Simple Relational Correctness Proofs for Static Analyses and Program Transformations

Nick Benton
Microsoft Research
7 J J Thomson Avenue
Cambridge CB3 0FB, UK

nick@microsoft.com

Abstract

We show how some classical static analyses for imperative programs, and the optimizing transformations which they enable, may be expressed and proved correct using elementary logical and denotational techniques. The key ingredients are an interpretation of program properties as relations, rather than predicates, and a realization that although many program analyses are traditionally formulated in very intensional terms, the associated transformations are actually enabled by more liberal extensional properties.

We illustrate our approach with formal systems for analysing and transforming while-programs. The first is a simple type system which tracks constancy and dependency information and can be used to perform dead-code elimination, constant propagation and program slicing as well as capturing a form of secure information flow. The second is a relational version of Hoare logic, which significantly generalizes our first type system and can also justify optimizations including hoisting loop invariants. Finally we show how a simple available expression analysis and redundancy elimination transformation may be justified by translation into relational Hoare logic.

Categories and Subject Descriptors: F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages – *Denotational semantics, Partial evaluation, Program analysis*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs – *Logics of programs, Invariants*; D.3.4 [Programming Languages]: Processors – *Compilers, Optimization*;

General Terms: Languages, Theory, Verification

Keywords: Program analysis, optimizing compilation, types, denotational semantics, partial equivalence relations, Hoare logic, dependency, information flow, security

1 Introduction

Although static program analyses are routinely proved correct, the soundness of the optimizing transformations which they enable is much less frequently addressed. Much of the work which *has* been done on formalizing and validating analysis-based transformations comes from the functional programming community (see Section 5 for related work) – the literature on optimizations for imperative languages contains few formal specifications of transformations, let alone proofs of their correctness. One might think that this is because the correctness of most imperative optimizations is entirely trivial, but what literature there is on the subject [22, 18, 19, 20] (not to mention the occasional behaviour of real optimizing compilers) indicates that this is not so.

Why is proving correctness of analysis-based transformations hard? We wish to establish that, given the results of a static analysis, the original program and the transformed program are observationally equivalent. The first problem is that many transformations involve locally replacing some part P of a larger program $C[P]$ with a new version P' which is *not* generally observationally equivalent to P (i.e. $P \not\sim P'$), though they *are* equivalent in that particular context: $C[P] \sim C[P']$. Simply having an analysis which (albeit correctly) deduces that certain predicates $\psi(P)$ hold of program fragments does not straightforwardly allow one to justify such transformations *unless* the predicates $\psi(\cdot)$ also somehow involve sets of contexts $C[\cdot]$, which is often not the case.

The second difficulty in proving correctness of optimizing transformations is that program analyses, especially for imperative languages, are often specified in a very intensional way. For example, "an assignment $[x := a]^l$ *may reach* a certain program point if there is an execution of the program where x was last assigned a value at l when the program point is reached". Notions such as 'program point' and 'where a variable was last assigned' are not present in natural operational or denotational semantics, so the correctness of these analyses is frequently formulated in terms of a new (and essentially bogus) *instrumented semantics* which tracks the extra information. Even where the instrumented semantics is related back to the original one, the relation is usually a rather weak form of adequacy which certainly does not help with establishing equivalences: the instrumented semantics will generally have a weaker equational theory than the original one.

This paper demonstrates that, at least in some simple cases, both of these difficulties can be overcome by use of elementary ideas which are commonplace in the semantics community, but which have not previously been fully exploited in the context of compiler analyses and transformations.

The first difficulty is approached by taking seriously the notion of the semantics of types as (special kinds of) relations, rather than predicates. Typed lambda calculi are routinely presented using judgements of the form

$$\Gamma \vdash M = M' : A$$

which does not assert "under assumptions Γ, M equals M' and they both have type A", but rather "under assumptions Γ, M and M' are equal *at* type A". Such calculi can be modelled by interpreting types as partial equivalence relations over some untyped universe such as D_∞. Many program analyses are presented as non-standard type systems, and partial equivalence relations have been used to give semantics to these non-standard types (equivalently, elements of abstract domains), at least in the cases of binding-time [15] and security analyses [26]. However, even in those cases, the emphasis has been on simple typing judgements rather than equational reasoning. Our approach is to treat all abstract properties as relations, including those which have naive interpretations as predicates (e.g. 'X is 5'), and to present transformations by giving rules for deriving (non-standard) typed equations in context.

The second difficulty, the apparently intensional nature of properties, is often something of a red herring, attributable to a confusion between certain analysis algorithms and the semantics of the information they produce. Of course, analyses relating to properties such as time or space usage can only be justified relative to semantic models which make those aspects of computation explicit. But many transformations performed by optimizing compilers can be justified using more extensional semantics, not only in the weak sense that every input program is provably equivalent to its transformed version, but also in the stronger sense that there is a generic correctness argument for all programs. The true preconditions for applying a transformation tend to be extensional ("this command does not change the value of X+Y") even if an analysis algorithm only discovers those properties if a stronger intensional property holds ("this command does not contain any assignments to X or Y").

As a facetious example of the difference between the intensional and extensional approaches, consider why the following transformation is correct:

```
X := 7;              X := 7;
Y := Y+1;    ==>     Y := Y+1;
Z := X;              Z := 7;
```

The extensional answer is "when X is evaluated on the last line, its value will always be 7". The intensional answer is something like "the only definition of X which reaches its use on line 3 is the one on line 1, and the right hand side of that definition does not contain any variable which is assigned to in lines 1 or 2". This may well be an accurate account of *how* an algorithm works, but it is not a good basis for thinking about *what* it establishes. Things get even worse if we consider a sequence like

```
X := 7;            X := 7;            X := 7;
Y := Y+1;  ==>     Y := Y+1;  ==>     Y := Y+1;
X := 7;            X := 7;
Z := X;            Z := 7;            Z := 7;
```

After the first transformation, the intensional justification for the change to line 4 refers to the definition of X on line 3. But after the second transformation, that definition has gone, which complicates proving the correctness of the combined transformation. Problems of this sort occur both in real compilers (keeping analysis results

sound during transformations is notoriously tricky) and in proofs (see for example the discussion of interference between 'forward' analyses and 'backward' transformation patterns in [20]).

