim \mathbf{Skip} : R \Rightarrow T \quad D \sim D' : T \Rightarrow S}{C; D \sim D' : R \Rightarrow S} \quad \text{UNARY} \frac{\triangleright C : A \quad \triangleright C' : A' \quad R = \{(\sigma, \sigma'). \forall \tau \tau'. A \sigma \tau \rightarrow A' \sigma' \tau' \rightarrow \tau S \tau'\}}{C \sim C' : R \Rightarrow S} \\
\text{SUB} \frac{C \sim C' : R \Rightarrow S \quad R' \subseteq R \quad S \subseteq S'}{C \sim C' : R' \Rightarrow S'} \quad \text{SETOP} \frac{C \sim C' : R \Rightarrow S \quad C \sim C' : T \Rightarrow U \quad R' = R \odot T \quad S' = S \odot U \quad \odot \in \{\cup, \cap\}}{C \sim C' : R' \Rightarrow S'}
\end{array}$$

Figure 4. Nonsynchronous RHL rules (excerpt)

2.3 Discussion: auxiliary state and self-composition

Self-composition [8] solves the certification problem for non-interference of C by reducing it to a one-execution property of the program $C; C'$ where C' arises from C by replacing all program variables x in C by fresh copies x' . Written in pre-/postcondition style, the specification for $C; C'$ then is $\{\sim_L\} C; C' \{\sim_L\}$ where for some fixed (“low”) subset L of nonprimed variables the predicate \sim_L is defined as $\{s. \forall x \in L. s \ x = s \ x'\}$. The validity of this approach follows from the observation that the execution of the first half of the self-composed program, i.e. the execution of C , neither modifies nor reads primed variables, whereas the execution of C' neither reads nor modifies nonprimed variables. Comparing self-composition with decomposition, we see that a witness ϕ amounts to an assertion applicable at the point of self-composition, i.e. at the state where C has executed but C' has not started yet. The orientation of the witness corresponds to the (in principle arbitrary) order in which the two copies are self-composed. The relations R and S are \sim_L , but the entire situation is formulated without duplicating variables or code, purely semantically. This interpretation also explains the inclusion property of witnesses, as equation (1) precisely captures the possible applications of the consequence rule at the point of self-composition.

Terauchi and Aiken observed that the efficiency of verification based on self-composition is improved if the self-composition operation is applied only to small program fragments [32]. They also demonstrated that type systems for non-interference yield appropriate transformation rules. These transformations push self-composition towards the leaves of the syntax tree, in a manner that resembles our decomposition rules for compound phrases.

Decomposition triangles resemble the use of universally quantified state in Kleymann’s formal treatment of auxiliary variables [21], an effect that is particularly apparent in the dead-code rule. Formally, our universal quantification concerns states in the opposite execution and occurs inside (relational) assertions whereas in Kleymann’s case the quantification concerns a single

execution and occurs in the semantic interpretation of judgements. Nevertheless, we will make use of this observation in the following section, where we model auxiliary state (of arbitrary type) by decomposition with respect to a silent instruction.

3. Parameterized simulation and noninterference

Like unary properties, instances of relational simulations often involve pre- and postrelations that are parameterized by some notion of auxiliary “state”. This section reduces parametric simulations to the decomposed nonparameterized style and then instantiates the development to noninterference for the language of objects.

3.1 Constructions on transition systems, parameterization

Parameterized relational security concerns the following property.

DEFINITION 2. For transition systems \mathcal{T} and \mathcal{T}' as before, type \mathcal{Z} of auxiliary “states”, and parameterized relations $R : \mathcal{Z} \Rightarrow (\mathcal{S} \times \mathcal{S}')$, we write $\models_{\text{par}}^{(\mathcal{T}, \mathcal{T}')} P \sim P' : R \Rightarrow S$ if for all z, s, s', t , and t' with $(s, P, t) \in \mathcal{T}$ and $(s', P', t') \in \mathcal{T}'$, $sR_z s'$ implies $tS_z t'$, where R_z denotes the application of R to parameter z .


Exploiting the above connection between simulation diagrams and auxiliary state, we reduce parameterized simulation to the nonparameterized form, using two constructions on transition systems. The product transition system $\mathcal{T} \times \mathcal{T}'$ is given by $\{((s, s'), (P, P'), (t, t')). (s, P, t) \in \mathcal{T} \wedge (s', P', t') \in \mathcal{T}'\}$ and thus internalizes the two-execution nature of simulations. The identity transition system for parameters \mathcal{Z} is denoted by $\mathcal{I}_{\mathcal{Z}}$ and given by $\{(z, *, y). z = y\}$ where $*$ is the unique value of some singleton set of labels. These constructions yield further laws regarding the extremal witnesses – see Figure 5.

For parameterized relation $R : \mathcal{Z} \Rightarrow (\mathcal{S} \times \mathcal{S}')$ we write \vec{R} for the relation $\{(s, (s', z)). (s, s') \in Rz\}$. The following currying lemma justifies the above constructions on transition systems by

$$\begin{aligned}
\phi_L^{T \times T'}(P, P') R &= \{((t, t'), t''). ((t, t'), t'') \in \phi_L^T P \{(s, (r', r'')). (r', r'') \in \phi_L^{T'} P' \{(s', s''). ((s, s'), s'') \in R\}\}\} \\
\phi_R^{T \times T'}(P, P') S &= \{(s'', (s, s')). ((s'', s), s') \in \phi_R^{T'} P' \{((t'', t), t'). (t'', t) \in \phi_R^T P \{(r'', r), (r'', (r, t')) \in S\}\}\} \\
\phi_L^{T \times \mathcal{I}Z}(P, *) R &= \{((t, z), t'). (t, (z, t')) \in \phi_L^T P \{(s, (y, r')). ((s, y), r') \in R\}\} \\
\phi_R^{T' \times \mathcal{I}Z}(P', *) S &= \{(s, (s', z)). (s, s') \in \phi_R^{T'} P' \{(t, t'). (t, (t', z)) \in S\}\} \\
\models^T P : Dec_L R \phi &= \models^{T \times \mathcal{I}Z}(P, *) : \{((s, s'), (t, t')). s' = t' \wedge (sR s' \rightarrow t\phi t')\} \\
\models^{T \times \mathcal{I}Z}(P, *) : Dec_R R \phi &= \models^T P : Dec_R \{((s, z), s'). sR(s', z)\} \{((s, z), s'). s\phi(s', z)\}
\end{aligned}$$

Figure 5. Product decomposition laws

relating \mathcal{Z} -parameterized behaviour over $T \times T'$ to nonparameterized behaviour over $T \times (T' \times \mathcal{I}Z)$.



LEMMA 5. For $R, S : \mathcal{Z} \Rightarrow (S \times S')$ we have $\models_{\text{Par}}^{(T, T')} P \sim P' : R \Rightarrow S$ exactly iff $\models_{\text{Par}}^{(T' \times \mathcal{I}Z)} P \sim (P', *) : \vec{R} \Rightarrow \vec{S}$.

The product $(P', *)$ is again in the form of Definition 1, resulting in a form of “internal notion of transitivity”. For general triples of transitions, we remark that the weakest liberal prerelation decomposes in the two following ways

$$\begin{aligned}
WLR_{P, (P', P'')}^{T, (T' \times T'')}(S) &= \{(s, x). x \in WLR_{P', P''}^{T', T''}(\{y. s \in WLP_P^T(\{t. tSy\})\})\} \\
&= \{(s, x). s \in WLP_P^T(\{t. x \in WLR_{P', P''}^{T', T''}(\{y. tSy\})\})\}.
\end{aligned}$$

Comparing the latter equation with Lemma 4(3) highlights the correspondence between program pairs as considered in Definition 1 and formal products of transition systems.

