

Thirty-seven years of relational Hoare logic: remarks on its principles and history (extended version)

David A. Naumann*

Stevens Institute of Technology

Abstract. Relational Hoare logics extend the applicability of modular, deductive verification to encompass important 2-run properties including dependency requirements such as confidentiality and program relations such as equivalence or similarity between program versions. A considerable number of recent works introduce different relational Hoare logics without yet converging on a core set of proof rules. This paper looks backwards to little known early work. This brings to light some principles that clarify and organize the rules as well as suggesting a new rule and a new notion of completeness.

A shorter version of this paper appears in *ISOLA*, the 9th International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (Springer LNCS 12477, pages 93–116). This version has additional material as endnotes, and minor revisions.

1 Introduction

Even in the archivally published part of the scientific literature, there are some gems known to few but deserving the attention of many. Such a gem is a paper by Nissim Francez published in 1983, around the time of Apt’s two-part paper “Ten Years of Hoare Logic” [2,3]. Relational Hoare Logic (RHL) formalizes reasoning about two programs. The term, and a version of the logic, are introduced in a well known gem by Nick Benton published in 2004 [22]. Relating two programs is far from new, and is important: it encompasses equivalence (as in compilation), refinement and conditional similarity (as in software development, evolution, and re-engineering), and properties of a single program (like determinacy of output) for which one must consider two executions. Reasoning about two executions inevitably leads to reasoning about two different programs—that is one of the principles already articulated in the paper by Francez titled “Product Properties and Their Direct Verification” [35] which introduces many rules of RHL.

The fundamental safety property is *partial correctness*: for each of the program’s runs, if the initial state satisfies the designated precondition, and the run terminates, then the final state satisfies the designated postcondition. The

* Partially supported by NSF CNS 1718713 and ONR N00014-17-1-2787.

fundamental liveness property is *termination*: for each of the program’s runs, if the initial state satisfies the precondition then the run is finite. Many interesting or desirable behavioral properties of a program are such *trace properties*, that is, a condition on runs is required to hold for all runs. Relations between programs involve two runs at a time, for example one notion of equivalence is that from the same initial state, runs of the two programs reach the same final state. One cannot expect this property if the programs are nondeterministic. What then is determinacy? It is the property of a program that from any initial state, its runs all diverge or all terminate in the same state. This can be defined more parsimoniously: from any initial state, any two runs either both diverge or both terminate in the same state. Behavioral program properties involving multiple runs have been dubbed *hyperproperties* and an important subclass are the *k-safety* properties which, for some fixed $k \geq 1$ can be defined by requiring all k -tuples of runs satisfy a condition on k -tuples [69,26]. Francez uses the term *power property* for what is now called k -safety, and *product property* for relations between programs.

The k -safety properties are an attractive object of study because they are amenable to reasoning techniques that are natural generalizations of those for safety properties. As is likely to occur to anyone familiar with programming or automata theory, to prove a 2-safety property one may construct a *product program* whose runs represent pairs of runs of the original. Francez points out that product programs can be expressed as ordinary programs, so that Hoare logic (HL) can be used to prove 2-safety: If C is a sequential program acting on variables, we can choose fresh names and obtain a renamed copy C' , and then $C;C'$ serves as a product. This particular product construction is often called *self-composition*, a term from Barthe et al [17,18] who rediscover the idea (and also consider other forms of product) for proving information flow security.

By now, scientific interest together with practical importance has led to exciting achievements. Related challenges are under active study by a number of research groups, often making use of some form of RHL. Quite a few papers have appeared with RHLs, some emphasizing frameworks meant to provide unifying principles [22,16,1,50], but there is considerable variety in the proof rules included or omitted. By contrast, the core rules of HL for imperative programs appear in many places with relatively little variation. There is a scientific explanation for this situation: the recipe for boiling a logic down to its essentials is to state and prove a completeness theorem that says true properties are provable. But, through product programs, relational reasoning is reduced to HL—so completeness in this sense is a trivial consequence of completeness for *unary* (i.e., 1-safety) properties, as Francez observes. His paper concludes with a problem that is still open: “It would be interesting to obtain a formal characterization of the situation in which the proposed method achieves actual proof simplification over the indirect proofs using Hoare’s (unary) logic.”

In this paper I describe, in as plain terms as I can, various reasoning principles and their embodiment in proof rules. One contribution is to systematize knowledge and in particular to emphasize the importance of program equiva-

lences to reduce the number of core rules. I also introduce a new rule to fill a gap that becomes evident when one considers one of the key motivations for relational reasoning. Finally, I introduce a new notion: alignment completeness. It appears naturally when one recapitulates, as Francez does and I do with slightly more generality, the development from Floyd to Hoare.

Scientists thrive on getting credit and good scientists take care in giving credit. But it is not always easy to determine the origin of ideas, in part because good ideas may be independently rediscovered several times before becoming firmly ensconced in the edifice of the known. My primary aim in this paper is to explain some principles as I have understood them, not to give a survey of prior work. I do point out ideas found in the paper by Francez, and cite some other work in passing. Other early works that may have been overlooked can be found in the first paragraph of Sec. 5.

Outline. Following background on the inductive assertion method and HL (Sec. 2), the method is extended to aligned pairs of runs in Sec. 3, as background for RHL which comprises Sec. 4. Sec. 5 discusses related work and Sec. 6 concludes. *In this long version of the paper, superscript numerals refer to end notes. These provide additional details including related work on completeness.*

2 Preliminaries

2.1 The inductive assertion method

We focus on the simple imperative or “while” language with assignments and possibly other primitive commands like nondeterministic assignment. The reader is expected to be familiar with transition semantics, in which the program acts on stores, where a store is a total mapping from variables to values.¹ The following abstraction of a program’s semantics is convenient.

An **automaton** is a tuple $(Ctrl, Sto, init, fin, \mapsto)$ where Sto is a set (the data stores), $Ctrl$ is a finite set that contains distinct elements $init$ and fin , and $\mapsto \subseteq (Ctrl \times Sto) \times (Ctrl \times Sto)$ is the transition relation. We require $(c, s) \mapsto (d, t)$ to imply $c \neq fin$ and $c \neq d$ and call these the **finality** and **non-stuttering** conditions respectively.² A pair (c, s) is called a **state**. Let β and γ range over states. A **trace** of an automaton is a non-empty sequence τ of states, consecutive under the transition relation, with $ctrl(\tau_0) = init$. It is **terminated** provided τ is finite and $ctrl(\tau_{-1}) = fin$, where τ_{-1} denotes the last state of τ .

In structural operational semantics, transitions act on configurations of the form (c, s) where c is a command, and **skip** by itself indicates termination. This fits our model: take $init$ to be the program of interest, $Ctrl$ to be all commands,³ and fin to be **skip**. Another instantiation treats $Ctrl$ as the points in the program’s control flow graph (CFG).

A partial correctness property is given by a pair of store predicates, P, Q , for which we write $P \rightsquigarrow Q$. In a formal logic, P and Q range over formulas in some assertion language, usually first order logic for a designated signature that includes types and operations used in the program. We write $s \models P$ to

say s satisfies predicate P , and define $(c, s) \models P$ iff $s \models P$. As a means to specify requirements, the notation is inadequate. The postcondition $y = x + 1$ can be achieved by changing x or by changing y . This problem is best solved by including a **frame condition** which for simple imperative programs is just a list \bar{x} of variables permitted to change, which we write as $P \rightsquigarrow Q[\bar{x}]$. Its meaning can be reduced to the simpler form provided we distinguish between program variables and spec-only variables not allowed to occur in programs.⁴ That being so, we focus on the form $P \rightsquigarrow Q$ and use for it the succinct term **spec**.

Let us spell out two semantics for specs in terms of an automaton A . The **basic semantics** is as follows. For a finite trace τ to satisfy $P \rightsquigarrow Q$ means that $\tau_0 \models P$ and $ctrl(\tau_{-1}) = fin$ imply $\tau_{-1} \models Q$, in which case we write $\tau \models P \rightsquigarrow Q$. Then A satisfies $P \rightsquigarrow Q$ just if all its finite traces do. The **non-stuck semantics** adds a second condition for τ to satisfy the spec: $ctrl(\tau_{-1}) \neq fin$ implies $\tau_{-1} \mapsto -$, where $\tau_{-1} \mapsto -$ means there is at least one successor state. Stuck states are often used to model runtime faults.

The inductive assertion method (IAM) of Floyd [34],⁵ is a way to establish that command C satisfies spec $P \rightsquigarrow Q$. The first idea is to generalize the problem: in addition to establishing that Q must hold in a final state, we establish additional conditions at intermediate steps, with the aim to reason by induction on steps of execution. The second idea is to designate which intermediate steps in terms of C 's CFG. An **assertion** is thus a formula R associated with a particular point in the CFG, designating the claim that in any run, R holds whenever control is at that point. This beautiful idea, called **annotation**, has a simple representation in syntax which has become commonplace: the assert statement. The third idea ensures that the claim is strong enough to be an induction hypothesis to prove the spec: (i) The entry point is annotated as P and the exit point is annotated as Q . (ii) Every cycle in the CFG is annotated with at least one assertion. The annotated points are said to form a **cutpoint set**. Such an annotation determines a finite set of acyclic paths through the CFG, each starting and ending with an annotation and having no intervening one—we call these **segments**.

Floyd shows, by induction on execution steps, that C satisfies $P \rightsquigarrow Q$ provided that the **verification conditions** (VCs) all hold [34]. Each segment determines the following VC: for any state that satisfies the initial assertion, and any execution along the segment, such that the branch conditions hold, the final assertion holds in the last state. In effect, the VCs are cases of the induction step in a proof that the assertions hold in any run from a state satisfying P .

Given a program and cutpoint set for its CFG, there is an automaton with $Ctrl$ the cutpoint set; the transitions $(c, s) \mapsto (d, t)$ are given by the semantics of a segment from c to d . An annotation assigns a store predicate $anno(c)$ to each cutpoint c . Define the state set $S \subseteq (Ctrl \times Sto)$ by $(c, s) \in S$ iff $s \models anno(c)$. Then the VCs amount to the condition that S is closed under \mapsto .