Another familiar example of the intensional/extensional distinction arises in optimizing compilation of lazy functional languages [9]. Some analysis algorithms aimed at detecting when CBN can be replaced by CBV look for functions which always evaluate their arguments ('neededness'), which is an intensional property. Their correctness (and that of the associated transformation) can be established in terms of the extensional property of strictness, which is much easier to reason about. Of course, since the extensional property is also weaker (holds of more functions) one is then naturally led to reformulate the analysis to establish the extensional property more directly.

2 The Language of `while`-Programs

The syntax and denotational semantics of the language of `while`-programs are entirely standard (see, e.g. [34]). To fix notation, they are briefly summarized in Figure 1. We sometimes use F_τ as a metavariable ranging over τ exp where $\tau \in \{\texttt{int}, \texttt{bool}\}$.

The denotational semantics is given in the category of ω-complete partial orders (predomains) and continuous functions. We write $\lceil \cdot \rceil : D \to D_\perp$ for the injection of a domain into its lift and $(\cdot)^* : (D \to D'_\perp) \to (D_\perp \to D'_\perp)$ for the associated extension operation. When $R \subseteq D \times D, R_\perp \subseteq D_\perp \times D_\perp$ is the relation defined by

$$R_\perp = \{(\lceil x \rceil, \lceil y \rceil) \mid (x, y) \in R\} \cup \{(\perp, \perp)\}$$

If $f : D \to E, x \in D$ and $y \in E$ then we define $f[x \mapsto y] : D \to E$ in the usual way:

$$(f[x \mapsto y])(z) = \begin{cases} y & \text{if } z = x \\ f(z) & \text{otherwise} \end{cases}$$

The denotational semantics is fully abstract with respect to the obvious operational semantics and definition of observational equivalence.

3 Dependency, Dead Code and Constants

In this section we present DDCC, a simple analysis and transformation system for `while`-programs which tracks dependency and constancy information, enabling optimizations such as constant-folding and dead-code elimination. As indicated in the introduction, the system is presented as a non-standard type system for deriving typed equalities between expressions and between commands.

3.1 DDCC Syntax and Semantics

3.1.1 Formulae

We begin by defining the syntax of some non-standard types for expressions. For $\tau \in \{\texttt{int}, \texttt{bool}\}, c \in [\![\tau]\!]$:

$$\phi_\tau \quad := \quad \mathbb{F}_\tau \mid \{c\}_\tau \mid \Delta_\tau \mid \mathbb{T}_\tau$$

Intuitively, $\{c\}_\tau$ is the type of τ-expressions equal to the constant c, Δ_τ is the type of τ-expressions whose value we do not know, whilst $\mathbb{T}_\tau$ is the type of τ-expressions whose value we do not care about. $\mathbb{F}_\tau$ is an empty expression type, which we have included

Syntax

$$\begin{aligned}
X \in \mathbb{V} &= \{\mathtt{X}, \mathtt{Y}, \ldots\} \\
n \in \mathbb{Z} \qquad b \in \mathbb{B} &= \{true, false\} \\
iop \quad \in \quad &\{+, \times, -, \ldots\} \subseteq \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z} \\
bop \quad \in \quad &\{<, =, \ldots\} \subseteq \mathbb{Z} \times \mathbb{Z} \to \mathbb{B} \\
lop \quad \in \quad &\{\vee, \wedge, \ldots\} \subseteq \mathbb{B} \times \mathbb{B} \to \mathbb{B} \\
\mathtt{int\ exp} \ni E \quad := \quad &\mathtt{n} \mid X \mid E \ iop \ E \\
\mathtt{bool\ exp} \ni B \quad := \quad &\mathtt{b} \mid E \ bop \ E \mid \mathtt{not}\ B \mid B \ \mathtt{lop}\ B \\
\mathtt{com} \ni C \quad := \quad &\mathtt{skip} \mid X\mathtt{:=}E \mid C\mathtt{;}C \mid \mathtt{if}\ B\ \mathtt{then}\ C\ \mathtt{else}\ C \mid \mathtt{while}\ B\ \mathtt{do}\ C
\end{aligned}$$

Denotational Semantics

$$S \in \mathbb{S} = \mathbb{V} \to \mathbb{Z}$$

$$\begin{aligned}
[\![E]\!] \in \mathbb{S} \to \mathbb{Z} \\
[\![\mathtt{n}]\!]S &= n \\
[\![X]\!]S &= S(X) \\
[\![E_1 \ \mathtt{iop}\ E_2]\!]S &= ([\![E_1]\!]S) \ iop \ ([\![E_2]\!]S)
\end{aligned}$$

$$\begin{aligned}
[\![B]\!] \in \mathbb{S} \to \mathbb{B} \\
[\![\mathtt{b}]\!]S &= b \\
[\![E_1 \ \mathtt{bop}\ E_2]\!]S &= ([\![E_1]\!]S) \ bop \ ([\![E_2]\!]S) \\
[\![B_1 \ \mathtt{lop}\ B_2]\!]S &= ([\![B_1]\!]S) \ lop \ ([\![B_2]\!]S) \\
[\![\mathtt{not}B]\!]S &= \neg([\![B]\!]S)
\end{aligned}$$