3.2 Application: auxiliary state for language with objects

Specializing Lemma 5 to $T' = \mathcal{T}_{\text{Obj}}$ yields identities such as

$$\begin{aligned}
\phi_R^{\mathcal{T}_{\text{Obj}} \times \mathcal{I}Z}(x' := e', *) S &= \{(\tau, ((s', h'), z)). \tau S((s'[x' \mapsto \llbracket e' \rrbracket_{s'}], h'), z)\} \\
\phi_R^{\mathcal{T}_{\text{Obj}} \times \mathcal{I}Z}(C; D, *) S &= \phi_R^{\mathcal{T}_{\text{Obj}} \times \mathcal{I}Z}(C, *) (\phi_R^{\mathcal{T}_{\text{Obj}} \times \mathcal{I}Z}(D, *) S)
\end{aligned}$$

and similarly for the other instruction forms. Together with the (unmodified) laws for $\phi_L^{\mathcal{T}_{\text{Obj}}} C R$, these may be used to derive inclusions $\phi_L^{\mathcal{T}_{\text{Obj}}} C \vec{R} \subseteq \phi_R^{\mathcal{T}_{\text{Obj}} \times \mathcal{I}Z}(C', *) \vec{S}$, i.e. properties in the shape of Lemma 4(1). Proceeding in parallel with the previous section, we may also derive a logic $\vdash_{\text{Par}} C \sim C' : R \Rightarrow S$ using the interpretation

$$\exists \phi. \triangleright C : Dec_L \vec{R} \phi \wedge \triangleright C' : Dec_R (\vec{S})^\sharp \phi^\sharp$$

where the operation $\psi^\sharp \equiv \{(x, z), x'\}. (x, (x', z)) \in \psi\}$ shifts the auxiliary value z to the left component, so that it is not affected by the execution of C' .

By construction, we have that $\vdash_{\text{Par}} C \sim C' : R \Rightarrow S$ implies $\models_{\text{Par}}^{(\mathcal{T}_{\text{Obj}}, \mathcal{T}_{\text{Obj}})} C \sim C' : R \Rightarrow S$. Moreover, the interpretation admits the derivation of similar proof rules as in the previous section, using essentially the same constructions of witnesses. Two exemplary rules are shown in Figure 6, highlighting the formal similarity to the

earlier rules. In both cases, the given prerelations R indeed coincide with $WLR_{C, (C', *)}^{\mathcal{T}_{\text{Obj}}, (\mathcal{T}_{\text{Obj}} \times \mathcal{I}Z)}(\vec{S})$.

3.3 Application: noninterference for simple objects

Following Banerjee-Naumann [5], our model of noninterference employs indistinguishability relations between states that are parameterized by partial bijections over locations, i.e. sets $\beta \subseteq \mathcal{L} \times \mathcal{L}$ satisfying $(\ell = \ell_1) \Leftrightarrow (\ell' = \ell'_1)$ for any $(\ell, \ell') \in \beta$ and $(\ell_1, \ell'_1) \in \beta$. These bijections serve two purposes: first, they allow both executions to allocate or otherwise use private objects, i.e. objects that have no correspondence in the opposite execution. Second, for those objects that do have a correspondent in the opposite execution, they hide the precise locations holding either object. Thus, locations behave anonymously.

The formal setup is based on indistinguishability relations for values, stores, and heaps, given fixed associations Γ and Ξ that associate observation levels $l \in \{\text{high}, \text{low}\}$ to variables and field names, respectively. The two levels correspond to the observation capability of private and public agents, respectively. For values, the indistinguishability relation $v \sim_\beta v'$ is defined by

$$\begin{array}{l}
i = i' \\
i \sim_\beta i'
\end{array}
\quad
\frac{}{\text{Null} \sim_\beta \text{Null}}
\quad
\frac{(\ell, \ell') \in \beta}{\ell \sim_\beta \ell'}$$

For stores and objects, we define

$$\begin{aligned}
s \sim_\beta s' &\equiv \forall x. \Gamma(x) = \text{low} \rightarrow (sx) \sim_\beta (s'x) \\
(c, F) \sim_\beta (c', F') &\equiv c = c' \wedge LFlds_\Xi(F) = LFlds_\Xi(F') \\
&\quad \wedge \forall f. \Xi(f) = \text{low} \rightarrow (Ff) \sim_\beta (F'f)
\end{aligned}$$

where $LFlds_\Xi(F) = \{f. \exists v. Ff = v \wedge \Xi(f) = \text{low}\}$ denotes the public nonempty fields of an object. Finally, indistinguishability of heaps and states is defined as

$$\begin{aligned}
h \sim_\beta h' &\equiv \text{dom } \beta \subseteq \text{dom } h \wedge \text{rng } \beta \subseteq \text{dom } h' \\
&\quad \wedge \forall (\ell, \ell') \in \beta. h\ell \sim_\beta h'\ell' \\
(s, h) \sim_\beta (s', h') &\equiv s \sim_\beta s' \wedge h \sim_\beta h'
\end{aligned}$$

Noninterference then requires that β -indistinguishable initial states yield γ -indistinguishable final states, for some $\gamma \supseteq \beta$:

DEFINITION 3. C is noninterferent if for any $\sigma \sim_\beta \sigma'$ with $\sigma \xrightarrow{C} \tau$ and $\sigma' \xrightarrow{C} \tau'$, there exists some $\gamma \supseteq \beta$ with $\tau \sim_\gamma \tau'$.

Instead of understanding this condition as requiring the terminal states to be in some postrelation that is indexed by γ , we employ the reading that they are again in a β -indexed relation that internally quantifies over γ . This allows the partial bijection β to be treated as the auxiliary parameter: for

$$\begin{aligned}
R_{\text{NI}} &= \lambda \beta. \{(\sigma, \sigma'). \sigma \sim_\beta \sigma'\} \\
S_{\text{NI}} &= \lambda \beta. \{(\sigma, \sigma'). \exists \gamma \supseteq \beta. \sigma \sim_\gamma \sigma'\}
\end{aligned}$$

$$\begin{array}{c}
\frac{R = \lambda z. \{((s, h), (s', h')). (s[x \mapsto \llbracket e \rrbracket_s], h)(S_z)(s'[x' \mapsto \llbracket e' \rrbracket_{s'}], h')\}}{\vdash_{\text{Par}} x := e \sim x' := e' : R \Rightarrow S} \\
\\
\frac{R = \lambda z. \left\{ \begin{array}{l} ((s, h), (s', h')). \forall \ell c F v \ell' c' F' v'. s y = \ell \rightarrow h \ell = (c, F) \rightarrow F f = v \rightarrow \\ s' y' = \ell' \rightarrow h' \ell' = (c', F') \rightarrow F' f' = v' \rightarrow (s[x \mapsto v], h)(S_z)(s'[x' \mapsto v'], h') \end{array} \right\}}{\vdash_{\text{Par}} x := y.f \sim x' := y'.f' : R \Rightarrow S}
\end{array}$$

Figure 6. Example proof rules for parameterized relational logic $\vdash_{\text{Par}} C \sim C' : R \Rightarrow S$

noninterference coincides with $\models_{\text{Par}}^{(\mathcal{T}_{\text{Obj}}, \mathcal{T}_{\text{Obj}})} C \sim C : R_{\text{NI}} \Rightarrow S_{\text{NI}}$ and, in fact, also with $\models_{\text{Par}}^{(\mathcal{T}_{\text{Obj}}, \mathcal{T}_{\text{Obj}})} C \sim C : S_{\text{NI}} \Rightarrow S_{\text{NI}}$.