The IAM requires us to reason about the semantics of straight-line program fragments, which is amenable to automation in a number of different ways, in-

$$\begin{array}{c}
 x := e : P_e^x \rightsquigarrow P \\
 \\
 \frac{C : P \wedge e \rightsquigarrow Q \quad D : P \wedge \neg e \rightsquigarrow Q}{\text{if } e \text{ then } C \text{ else } D : P \rightsquigarrow Q} \quad \frac{C : P \rightsquigarrow R \quad D : R \rightsquigarrow Q}{C ; D : P \rightsquigarrow Q} \\
 \\
 \frac{C : P \rightsquigarrow Q \quad D : P \rightsquigarrow R}{C + D : P \rightsquigarrow Q \vee R} \\
 \\
 \frac{C : P \wedge b \rightsquigarrow P}{\text{while } b \text{ do } C \text{ od} : P \rightsquigarrow P \wedge \neg b}
 \end{array}$$

Fig. 1. Syntax-directed rules of HL for simple imperative programs.

$$\begin{array}{c}
 \frac{P \Rightarrow R \quad C : R \rightsquigarrow S \quad S \Rightarrow Q}{C : P \rightsquigarrow Q} \text{ (CONSEQ)} \\
 \\
 \frac{C : P \rightsquigarrow Q \quad C : P \rightsquigarrow R}{C : P \rightsquigarrow Q \wedge R} \text{ (CONJ)} \quad \frac{C : P \rightsquigarrow R \quad C : Q \rightsquigarrow R}{C : P \vee Q \rightsquigarrow R} \text{ (DISJ)} \\
 \\
 \frac{C : P \rightsquigarrow Q \quad FV(R) \cap Vars(C) = \emptyset}{C : P \wedge R \rightsquigarrow Q \wedge R} \text{ (FRAME)}
 \end{array}$$

Fig. 2. Rules to manipulate specs in HL.

cluding the direct use of operational semantics [51]. What makes program verification difficult is finding inductive intermediate assertions.

2.2 Hoare logic

Hoare showed that the IAM can be presented as a deductive system, in which inference rules capture the semantics of the constructs of the programming language and verification conditions are, to some degree, compositional in terms of program syntax. The system derives what are known variously as *partial correctness assertions*, Hoare triples, etc., and which ascribe a spec to a program. Hoare wrote $P\{C\}Q$ but it has become standard to write $\{P\}C\{Q\}$. We write $C : P \rightsquigarrow Q$ and call it a **correctness judgment**.

There are many Hoare logics, because a deductive system is defined for a particular language, i.e., set of program constructs. The textbook by Apt et al [5] has logics encompassing procedures, concurrency, etc. In this paper we focus on sequential programs but the principles apply broadly.

Rules for some program constructs can be found in Fig. 1. The axiom for assignment involves capture-avoiding substitution of an expression for a variable, written P_e^x , wherein we see that the system treats program expressions and variables as mathematical ones, a slightly delicate topic that does not obtrude in the sequel. These rules transparently embody the reasoning that underlies the IAM.

The sequence rule adds an intermediate assertion, changing one verification condition into two (typically simpler) ones. The rule for conditional alternatives has two premises, corresponding to the two paths through the CFG. Nondeterministic choice (notation $+$) again gives rise to two paths. The rules in Fig. 1 provide for deductive proof following the program structure but are incomplete; e.g., they give no way to prove the judgment $x := x + 1 : (x = y) \rightsquigarrow (x > y)$. Such gaps are bridged by rules like those in Fig. 2. For HL to be a self-contained deductive system it needs to include means to infer valid formulas, such as the first and third premises of rule CONSEQ.

For while programs, the syntax-directed rules together with CONSEQ are complete in the sense that any true judgment can be proved.⁶ The other rules embody useful reasoning principles. DISJ provides proof by cases, which is useful when a program’s behavior has two quite different cases. In terms of IAM, one might apply the method twice, to prove $C : P \rightsquigarrow R$ and $C : Q \rightsquigarrow R$ with entirely different annotations, and then conclude $C : P \vee Q \rightsquigarrow R$ by some argument about the meaning of specs. This principle is expressed directly in Hoare logic, as is the oft-used principle of establishing conjuncts of a postcondition separately.

Modular reasoning. HL easily admits procedure-modular reasoning, sometimes formalized by judgments of the form $H \vdash C : P \rightsquigarrow Q$ where hypothesis H comprises procedure signatures and their specs [61,10]. With the addition of procedures, CONSEQ is not sufficient for completeness. Other rules are needed to manipulate specs, such as substitution rules to adapt procedure specs to their calling contexts [40,62,5]. We use the name FRAME for a rule Hoare called Invariance [40], with a nod towards a similar rule in separation logic [60] where disjointness of heap locations is expressed by using the separating conjunction in place of \wedge . With explicit frame conditions the rule can be phrased like this: From $C : P \rightsquigarrow Q[\bar{x}]$ and $FV(R) \cap \bar{x} = \emptyset$ infer $C : P \wedge R \rightsquigarrow Q \wedge R[\bar{x}]$. The principle here is to reason “locally” with assertions P, Q pertinent to the effect of C , and then infer a spec $P \wedge R \rightsquigarrow Q \wedge R$ needed to reason about a larger program of which C is a part. Locality is important for effective reasoning about programs involving the heap. (Explicit frame conditions for the heap can be found, for example in the Dafny language [49] and in the variation of HL dubbed region logic [10].) The notion of *adaptation completeness* characterizes the extent to which a HL has sufficient rules for reasoning about specs [62,43,58,6].⁷

Refinement. Validity of assertions is a separate concern from program correctness and CONSEQ brings the two together in a simple way—but it has nothing to do with a specific command. It connects two specs, in a way that can be made precise by defining the *intrinsic refinement order* (\sqsubseteq , “refined by”) on specs. Fixing a class of programs, we define $P \rightsquigarrow Q \sqsubseteq R \rightsquigarrow S$ iff $C : R \rightsquigarrow S$ implies $C : P \rightsquigarrow Q$ for all C . This relation can itself be given a deductive system, with rules including that $P \rightsquigarrow Q \sqsubseteq R \rightsquigarrow S$ can be inferred from $P \Rightarrow R$ and $S \Rightarrow Q$. The program correctness rule infers $C : spec_1$ from $C : spec_0$ and $spec_1 \sqsubseteq spec_0$. Using frame conditions, the FRAME rule can also be phrased as a spec refinement: $P \wedge R \rightsquigarrow Q \wedge R[\bar{x}] \sqsubseteq P \rightsquigarrow Q[\bar{x}]$ provided $FV(R) \cap \bar{x} = \emptyset$.

Disentangling spec reasoning from reasoning about the correctness judgment helps clarify that that adaptation completeness is about spec refinement [58]. But it does come at a cost: To account for the `CONJ` and `DISJ` rules one needs not only the relation \sqsubseteq on specs but also meet/join operators.⁸ Explicit formalization of spec refinement could be useful for relational specs, owing to the additional manipulations that exist owing to the additional dimension. But I do not develop the topic further in this paper.

Program transformation. Verification tools employ semantics-preserving transformations as part of the process of generating VCs.⁹ Less commonly, transformations are an ingredient in a Hoare logic. An instance of this is the logic for distributed programs of Apt et al [4,5], where the main rule for a distributed program has as premise the correctness of a derived sequential program. Another instance is rules to justify the use of auxiliary or *ghost* state in reasoning [63]. One such rule uses variable blocks `var \bar{x} in C` . The variables \bar{x} are called *auxiliary in C* provided their only occurrences are in assignments to variables in \bar{x} . Writing $C \setminus \bar{x}$ for the command obtained by replacing all such assignments with `skip`, the rule is

$$\frac{\bar{x} \notin FV(P, Q) \quad \bar{x} \text{ auxiliary in } C \quad \text{var } \bar{x} \text{ in } C : P \rightsquigarrow Q}{C \setminus \bar{x} : P \rightsquigarrow Q} \text{ (AUXVAR)}$$

It is sound because the auxiliary variables cannot influence values or branch conditions, and thus have no effect on the variables in P or Q , nor on termination.¹⁰ As relational correctness judgments can express both dependency and program equivalences, we should be able to bring both the condition “auxiliary in” and the transformation $C \setminus \bar{x}$ into the logic, making the above rule admissible.

3 Relational properties, alignment, and program products

Here are some example relational properties of a single program.

- (determinacy)** For all terminated traces τ, v from the same initial state, the final states are the same: $\tau_0 = v_0$ implies $\tau_{-1} = v_{-1}$.
- (monotonicity)** For all terminated traces τ, v , if $\tau_0(x) \leq v_0(x)$ then $\tau_{-1}(z) \leq v_{-1}(z)$. Here x, z are integer variables.
- (dependence, non-interference)** (“ z depends on nothing except possibly x ”) For all terminated traces τ, v , if $\tau_0(x) = v_0(x)$ then $\tau_{-1}(z) = v_{-1}(z)$.

Here are some example relations between programs C and D .

- (equivalence)** For all terminated traces τ of C and v of D , if $\tau_0 = v_0$ then $\tau_{-1} = v_{-1}$. Determinacy is self-equivalence in this sense.
- (majorization)** For all terminated traces τ of C and v of D , if $\tau_0(x) = v_0(x)$ then $\tau_{-1}(z) > v_{-1}(z)$.
- (refinement)** For all terminated traces τ of C , there is a terminated trace v of D with $\tau_0 = v_0$ and $\tau_{-1} = v_{-1}$.

(relative termination) (For a given relation \mathcal{R} .) For all initial states β, γ that satisfy \mathcal{R} , if C has a terminated trace from β then D has a terminated trace from γ [39].

(mutual termination) For all initial β, γ that satisfy \mathcal{R} , C can diverge from β iff D can diverge from γ [37].

Refinement and relative termination involve existential quantification over traces, as do generalizations of refinement such as simulation and also dependence for nondeterministic programs (if τ is terminated and $\tau_0(x) = \gamma(x)$ then there is terminated v with $v_0 = \gamma$ and $v_{-1}(z) = \tau_{-1}(z)$). We refer to these as $\forall\exists$ properties, by contrast with the preceding items which universally quantify traces (denoted $\forall\forall$). The $\forall\forall$ properties above are also *termination-insensitive* in the sense that they only constrain terminating traces. In this paper we focus on termination-insensitive $\forall\forall$ properties while discussing some $\forall\exists$ properties (which are hyperliveness [26], not 2-safety) in passing. Mutual termination also involves existentials, unless programs are deterministic as they are in Benton [22] where mutual termination is used.