$$\begin{aligned}
[\![C]\!] \in \mathbb{S} \to \mathbb{S}_\perp \\
[\![\mathtt{skip}]\!] &= \lambda S. \lceil S \rceil \\
[\![X\mathtt{:=}E]\!] &= \lambda S. \lceil S[X \mapsto [\![E]\!]S] \rceil \\
[\![C_1\mathtt{;}C_2]\!] &= [\![C_2]\!]^* \circ [\![C_1]\!] \\
[\![\mathtt{if}\ B\ \mathtt{then}\ C_1\ \mathtt{else}\ C_2]\!] &= \lambda S. [\![B]\!]S \Longrightarrow [\![C_1]\!]S \mid [\![C_2]\!]S \\
[\![\mathtt{while}\ B\ \mathtt{do}\ C]\!] &= fix\ f. \lambda S. [\![B]\!]S \Longrightarrow f^*([\![C]\!]S) \mid \lceil S \rceil
\end{aligned}$$

Figure 1. Syntax and Semantics of `while` Programs

for completeness.[1] Semantically, the denotation of ϕ_τ is a binary relation on $[\![\tau]\!]$:

$$\begin{aligned}
[\![\mathbb{F}_\tau]\!] &= \emptyset \\
[\![\{c\}_\tau]\!] &= \{(c,c)\} \\
[\![\Delta_\tau]\!] &= \{(x,x) \mid x \in [\![\tau]\!]\} \\
[\![\mathbb{T}_\tau]\!] &= [\![\tau]\!] \times [\![\tau]\!]
\end{aligned}$$

Types for states are then finite maps from variables to types for `int` exps, written as lists with the usual conventions. In particular, writing $\Phi, X : \phi_{\mathtt{int}}$ implies that X does not occur in Φ.

$$\Phi \quad := \quad - \mid \Phi, X : \phi_{\mathtt{int}}$$

State types are interpreted as binary relations on $\mathbb{S}$:

$$\begin{aligned}
[\![-]\!] &= \mathbb{S} \times \mathbb{S} \\
[\![\Phi, X : \phi_{\mathtt{int}}]\!] &= [\![\Phi]\!] \cap \{(S,S') \mid (S(X), S'(X)) \in [\![\phi_{\mathtt{int}}]\!]\}
\end{aligned}$$

3.1.2 Entailment

There is a subtyping relation $\leq$ on expression types, which is axiomatised as follows:

$$\mathbb{F}_\tau \leq \phi_\tau \qquad \{c\}_\tau \leq \Delta_\tau$$

$$\phi_\tau \leq \mathbb{T}_\tau \qquad \phi_\tau \leq \phi_\tau$$

$$\frac{\phi_\tau \leq \phi'_\tau \quad \phi'_\tau \leq \phi''_\tau}{\phi_\tau \leq \phi''_\tau}$$

The above induces a depth- and width-subtyping relation on state types:

$$\Phi \leq - \qquad\qquad \Phi, X : \phi_{\mathtt{int}} \leq \Phi$$

$$\Phi, X : \mathbb{F}_{\mathtt{int}} \leq \Phi'$$

$$\frac{\Phi \leq \Phi'}{\Phi \leq \Phi', X : \mathbb{T}_{\mathtt{int}}} \qquad \frac{\phi_{\mathtt{int}} \leq \phi'_{\mathtt{int}}}{\Phi, X : \phi_{\mathtt{int}} \leq \Phi, X : \phi'_{\mathtt{int}}}$$

Because $\Phi, X : \mathbb{T}_{\mathtt{int}} \leq \Phi$ and $\Phi \leq \Phi, X : \mathbb{T}_{\mathtt{int}}$, absence of a variable from a state type is equivalent to it being present with type $\mathbb{T}_{\mathtt{int}}$.

LEMMA 1.

1. *For all ϕ_τ and Φ, $[\![\phi_\tau]\!]$ and $[\![\Phi]\!]$ are partial equivalence relations.*

2. *The $\leq$ relation on state types is reflexive and transitive.*

3. *If $\phi_\tau \leq \phi'_\tau$ then $[\![\phi_\tau]\!] \subseteq [\![\phi'_\tau]\!]$.*

4. *If $\Phi \leq \Phi'$ then $[\![\Phi]\!] \subseteq [\![\Phi']\!]$.*

3.1.3 Judgements

DDCC has two basic forms of judgement. For expressions, with $F, F' \in \tau$ exp, we have judgements of the form

$$\vdash F \sim F' : \Phi \Rightarrow \phi_\tau$$

whilst for commands, $C, C' \in \mathtt{com}$, there are judgements of the form

$$\vdash C \sim C' : \Phi \Rightarrow \Phi'$$

---

[1]This is really just a matter of taste. $\mathbb{F}_\tau$ does not appear in many interesting derivations.