This motivates the definition of the judgement forms

$$\begin{aligned}
\text{LOW}(C) &\equiv \vdash_{\text{Par}} C \sim C : S_{\text{NI}} \Rightarrow S_{\text{NI}} \\
\text{HIGH}(C) &\equiv \vdash_{\text{Par}} C \sim \text{Skip} : R_{\text{NI}} \Rightarrow R_{\text{NI}}
\end{aligned}$$

which interpret the type judgements $[low] \vdash C$ and $[high] \vdash C$, respectively.

Figure 7 shows the (flow-insensitive) typing rules that have been derived in this fashion. In the rules for assignment and field modification, we write $\models e : low$ if $s \sim_{\beta} s'$ implies $\llbracket e \rrbracket_s \sim_{\beta} \llbracket e \rrbracket_{s'}$. The similar notation for boolean expressions used in the low rules for conditionals and loops additionally requires that the resulting values are identical. These semantic guarantees are implied by the auxiliary derivation systems for expressions that are typically part of type systems for noninterference (see [5, 29, 33]).

3.4 Discussion: ghost variables

JML and similar specification methodologies support ghost variables, i.e. variables that may be assigned language- or specification-level expressions but which influence neither the content of ordinary program variables nor the control flow. Ghost variables (and extensions of this concept such as ghost methods, fields, classes, ...) are used to simplify the verification process and also as an instrumentation technique for specifying extra-functional behaviour such as resource consumption. The use of ghost variables is a typical instance of simulation. It may be formulated either as the comparison of programs from different languages over different state spaces (instrumented language over extended states versus erased language), or as the comparison of two executions of the same annotated program, but for pre- and postrelations that ignore the content of ghost variables. The latter view suggests that ghost variables conceptually coincide with high-security variables - a view that is consistent with the admitted flow direction from ordinary into ghost variables and the independence of control flow from ghost variables. Under this view, the erasure of ghost variables finds its counterpart in the intimate relationship between noninterference and program slicing. The former view corresponds arguably more closely to Hofmann and Pavlova's technique for formally eliminating ghost variables, where derivations in a program logic for the ghost-annotated language are transformed into derivations for equivalent specifications in a logic for the erased language [17].

4. Error states and separation logic

Occasionally, verification formalisms are formulated for operational models that include explicit error states. An important aspect of these formalisms then is their capability to demonstrate that error states are in fact unreachable. Well-known instances of this approach are Reynolds' and Yang's presentations of unary [27] and relational [36] separation logics (SL). The operational semantics underlying these works contain explicit error states called **abort** and **fault**, respectively. These are entered whenever an attempt is made to dereference a dangling pointer, and their absence from executions of well-specified programs is a core aspect of SL.

In this section, we extend decomposition to models with error states, based on a unary logic that exposes error states in assertions. We then derive variants of unary and relational separation logics. The choice of operational model, and the use of separate pre- and postconditions in SL judgements, were guided by the desire to stay close to the original presentations.

4.1 Decomposition of language with explicit error states

We consider syntactically the same language as before, but employ an operational semantics with states of type $\Sigma \text{ option}$, ranged over by μ, ν, \dots . We write \perp for the error state and $\lfloor \sigma \rfloor$ for the injection of $\sigma \in \Sigma$. The operational semantics \mathcal{T}_{Err} is obtained by

- injecting the rules of \mathcal{T}_{Obj} , lifting initial states, and lifting final states as appropriate;
- adding behaviour to the object-manipulating instructions such that attempts to dereference a dangling pointer yield \perp , and likewise attempts to read from a field that is not in an (allocated) object's field map. We also add erroneous behaviour if the object variable contains an integer value or³ Null. Figure 8 shows the resulting rules for **getfield** and **putfield**.
- adding the rule $\frac{\perp \xrightarrow{C} \perp}{\perp \xrightarrow{C} \perp}$. Equivalently, one could propagate error states by appropriately modifying the rules for sequential composition and loops.

Traditionally, program logics for languages with error states build the avoidance of errors into the interpretation of judgements: the satisfaction of the precondition by an initial state guarantees that the execution will not end in \perp . In contrast, our unary logics exposes error states explicitly in the assertions, by employing assertions of type $\mathcal{A} \equiv \Sigma \text{ option} \Rightarrow \Sigma \text{ option} \Rightarrow \mathbb{T}$. Both the format of judgements $\vdash^{\text{Err}} C : A$, and their partial-correctness interpretation $\models^{\text{Err}} C : A \equiv \forall \mu \nu. (\mu, C, \nu) \in \mathcal{T}_{\text{Err}} \rightarrow A \mu \nu$ are in accordance with the generic setup of Section 1.1. Once the programmer has been exposed to the existence of error states by their occurrence in the operational model, this approach appears more flexible, as the programmer may explicitly reason about the presence or absence of error states. In the relational setting, coincident occurrence of error states can be expressed directly, as can asymmetric variations of this discipline. Preconditions remain redundant in a termination-insensitive setting, and their use in total-correctness logics remains orthogonal to the issue of error states.

The proof rules (see Figure 9) reflect the changes in the operational semantics but are otherwise similar to the earlier rules. They allow us to specify situations where, for example, the branches of a conditional differ in their error behaviour. Soundness and completeness, i.e. the equivalence between $\vdash^{\text{Err}} C : A$ and $\models^{\text{Err}} C : A$, are proven in the same way as Theorem 1.

Some decomposition laws for $\phi_{\perp}^{\text{Err}} C R$ are given in Figure 10. Subsequently, the proof rules for relational logics without (cf. Figures 3 and 4) and with (cf. Figure 6) auxiliary state can be derived

³The fact that type errors should arguably be kept separate from dangling-pointer errors is orthogonal to our argument.

$$\begin{array}{c}
\frac{}{[l] \vdash \text{Skip}} \quad \frac{\models e : \text{low}}{[low] \vdash x := e} \quad \frac{\Gamma(x) = \text{high}}{[high] \vdash x := e} \quad \frac{\Gamma(x) = \text{low} \quad \forall (f, e) \in \iota. \models e : \text{low}}{[low] \vdash x := \text{new } c \iota} \quad \frac{\Gamma(x) = \text{high}}{[high] \vdash x := \text{new } c \iota} \quad \frac{\Gamma(y) = \text{low} \quad \Xi(f) = \text{low}}{[low] \vdash x := y.f} \\
\frac{\Gamma(y) = \text{high}}{[high] \vdash x := y.f} \quad \frac{\Gamma(x) = \text{low} \quad \Xi(f) = \text{low} \quad \models e : \text{low}}{[low] \vdash x.f := e} \quad \frac{\Xi(f) = \text{high}}{[high] \vdash x.f := e} \quad \frac{[low] \vdash C \quad [low] \vdash D \quad \models b : \text{low}}{[low] \vdash \text{If } b \text{ then } C \text{ else } D} \\
\frac{[high] \vdash C \quad [high] \vdash D}{[high] \vdash \text{If } b \text{ then } C \text{ else } D} \quad \frac{[l] \vdash C \quad [l] \vdash D}{[l] \vdash C; D} \quad \frac{[low] \vdash C \quad \models b : \text{low}}{[low] \vdash \text{While } b \text{ do } C} \quad \frac{[high] \vdash C}{[high] \vdash \text{While } b \text{ do } C} \quad \frac{[high] \vdash C}{[low] \vdash C}
\end{array}$$