Let $A' = (Ctrl', Sto', init', fin', \mapsto')$ be an automaton. A relational spec $\mathcal{R} \approx \mathcal{S}$ is comprised of relations \mathcal{R} and \mathcal{S} from Sto to Sto' . We write $(c, s), (c', s') \models \mathcal{R}$ to mean $s, s' \models \mathcal{R}$. Finite traces τ of A and τ' of A' satisfy $\mathcal{R} \approx \mathcal{S}$, written $\tau, \tau' \models \mathcal{R} \approx \mathcal{S}$, just if $\tau_0, \tau'_0 \models \mathcal{R}$, $ctrl(\tau_{-1}) = fin$, and $ctrl(\tau'_{-1}) = fin$ imply $\tau_{-1}, \tau'_{-1} \models \mathcal{S}$. The non-stuck semantics of relational specs requires, in addition, that $ctrl(\tau_{-1}) \neq fin$ implies $\tau_{-1} \mapsto -$ and $ctrl(\tau'_{-1}) \neq fin'$ implies $\tau'_{-1} \mapsto -$. Finally, the pair A, A' satisfies $\mathcal{R} \approx \mathcal{S}$ just if all pairs of finite traces do, and we write $A|A' : \mathcal{R} \approx \mathcal{S}$ for satisfaction. (Where I write $A|A'$, as in [55,11], Francez writes $A \times A'$, and Benton's $A \sim A'$ is popular.)

A key idea (in [35] and elsewhere) is to form a single automaton, runs of which encode pairs of runs of the considered programs, and to which IAM can be applied. For a single program there is not much flexibility in how it is represented as an automaton or CFG but there are many product automata for a given pair of programs—these represent different ways of aligning the steps of the two programs.¹¹ This flexibility is crucial for the effectiveness of the IAM, specifically on the simplicity of annotations and thus the ease of finding them and proving the VCs. To discuss this we consider the four examples in Fig. 3.

Consider proving monotonicity of $P0$. To express relations we use dashed ($'$) identifiers for the second run, so the spec can be written $x \leq x' \approx z \leq z'$. One can prove the functional property that $P0$ computes factorial ($x!$) and then prove monotonicity for the recursive definition of $!$. But, as pointed out in [35], one can also consider two runs from initial values x, x' with $x \leq x'$, aligning their iterations in lockstep with invariant $y \leq y' \wedge z \leq z'$ and no use of $!$.

Consider proving that $P2$ is equivalent to $P0$, which we again specify just using the relevant variables: $P0|P2 : x = x' \approx z = z'$. Lockstep alignment of their iterations is not helpful; we would like to align each iteration of $P0$ with two iterations of $P2$ in order to use simple annotations like $y = y' \wedge z = z'$.


```

P0: (* z := x! *) y := x; z := 1; while y ≠ 0 do z := z*y; y := y-1 od

P1: (* z := 2x *) y := x; z := 1; while y ≠ 0 do z := z*2; y := y-1 od

P2: (* z := x!, half as fast *)
y := x; z := 1; w := 0;
while y ≠ 0 do if w mod 2 = 0 then z := z*y; y := y-1 fi; w := w+1 od

P3: (* z := 2x, a third as fast *)
y := x; z := 1; w := 0;
while y ≠ 0 do if w mod 3 = 0 then z := z*2; y := y-1 fi; w := w+1 od
    
```

Fig. 3. Example programs. $P0$ and $P1$ are from [35].

3.1 Product automata represent alignments

Let \otimes denote the cartesian product of relations,¹² so $\mapsto \otimes \mapsto'$ is a relation on $(Ctrl \times Sto) \times (Ctrl' \times Sto')$, i.e., on state pairs. Let id_A be the identity relation on states of A . A **pre-product** of A and A' is an automaton $P_{A,A'}$ of the form $((Ctrl \times Ctrl'), (Sto \times Sto'), (init, init'), (fin, fin'), \Rightarrow)$ such that we have $\Rightarrow \subseteq (\mapsto \otimes \mapsto') \cup (\mapsto \otimes id_{A'}) \cup (id_A \otimes \mapsto')$. The union is disjoint, owing to non-stuttering of A and A' . Each transition of $P_{A,A'}$ corresponds to one of both A and A' , or else one of A or A' leaving the other side unchanged. Such \Rightarrow satisfies the requirements of finality and non-stuttering.

Let T be a trace of a pre-product of A, A' . Mapping the first projection (fst) over T does not necessarily yield a trace of A , as it may include stuttering steps (related by id_A). So we define $left(T)$ to be $destutter(map(fst, T))$ where $destutter$ removes stuttering transitions.¹³ Observe that $left(T)$ is a trace of A , and we obtain *mutatis mutandis* a trace, $right(T)$, of A' . A **pre-product** is **adequate** if it covers all finite traces: For all finite traces τ of A and τ' of A' there is a trace T of $P_{A,A'}$ with $\tau \preceq left(T)$ and $\tau' \preceq right(T)$, where \preceq means prefix. It is **weakly adequate** if it covers all finite prefixes τ, τ' of terminated traces. (To see that equality $\tau = left(T)$ and $\tau' = right(T)$ would be too restrictive, consider lockstep alignment with τ strictly shorter or longer than τ' .)

Owing to the definition of states of a pre-product, a relational spec $\mathcal{R} \approx \mathcal{S}$ for A, A' can be seen as a unary spec $\mathcal{R} \rightsquigarrow \mathcal{S}$ for $P_{A,A'}$. For a trace T of $P_{A,A'}$ we have $T \models \mathcal{R} \rightsquigarrow \mathcal{S}$ iff $left(T), right(T) \models \mathcal{R} \approx \mathcal{S}$ by definitions. We obtain the following by definitions.

Theorem 1. *For the basic semantics of specs, if $P_{A,A'}$ is a weakly adequate pre-product of A, A' then $P_{A,A'}$ satisfies $\mathcal{R} \rightsquigarrow \mathcal{S}$ iff the pair A, A' satisfies $\mathcal{R} \approx \mathcal{S}$.*

This confirms that a relational spec can be proved using the IAM and a pre-product. The challenge is to construct one that admits a simple annotation and is at least weakly adequate. Adequacy is a more robust condition that holds for several forms of product. The number of cutpoints needed for a product may be

on the order of the product of the number for the underlying automata, but a good alignment makes many unreachable; those can be annotated as false so the corresponding VCs are vacuous.

Apropos stuck states, an adequate pre-product may have stuck states that do not correspond to stuck states of A or A' ; this is one way for a pre-product to be helpful, rendering unreachable states such as those where the guards of a conditional are not in agreement. However, if $P_{A,A'}$ is an adequate pre-product and satisfies $\mathcal{R} \rightsquigarrow \mathcal{S}$ in non-stuck semantics, it does not follow that A, A' satisfies $\mathcal{R} \approx \mathcal{S}$ in non-stuck semantics —contrary to a misstatement in a previous version of this paper. The problem is one-sided divergence. For example, suppose trace τ of A is not terminated but τ_{-1} is stuck. If A' can diverge then it is possible for a $P_{A,A'}$ to be adequate yet have an infinite sequence of traces T_i such that $i < j \Rightarrow T_i < T_j$ (so that it is not stuck), yet $\text{fst}(T_i) = \tau$ for all i . That problem can be solved by requiring A' to be terminating (and A as well, to prevent divergence on the left), an approach taken in [12]. Alternatively, products can be required not to have one-sided divergence, an approach taken in [8].

Here are some pre-products defined for arbitrary A, A' . For brevity, we express product states as pairs of A - and A' -states, as if the product's state had type $(Ctrl \times Sto) \times (Ctrl' \times Sto')$.

only-lockstep. $(\gamma, \gamma') \Rightarrow_{olck} (\beta, \beta')$ iff $\gamma \mapsto \beta$ and $\gamma' \mapsto' \beta'$.

eager-lockstep. $(\gamma, \gamma') \Rightarrow_{elck} (\beta, \beta')$ iff $(\gamma, \gamma') \Rightarrow_{olck} (\beta, \beta')$, or $\text{ctrl}(\gamma) = \text{fin}$ and $\gamma' \mapsto' \beta'$ and $\gamma = \beta$, or $\text{ctrl}(\gamma') = \text{fin}'$ and $\gamma \mapsto \beta$ and $\gamma' = \beta'$.

interleaved. $(\gamma, \gamma') \Rightarrow_{int} (\beta, \beta')$ iff $\gamma \mapsto \beta$ and $\gamma' = \beta'$ or $\gamma' \mapsto' \beta'$ and $\gamma = \beta$.

maximal. The union $\Rightarrow_{olck} \cup \Rightarrow_{int}$.

sequenced. $(\gamma, \gamma') \Rightarrow_{seq} (\beta, \beta')$ iff $\gamma \mapsto \beta$ and $\text{ctrl}(\gamma') = \text{init}'$ and $\gamma' = \beta'$ or $\text{ctrl}(\gamma) = \text{fin}$ and $\gamma = \beta$ and $\gamma' \mapsto' \beta'$.

simple-condition. Given “alignment condition” $ac \subseteq (Ctrl \times Sto) \times (Ctrl' \times Sto')$, define \Rightarrow_{scnd} by $(\gamma, \gamma') \Rightarrow_{scnd} (\beta, \beta')$ iff either $(\gamma, \gamma') \in ac$ and $(\gamma, \gamma') \Rightarrow_{olck} (\beta, \beta')$ or $(\gamma, \gamma') \notin ac$ and $(\gamma, \gamma') \Rightarrow_{int} (\beta, \beta')$.

As Francez observes, interleaved has a relatively large reachable state space, making it more difficult to find inductive invariants.

The only-lockstep form is not adequate, in general, because a terminated state or stuck state can be reached on one side before it is on the other. The eager-lockstep, interleaved, and maximal pre-products are all adequate. The sequenced form is not adequate in general: a stuck or divergent state on the left prevents coverage on the right. Sequenced is weakly adequate if A, A' have no stuck states.

The simple-condition product can also fail to be adequate: if ac holds, both sides are required to take a step, which may be impossible if one side is stuck or terminated. It is also insufficiently general: as we show later, it may be most convenient to designate that steps should be taken on one side or the other. This suggests the following, which subsumes the preceding constructions.

3-condition. Given state conditions l, r, b , define $(\gamma, \gamma') \Rightarrow_{3cnd} (\beta, \beta')$ iff either $(\gamma, \gamma') \in l$ and $\gamma \mapsto \beta$ and $\gamma' = \beta'$, or $(\gamma, \gamma') \in b$ and $(\gamma, \gamma') \Rightarrow_{olck} (\beta, \beta')$, or $(\gamma, \gamma') \in r$ and $\gamma' \mapsto' \beta'$ and $\gamma = \beta$.