We write $\vdash C : \Phi \Rightarrow \Phi'$ as shorthand for $\vdash C \sim C : \Phi \Rightarrow \Phi'$ and similarly for single-subject expression judgements. If we define

$$
\begin{aligned}
[\![\Phi \Rightarrow \phi_\tau]\!] &\subseteq (\mathbb{S} \to [\![\tau]\!]) \times (\mathbb{S} \to [\![\tau]\!]) \\
&\equiv \{(f, f') \mid \forall (S, S') \in [\![\Phi]\!]. \ (fS, f'S') \in [\![\phi_\tau]\!]\}
\end{aligned}
$$

$$
\begin{aligned}
[\![\Phi \Rightarrow \Phi']\!] &\subseteq (\mathbb{S} \to \mathbb{S}_\perp) \times (\mathbb{S} \to \mathbb{S}_\perp) \\
&\equiv \{(f, f') \mid \forall (S, S') \in [\![\Phi]\!]. \ (fS, f'S') \in [\![\Phi]\!]_\perp\}
\end{aligned}
$$

then the intended meanings of the judgements are:

$$
\begin{aligned}
\models F_\tau \sim F'_\tau : \Phi \Rightarrow \phi_\tau &\equiv ([\![F_\tau]\!], [\![F'_\tau]\!]) \in [\![\Phi \Rightarrow \phi_\tau]\!] \\
\models C \sim C' : \Phi \Rightarrow \Phi' &\equiv ([\![C]\!], [\![C']\!]) \in [\![\Phi \Rightarrow \Phi']\!]
\end{aligned}
$$

LEMMA 2. $[\![\Phi \Rightarrow \phi_\tau]\!]$ and $[\![\Phi \Rightarrow \Phi']\!]$ *are admissible partial equivalence relations.*

Some basic rules for deriving DDCC judgements are shown in Figure 2. The rules for expressions refer to abstract versions $\widehat{op}$ of each primitive binary operator op in the language. A typical definition is that for multiplication:

$\widehat{\times}$	$\mathbb{F}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\{n\}_{\texttt{int}}$	$\Delta_{\texttt{int}}$	$\mathbb{T}_{\texttt{int}}$
$\mathbb{F}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$
$\{0\}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$
$\{m\}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\{m \times n\}_{\texttt{int}}$	$\Delta_{\texttt{int}}$	$\mathbb{T}_{\texttt{int}}$
$\Delta_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\Delta_{\texttt{int}}$	$\Delta_{\texttt{int}}$	$\mathbb{T}_{\texttt{int}}$
$\mathbb{T}_{\texttt{int}}$	$\mathbb{F}_{\texttt{int}}$	$\{0\}_{\texttt{int}}$	$\mathbb{T}_{\texttt{int}}$	$\mathbb{T}_{\texttt{int}}$	$\mathbb{T}_{\texttt{int}}$

The general correctness condition for abstract operations is familiar from abstract interpretation:

Definition 1. We say $\widehat{op}$ soundly abstracts the operation op if

$$\forall (x, x') \in [\![\phi_\tau]\!], (y, y') \in [\![\phi'_\tau]\!]. \ (x \ op \ y, x' \ op \ y') \in [\![\phi_\tau \ \widehat{op} \ \phi'_\tau]\!].$$

The most interesting of the rules in Figure 2 are those for conditionals and while-loops. Observe that for two conditionals to be related, not only do their true and false branches have to be pairwise related, but they also have to agree on which branch is taken; this is expressed by the use of $\Delta_{\texttt{bool}}$ in the premises of the rule. Similar considerations apply to the rule for while-loops, which ensures that related loops execute in lockstep.

3.2 Equations

Using only the rules in Figure 2, most of the interesting judgements one can prove relate a phrase to itself at some type. In other words, they constitute an analysis system but not yet a program transformation system. However, the advantage of our formulation is that program transformations can now be specified and justified simply by adding new inference rules whose soundness may be straightforwardly and independently checked in terms of the semantics.

3.2.1 Basic equations

Our first set of transformation rules express universally applicable structural equivalences for while-programs, without requiring any of the extra information gathered by the analysis.

Sequential unit laws:

$$
\frac{\vdash C : \Phi \Rightarrow \Phi'}{\vdash (\texttt{skip};C) \sim C : \Phi \Rightarrow \Phi'} \ \text{[D-SU1]}
$$

$$
\frac{\vdash C : \Phi \Rightarrow \Phi'}{\vdash (C;\texttt{skip}) \sim C : \Phi \Rightarrow \Phi'} \ \text{[D-SU2]}
$$

Associativity:

$$
\frac{\vdash (C_1;C_2);C_3 : \Phi \Rightarrow \Phi'}{\vdash ((C_1;C_2);C_3) \sim (C_1;(C_2;C_3)) : \Phi \Rightarrow \Phi'}
$$

In practice, one usually identifies programs up to associativity of sequential composition, rather than making explicit use of the rule above.

Commuting conversion for conditional:

$$
\frac{\vdash \texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2 : \Phi \Rightarrow \Phi' \quad \vdash C_3 : \Phi' \Rightarrow \Phi''}{\begin{array}{l} \vdash \ \ (\texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2);C_3 \\ \sim \ \ \texttt{if } B \texttt{ then } (C_1;C_3) \texttt{ else } (C_2;C_3) : \Phi \Rightarrow \Phi'' \end{array}} \ \text{[D-CC]}
$$

Equivalent branches for conditional:

$$
\frac{\vdash C_1 \sim C_2 : \Phi \Rightarrow \Phi'}{\vdash \texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2 \sim C_1 : \Phi \Rightarrow \Phi'} \ \text{[D-BrE]}
$$

Loop unrolling:

$$
\frac{\vdash \texttt{while } B \texttt{ do } C : \Phi \Rightarrow \Phi'}{\begin{array}{l} \vdash \ \ \texttt{while } B \texttt{ do } C \\ \sim \ \ \texttt{if } B \texttt{ then } C;(\texttt{while } B \texttt{ do } C) \texttt{ else skip} : \Phi \Rightarrow \Phi' \end{array}} \ \text{[D-LU1]}
$$

$$
\frac{\vdash \texttt{while } B \texttt{ do } C : \Phi \Rightarrow \Phi'}{\begin{array}{l} \vdash \ \ \texttt{while } B \texttt{ do } C \\ \sim \ \ \texttt{while } B \texttt{ do } (C;\texttt{if } B \texttt{ then } C \texttt{ else skip}) : \Phi \Rightarrow \Phi' \end{array}} \ \text{[D-LU2]}
$$