Figure 7. Derived typing rules for noninterference

$$\begin{array}{c}
\frac{s y = \ell \quad h \ell = (c, F) \quad F f = v \quad (\exists i. s y = i) \vee s y = \text{Null} \vee (\exists \ell. s y = \ell \wedge (\ell \notin \text{dom } h \vee (\exists c F. h \ell = (c, F) \wedge f \notin \text{dom } F)))}{[(s, h)] \xrightarrow{x:=y.f} [(s[x \mapsto v], h)]} \quad \frac{}{[(s, h)] \xrightarrow{x:=y.f} \perp} \\
\frac{s x = \ell \quad h \ell = (c, F)}{[(s, h)] \xrightarrow{x.f:=e} [(s, h[\ell \mapsto (c, F[f \mapsto [e]]_s])]} \quad \frac{(\exists i. s x = i) \vee s x = \text{Null} \vee (\exists \ell. s x = \ell \wedge \ell \notin \text{dom } h)}{[(s, h)] \xrightarrow{x.f:=e} \perp}
\end{array}$$

Figure 8. Operational semantics with error states (excerpt)

$$\begin{array}{c}
\frac{}{\vdash^{\text{Err}} x := e : \lambda \mu \nu. (\exists s h. \mu = [(s, h)] \wedge \nu = [(s[x \mapsto [e]]_s], h)) \vee (\mu = \nu = \perp)} \quad \frac{\vdash^{\text{Err}} C : A \quad \vdash^{\text{Err}} D : B}{\vdash^{\text{Err}} C; D : \lambda \mu \nu. \exists \rho. A \mu \rho \wedge B \rho \nu} \\
\frac{}{\vdash^{\text{Err}} x := \text{new } c \iota : \lambda \mu \nu. (\exists s h \ell. \mu = [(s, h)] \wedge \ell \notin \text{locs}(s, h) \wedge \nu = (s[x \mapsto \ell], h[\ell \mapsto (c, [l]_s)])) \vee (\mu = \nu = \perp)} \\
A = \lambda \mu \nu. \left(\begin{array}{c} \exists s h. \mu = [(s, h)] \wedge \left(\begin{array}{c} (\exists i. s y = i \wedge \nu = \perp) \vee (s y = \text{Null} \wedge \nu = \perp) \vee \\ (\ell \notin \text{dom } h \wedge \nu = \perp) \vee \\ (\exists \ell. s y = \ell \wedge \left(\begin{array}{c} \exists c F. h \ell = (c, F) \wedge \left(\begin{array}{c} \exists v. F f = v \wedge \nu = (s[x \mapsto v], h) \\ \vee (f \notin \text{dom } F \wedge \nu = \perp) \end{array} \right) \end{array} \right) \end{array} \right) \end{array} \right) \vee (\mu = \nu = \perp) \end{array} \right) \\
\frac{}{\vdash^{\text{Err}} x := y.f : A} \\
\frac{}{\vdash^{\text{Err}} C : A \quad \vdash^{\text{Err}} D : B} \\
\frac{}{\vdash^{\text{Err}} \text{If } b \text{ then } C \text{ else } D : \lambda \mu \nu. (\exists s h. \mu = [(s, h)] \wedge (([b]_s \rightarrow A \mu \nu) \wedge (\neg [b]_s \rightarrow B \mu \nu))) \vee (\mu = \nu = \perp)} \\
\frac{\vdash^{\text{Err}} C : B \quad \forall s h. \neg [b]_s \rightarrow A [(s, h)] [(s, h)] \quad \forall s h. [b]_s \rightarrow ((B [(s, h)] \perp) \rightarrow A [(s, h)] \perp) \wedge (\forall \sigma \nu. B [(s, h)] [\sigma] \rightarrow A [\sigma] \nu \rightarrow A [(s, h)] \nu)}{\vdash^{\text{Err}} \text{While } b \text{ do } C : \lambda \mu \nu. (\exists s h. \mu = [(s, h)] \wedge A \mu \nu \wedge (\forall t k. \nu = [(t, k)] \rightarrow \neg [b]_t)) \vee (\mu = \nu = \perp)}
\end{array}$$

Figure 9. Unary program logic with error states (excerpt).

$$\begin{array}{l}
\phi_{\perp}^{\text{Err}} x := e R = \{(\nu, \mu'). (\exists s h. \nu = [(s[x \mapsto [e]]_s], h) \wedge [(s, h)] R \mu' \vee (\nu = \perp \wedge \perp R \mu')\} \\
\phi_{\perp}^{\text{Err}} x := \text{new } c \iota R = \{(\nu, \mu'). (\exists s h \ell. \nu = [(s[x \mapsto \ell], h[\ell \mapsto (c, [l]_s)])] \wedge \ell \notin \text{locs}(s, h) \wedge [(s, h)] R \mu' \vee (\nu = \perp \wedge \perp R \mu')\} \\
\phi_{\perp}^{\text{Err}} \text{While } b \text{ do } C R = \{(\nu, \mu'). (b, C, R, \nu, \mu') \in \Phi_{\perp}^{\text{while}}\}
\end{array}$$

where $\Phi_{\perp}^{\text{while}}$ is inductively defined by (i.e. the smallest set closed under) the three rules

$$\begin{array}{c}
\frac{(\perp, \mu') \in \phi_{\perp}^{\text{Err}} C (R \cap \{(\mu, \mu'). \forall s h. \mu = [(s, h)] \rightarrow [b]_s\})}{(b, C, R, \perp, \mu') \in \Phi_{\perp}^{\text{while}}} \quad \frac{[(t, k)] R \mu' \quad \neg [b]_t}{(b, C, R, [(t, k)], \mu') \in \Phi_{\perp}^{\text{while}}} \\
\frac{(b, C, \phi_{\perp}^{\text{Err}} C (R \cap \{(\mu, \mu'). \forall s h. \mu = [(s, h)] \rightarrow [b]_s\}), \nu, \mu') \in \Phi_{\perp}^{\text{while}}}{(b, C, R, \nu, \mu') \in \Phi_{\perp}^{\text{while}}}
\end{array}$$

Figure 10. Laws for $\phi_{\perp}^{\text{Err}} C R$ (excerpt)

$$\frac{R = \lambda z. \{(\mu, \mu'). \forall \nu \nu'. \left(\begin{array}{c} (\forall s h. \mu = [(s, h)] \rightarrow \nu = [(s[x \mapsto [e]]_s], h) \wedge (\mu = \perp \rightarrow \nu = \perp) \wedge \\ (\forall s' h'. \mu' = [(s', h')] \rightarrow \nu' = [(s'[x' \mapsto [e']_{s'}], h')) \wedge (\mu' = \perp \rightarrow \nu' = \perp) \end{array} \right) \rightarrow \nu(S_z) \nu'\}}{\vdash_{\text{Par}}^{\text{Err}} x := e \sim x' := e' : R \Rightarrow S}$$

Figure 11. Error-aware parameterized relational logic (Assign-Assign rule)

in the same style as in the previous section. Figure 11 shows an exemplary rule from the latter. The executions may exhibit different behaviour with respect to the error-status of terminal states.

4.2 Application: derivation of unary separation logic

From the logic $\vdash_{\text{Par}}^{\text{Err}} C \sim C' : R \Rightarrow S$ we derive a unary separation logic. In order to stay close to Reynolds' original presentation [27], we employ judgements with separate pre- and postconditions and explicit auxiliary variables. However, locations remain anonymous, and we do not support deallocation.