3.2 Examples

Consider proving that $P0$ majorizes $P1$, for inputs $x > 3$, that is, $P0|P1 : x = x' \wedge x > 3 \approx z > z'$.¹⁴ Francez observes that using sequenced product would require reasoning about $z = x!$ and $z' = 2^{x'}$, and suggests aligning the iterations in lockstep and using this relational invariant: $y = y' \wedge (z = z' = 1 \vee z > z')$. This condition is not preserved by the loop bodies under guard condition $y > 0$, for example in the state $y = 2, z = 6, z' = 4$ reached when $x = x' = 3$, but here we are concerned with the case $x > 3$. If we add $x > 3$ as a conjunct we get a condition that is indeed invariant for lockstep executions, but it is not inductive—that is, the verification condition for the loop body is not valid. But there is a simple invariant with which the relation can be proved:

$$y = y' \wedge ((y > 4 \wedge z = z' = 1) \vee (y > 0 \wedge z > 2 * z') \vee (y = 0 \wedge z > z')) \quad (1)$$

This is not established by the initialization, in case $x = 4$. Instead we use this invariant to prove correctness under precondition $x > 4$ and separately prove correctness under the very specific precondition $x = 4$ which can be proved, for example, by unrolling the loops. In short, we do case analysis, as in rule DISJ.¹⁵

Program $P2$ is equivalent to $P0$, and $P3$ to $P1$, but neither fact is easily proved using lockstep alignment. For the simplest invariants in proving $P0$ equivalent to $P2$ we should let $P2$ take two iterations for each one of $P0$. The question is how to formulate that nicely.

As another example, $P2$ majorizes $P3$, for $x > 4$, but again this is not easily proved by reasoning about lockstep alignment of the loops. Both programs have gratuitous iterations in which y and z are not changed. We would like to align the computations so that when $w = w' = 0$ we can assert (1). Indeed, when $w \neq 0$ (respectively $w' \neq 0$), an iteration on the left (resp. right) has no effect on the other variables and thus maintains (1). For this proof we may try a simple-condition product so joint steps are taken when $(w \bmod 2) = 0 = (w' \bmod 3)$. But this is insufficient: it allows one side to run ahead in states where the condition does not require both sides to step together, precluding a simple invariant. What we need is a 3-condition product. The left may advance independently when $w \bmod 2 \neq 0$ and $w/2 = w'/3$; the right when $w' \bmod 3 \neq 0$ and $w/2 = w'/3$. Then (1) is invariant.

The examples only scratch the surface. Compilation, refactoring, and program revision need less obvious alignments, but often do admit alignments for which simple and even inferrable invariants suffice.

In examples like equivalence of $P0$ and $P2$ there is a fixed correspondence between loops of one versus the other program, a pattern that arises in some loop transformations used in compilers (e.g., to introduce vector operations). For majorization of $P3$ by $P2$ our alignment is more data-dependent, although it is not surprising that it can be described succinctly since the iterations have a regular pattern. Here is a less regular example (from [11]): the program uses a loop to sum the integers in a list, where list elements have a boolean flag that indicates an element should be ignored. The property is that two runs yield the

same sum, provided the two lists have the same non-deleted elements in the same order. This can be handled nicely using a 3-condition product.¹⁶

One can imagine more elaborate product automata using ghost state to track alignment conditions, but it seems that in any case what is needed is to designate when to advance on one side, the other side, or both.

4 Rules of relational program logic

As has been rediscovered and noted several times, it is not difficult to use program syntax to make a program that behaves as a product of programs. A simple case, mentioned earlier, is the sequence $C; C'$ where C' has no variables in common with C , and which corresponds to the sequenced product automaton. But it is also natural to interleave code from such disjoint programs, so as to align intermediate points in control flow. For a deductive system one also needs to account for the connection between such a product and the original program (or programs), the primary objects of interest. It is also desirable to disentangle reasoning principles, such as various alignments, from details of encoding. Furthermore, although disjoint variables suffice to reduce relational reasoning to standard HL for simple imperative programs, this is no longer the case for languages with more elaborate notions of state. For example, many languages feature a single heap and it is not trivial to use it to encode two disjoint heaps (see [59,23]). Another example is assembly language for a conventional architecture with a fixed set of registers. In such situations it may be preferable to work more directly with the relational correctness judgment, suitably interpreted, rather than depending entirely on products encoded as single programs.

We have reached the main topic of this paper, deductive systems for the relational judgment $C|C' : \mathcal{R} \approx \mathcal{S}$, in which various principles of reasoning are manifest in proof rules. With HL in mind we may expect syntax-directed rules that embody program semantics, rules for manipulation of specs, and rules for program transformation. In addition, relational reasoning benefits from judicious alignment of program fragments. For lockstep automata, the corresponding rules are dubbed “diagonal” [35] and relate programs with the same control structure. The sequenced and interleaved automata involve one-sided steps, corresponding to proof rules syntax-directed on one side. The 3-condition product is manifest in a three-premise rule for relating two loops. There are also rules that involve both relational and unary judgments.

Good alignment not only enables use of simple assertions, it is also essential to enable the use of relational specs for procedure calls. For lack of space we do not delve into this topic.

We refrain from formalizing relational formulas but we do assume they are closed under the propositional connectives with classical semantics. Usual formulations of HL rely on the use of program variables and expressions both as part of programs and as terms in formulas; in relational formulas we need to designate whether they refer to the left or right execution. As an alternative to the dashed/undashed convention used in Sec. 3, we use the notation $\langle e \rangle$ (resp.

$$\begin{array}{c}
 x := e \mid x' := e' : \mathcal{R}_{e|e'}^{x|x'} \approx \mathcal{R} \quad \frac{C|C' : \mathcal{R} \approx \mathcal{Q} \quad D|D' : \mathcal{Q} \approx \mathcal{S}}{C; D \mid C'; D' : \mathcal{R} \approx \mathcal{S}} \\
 \\
 \frac{C|C' : \mathcal{R} \wedge \{e\} \wedge \{e'\} \approx \mathcal{S} \quad D|D' : \mathcal{R} \wedge \{-e\} \wedge \{-e'\} \approx \mathcal{S} \\
 C|D' : \mathcal{R} \wedge \{e\} \wedge \{-e'\} \approx \mathcal{S} \quad D|C' : \mathcal{R} \wedge \{-e\} \wedge \{e'\} \approx \mathcal{S}}{\text{if } e \text{ then } C \text{ else } D \mid \text{if } e' \text{ then } C' \text{ else } D' : \mathcal{R} \approx \mathcal{S}} \\
 \\
 \frac{\mathcal{R} \Rightarrow \{e\} = \{e'\} \\
 C|C' : \mathcal{R} \wedge \{e\} \wedge \{e'\} \approx \mathcal{S} \quad D|D' : \mathcal{R} \wedge \{-e\} \wedge \{-e'\} \approx \mathcal{S}}{\text{if } e \text{ then } C \text{ else } D \mid \text{if } e' \text{ then } C' \text{ else } D' : \mathcal{R} \approx \mathcal{S}} \text{ (ALTAGREE)} \\
 \\
 \frac{\mathcal{Q} \Rightarrow \{e\} = \{e'\} \quad C|C' : \mathcal{Q} \wedge \{e\} \wedge \{e'\} \approx \mathcal{Q}}{\text{while } e \text{ do } C \text{ od} \mid \text{while } e' \text{ do } C' \text{ od} : \mathcal{Q} \approx \mathcal{Q} \wedge \{-e\} \wedge \{-e'\}} \text{ (ITERAGREE)} \\
 \\
 \frac{\mathcal{Q} \Rightarrow \{e\} = \{e'\} \vee (\mathcal{L} \wedge \{e\}) \vee (\mathcal{R} \wedge \{e'\}) \quad C|C' : \mathcal{Q} \wedge \{e\} \wedge \{e'\} \wedge \neg \mathcal{L} \wedge \neg \mathcal{R} \approx \mathcal{Q} \\
 C|\text{skip} : \mathcal{Q} \wedge \mathcal{L} \wedge \{e\} \approx \mathcal{Q} \quad \text{skip}|C' : \mathcal{Q} \wedge \mathcal{R} \wedge \{e'\} \approx \mathcal{Q}}{\text{while } e \text{ do } C \text{ od} \mid \text{while } e' \text{ do } C' \text{ od} : \mathcal{Q} \approx \mathcal{Q} \wedge \{-e\} \wedge \{-e'\}}
 \end{array}$$

Fig. 4. Diagonal syntax-directed rules.

$\{e\}$) for the value of expression e on the left (resp. right) side. As naming convention we tend to use dashed names for commands on the right side, but this does not imply renaming of variables or anything of the sort. In the logic, the programs are considered to act on distinct states which may or may not have the same variables. For example, we can write $\{x\} \leq \{x\}$ rather than $x \leq x'$.

4.1 Diagonal and one-side rules

The rules in Fig. 4 relate programs with the same control structure. Such rules are found in [35,22,73] and many other papers. In the assignment rule, the notation $\mathcal{R}_{e|e'}^{x|x'}$ is meant to be the formula \mathcal{R} in which left-side occurrences of x are replaced by e and right-side occurrences of x' by e' . For example, $(\{x\} = \{x\})_{x+1|y}^{x|x}$ is $\{x+1\} = \{y\}$. The first rule for if/else is general, covering the possible control flows, whereas ALTAGREE is applicable when the guard conditions are in agreement (and can be understood in terms of simple-condition pre-product with a condition to ensure adequacy). ALTAGREE can be derived from the first rule, using that $C|C' : \text{false} \approx \mathcal{S}$ and RELCONSEQ (Fig. 6).

The ITERAGREE rule (e.g., [22,73]) is applicable when the loop conditions remain in agreement under lockstep alignment; it uses a single invariant relation \mathcal{Q} much like the unary loop rule. The rule can be use to prove example $P0$ majorizes $P1$, for $x > 4$, using (1) as invariant. Francez gives a loop rule that corresponds to the eager-lockstep product:¹⁷ with a single invariant like in ITERAGREE but with additional premises $C|\text{skip} : \mathcal{Q} \wedge \{e\} \wedge \{-e'\} \approx \mathcal{Q}$ and $\text{skip}|C' : \mathcal{Q} \wedge \{-e\} \wedge \{e'\} \approx \mathcal{Q}$ to handle the situation that one loop continues

$$\begin{array}{c}
x := e | \text{skip} : \mathcal{R}_{e^x} \approx \mathcal{R} \qquad \frac{C | \text{skip} : \mathcal{R} \approx \mathcal{Q} \quad D | D' : \mathcal{Q} \approx \mathcal{S}}{C; D | D' : \mathcal{R} \approx \mathcal{S}} \text{ (LEFTSEQ)} \\
\\
\frac{B | C : \mathcal{R} \wedge \{e\} \approx \mathcal{S} \quad D | C : \mathcal{R} \wedge \{-e\} \approx \mathcal{S}}{\text{if } e \text{ then } B \text{ else } D | C : \mathcal{R} \approx \mathcal{S}} \\
\\
\frac{\text{while } e \wedge b \text{ do } B \text{ od} | C : \mathcal{P} \approx \mathcal{Q} \quad \text{while } e \text{ do } B \text{ od} | D : \mathcal{Q} \approx \mathcal{R} \quad \mathcal{Q} \wedge \{-e\} \Rightarrow \mathcal{R}}{\text{while } e \text{ do } B \text{ od} | C; D : \mathcal{P} \approx \mathcal{R}} \text{ (WHSEQ)}
\end{array}$$

Fig. 5. Some left side and mixed structure rules.

while the other has terminated; it is seldom helpful. Our second loop rule, from Beringer [23], corresponds to the 3-condition product: It augments the invariant \mathcal{Q} with two other relations: \mathcal{L} is precondition for an iteration on the left while the right side remains stationary; *mutatis mutandis* for \mathcal{R} . The side condition $\mathcal{Q} \Rightarrow ((\{e\} = \{e'\}) \vee (\mathcal{L} \wedge \{e\}) \vee (\mathcal{R} \wedge \{e'\}))$ ensures adequacy, i.e., covering all pairs of unary traces.