Self-assignment elimination:

$$\vdash X \texttt{:=} X \sim \texttt{skip} : \Phi, X : \phi_{\texttt{int}} \Rightarrow \Phi, X : \phi_{\texttt{int}} \ \text{[D-SAs]}$$

In conjunction with the core rules, the rules above can be used to derive many of the basic equalities one might expect.[2] From a pragmatic point of view, however, they are somewhat unwieldy: even very simple proofs get quite large, with many applications of the symmetry and transitivity rules and many repeated sub-derivations. Reformulating the rules as logically equivalent versions which can be applied in more general contexts helps immensely. For example, a better formulation of one of the skip rules is the following:

$$
\frac{\vdash C \sim \texttt{skip} : \Phi \Rightarrow \Phi' \quad \vdash C' \sim C'' : \Phi' \Rightarrow \Phi''}{\vdash (C;C') \sim C'' : \Phi \Rightarrow \Phi''} \ \text{[D-SU1']}
$$

Presenting rules in this style is essentially trying to produce a system with a kind of cut-elimination property, but we leave serious consideration of proof-theoretic matters to future work.

3.2.2 Optimizing Transformations

In this section we consider some more interesting rules, in which equations are predicated on information in the type system.

---

[2]Though the rules presented are in no sense complete. There are sound rules (arithmetic identities and equivalences for nested conditionals, for example) which are not consequences of the ones we have given.

Subtyping and Structural

$$\vdash C \sim C' : \Phi \Rightarrow - \text{ [D-CT1]} \qquad \vdash C \sim C' : \Phi, X : \mathbb{F}_{\text{int}} \Rightarrow \Phi' \text{ [D-CT2]} \qquad \vdash F_\tau \sim F'_\tau : \Phi \Rightarrow \mathbb{T}_\tau \text{ [D-ET1]}$$

$$\vdash F_\tau \sim F'_\tau : \Phi, X : \mathbb{F}_{\text{int}} \Rightarrow \phi_\tau \text{ [D-ET2]} \qquad \frac{\vdash F_\tau \sim F'_\tau : \Phi \Rightarrow \phi_\tau}{\vdash F'_\tau \sim F_\tau : \Phi \Rightarrow \phi_\tau} \text{ [D-ESym]} \qquad \frac{\vdash F_\tau \sim F'_\tau : \Phi \Rightarrow \phi_\tau \quad \Phi' \leq \Phi \quad \phi_\tau \leq \phi'_\tau}{\vdash F_\tau \sim F'_\tau : \Phi' \Rightarrow \phi'_\tau} \text{ [D-ESub]}$$

$$\frac{\vdash C \sim C' : \Phi_1 \Rightarrow \Phi_2 \quad \Phi'_1 \leq \Phi_1 \quad \Phi_2 \leq \Phi'_2}{\vdash C \sim C' : \Phi'_1 \Rightarrow \Phi'_2} \text{ [D-CSub]} \qquad \frac{\vdash F_\tau \sim F'_\tau : \Phi \Rightarrow \phi_\tau \quad \vdash F'_\tau \sim F''_\tau : \Phi \Rightarrow \phi_\tau}{\vdash F_\tau \sim F''_\tau : \Phi \Rightarrow \phi_\tau} \text{ [D-ETr]}$$

$$\frac{\vdash C \sim C' : \Phi \Rightarrow \Phi'}{\vdash C' \sim C : \Phi \Rightarrow \Phi'} \text{ [D-CSym]} \qquad \frac{\vdash C \sim C' : \Phi \Rightarrow \Phi' \quad \vdash C' \sim C'' : \Phi \Rightarrow \Phi'}{\vdash C \sim C'' : \Phi \Rightarrow \Phi'} \text{ [D-CTr]}$$

Expressions

$$\vdash X \sim X : \Phi, X : \phi_{\text{int}} \Rightarrow \phi_{\text{int}} \text{ [D-V]} \qquad \vdash \texttt{n} \sim \texttt{n} : \Phi \Rightarrow \{n\}_{\text{int}} \text{ [D-N]}$$

$$\vdash \texttt{b} \sim \texttt{b} : \Phi \Rightarrow \{b\}_{\text{bool}} \text{ [D-B]} \qquad \frac{\vdash F_\tau \sim G_\tau : \Phi \Rightarrow \phi_\tau \quad \vdash F'_\tau \sim G'_\tau : \Phi \Rightarrow \phi'_\tau}{\vdash F_\tau \texttt{ op } F'_\tau \sim G_\tau \texttt{ op } G'_\tau : \Phi \Rightarrow (\phi_\tau \, \widehat{op} \, \phi'_\tau)} \text{ [D-}op\text{]}$$

Commands

$$\vdash \texttt{skip} \sim \texttt{skip} : \Phi \Rightarrow \Phi \text{ [D-Skip]} \qquad \frac{\vdash C_1 \sim C'_1 : \Phi \Rightarrow \Phi' \quad \vdash C_2 \sim C'_2 : \Phi' \Rightarrow \Phi''}{\vdash (C_1 \,;\, C_2) \sim (C'_1 \,;\, C'_2) : \Phi \Rightarrow \Phi''} \text{ [D-Seq]}$$

$$\frac{\vdash E \sim E' : \Phi, X : \phi_{\text{int}} \Rightarrow \phi'_{\text{int}}}{\vdash X \texttt{:=} E \sim X \texttt{:=} E' : \Phi, X : \phi_{\text{int}} \Rightarrow \Phi, X : \phi'_{\text{int}}} \text{ [D-Ass]} \qquad \frac{\vdash B \sim B' : \Phi \Rightarrow \Delta_{\text{bool}} \quad \vdash C \sim C' : \Phi \Rightarrow \Phi}{\vdash (\texttt{while } B \texttt{ do } C) \sim (\texttt{while } B' \texttt{ do } C') : \Phi \Rightarrow \Phi} \text{ [D-Whl]}$$

$$\frac{\vdash B \sim B' : \Phi \Rightarrow \Delta_{\text{bool}} \quad \vdash C_1 \sim C'_1 : \Phi \Rightarrow \Phi' \quad \vdash C_2 \sim C'_2 : \Phi \Rightarrow \Phi'}{\vdash (\texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2) \sim (\texttt{if } B' \texttt{ then } C'_1 \texttt{ else } C'_2) : \Phi \Rightarrow \Phi'} \text{ [D-If]}$$