In addition to serving as a test case for the decomposed relational logics, our encoding emphasizes the relational nature of (even unary) separation logic: similar to the embedding of low proof rules for noninterference, we interpret a separation logic judgement for C as a relational specification $\vdash_{\text{Par}}^{\text{Err}} C \sim C : R \Rightarrow S$ i.e. as an assertion about two executions of C . The relations R and S arise from the separation logic pre- and postconditions by incorporating the effect of future applications of the frame rule, in a style that is reminiscent of Birkedal et al.'s Kripke-resource extension [13] (see also [11],[4]). Thus, the encoding highlights that separation logic, like noninterference, concerns indistinguishability where states are partitioned into public components on which states are equal (the footprint) and nonpublic components (the frames)⁴.

We first outline a simplified encoding for the allocation-free language fragment, before integrating allocation in Section 4.2.2.

4.2.1 Unary separation-logic for allocation-free language

SL assertions concern lifted states that have been extended by an additional store holding the values of auxiliary variables (denoted as upper-case letters X, Y, \dots). In order to admit a formulation of separating conjunction at the granularity level of fields, we follow Parkinson [26] and employ *flat* heaps $\kappa \in (\mathcal{L} \times \mathcal{F}) \rightarrow (\mathcal{C} \times \mathcal{V})$. The flattening of heap h , denoted h^\flat , is given pointwise by

$$h^\flat(\ell, f) = (c, v) \iff \exists F. h \ell = (c, F) \wedge F f = v$$

and is injective up to the identification of unallocated objects with objects without fields. Assertions M, N, \dots are sets of triples (r, s, κ) where $r \in \mathcal{R}$ is the auxiliary store and s is the ordinary store. Thus, assertions in separation logic cannot talk about error states explicitly, in accordance with Reynolds and Parkinson.

The interpretation of a triple mediates between error-free states in assertions and error-aware states in the operational model, and also between flat and nonflat heaps. A separation logic specification is *valid*, notation $\models_{\text{SL}} \{M\} C \{N\}$, if whenever $(r, s, h^\flat) \in M$ and $\llbracket (s, h) \rrbracket \xrightarrow{C} \nu$ then there exist some t and k such that $\nu = \llbracket (t, k) \rrbracket$ and $(r, t, k^\flat) \in N$. Thus, any terminating execution starting in a state satisfying M yields a lifted state satisfying N .

Some simple operators of separation logic are given in Figure 12. Contrary to Parkinson we do not model Java-style garbage collection and consequently use classic rather than intuitionistic separation logic. The singleton mapsto assertions for ordinary and auxiliary variables are formulated on a per-field basis.

We say that assertion M is independent from the variables V , notation $\text{indep}_V M$, if for all $x \in V$ and all $v, (r, s, \kappa) \in M$ implies $(r, s[x \mapsto v], \kappa) \in M$.

The relational interpretation of assertion N with respect to variable set V yields a relation over error-aware states that is parameterized in the sense of the previous section by triples (M, r, M') . The assertions M and M' are separately conjoined with N in the

⁴The analogy is not perfect, however: in separation logic, the program operating on the footprint may neither influence nor be influenced by the private component of states, whereas noninterference allows public data to flow into private variables (but no flow in the opposite direction).

left and right execution, respectively, while r is an auxiliary store. The definition

$$R_V^N(M, r, M') \equiv \left\{ \begin{array}{l} (\llbracket (s, h) \rrbracket, \llbracket (s', h') \rrbracket). s = s' \wedge \\ \text{indep}_V M \wedge (r, s, h^\flat) \in N * M \wedge \\ \text{indep}_V M' \wedge (r, s, h^\flat) \in N * M' \end{array} \right\}$$

yields pairs of lifted states with identical stores but possibly different heaps, such that the appropriately r -extended states satisfy $N * M$ (left) and $N * M'$ (right), provided that the frames M and M' are independent from V . Frames that depend on V yield the empty relation, and no error state occurs in the relation.

With these definitions in place, we introduce the judgement form $\models_{\text{SL}} \{L\} C \{N\} [V]$ as an abbreviation of the self-simulation statement $\vdash_{\text{Par}}^{\text{Err}} C \sim C : R_V^L \Rightarrow R_V^N$.

We thus specify that the terminal states of runs starting in R_V^L -related states are related by R_V^N . In particular, both final states are lifted, and the parameterization mechanism propagates the frames from initial to final states.

The following result proves the faithfulness of our embedding with respect to the semantic guarantee and follows from the soundness of the logic $\vdash_{\text{Par}}^{\text{Err}} C \sim C : R \Rightarrow S$ by instantiating both frame parameters by emp.

THEOREM 3. *If $\models_{\text{SL}} \{L\} C \{N\} [V]$ then $\models_{\text{SL}} \{L\} C \{N\}$.*

Figure 13 collects the derived proof rules, where program variables possibly modified by C are collected in set V . Due to the Kripke-resource-style relational interpretation of judgements, the proof of the frame rule essentially amounts to transferring the frame L between the public part (the assertions in the concluding judgement) and the private part (the parameter of the hypothetical judgement). The legality of this operation follows from the associativity of $*$, the compatibility of $\text{indep}_V M$ with $*$, and the contravariance of $\text{indep}_V M$ with respect to V .

4.2.2 Allocation

The interpretation R_V^N does not admit the derivation of a proof rule for object allocation: the allocation nondeterminism results in the conjunct $s = s'$ of the invariant not being preserved. Also, the location chosen in one execution may be already in use in the opposite execution, either in the heap region specified by the footprint assertion or in the invisible frame assertion. Motivated by Section 3.3, we outline a remedy using partial bijections. The modified interpretation

$$S_V^N(M, r, r', M') \equiv \left\{ \begin{array}{l} (\llbracket (s, h) \rrbracket, \llbracket (s', h') \rrbracket). \exists \beta. \\ s \sim_\beta s' \wedge \text{locs } r \subseteq \text{dom } h \wedge \\ r \sim_\beta r' \wedge \text{locs } r' \subseteq \text{dom } h' \wedge \\ \text{indep}_V M \wedge (r, s, h^\flat) \in N * M \wedge \\ \text{indep}_V M' \wedge (r', s', h'^\flat) \in N * M' \end{array} \right\}$$

replaces $s = s'$ with $s \sim_\beta s'$ for some partial bijection β , and extends this condition to the (now) two versions of the auxiliary store. The duplication of the auxiliary store is necessitated by the introduction of $s \sim_\beta s'$, since otherwise assertions like $X = x$ which relate auxiliary to ordinary variables could never hold in related states. In addition, auxiliary variables should only mention allocated locations, for either execution. All indistinguishability relations are for the context $\Gamma = \lambda x. \text{low}$: we designate all (including auxiliary) variables to be relevant. Our interpretation does not require the heaps h and h' to be related by β .

$$\begin{aligned}
M * N &\equiv \{(r, s, \kappa). \exists d_1 d_2. d_1 \uplus d_2 = \text{dom } \kappa \wedge (r, s, \kappa|_{d_1}) \in M \wedge (r, s, \kappa|_{d_2}) \in N \} \\
\text{emp} &\equiv \{(r, s, \kappa). \text{dom } \kappa = \emptyset\} \\
X.f \mapsto (c, v) &\equiv \{(r, s, \kappa). \exists \ell. r X = \ell \wedge \text{dom } \kappa = \{(\ell, f)\} \wedge \kappa(\ell, f) = (c, v)\} \\
x.f \mapsto (c, v) &\equiv \{(r, s, \kappa). \exists \ell. s x = \ell \wedge \text{dom } \kappa = \{(\ell, f)\} \wedge \kappa(\ell, f) = (c, v)\}
\end{aligned}$$