To relate differing programs, a natural idea is one-side rules, some of which we give in Fig. 5. The assignment rule is from Francez, where several one-side rules are given with `skip` on the other side, corresponding to interleaved product. The alternation rule is given in the more general form found in Barthe et al [14,16] and in Beringer [23] which also gives LEFTSEQ. If we identify D' with `skip`; D' (see Sec. 4.5), rule LEFTSEQ can be derived from sequence rule in Fig. 4 by replacing $C; D | D'$ with $C; D | \text{skip}; D'$. Right-side rules can be derived using rule SWAP (Sec. 4.4).

In addition to one-side rules that relate a structured program with an arbitrary one, Francez considers rules for relating different program structures, for example WHSEQ. The rule is unusual in that the premises are not judgments for subprograms of the one in the conclusion. The rule is derivable provided there are rules to rewrite programs to equivalent ones (see Sec. 4.5). Since `while` e `do` B `od` is unconditionally equivalent to the sequence (`while` $e \wedge b$ `do` B `od`); `while` e `do` B `od`, rewriting the conclusion results in a relation between two sequences.

4.2 From unary correctness to relational

If the variables of C' are disjoint from those of C then the semantics of command $C; C'$ amounts to the sequenced product of the corresponding automata, suggesting:

$$\frac{C \text{ and } C' \text{ have disjoint variables} \quad C; C' : \mathcal{R} \rightsquigarrow \mathcal{S}}{C | C' : \mathcal{R} \approx \mathcal{S}} \text{ (SEQPROD)}$$

For programs that cannot get stuck, it is sound in basic semantics according to Theorem 1 and the weak adequacy of sequenced product. Stuckness can be addressed using additional unary premises.

SEQPROD is useful as means to obtain relational judgments for small sub-programs such as assignments and basic blocks where a functional spec is not difficult to prove. An alternative way to get relational correctness from unary is by this rule, essentially introduced by Yang [73].

$$\frac{C : P \rightsquigarrow Q \quad D : R \rightsquigarrow S}{C \mid D : \langle P \rangle \wedge \langle R \rangle \approx \langle Q \rangle \wedge \langle S \rangle} \text{ (EMBED)}$$

It is sound in both basic and non-stuck semantics.

Typically, the relational assertion language does not express equality of entire states, but rather of specific variables and sometimes of partial heaps [73,11]. Equivalence of two programs can be specified as $C \mid C' : \mathcal{E} \approx \mathcal{F}$ where \mathcal{E} (resp. \mathcal{F}) expresses agreement on whatever parts of the state are read (resp. written) by C or C' . In a unary logic with frame conditions, suitable \mathcal{E}, \mathcal{F} can be derived from the frame condition [8] but I leave this informal in the following rule which yields a relational judgment from a unary one.

$$\frac{C : P \rightsquigarrow Q}{C \mid C : \mathcal{E} \wedge \mathbb{B}P \approx \mathcal{F}} \text{ (EREFL)}$$

Here $\mathbb{B}P$ abbreviates $\langle P \rangle \wedge \langle P \rangle$. One can add postcondition $\mathbb{B}Q$ by means of EMBED and RELCONSEQ. Further agreements can be added using RELFRAME (Fig. 6).

4.3 From relational correctness to unary

Preservation of unary correctness by equivalence transformation can be expressed as follows, where \mathcal{E}, \mathcal{F} are suitable agreements as in EREFL.

$$\frac{C : P \rightsquigarrow Q \quad C \mid D : \mathcal{E} \wedge \mathbb{B}P \approx \mathcal{F}}{D : P \rightsquigarrow Q} \text{ (ECORR)}$$

Whereas using unary judgments to infer relational ones allows for a deductive system in which the unary judgment stands on its own, this rule makes a dependency in reverse. We now take a further step which entangles assertion reasoning with correctness judgments.

Francez [35] motivates interest in the property of monotonicity by considering that it could be a requirement on a procedure passed to a routine for numeric integration. Similarly, a sorting routine requires that the comparator passed to it computes a transitive relation, and collections libraries require that the *equals* method compute a symmetric relation (at least) [68]. Evidently the functional correctness of such routines relies on these k -safety properties, but the cited papers do not even sketch such reasoning. Let us do so, glossing over details about parameter passing.

Consider a sorting routine that uses comparator $comp$ with inputs x, y and output z . Suppose in the proof of $sort(a, comp) : \text{true} \rightsquigarrow sorted(a)$ we rely on symmetry.¹⁸ That is, some use of CONSEQ is for an entailment that is valid owing to symmetry of comparison. Symmetry can be expressed as the relational judgment $comp|comp : x = y' \wedge y = x' \approx \triangleright z = z'$. But we need to connect this with reasoning about unary assertions, within the confines of a logic of relational and unary correctness judgments.

Such a connection is made in tools and theories that allow “pure methods” to be used in assertions while avoiding illogical nonsense using arbitrary program functions as mathematical ones [28,13]. Let C be some command meant to compute a function of input variables \bar{x} as output z . Let f be an uninterpreted (and fresh) name which we will use to represent that function. We have already seen how to express that z depends only on \bar{x} , deterministically: $C|C : \bar{x} = \bar{x}' \approx \triangleright z = z'$. A property such as symmetry or monotonicity has the form $C|C : \mathcal{R}(\bar{x}, \bar{x}') \approx \triangleright \mathcal{S}(z, z')$. To express that f is the function computed in z we use a unary spec, thus $C : \text{true} \rightsquigarrow z = f(\bar{x})$. Finally, we express the relational property of f as a first order (unary) formula: $\forall \bar{x}, \bar{x}'. \mathcal{R}(\bar{x}, \bar{x}') \Rightarrow \mathcal{S}(f(\bar{x}), f(\bar{x}'))$. With these ingredients we can state a rule.

$$\frac{C|C : \bar{x} = \bar{x}' \approx \triangleright z = z' \quad C|C : \mathcal{R}(\bar{x}, \bar{x}') \approx \triangleright \mathcal{S}(z, z') \quad f \text{ fresh} \quad C : \text{true} \rightsquigarrow z = f(\bar{x}) ; (\forall \bar{x}, \bar{x}'. \mathcal{R}(\bar{x}, \bar{x}') \Rightarrow \mathcal{S}(f(\bar{x}), f(\bar{x}')))) \vdash D : P \rightsquigarrow Q}{\text{“link } D \text{ with } C\text{”} : P \rightsquigarrow Q} \text{ (CMDFUN)}$$

We are glossing over procedures and parameter passing, and termination of C . The last premise, for D , is meant to indicate reasoning under a hypothesis. The hypothesis includes a unary judgment, as in formalizations of HL with procedures. It also includes the axiom about f for reasoning about assertions. The rule does not require C to be entirely deterministic and have no effects on other variables besides z , but we should disallow writes to \bar{x} , so $z = f(\bar{x})$ means what we want.

From $C : \text{true} \rightsquigarrow z = f(\bar{x})$ one can derive $C|C : \bar{x} = \bar{x}' \approx \triangleright z = z'$ by EMBED and RELCONSEQ. But CMDFUN does not require proof of $C : \text{true} \rightsquigarrow z = f(\bar{x})$. Instead, that spec is used to define f in terms of C , in reasoning about D .

4.4 Reasoning about specs

The reasoning embodied by CONSEQ and other spec rules in HL is also needed in RHL, e.g., in Sec. 3.2 we suggested an appeal to the relational disjunction rule. Some of these rules are in Fig. 6. In addition to logical connectives, it is natural to consider formulas with converse and relational composition, for which I write \mathcal{R}^\sim and $\mathcal{R};\mathcal{S}$ respectively. Rule SWAP is sound in basic and non-stuck semantics (but not for relative termination). Rule COMP is not sound in basic or non-stuck semantics, owing to possible divergences of C_1 ; these are precluded under relative termination and mutual termination semantics. Soundness of COMP can also be achieved using an additional premise for termination.

$$\begin{array}{c}
 \frac{\mathcal{P} \Rightarrow \mathcal{R} \quad C|D : \mathcal{R} \approx \mathcal{S} \quad \mathcal{S} \Rightarrow \mathcal{Q}}{C|D : \mathcal{P} \approx \mathcal{Q}} \text{ (RELCONSEQ)} \\
 \\
 \frac{C|C' : \mathcal{P} \approx \mathcal{Q} \quad FV(\mathcal{R}) \text{ disjoint from } Vars(C, C')}{C|C' : \mathcal{P} \wedge \mathcal{R} \approx \mathcal{Q} \wedge \mathcal{R}} \text{ (RELFRAME)} \\
 \\
 \frac{C|C' : \mathcal{P} \approx \mathcal{Q}}{C'|C : \mathcal{P} \approx \mathcal{Q}} \text{ (SWAP)} \quad \frac{C_0|C_1 : \mathcal{P} \approx \mathcal{Q} \quad C_1|C_2 : \mathcal{R} \approx \mathcal{S}}{C_0|C_2 : \mathcal{P}; \mathcal{R} \approx \mathcal{Q}; \mathcal{S}} \text{ (COMP)}
 \end{array}$$

Fig. 6. Some rules that manipulate specs.