Figure 2. Core DDCC System

Dead assignment elimination:

$$\vdash (X \texttt{:=} E) \sim \texttt{skip} : \Phi, X : \phi_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}} \text{ [D-DAs]}$$

Intuitively, the dead assignment rule says that an assignment to a variable is equivalent to skip *if* we are in a context in which the value of that variable does not matter.

Constant folding:

$$\frac{\vdash F_\tau : \Phi \Rightarrow \{c\}_\tau}{\vdash F_\tau \sim \texttt{c} : \Phi \Rightarrow \{c\}_\tau} \text{ [D-CF]}$$

Known branch:

$$\frac{\vdash B : \Phi \Rightarrow \{true\} \quad \vdash C_1 \sim C' : \Phi \Rightarrow \Phi'}{\vdash (\texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2) \sim C' : \Phi \Rightarrow \Phi''} \text{ [D-KBT]}$$

$$\frac{\vdash B : \Phi \Rightarrow \{false\} \quad \vdash C_2 \sim C' : \Phi \Rightarrow \Phi'}{\vdash (\texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2) \sim C' : \Phi \Rightarrow \Phi''} \text{ [D-KBF]}$$

Dead while:

$$\frac{\vdash B : \Phi \Rightarrow \{false\}}{\vdash (\texttt{while } B \texttt{ do } C) \sim \texttt{skip} : \Phi \Rightarrow \Phi} \text{ [D-DWh]}$$

Divergence:

$$\frac{\vdash B : \Phi \Rightarrow \{true\} \quad \vdash C : \Phi \Rightarrow \Phi}{\vdash (\texttt{while } B \texttt{ do } C) : \Phi \Rightarrow \Phi'} \text{ [D-Div]}$$

The type Φ' in the conclusion of the rule above is arbitrary because the loop will diverge when executed in any state in the domain of Φ.

The following is an easy induction, relying on Lemmas 1 and 2:

THEOREM 1. *Assuming the abstract operations satisfy the correctness condition given in Definition 1, the core DDCC rules of Figure 2 and the additional rules of Section 3.2 are all sound:*

$$\vdash F_\tau \sim F'_\tau : \Phi \Rightarrow \phi_\tau \quad \Longrightarrow \quad \models F_\tau \sim F'_\tau : \Phi \Rightarrow \phi_\tau$$
$$\vdash C \sim C' : \Phi \Rightarrow \Phi' \quad \Longrightarrow \quad \models C \sim C' : \Phi \Rightarrow \Phi'$$

3.3 Example Transformations

These rules are sufficient to capture some non-trivial transformations, including constant propagation, dead-code elimination and program slicing [33]. Some example derivations are shown in Figure 3. We leave it as an exercise to prove larger examples, such as the slicing transformation:

18

Constants, known branches and dead code:

D_1 :

$$\dfrac{\dfrac{}{\vdash X : \Phi, X : \{3\} \Rightarrow \{3\}}\text{[D-V]} \quad \dfrac{}{\vdash 3 : \Phi, X : \{3\} \Rightarrow \{3\}}\text{[D-N]}}{\vdash X = 3 : \Phi, X : \{3\} \Rightarrow \{true\}}\text{[D-=]}$$

$$\dfrac{\dfrac{}{\vdash 7 : \Phi, X : \{3\} \Rightarrow \{7\}}\text{[D-N]}}{\vdash X \texttt{:=7} : \Phi, X : \{3\} \Rightarrow \Phi, X : \{7\}}\text{[D-Ass]}$$

$$\dfrac{}{\vdash (\texttt{if } X = 3 \texttt{ then } X\texttt{:=7 else skip}) \sim (X\texttt{:=7}) : \Phi, X : \{3\} \Rightarrow \Phi, X : \{7\}}\text{[D-KBT]}$$

D_2 :

$$\dfrac{\dfrac{}{\vdash X : \Phi, X : \{7\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \{7\}}\text{[D-V]} \quad \dfrac{}{\vdash 1 : \Phi, X : \{7\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \{1\}}\text{[D-N]}}{\vdash X + 1 : \Phi, X : \{7\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \{8\}}\text{[D-+]}$$

$$\dfrac{\vdash X + 1 \sim 8 : \Phi, X : \{7\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \{8\}}{}\text{[D-CF]}$$

$$\dfrac{\dfrac{\vdash {\begin{smallmatrix}Z\texttt{:=}X+1\\\sim Z\texttt{:=8}\end{smallmatrix}} : \Phi, X : \{7\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \{7\}, Z : \{8\}}{}\text{[D-Ass]} \quad \Phi, X : \{7\}, Z : \{8\} \leq \Phi, X : \mathbb{T}_{\text{int}}, Z : \{8\}}{\vdash {\begin{smallmatrix}Z\texttt{:=}X+1\\\sim Z\texttt{:=8}\end{smallmatrix}} : \Phi, X : \{7\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}}, Z : \{8\}}\text{[D-CSub]}$$