Figure 12. Separation logic operators

$$\begin{array}{c}
\frac{x \in V}{\vdash_{\text{SL}} \{Y = \llbracket e \rrbracket_s\} x := e \{x = Y\} [V]} \quad \frac{y \in V}{\vdash_{\text{SL}} \{X = x \wedge X.f \mapsto (c, Y)\} y := x.f \{Y = y \wedge X.f \mapsto (c, Y)\} [V]} \\
\hline
\frac{}{\vdash_{\text{SL}} \{M\} \text{Skip} \{M\} [V]} \quad \frac{}{\vdash_{\text{SL}} \{\exists v. Y = \llbracket e \rrbracket_s \wedge x.f \mapsto (c, v)\} x.f := e \{x.f \mapsto (c, Y)\} [V]} \quad \frac{\vdash_{\text{SL}} \{M \wedge \llbracket b \rrbracket_s\} C \{M\} [V]}{\vdash_{\text{SL}} \{M\} \text{While } b \text{ do } C \{M \wedge \neg \llbracket b \rrbracket_s\} [V]} \\
\frac{\vdash_{\text{SL}} \{M\} C \{L\} [V] \quad \vdash_{\text{SL}} \{L\} D \{N\} [V]}{\vdash_{\text{SL}} \{M\} C; D \{N\} [V]} \quad \frac{\vdash_{\text{SL}} \{M \wedge \llbracket b \rrbracket_s\} C \{N\} [V] \quad \vdash_{\text{SL}} \{M \wedge \neg \llbracket b \rrbracket_s\} D \{N\} [V]}{\vdash_{\text{SL}} \{M\} \text{If } b \text{ then } C \text{ else } D \{N\} [V]} \quad \frac{M \subseteq M' \quad N' \subseteq N}{\vdash_{\text{SL}} \{M'\} C \{N'\} [V]} \quad \frac{V \subseteq W \quad \text{indep}_{W} L}{\vdash_{\text{SL}} \{M\} C \{N\} [V]} \\
\frac{}{\vdash_{\text{SL}} \{M\} C \{N\} [V]} \quad \frac{}{\vdash_{\text{SL}} \{M * L\} C \{N * L\} [W]}
\end{array}$$

Figure 13. Derived rules for unary separation logic. All assertions have formally the form $\{(r, s, h). \dots\}$. Upper-case letters denote auxiliary variables and are interpreted in r . Lower-case variables x, y, \dots are interpreted in s .

The judgement for unary SL with allocation is defined in a similar way as before, by introducing $\vdash_{\text{SLA}} \{L\} C \{N\} [V]$ as an abbreviation for $\vdash_{\text{Par}}^{\text{Err}} C \sim C : S_V^L \Rightarrow S_V^N$. In contrast to Section 3.3, the partial bijection in the postrelation is not required to be related to that in the prerelation, hence we do not propagate β via the parameter mechanism. This may be understood as expressing that our SL concerns in-space rather than in-place properties.

The soundness result of this logic guarantees the same property as the allocation-free logic, but requires that auxiliary variables in the precondition concern at most allocated locations.

THEOREM 4. *If $\vdash_{\text{SLA}} \{L\} C \{N\} [V]$ and all $(r, s, h^b) \in L$ satisfy $\text{locs } r \subseteq \text{dom } h$ then $\models_{\text{SL}} \{L\} C \{N\}$.*

The additional condition enforces the reasonable discipline that an auxiliary variable cannot specify objects outside of the footprint, in particular precluding the possibility that an auxiliary variable freezes the initial value of a location that is in fact allocated by C . The condition is, however, an artifact of the use of partial bijections, and will be avoided by a third interpretation below.

The proof rules from Figure 13 largely remain sound, but the rules for conditionals and loops require the condition $\models b : \text{low}$. By enforcing that these boolean expressions do not involve any constant locations, we ensure that the two executions follow the same control flow path. In addition, the allocation rule

$$\frac{x \in V \quad N = *_{\iota(f_i)=e_i} x.f_i \mapsto (c, Y_i)}{\vdash_{\text{SLA}} \{\text{emp} \wedge \bigwedge_{\iota(f_i)=e_i} Y_i = \llbracket e_i \rrbracket_s\} x := \text{new } c \iota \{N\} [V]}$$

can now be derived, where the Y_i are assumed to be distinct.

4.3 A relational separation logic

Turning to the relational case, we employ assertions P, Q, \dots that are relations over flat states. We say that a quadruple is valid, notation $\models_{\text{RSL}} \{P\} C \sim C' \{Q\}$, if whenever $((r, s, h^b), (r', s', h'^b)) \in P$ and $\llbracket (s, h) \rrbracket \xrightarrow{C} \nu$ and $\llbracket (s', h') \rrbracket \xrightarrow{C'} \nu'$, then $\nu = \llbracket (t, k) \rrbracket$ and $\nu' = \llbracket (t', k') \rrbracket$ for some t, k, t', k' with $((r, t, k^b), (r', t', k'^b)) \in Q$. We thus continue to work in the termination-insensitive setting but require that both executions are error-avoiding.

Figure 14 presents the adaptation of the basic SL primitives, the modified definition of relational independence, and an interpretation of assertions that again yields sets of pairs of non-error

states. A parameter contains two (unrelated) auxiliary stores and a relational frame. Our interpretation differs from Yang's in two respects. First, we incorporate (relational) Kripke-resource extension, thus preparing for the integration of procedures and admitting a similarly simple proof of the frame rule as in the unary case. In contrast, Yang's interpretation quantifies universally over pairs of heap extensions. Second, we leave the stores of related states unconstrained rather than interpreting both programs over a joint store (as Yang) or imposing indistinguishability. As the executions of C and C' are operationally unrelated, Yang's condition that requires related phrases to employ distinct variables becomes redundant.

We introduce the judgement $\vdash_{\text{RSL}} \{P\} C \sim C' \{Q\} (V, V')$ as a shorthand for $\vdash_{\text{Par}}^{\text{Err}} C \sim C' : \mathbb{T}_{(V, V')}^P \Rightarrow \mathbb{T}_{(V, V')}^Q$.

THEOREM 5. *If $\vdash_{\text{RSL}} \{P\} C \sim C' \{Q\} (V, V')$ then $\models_{\text{RSL}} \{P\} C \sim C' \{Q\}$.*

Figure 15 shows the resulting one-sided rules for object creation, field access and modification, and the frame rule. The expected rules for pairs of e.g. field updates may either be derived separately or may be obtained from the given rules, their symmetric counterparts for the phrases in the right phrase, and the **Skip**-elimination rules – see also [36]. The rules for composition, conditionals, loops, etc. are inherited from the host logic $\vdash_{\text{Par}}^{\text{Err}} C \sim C' : R \Rightarrow S$, modulo the addition of variable sets V and V' , and are hence omitted.

Regarding logical rules, Yang includes an embedding rule for pairs of (total-correctness) unary SL judgements. However, a proof attempt for the termination-insensitive variant

$$\frac{\vdash_{\text{SL}} \{M\} C \{N\} [V] \quad \vdash_{\text{SL}} \{M'\} C' \{N'\} [V']}{\vdash_{\text{RSL}} \{M \times M'\} C \sim C' \{N \times N'\} (V, V')}$$

fails, and likewise an attempt for the alternative unary judgement form from Section 4.2.2. Both unary interpretations use Kripke-extension over unary predicates, whereas the interpretation of the concluding judgement requires us to prove a property for *relational* Kripke-extensions. As not every relational frame splits into a pair of unary frames, the hypotheses do not apply.