Let us abbreviate the agreement $\{x\} = \{x\}$ by $\mathbb{A}x$. We have focused on local agreements like $\mathbb{A}x$, but one may wish to include a global identity relation, for which we write \mathcal{I} . As Benton shows, partial equivalences (symmetric and transitive relations, *per* for short) are particularly important, and relation operations let us express such properties as valid implications: $\mathcal{R} \sim \Rightarrow \mathcal{R}$ (symmetry) and $\mathcal{R}; \mathcal{R} \Rightarrow \mathcal{R}$ (transitivity). Several works use relational specs to express partial declassification of secrets (e.g., [57]). To declassify the value of expression e , a typical precondition has the form $\mathbb{A}e \wedge \mathbb{B}P$ which is a per but not reflexive. Apropos rule COMP instantiated in the form $C_0|C_2 : \mathcal{R}; \mathcal{R} \approx \mathcal{S}; \mathcal{S}$, if \mathcal{S} is transitive we obtain $C_0|C_2 : \mathcal{R}; \mathcal{R} \approx \mathcal{S}$ using RELCONSEQ. Then if \mathcal{R} is reflexive ($\mathcal{I} \Rightarrow \mathcal{R}$) we obtain $C_0|C_2 : \mathcal{R} \approx \mathcal{S}$, as $\mathcal{I}; \mathcal{R}$ is equivalent to \mathcal{R} .

By analogy with rule ECORR we would like to reason about preservation of a relational property by equivalence transformation. Consider the relation $C|C' : \mathcal{R} \approx \mathcal{S}$ together with equivalences $D|C : \mathcal{E} \approx \mathcal{F}$ and $C'|D' : \mathcal{E} \approx \mathcal{F}$ where \mathcal{E}, \mathcal{F} are suitable agreements. By COMP we get $D|D' : \mathcal{E}; \mathcal{R}; \mathcal{E} \approx \mathcal{F}; \mathcal{S}; \mathcal{F}$. If \mathcal{E} is a conjunction of agreements including variables of \mathcal{R} , then \mathcal{R} is equivalent to $\mathcal{E}; \mathcal{R}; \mathcal{E}$ and likewise for \mathcal{S} so by RELCONSEQ we obtain $D|D' : \mathcal{R} \approx \mathcal{S}$. Besides enabling derivation of right-side rules from left-side rules, rule SWAP facilitates instantiating the preceding reasoning in case $C = C'$ and $D = D'$, to show a security property of C is preserved by the equivalence. (Take \mathcal{R}, \mathcal{S} to be agreement on non-secret variables.)

Benton [22] makes the beautiful observation that just as the relational spec $\mathbb{A}x \approx \mathbb{A}z$ characterizes a dependency property of a single program, it also captures that two programs are equivalent with respect to their effect on z , e.g. $z := x; y := z \mid z := x : \mathbb{A}x \approx \mathbb{A}z$ captures a dead-code elimination transform, for a context where the subsequent code does not use y and therefore requires no agreement on it.

With this in mind, consider programs in which atomic actions happen in different orders, for example $z := x + 1; w := y$ versus $w := y; z := w + 1$, the equivalence of which can be expressed by the spec $\mathbb{A}x \wedge \mathbb{A}y \approx \mathbb{A}z \wedge \mathbb{A}w$. A general rule for commuting assignments can be formulated requiring disjointness of the variables read in the assignments. Moreover, one can express such a rule

for assignments involving heap locations, given means to express agreements thereof.

Heap agreements are often needed up to bijective renaming of pointers [9,23], which can be encoded in ghost state. Such specs can be localized to the locations read and written by a given command, since preservation of additional agreements can be derived by RELFRAME. Yang’s logic [73] features a frame rule taking advantage of separating conjunction of relations. It is also possible to formulate a frame rule based on relational specs with frame conditions, as in the work of Banerjee et al [8] which features local equivalence specs derived from frame conditions.

4.5 Transformations

The diagonal and one-side rules enable reasoning in terms of convenient alignments but apply only to specific control structure patterns. Programs that do not exactly match the patterns can be rewritten by equivalences such as $\text{skip}; C \cong C$, $C; \text{skip} \cong C$, and the following:

while e **do** C **od** \cong **while** e **do** C ; **while** $e \wedge e\theta$ **do** C **od od**
while e **do** C **od** \cong **if** e **do** C **fi**; **while** e **do** C **od**

Commands C, C' are *unconditionally equivalent*, written $C \cong C'$, if they have exactly the same store traces.¹⁹ The relation can be formalized using laws like these together with congruence rules. Such equivalences can be used to desugar fancy control structures, as done in some verification tools; the justification is that $C : P \rightsquigarrow Q$ and $C \cong D$ implies $D : P \rightsquigarrow Q$ (cf. rule ECORR in Sec. 4.4). The relational logic of Banerjee et al [11] features a rule like this: from $C|C' : \mathcal{R} \approx \mathcal{S}$, $D \cong C$, and $C' \cong D'$, infer $D|D' : \mathcal{R} \approx \mathcal{S}$. The rule is applied in proving a loop tiling transformation, using the above rewrites to enable application of diagonal rules. Transformations are used similarly in [19,42]. To enable use of sequenced product one may use the equivalence $\text{var } x \text{ in } C \cong \text{var } x' \text{ in } C_{x'}^x$ for fresh x' .

It seems unparsimonious to rely on an additional program relation (\cong) for which axioms and rules must be provided and proved sound, in a setting where we already consider a form of program relation.²⁰ On the other hand, we have seen in Sec. 4.4 that there are limitations on the use of equivalence judgments for reasons of termination. Having a separate judgment of unconditional equivalence is one way to address termination in connection with the basic or non-stuck semantics of relational judgments.

4.6 Alignment completeness

The usual notion of completeness is that true judgments are provable. Suppose the relational judgment $C|C' : \mathcal{R} \approx \mathcal{S}$ is true. In a setting where \mathcal{R}, \mathcal{S} can be expressed as, or considered to be, unary formulas, one can prove it by application of SEQPROD. In turn, the sequence can be reduced to true judgments $C : \mathcal{R} \rightsquigarrow \mathcal{Q}$ and $C' : \mathcal{Q} \rightsquigarrow \mathcal{S}$. What matters is not that an explicit product

$C; C'$ can be formed but rather that store relations can be expressed as store predicates [35,17,24,23]. If so, the judgment is provable provided the unary HL is complete. Then a single rule for relational judgments (SEQPROD or EMBED) is complete on its own!²¹ A different notion is needed.

Suppose $C : P \rightsquigarrow Q$ can be proved using IAM with a particular annotation. Then there is a HL proof using that annotation, in the sense that at least the loop rule is instantiated according to the annotation (assuming that loops are cut at loop headers). Why? Because the VCs will be provable, by completeness of HL, and the syntax-directed rules suffice to compose the VCs. In this sense, HL is complete with respect to IAM for unary correctness.

A natural measure of completeness for RHL is whether any proof of $C|C' : \mathcal{R} \approx \mathcal{S}$ using IAM with a product automaton can be represented by an RHL proof *using the same annotation and alignment*. Turning this into a precise definition requires, first, a convincing general definition of product automaton; our 3-condition form is relatively general but does not encompass the use of ghost state for alignment conditions or store relations. Second, the correspondence between proof rules and aligned products, discussed informally throughout Secs. 4.1–4.5, needs to be made precise. To this end it may help to limit attention to annotations in which all branch points are cutpoints. We leave this to future work but note that formal proof outlines [5] may be a convenient intermediary.

It is straightforward to add ghost state to our notions of pre-product and adequacy, to express store relations and alignments. But some program transformations used in optimizing compilers reorder an unbounded number of atomic actions. These do not have an obvious representation by pre-product and they have not been formalized using RHL rules [56].

5 Selected additional related work

The idea of relating C to C' by unary reasoning about a program that represents their product goes back at least to the 1970s. In Reynolds' book [65] we find stepwise refinement from an algorithm C using local variables of abstract mathematical types to C' acting on concrete data structures, expressed by augmenting C with parts of C' interwoven in such a way that assertions can express the coupling relation between abstract and concrete data. DeRoeveer and Engelhardt call this Reynolds' method and devote a chapter to it, citing work by Susan Geary as precursor [66]. Morgan [52] formalizes the idea in terms of auxiliary variables, cf. rule AUXVAR. The idea of encoding two runs as a sequence of disjoint copies, and specifying determinacy as a Hoare triple, appears (in passing) in a 1986 paper by Csirmaz and Hart [30].

The influential papers by Benton [22] and Barthe et al [17] have been followed by many works. The rest of this section gives selected highlights.

Barthe, Crespo and Kunz [16] give several ways of formulating deductive reasoning about relational properties, including deductive systems for product programs in conjunction with unary HL. They formalize a judgment that con-

nects two commands with a command that represents their product. Products include assertions which must be verified to ensure what we call adequacy.

Beringer [23] considers partial correctness specs in “VDM style” i.e., as relations from pre- to post-state, so partial correctness means the relational semantics of the program is a subset of the spec. He defines relational decompositions, essentially the relations that hold at the semicolon of a product $C; C'$ (as in rule SEQPROD), and observes that given such an “interpolant” one can derive VCs for C and C' as quotients in the sense of relational calculus (also known as weakest prespecification [41]). This is used to derive a collection of RHL rules including diagonal and one-side rules as well as relational DISJ/CONJ, for imperative commands including the heap.

Beckert and Ulbrich [21] survey some of the main ideas in relational verification and describe a range of applications and works on verification. Maillard et al [50] introduce a general framework for relational logics, applicable to a range of computational effects such as exceptions. Aguirre et al [1] develop a logic based on relational refinement types, for terminating higher order functional programs, and provide an extensive discussion of work on relational logics. Recent proceedings of CAV include quite a few papers on relational verification, and further perspectives can be found in the report from a recent Dagstuhl Seminar on program equivalence [47].

Numerous works develop variations and extensions of the ideas in this paper. Terauchi and Aiken [69] observe that sequenced product necessitates use of strong intermediate assertions, and use a dependency type system to guide the construction of more effective products. They also coin the term 2-safety. Several works focus on modular reasoning and product constructions that enable use of relational specs for procedures [37,74,45,39,46,11,72,32]. Sousa and Dillig [68] formulate a logic for k -safety, with notation that stands for “any product” and may be understood as providing for lazy product construction. Eilers et al [32] give a k -product encoding that lessens code duplication. Whereas many works handle only lockstep alignment of loops, some cover the 3-condition automata [23,19]; Shemer et al [67] provide for more general alignment and infer state-dependent alignment conditions. Other works on inferring or expressing effective alignments include [36,64,25]. Product constructions for $\forall\exists$ properties appear in [15,27].