D_3 :

$$\dfrac{\dfrac{}{\vdash (X\texttt{:=7}) \sim \texttt{skip} : \Phi, X : \{3\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}}, Z : \mathbb{T}_{\text{int}}}\text{[D-DAss]} \quad \dfrac{\dfrac{}{\vdash (8) : \Phi, X : \mathbb{T}_{\text{int}}, Z : \mathbb{T}_{\text{int}} \Rightarrow \{8\}}}{\vdash Z\texttt{:=8} : \Phi, X : \mathbb{T}_{\text{int}}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}}, Z : \{8\}}\text{[D-SU1']}}{\vdash (X\texttt{:=7};Z\texttt{:=8}) \sim Z\texttt{:=8} : \Phi, X : \{3\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}}, Z : \{8\}}$$

$$\dfrac{\dfrac{D_1 \quad D_2}{\vdash {\begin{smallmatrix}(\texttt{if } X = 3 \texttt{ then } X\texttt{:=7 else skip};Z\texttt{:=}X+1)\\\sim (X\texttt{:=7};Z\texttt{:=8})\end{smallmatrix}} : \Phi, X : \{3\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}}, Z : \{8\}}\text{[D-Seq]} \quad D_3}{\vdash {\begin{smallmatrix}(\texttt{if } X = 3 \texttt{ then } X\texttt{:=7 else skip};Z\texttt{:=}X+1)\\\sim (Z\texttt{:=8})\end{smallmatrix}} : \Phi, X : \{3\}, Z : \mathbb{T}_{\text{int}} \Rightarrow \Phi, X : \mathbb{T}_{\text{int}}, Z : \{8\}}\text{[D-CTr]}$$

Figure 3. Examples of DDCC Transformations

```
I := 1;                  I := 1;
S := 0;
P := 1;                  P := 1
while I<N do (     ==>    while I<N do (
 S := S+I;
 P := P*I;                P := P*I;
 I := I+1;)               I := I+1;)
```

at type $N : \Delta_{\text{int}} \Rightarrow P : \Delta_{\text{int}}$. Here we expressed the fact that we were only interested in the final value of P simply by transforming it at a result type which only constrains the value of that variable to be preserved – all the others (in particular S) are typed at $\mathbb{T}_{\text{int}}$ and so are allowed to take any value.

3.4 Secure Information Flow

It is worth observing that the $\mathbb{T}, \Delta$ fragment of our calculus can be seen as a non-interference type system. Figure 4 presents a version of a type system for secure information flow due to Smith and Volpano [27]. In this system, a *security level*, σ, is either low (*L*) or high (*H*). A context γ is then a finite map from variables to security levels:

$$\gamma \quad := \quad - \mid \gamma, X : \sigma_{\text{int}}$$

Given such a context, the type system assigns a security level (σ_{int} or σ_{bool}) to each expression and (σ_{com}) to each command. The

property which the type system ensures is that any typeable command does not allow information to flow (either directly, via assignment, or indirectly, via control flow) from high security variables to low security ones. We define a translation $(\cdot)^*$ from the Smith/Volpano system into DDCC as follows:

Expression types: $L_\tau^* = \Delta_\tau$ and $H_\tau^* = \mathbb{T}_\tau$.

Contexts: $-^* = -$ and $(\gamma, X : \sigma_{\text{int}})^* = \gamma^*, X : \sigma_{\text{int}}^*$.

Judgements:

$$\begin{aligned}
(\gamma \vdash F : \sigma_\tau)^* &= \vdash F \sim F : \gamma^* \Rightarrow \sigma_\tau^* \\
(\gamma \vdash C : L_{\text{com}})^* &= \vdash C \sim C : \gamma^* \Rightarrow \gamma^* \\
(\gamma \vdash C : H_{\text{com}})^* &= \vdash C \sim \texttt{skip} : \gamma^* \Rightarrow \gamma^*
\end{aligned}$$

THEOREM 2. *For any judgement J derivable in the Smith/Volpano system, J^* is derivable in DDCC*

PROOF. This is a simple induction, relying on the dead assignment axiom in the case of high assignment statements, sequential unit for high sequential compositions and the equivalent branch rule for high conditionals. □

Definition 2. In the context of a security type assignment γ, a command C satisfies *strong sequential noninterference* if $\models C \sim C : \gamma^* \Rightarrow \gamma^*$.

$$\gamma, X : \sigma_{\text{int}} \vdash X : \sigma_{\text{int}} \qquad \gamma \vdash \mathtt{n} : \sigma_{\text{int}} \qquad \gamma \vdash \mathtt{b} : \sigma_{\text{bool}}$$

$$\frac{\gamma \vdash E : \sigma_{\text{int}} \quad \gamma \vdash E' : \sigma_{\text{int}}}{\gamma \vdash E \,\mathtt{iop}\, E' : \sigma_{\text{int}}} + \text{similar for } bop \text{ and } lop \qquad \frac{\gamma, X : \sigma_{\text{int}} \vdash E : \sigma_{\text{int}}}{\gamma, X : \sigma_{\text{int}} \vdash X\mathtt{:=}E : \sigma_{\text{com}}}$$

$$\frac{\gamma \vdash C : \sigma_{\text{com}} \quad \gamma \vdash C' : \sigma_{\text{com}}}{\gamma \vdash C\mathtt{;}C' : \sigma_{\text{com}}} \qquad \frac{\gamma \vdash B : \sigma_{\text{bool}} \quad \gamma \vdash C : \sigma_{\text{com}} \quad \gamma \vdash C' : \sigma_{\text{com}}}{\gamma \vdash \mathtt{if}\, B \,\mathtt{then}\, C \,\mathtt{else}\, C' : \sigma_{\text{com}}}$$

$$\frac{\gamma \vdash B : L_{\text{bool}} \quad \gamma \vdash C : L_{\text{com}}}{\gamma \vdash \mathtt{while}\, B \,\mathtt{do}\, C : L_{\text{com}}} \qquad \frac{\gamma \vdash F : L_{\tau}}{\gamma \vdash F : H_{\tau}} \qquad \frac{\gamma \vdash C : H_{\text{com}}}{\gamma \vdash C : L_{\text{com}}}$$