We are not aware of any discussion of this issue in the literature, but offer the following solution. Motivated by the shape of the rules in Figure 15, we interpret unary SL judgements in terms of the RSL-judgement form and thus formally quantify over relational

$$\begin{aligned}
\text{emp} &\equiv \{(r, s, \kappa), (r', s', \kappa')\}. \text{dom } \kappa = \emptyset \wedge \text{dom } \kappa' = \emptyset\} \\
P * Q &\equiv \left\{ ((r, s, \kappa), (r', s', \kappa')). \exists d_1 d_2 d'_1 d'_2. d_1 \uplus d_2 = \text{dom } \kappa \wedge d'_1 \uplus d'_2 = \text{dom } \kappa' \wedge \right. \\
&\quad \left. ((r, s, \kappa \upharpoonright_{d_1}), (r', s', \kappa' \upharpoonright_{d'_1})) \in P \wedge ((r, s, \kappa \upharpoonright_{d_2}), (r', s', \kappa' \upharpoonright_{d'_2})) \in Q \right\} \\
\text{rindp}_{V'}^{V'} P &\equiv \forall r s \kappa r' s' \kappa'. ((r, s, \kappa), (r', s', \kappa')) \in P \longrightarrow \left(\begin{aligned} &(\forall x v. x \in V \longrightarrow ((r, s[x \mapsto v], \kappa), (r', s', \kappa')) \in P) \wedge \\ &(\forall x' v'. x' \in V' \longrightarrow ((r, s, \kappa), (r', s'[x' \mapsto v'], \kappa')) \in P) \end{aligned} \right) \\
\mathbb{T}_{(V, V')}^P(r, Q, r') &\equiv \{[(s, h)], [(s', h')]\}. \text{rindp}_{V'}^{V'} Q \wedge ((r, s, h^b), (r', s', h'^b)) \in P * Q\}
\end{aligned}$$

Figure 14. Relational separation logic operators and interpretation of relational assertions

$$\begin{array}{c}
\frac{M = (X = x \wedge X.f \mapsto (c, Y)) \quad y \in V \quad N = (Y = y \wedge X.f \mapsto (c, Y))}{\vdash_{\text{RSL}} \{M \times A'\} y := x.f \sim \mathbf{Skip} \{N \times A'\} (V, V')} \quad \frac{M = (\exists v. Y = \llbracket e \rrbracket_s \wedge x.f \mapsto (c, v))}{\vdash_{\text{RSL}} \{M \times A'\} x.f := e \sim \mathbf{Skip} \{x.f \mapsto (c, Y) \times A'\} (V, V')} \\
\frac{M = \text{emp} \wedge \bigwedge_{i(f_i)=e_i} Y_i = \llbracket e_i \rrbracket_s \quad N = *_{i(f_i)=e_i} x.f_i \mapsto (c, Y_i) \quad x \in V}{\vdash_{\text{RSL}} \{M \times A'\} x := \mathbf{new } c \iota \sim \mathbf{Skip} \{N \times A'\} (V, V')} \quad \frac{V \subseteq W \quad \text{rindp}_{W'}^{W'} R \quad V' \subseteq W'}{\vdash_{\text{RSL}} \{P\} C \sim C' \{Q\} (V, V')} \\
\vdash_{\text{RSL}} \{P * R\} C \sim C' \{Q * R\} (W, W')
\end{array}$$

Figure 15. Selected rules of relational separation logic: object-related instructions, FRAME

frames, but use relations that contain emp in the appropriate components of assertions and compare the subject phrase with \mathbf{Skip} rather than itself. With $\vdash_{\text{USL}} \{M\} C \{N\} [V]$ as a shorthand for⁵

$$\begin{aligned}
&\vdash_{\text{RSL}} \{M \times \text{emp}\} C \sim \mathbf{Skip} \{N \times \text{emp}\} (V, \emptyset) \\
&\wedge \vdash_{\text{RSL}} \{\text{emp} \times M\} \mathbf{Skip} \sim C \{\text{emp} \times N\} (\emptyset, V)
\end{aligned} \quad (4)$$

we obtain a proof system that satisfies the following properties. First, the above embedding rule becomes derivable - in particular, $\vdash_{\text{USL}} \{M\} C \{N\} [V]$ implies

$$\vdash_{\text{RSL}} \{M \times M\} C \sim C \{N \times N\} (V, V)$$

and thus indeed guarantees a self-simulation property for C . Second, soundness holds in the sense of Theorem 3. Third, all rules in Figure 13 are validated, and so is the allocation rule from Section 4.2.2. As the interpretation does not involve any partial bijections the proof rules do not need the additional side conditions from Section 4.2.2, and soundness formally holds without a restriction on the locations held in auxiliary variables.

5. Discussion

Relational decomposition is a technique for integrating relational logics into unary verification frameworks. We established soundness and completeness of decomposition for general simulations, introduced relational variants of predicate transformers, and studied their relationship to unary transformers. We applied our findings to derive a number of analyses, including separation logics. Our development is backed up by a formalization in Isabelle/HOL.

While all our instantiations concerned a simple imperative language of objects, we pointed out the conceptual applicability of decomposition across language boundaries. We are currently extending our formalization to a bytecode-level variant of the language, aiming to formally relate the two language levels in a subsequent step. Both language levels should also be extended with language features such as arrays, exceptions, and methods. Our previous treatment of decomposition for noninterference [12] supports a restricted notion of (possibly recursive) procedures, but the

transfer of this development to virtual methods and non-lock-step occurrences of method invocations is an open topic.

A long-term goal is the transfer of our techniques to other mainstream languages such as C, and their integration into appropriate verification infrastructures and certified compilation environments. While the integration into production-strength infrastructures for the entire language – as developed in the Concurrent Cminor and CompCert projects – appears a daunting task, language fragments such as Spark/Ada may represent a realistic testbed for relational concepts that is both industrially relevant and formally tractable.

Clearly, even the applications that we touched upon provide ample room for further development, for example concerning extensional notions of declassification [6] and conditional information flow [2], the integration of separation disciplines and noninterference along the lines of Amtoft et al.’s relational logic [1], and a principled treatment of ghost variables and program instrumentation. Magill et al.’s two-step isolation of numeric program abstractions for reasoning about data structures may provide orientation how the latter topics interact with separation aspects [24].

While some of these applications (including program instrumentation) cope well with the condition that loops proceed in lock-step, this aspect is probably the most severe limitation of our development, and also of the logics by Benton and Yang. The necessity to appeal to the rules for injecting pairs of unary judgements (rule **Unary**, or the embedding rule of relational SL) arguably weakens the appeal of relational logics, and hampers the systematic verification of loop-reordering transformations typically found in modern compilers. A topic for future work is thus to investigate how formal arguments that support such transformations, for example in the setting of translation validation [7], can be rephrased as suitable relational invariants, possibly based on data structures such as Tate et al.’s program equivalence graphs [31] or by embedding aspects of Rhodium’s transformation rules [22].

Saabas and Uustalu show how type derivations yield semantics-preserving proof transformations between pairs of judgements of unary Hoare logics [28].

Our interpretations of SL have been rather elementary, despite the integration of Kripke-extensions: concepts such as permissions, separation algebras, and local actions have not been considered [14]. Local actions have an immediate relational reading, which suggests that it might be possible to consider a more fine-

⁵By the relational frame rule, the definition coincides with

$$\begin{aligned}
&\forall A' V'. \vdash_{\text{RSL}} \{M \times A'\} C \sim \mathbf{Skip} \{N \times A'\} (V, V') \\
&\wedge \forall A V'. \vdash_{\text{RSL}} \{A \times M\} \mathbf{Skip} \sim C \{A \times N\} (V', V).
\end{aligned}$$

grained decomposition of local reasoning, for example by interpreting them as preservation conditions on relations [10].