Richer formalisms like Dynamic Logic [31,20] and embedding in higher order logic [38,1,50] have their advantages and can address reasoning like rule CMD-FUN and the linking of procedures to their implementations which is often left semi-formal. But such embeddings, in particular, are far from providing the level of automation (and teachability!) that more direct implementations of HL/RHL can provide. Completeness results show how HL/RHL suffice for proving correctness judgments.

6 Conclusion

I spelled out a number of patterns of reasoning for program relations and relational properties of programs, in terms of product automata that model pairs of executions, and also as rules of relational program logic. Almost all the rules can be found in at least one prior publication but some “obvious” and useful rules are missing in several papers. Spelling out the inductive assertion method for relational properties, as Francez [35] does, makes explicit the alignment principles that should be embodied in deductive rules, guiding the design of such rules. On this basis I introduced the notion of alignment completeness, leaving its formalization to future work; it should be done for a more general form of product than the one I chose for expository purposes.

To streamline notation I focused on 2-run properties but there is strong motivation for some 3-run (e.g., transitivity). I am not aware of fundamentally different techniques or principles for k -run that are not at hand for 2-run.²²

Although several papers have described the need for k -safety properties in order to reason about unary correctness, to my knowledge this pattern of reasoning has not been provided by relational logics (aside from those embedded in expressive higher order logics). I present a new rule for this (CMDFUN) that stays within the limited resources of RHL, i.e., assertions, unary correctness, and relational correctness judgments.

A couple of years ago I moved to a smaller office. While winnowing paper files I came across the paper by Francez, which I had acquired but not fully appreciated when working full time at IBM as a programmer in the '80s. The dearth of citations shows I am not alone in not finding it when I searched online for relevant work. My copy is a publisher's reprint, affixed with stickers that indicate IBM paid a fee. Such stickers became obsolete but the flow of scientific knowledge is still too tangled with commerce.

Acknowledgments. The paper was improved thanks to comments from Krzysztof Apt, Anindya Banerjee, Gilles Barthe, Ramana Nagasamudram, and anonymous reviewers. The research was partially supported by NSF CNS 1718713 and ONR N00014-17-1-2787.

Notes

¹ Consult, for example, the book [5]. There may be infinitely many variables but any program has the *frame property*: there is a finite set of variables (or memory locations) such that the program acts on, and depends on, only those, leaving the rest unchanged.

² This loses no generality and facilitates the definition, later, of *destutter*.

- 3 For some variations of structural operational semantics it suffices for *Ctrl* to be something like finite sequences of subprograms of the program of interest.
- 4 Because then we can write $\hat{P} \rightsquigarrow \hat{Q}$ where \hat{P} conjoins to P some equations of the form $x = \hat{x}$, one for each program variable y not in \bar{x} , with each \hat{x} a distinct spec-only variable not free in P or Q . Conjoining the same equations to Q gives the property that program variables $y \notin \bar{x}$ are unchanged, given the stipulation that spec-only variables are never changed.
- 5 A detailed history of these ideas and terminology, which go back to Turing [70], is provided by Apt and Olderog in their article on Hoare’s logic [6].
- 6 See for example Apt et al [5] where completeness is defined relative to completeness of reasoning about assertions and with the requirement that the assertion language be expressive enough to capture loop invariants, which I gloss over in this paper.
- 7 There seems to be only a little room for variation in formulating the syntax-directed rules, but more room for the rules that manipulate correctness judgments. For example, consider this sound rule which is akin to the DISJ rule:

$$\frac{C : P \rightsquigarrow Q \quad x \notin FV(Q) \quad x \notin Vars(C)}{C : \exists x. P \rightsquigarrow Q}$$

In some sources one finds a variation with conclusion of this form: $C : \exists x. P \rightsquigarrow \exists x. Q$. The variation can be derived from the displayed rule using CONSEQ and the tautology $Q \Rightarrow \exists x. Q$.

- 8 Such operators are available in JML and are useful in connection with behavioral subtyping [48]. One can also define an intrinsic refinement order on programs in terms of the specs they satisfy. Tony Hoare long ago shifted his attention towards algebraic approaches in which programs and specs are freely combined and related by refinement. In the variation called refinement calculus [7,53,54], specs are considered as a kind of atomic command the central judgment is refinement, of programs, specs, and combinations thereof, which subsumes the correctness judgment $C : P \rightsquigarrow Q$ as $P \rightsquigarrow Q \sqsubseteq C$.
- 9 For example, rewriting potentially faulting or diverging expressions into primitive commands, so other expressions are pure and have mathematical semantics as needed by the “pun” of program expressions in formulas. Another example is assigning subexpressions to temporary variables in order for atomic commands to make at most one heap access.
- 10 We give the rule AUXVAR in a simple form, for clarity. But one may want other ghost computation besides assignments. For soundness of such a rule the ghost code must be terminating [33].

- 11 Francez and others notice the analogy with disjoint parallelism and use terms like “synchronization”. I prefer the term “alignment” [44] which is apt and not confusing even if relations are considered between concurrent programs.
- 12 That is, $(a, b)(R \otimes S)(c, d)$ iff aRc and bSd .
- 13 For trace τ define $destutter(\tau) = \tau$ if $len(\tau) = 1$, $destutter(\gamma_0 :: \gamma_1 :: \tau) = destutter(\gamma_1 :: \tau)$ if $\gamma_0 = \gamma_1$, and $destutter(\gamma_0 :: \gamma_1 :: \tau) = \gamma_0 :: destutter(\gamma_1 :: \tau)$ otherwise.
- 14 In [35] this is claimed for $x > 2$, but that is not true of these programs, as we have $3! < 2^3$.
- 15 In the case we have precondition $x > 4$, the initializations establish (1), in particular the first of the disjuncts. Note that following the loop we have both (1) and the negated guard, i.e. $y = 0$, whence the postcondition $z > z'$. It remains to show that (1) is preserved by the loop body when $y \neq 0$. To this end we reason forward, considering each disjunct in turn, and writing y_0, z_0, y'_0, z'_0 for the initial values. so we have $y = y_0 - 1, y' = y'_0 - 1, z = z_0 * y_0$, and $z' = z_0 * 2$.
- If the first disjunct holds initially, i.e., $y_0 > 4 \wedge z_0 = z'_0 = 1$, we get $z = y_0 > 4 = 2 * 2 = z' * 2$ and $y > 0$, whence the second disjunct in (1).
 - If the second disjunct holds initially, i.e., $y_0 > 0 \wedge z_0 > 2 * z'_0$, we make a further case split:
 - if $y_0 = 1 \wedge z_0 > 2 * z'_0$ then $y = 0$ and we get the rest of the third disjunct by $z = z_0 * 1 > 2 * z'_0 = z'$.
 - if $y_0 > 1 \wedge z_0 > 2 * z'_0$ then $y > 0$ and we get the rest of the second disjunct by $z = z_0 * y_0 \geq z_0 * 2 > 2 * z'_0 * 2 = 2 * z'$
 - The third disjunct does not hold initially, given the guard conditions $y \neq 0$.
- 16 Advance on the left (resp. right) if the next element on the left (resp. right) is deleted; if neither are deleted then both sides advance together.
- 17 As does the rule Fusion 2 of Sousa and Dillig [68].
- 18 Transitivity is certain to be needed, but it is 3-safety which is inconvenient for expository purposes.
- 19 That is, projecting out just the stores from their state traces.
- 20 Indeed, Benton formulates some unconditional equivalences within his relational calculus DDCC [22, sec 3].
- 21 The fact that the technique of SEQPROD is complete relative to unary HL is observed by Francez [35] but not worked out in detail. The fact is also

mentioned in Barthe et al [17], for the special case of relating a program to itself; it is evident that it holds more generally as noted in [14]. Semantic completeness of the technique, for general relational properties, is proved by Beringer [23]. A RHL is proved complete for deterministic programs, on this basis, in [19]. The logic of Sousa and Dillig [68] includes a rule like SEQPROD for k -products, and their Theorem 2 is completeness relative to completeness of an underlying HL; Wang et al [71] prove a similar result specialized to program equivalence. The crux of these completeness results is that product programs are complete in the sense of representing all pairs or k -tuples of unary executions (called adequacy in this paper). Francez gives a semantic completeness result of this sort, for eager-lockstep product, as do Eilers et al [32] for a more general form of product. For (higher order) functional programs, Aguirre et al [1] prove completeness for a RHL via embedding in a unary logic.

- 22 Another topic omitted for lack of space is the soundness of the relational rules under $\forall\exists$ interpretations. As in the case of partial versus total correctness interpretation of unary specs, most of the rules are sound but loop rules must be changed. Different notions of adequacy are needed for product automata.

References

1. Aguirre, A., Barthe, G., Gaboardi, M., Garg, D., Strub, P.: A relational logic for higher-order programs. *J. Funct. Program.* **29** (2019), <https://doi.org/10.1017/S0956796819000145>
2. Apt, K.: Ten years of Hoare’s logic, a survey, part I. *ACM Trans. Progr. Lang. Syst.* **3**(4), 431–483 (1981)
3. Apt, K.: Ten years of Hoare’s logic, a survey, part II: nondeterminism. *Theor. Comput. Sci.* **28** (1984)
4. Apt, K.R.: Correctness proofs of distributed termination algorithms. *ACM Trans. Progr. Lang. Syst.* **8**, 388–405 (1986)
5. Apt, K.R., de Boer, F.S., Olderog, E.R.: *Verification of Sequential and Concurrent Programs*. Springer, 3 edn. (2009)
6. Apt, K.R., Olderog, E.: Fifty years of Hoare’s logic. *Formal Asp. Comput.* **31**(6), 751–807 (2019)
7. Back, R.J., von Wright, J.: *Refinement Calculus: A Systematic Introduction*. Springer-Verlag (1998)
8. Banerjee, A., Nagasamudram, R., Nikouei, M., Naumann, D.A.: A relational program logic with data abstraction and dynamic framing. *CoRR* **abs/1910.14560** (2019), <http://arxiv.org/abs/1910.14560>
9. Banerjee, A., Naumann, D.A.: Ownership confinement ensures representation independence for object-oriented programs. *Journal of the ACM* **52**(6), 894–960 (2005)
10. Banerjee, A., Naumann, D.A.: Local reasoning for global invariants, part II: Dynamic boundaries. *Journal of the ACM* **60**(3), 19:1–19:73 (2013)
11. Banerjee, A., Naumann, D.A., Nikouei, M.: Relational logic with framing and hypotheses. In: *Foundations of Software Tech. and Theoretical Comp. Sci.* pp. 11:1–11:16 (2016), technical report at <http://arxiv.org/abs/1611.08992>