Separation logic is also treated as a further instance of self-composition by Barthe et al. [8], but in a termination- and error-insensitive manner. Indeed, breaking down error-avoidance of self-composed programs into *equi*-error-avoidance of its constituents would require redoing the soundness analysis of the syntactic manipulations, including the analysis of the Terauchi-Aiken optimizations. On the other hand, the application of self-composition to concurrent systems is at present beyond the capability of our logic-based analysis, and we are unaware of previous attempts to develop relational logics for concurrent languages.

References

- [1] T. Amtoft, S. Bandhakavi, and A. Banerjee. A logic for information flow in object-oriented programs. In *Principles of Programming Languages*, pages 91–102. ACM Press, 2006.
- [2] T. Amtoft, J. Hatcliff, E. Rodriguez, Robby, J. Hoag, and D. Greve. Specification and checking of software contracts for conditional information flow. In *15th International Symposium on Formal Methods (FM'08)*, volume 5014 of *LNCS*, pages 229–245. Springer, May 2008.
- [3] A. W. Appel. Foundational proof-carrying code. In J. Halpern, editor, *Logic in Computer Science*, pages 247–258. IEEE Press, June 2001.
- [4] A. W. Appel and S. Blazy. Separation logic for small-step cminor. In K. Schneider and J. Brandt, editors, *Theorem Proving in Higher Order Logics, 20th International Conference (TPHOLs 2007), Proceedings*, volume 4732 of *LNCS*, pages 5–21. Springer, 2007.
- [5] A. Banerjee and D. Naumann. Stack-based access control for secure information flow. *Journal of Functional Programming*, 15:131–177, Mar. 2005. Special Issue on Language-Based Security.
- [6] A. Banerjee, D. A. Naumann, and S. Rosenberg. Towards a logical account of declassification. In M. W. Hicks, editor, *Proceedings of the 2007 Workshop on Programming Languages and Analysis for Security (PLAS 2007)*, pages 61–66. ACM Press, 2007.
- [7] C. W. Barrett, Y. Fang, B. Goldberg, Y. Hu, A. Pnueli, and L. D. Zuck. Tvoc: A translation validator for optimizing compilers. In K. Etesami and S. K. Rajamani, editors, *Computer Aided Verification, 17th International Conference, (CAV 2005), Proceedings*, volume 3576 of *LNCS*, pages 291–295. Springer, 2005.
- [8] G. Barthe, P. D’Argenio, and T. Rezk. Secure Information Flow by Self-Composition. In R. Foccardi, editor, *Computer Security Foundations Workshop*, pages 100–114. IEEE Press, 2004.
- [9] N. Benton. Simple relational correctness proofs for static analyses and program transformations. In N. D. Jones and X. Leroy, editors, *Proceedings of the 31st ACM Symposium on Principles of Programming Languages, POPL’04, Venice, Italy*, pages 14–25. ACM Press, 2004.
- [10] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations with dynamic allocation. In M. Leuschel and A. Podelski, editors, *Proceedings of the 9th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP 2007)*, pages 87–96. ACM Press, 2007.
- [11] L. Beringer. Relational bytecode correlations. *Journal of Logic and Algebraic Programming*. To appear.
- [12] L. Beringer and M. Hofmann. Secure information flow and program logics. In *IEEE Computer Security Foundations Symposium*, pages 233–248. IEEE Press, 2007.
- [13] L. Birkedal, N. Torp-Smith, and H. Yang. Semantics of separation-logic typing and higher-order frame rules. In *Proceedings of the 20th IEEE Symposium on Logic in Computer Science (LICS 2005)*, pages 260–269. IEEE Press, 2005.
- [14] C. Calcagno, P. W. O’Hearn, and H. Yang. Local action and abstract separation logic. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007), Proceedings*, pages 366–378. IEEE Press, 2007.
- [15] A. Darvas, R. Hähnle, and D. Sands. A theorem proving approach to analysis of secure information flow. In R. Gorrieri, editor, *Workshop on Issues in the Theory of Security*. IFIP WG 1.7, ACM SIGPLAN and GI FoMSESS, 2003.
- [16] M. V. Hermenegildo and J. Palsberg, editors. *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2010)*. ACM Press, 2010.
- [17] M. Hofmann and M. Pavlova. Elimination of ghost variables in program logics. In *Trustworthy Global Computing: Revised Selected Papers from the Third Symposium TGC 2007*, number 4912 in *LNCS*, pages 1–20. Springer, 2008.
- [18] S. Hunt and D. Sands. On flow-sensitive security types. In J. G. Morrisett and S. L. P. Jones, editors, *Principles of Programming Languages*, pages 79–90. ACM Press, 2006.
- [19] C. B. Jones. *Systematic Software Development Using VDM*. Prentice-Hall International, Englewood Cliffs, New Jersey, second edition, 1990.
- [20] R. Joshi and K. R. M. Leino. A semantic approach to secure information flow. *Science of Computer Programming*, 37(1-3):113–138, 2000.
- [21] T. Kleymann. *Hoare Logic and VDM: Machine-Checked Soundness and Completeness Proofs*. PhD thesis, LFCS, University of Edinburgh, 1998.
- [22] S. Lerner, T. D. Millstein, E. Rice, and C. Chambers. Automated soundness proofs for dataflow analyses and transformations via local rules. In J. Palsberg and M. Abadi, editors, *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2005)*, pages 364–377. ACM Press, 2005.
- [23] X. Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [24] S. Magill, M.-H. Tsai, P. Lee, and Y.-K. Tsay. Automatic numeric abstractions for heap-manipulating programs. In Hermenegildo and Palsberg [16], pages 211–222.
- [25] T. Nipkow. Hoare logics for recursive procedures and unbounded nondeterminism. In J. Bradfield, editor, *Computer Science Logic*, volume 2471 of *LNCS*, pages 103–119. Springer, 2002.
- [26] M. Parkinson. *Local Reasoning for Java*. PhD thesis, University of Cambridge, 2005.
- [27] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science*, Copenhagen, Denmark, July 2002. IEEE Press.
- [28] A. Saabas and T. Uustalu. Program and proof optimizations with type systems. *Journal of Logic and Algebraic Programming*, 77(1-2):131–154, 2008.
- [29] A. Sabelfeld and A. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communication*, 21:5–19, Jan. 2003.
- [30] D. A. Schmidt. Structure-preserving binary relations for program abstraction. In T. Æ. Mogensen, D. A. Schmidt, and I. H. Sudborough, editors, *The Essence of Computation*, volume 2566 of *LNCS*, pages 245–268. Springer, 2002.
- [31] R. Tate, M. Stepp, and S. Lerner. Generating compiler optimizations from proofs. In Hermenegildo and Palsberg [16], pages 389–402.
- [32] T. Terauchi and A. Aiken. Secure information flow as a safety problem. In C. Hankin and I. Siveroni, editors, *Static Analysis Symposium*, volume 3672 of *LNCS*, pages 352–367. Springer, 2005.
- [33] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):167–187, 1996.
- [34] G. Winskel. *The Formal Semantics of Programming Languages: An Introduction*. MIT Press, 1993.
- [35] www.cs.princeton.edu/~eberinge.
- [36] H. Yang. Relational separation logic. *Theoretical Computer Science*, 375(1-3):308–334, 2007.