12. Banerjee, A., Naumann, D.A., Nikouei, M.: Relational logic with framing and hypotheses. In: 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. pp. 11:1–11:16 (2016), technical report at <http://arxiv.org/abs/1611.08992>
13. Banerjee, A., Naumann, D.A., Nikouei, M.: A logical analysis of framing for specifications with pure method calls. *ACM Trans. Progr. Lang. Syst.* **40**(2), 6:1–6:90 (2018)
14. Barthe, G., Crespo, J.M., Kunz, C.: Relational verification using product programs. In: *Formal Methods. LNCS*, vol. 6664, pp. 200–214 (2011)
15. Barthe, G., Crespo, J.M., Kunz, C.: Beyond 2-safety: Asymmetric product programs for relational program verification. In: *LFCS*. pp. 29–43 (2013)
16. Barthe, G., Crespo, J.M., Kunz, C.: Product programs and relational program logics. *J. Logical and Algebraic Methods in Programming* **85**(5), 847–859 (2016)
17. Barthe, G., D’Argenio, P.R., Rezk, T.: Secure information flow by self-composition. In: *IEEE CSFW*. pp. 100–114 (2004), see extended version [18].
18. Barthe, G., D’Argenio, P.R., Rezk, T.: Secure information flow by self-composition. *Math. Struct. Comput. Sci.* **21**(6), 1207–1252 (2011)
19. Barthe, G., Grégoire, B., Hsu, J., Strub, P.: Coupling proofs are probabilistic product programs. In: *POPL*. pp. 161–174 (2017)
20. Beckert, B., Hähnle, R., Schmitt, P.H. (eds.): *Verification of Object-Oriented Software: The KeY Approach, LNCS*, vol. 4334. Springer-Verlag (2007)
21. Beckert, B., Ulbrich, M.: Trends in relational program verification. In: Müller, P., Schaefer, I. (eds.) *Principled Software Development*, pp. 41–58. Springer (2018)
22. Benton, N.: Simple relational correctness proofs for static analyses and program transformations. In: *POPL*. pp. 14–25 (2004)
23. Beringer, L.: Relational decomposition. In: *Interactive Theorem Proving (ITP)*. LNCS, vol. 6898 (2011)
24. Beringer, L., Hofmann, M.: Secure information flow and program logics. In: *IEEE CSF*. pp. 233–248 (2007)
25. Churchill, B.R., Padon, O., Sharma, R., Aiken, A.: Semantic program alignment for equivalence checking. In: *PLDI*. pp. 1027–1040 (2019)
26. Clarkson, M.R., Schneider, F.B.: Hyperproperties. *Journal of Computer Security* **18**(6), 1157–1210 (2010)
27. Clochard, M., Marché, C., Paskevich, A.: Deductive verification with ghost monitors. *Proc. ACM Program. Lang.* **4**(POPL), 2:1–2:26 (2020)
28. Cok, D.R.: Reasoning with specifications containing method calls and model fields. *Journal of Object Technology* **4**(8), 77–103 (2005)
29. Csirmaz, L.: Program correctness on finite fields. *Periodica Mathematica Hungarica* **33**(1), 23–33 (1996)
30. Csirmaz, L., Hart, B.: Program correctness on finite fields. In: *IEEE Symp. on Logic in Computer Science (LICS)*. pp. 4–10 (1986), see also [29].
31. Darvas, A., Hähnle, R., Sands, D.: A theorem proving approach to analysis of secure information flow. In: *Security in Pervasive Computing*. pp. 193–209 (2005)
32. Eilers, M., Müller, P., Hitz, S.: Modular product programs. *ACM Trans. Program. Lang. Syst.* **42**(1), 3:1–3:37 (2020)
33. Filliâtre, J., Gondelman, L., Paskevich, A.: The spirit of ghost code. *Formal Methods in System Design* **48**(3), 152–174 (2016)
34. Floyd, R.: Assigning meaning to programs. In: *Symp. on Applied Math.* 19, *Math. Aspects of Comp. Sci.* pp. 19–32. Amer. Math. Soc. (1967)
35. Francez, N.: Product properties and their direct verification. *Acta Informatica* **20**, 329–344 (1983)

36. Girka, T., Mentré, D., Régis-Gianas, Y.: Verifiable semantic difference languages. In: *Principles and Practice of Declarative Programming (PPDP)* (2017)
37. Godlin, B., Strichman, O.: Inference rules for proving the equivalence of recursive procedures. *Acta Inf.* **45**(6), 403–439 (2008)
38. Grimm, N., Maillard, K., Fournet, C., Hritcu, C., Maffei, M., Protzenko, J., Ramananandro, T., Rastogi, A., Swamy, N., Béguelin, S.Z.: A monadic framework for relational verification: applied to information security, program equivalence, and optimizations. In: *CPP* (2018)
39. Hawblitzel, C., Kawaguchi, M., Lahiri, S.K., Rebêlo, H.: Towards modularly comparing programs using automated theorem provers. In: *CADE*. pp. 282–299 (2013)
40. Hoare, C.A.R.: Procedures and parameters: an axiomatic approach. In: Engler, E. (ed.) *Symposium on the Semantics of Algorithmic Languages* (1971)
41. Hoare, C.A.R., He, J.: The weakest prespecification. *Inf. Process. Lett.* **24**(2), 127–132 (1987)
42. Kiefer, M., Klebanov, V., Ulbrich, M.: Relational program reasoning using compiler IR: Combining static verification and dynamic analysis. *J. Automated Reasoning* **60**, 337–363 (2018)
43. Kleymann, T.: Hoare logic and auxiliary variables. *Formal Aspects of Computing* **11**, 541–566 (1999)
44. Kovács, M., Seidl, H., Finkbeiner, B.: Relational abstract interpretation for the verification of 2-hypersafety properties. In: *ACM CCS* (2013)
45. Lahiri, S.K., Hawblitzel, C., Kawaguchi, M., Rebêlo, H.: SYMDIFF: A language-agnostic semantic diff tool for imperative programs. In: *CAV*. pp. 712–717 (2012)
46. Lahiri, S.K., McMillan, K.L., Sharma, R., Hawblitzel, C.: Differential assertion checking. In: *Joint Meeting of the European Software Engineering Conference and the ACM Symposium on the Foundations of Software Engineering* (2013)
47. Lahiri, S.K., Murawski, A.S., Strichman, O., Ulbrich, M.: Program equivalence (Dagstuhl Seminar 18151). *Dagstuhl Reports* **8**(4), 1–19 (2018)
48. Leavens, G.T., Naumann, D.A.: Behavioral subtyping, specification inheritance, and modular reasoning. *ACM Trans. Progr. Lang. Syst.* **37**, 13:1–13:88 (2015)
49. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: *LPAR*. pp. 348–370 (2010)
50. Maillard, K., Hritcu, C., Rivas, E., Muylder, A.V.: The next 700 relational program logics. *Proc. ACM Program. Lang.* **4**(POPL), 4:1–4:33 (2020)
51. Moore, J.S.: Inductive assertions and operational semantics. *Int. J. Softw. Tools Technol. Transf.* **8**(4-5), 359–371 (2006)
52. Morgan, C.: Auxiliary variables in data refinement. *Inf. Process. Lett.* **29**(6), 293–296 (1988)
53. Morgan, C.: *Programming from Specifications*, second edition. Prentice Hall (1994)
54. Morris, J.M.: A theoretical basis for stepwise refinement and the programming calculus. *Sci. Comput. Program.* **9**, 287–306 (1987)
55. Müller, C., Kovács, M., Seidl, H.: An analysis of universal information flow based on self-composition. In: *IEEE CSF* (2015)
56. Namjoshi, K.S., Singhania, N.: Loopy: Programmable and formally verified loop transformations. In: *Static Analysis Symposium (SAS)*. pp. 383–402 (2016)
57. Nanevski, A., Banerjee, A., Garg, D.: Verification of information flow and access control policies with dependent types. In: *IEEE Symp. on Sec. and Priv.* (2011)
58. Naumann, D.A.: Calculating sharp adaptation rules. *Inf. Process. Lett.* **77**, 201–208 (2001)
59. Naumann, D.A.: From coupling relations to mated invariants for secure information flow. In: *ESORICS. LNCS*, vol. 4189, pp. 279–296 (2006)

60. O’Hearn, P.W., Reynolds, J.C., Yang, H.: Local reasoning about programs that alter data structures. In: CSL. LNCS, vol. 2142, pp. 1–19. Springer-Verlag (2001)
61. O’Hearn, P.W., Yang, H., Reynolds, J.C.: Separation and information hiding. *ACM Trans. Program. Lang. Syst.* **31**(3), 1–50 (2009)
62. Olderog, E.R.: On the notion of expressiveness and the rule of adaptation. *Theoretical Computer Science* **30**, 337–347 (1983)
63. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs i. *Acta Inf.* **6** (1976)
64. Pick, L., Fedyukovich, G., Gupta, A.: Exploiting synchrony and symmetry in relational verification. In: *CAV*. pp. 164–182 (2018)
65. Reynolds, J.C.: *The Craft of Programming*. Prentice-Hall (1981)
66. de Roever, W.P., Engelhardt, K.: *Data Refinement: Model-Oriented Proof Methods and their Comparison*. Cambridge University Press (1998)
67. Shemer, R., Gurfinkel, A., Shoham, S., Vizel, Y.: Property directed self composition. In: *CAV*. pp. 161–179 (2019)
68. Sousa, M., Dillig, I.: Cartesian Hoare Logic for verifying k-safety properties. In: *PLDI*. pp. 57–69 (2016)
69. Terauchi, T., Aiken, A.: Secure information flow as a safety problem. In: *Static Analysis Symposium (SAS)*. LNCS, vol. 3672, pp. 352–367 (2005)
70. Turing, A.: On checking a large routine. In: *Report of a Conference on High Speed Automatic Calculating Machines*. pp. 67–69 (1949), univ. Math. Lab., Cambridge
71. Wang, Y., Dillig, I., Lahiri, S.K., Cook, W.R.: Verifying equivalence of database-driven applications. *Proc. ACM Program. Lang.* **2**(POPL), 56:1–56:29 (2018)
72. Wood, T., Drossopoulou, S., Lahiri, S.K., Eisenbach, S.: Modular verification of procedure equivalence in the presence of memory allocation. In: *ESOP* (2017)
73. Yang, H.: Relational separation logic. *Theor. Comput. Sci.* **375**, 308–334 (2007)
74. Zaks, A., Pnueli, A.: CoVaC: Compiler validation by program analysis of the cross-product. In: *Formal Methods*. LNCS, vol. 5014, pp. 35–51 (2008)