Figure 4. Smith/Volpano Type System

This version of non-interference is the semantic security property intended by Smith and Volpano, though the actual property established by the soundness proof in [27] is more syntactic and intensional, as it is defined in terms of their particular typing rules. Our notion of intereference is *strong* because it is termination-sensitive: varying the high-security inputs affects neither the low-security outputs *nor* the termination behaviour. In the absence of any termination analysis, this is enforced by the rather brutal approach of making all high-security commands total. The weaker notion of non-interference that is achieved by the earlier system of Volpano, Smith and Irvine [29] does not seem to translate directly into DDCC.

Even without constant tracking, DDCC is marginally more powerful than the Smith/Volpano system. For example, if H is a high-security variable, and L is low-security then the following are easily shown to satisfy non-interference in DDCC, but would be rejected by the Smith/Volpano system:

1. if H > 3 then H := L ; L := 1 else L := 1
2. L := H ; L := 3

4 Relational Hoare Logic

There are many common optimizing transformations which are not captured by DDCC. In particular:

- It does not capture any transformations that take advantage of the fact that one knows statically which way a boolean test must have evaluated if one is within a particular branch of a conditional, or either in the body of or have just left a while-loop. For example, the judgement

$$\vdash (\mathtt{if}\, X = 3 \,\mathtt{then}\, Y\mathtt{:=}X \,\mathtt{else}\, Y\mathtt{:=}3)$$
$$\sim (Y\mathtt{:=}3) : X : \Delta \Rightarrow Y : \{3\}$$

is semantically valid but not derivable.

- It cannot express the preservation of the values of expressions, except where they are statically known to be a particular constant. These means even trivial code-motion transformations cannot be derived.

We can address these shortcomings by making piecemeal additions to the system, such as quantification over variables ranging over integers or PERs. However, there is a simple and elegant system, which we call Relational Hoare Logic (RHL), into which many of these extensions or alternative type systems can be embedded.

Unlike DDCC, RHL does not look like a conventional type-based analysis system – it has a rather general syntax for relations and is parameterized on some system for deciding the entailment relation between them. The intention is that more specific analyses and transformations can formulated as subsystems of RHL by restricting the syntax of assertions and providing particular approximations to the entailment relation. Another way in which RHL goes beyond DDCC is that it is not restricted to partial equivalence relations, which deserves some comment.

PERs are certainly privileged: they are the basis of equational reasoning, and we will nearly always be trying to prove that one program phrase is related to another by a PER so that we can perform a rewrite in some context. However, in order to establish that a two phrases are related by a PER, we often have to do some local reasoning using more general relations. This is familiar in the semantics of polymorphic type theories: types are interpreted by PERs, and polymorphism by quantification over PERs, but parametricity theorems and equivalence results for implementations of abstract dataypes arise from substituting more general relations. To give some intuition for why this might be so, consider proving the equivalence

```
X := -Y;              X := Y;
Z := Z-X;      ==>    Z := Z+X;
X := -X;
```

at, say, $Y : \Delta_{\text{int}}, Z : \Delta_{\text{int}} \Rightarrow X : \Delta_{\text{int}}, Y : \Delta_{\text{int}}, Z : \Delta_{\text{int}}$. If we try to to establish that the two commands are related by this PER by relating their intermediate states (though this is not the only approach one could take), we will need to use the relation that the value of X in one state is the negation of that in the other, which is not a PER.

RHL is an extremely simple variation on traditional Floyd-Hoare logic [13]. Instead of assertions which denote predicates on states and judgements which say that terminating execution of a command in a state satisfying a precondition will yield a state satisfying a postcondition, we directly axiomatise when a *pair* of commands map a given pre-*relation* into a given post-*relation*. Binary relations on states are simply specified by boolean expessions of the language over variables tagged with an indication of which of the two states they refer to. At first sight, this may seem frighteningly simple-minded, but it actually works rather nicely. In this presentation we do not consider quantification over metavariables ("ghost variables") denoting integers: their addition is straighforward, but simple global analyses seem to be expressible without them.

20

$$\vdash \mathtt{skip} \sim \mathtt{skip} : \Phi \Rightarrow \Phi \text{ [R-Skip]}$$

$$\frac{\vdash C \sim C' : \Phi \wedge (B\langle 1\rangle \wedge B'\langle 2\rangle) \Rightarrow \Phi' \quad \vdash D \sim D' : \Phi \wedge \mathtt{not}(B\langle 1\rangle \vee B'\langle 2\rangle) \Rightarrow \Phi'}{\vdash \mathtt{if}\ B\ \mathtt{then}\ C\ \mathtt{else}\ D \sim \mathtt{if}\ B'\ \mathtt{then}\ C'\ \mathtt{else}\ D' : \Phi \wedge (B\langle 1\rangle = B'\langle 2\rangle) \Rightarrow \Phi'} \text{ [R-If]}$$

$$\frac{\vdash C \sim C' : \Phi \Rightarrow \Phi' \quad \vdash D \sim D' : \Phi' \Rightarrow \Phi''}{\vdash C \mathbin{;} D \sim C' \mathbin{;} D' : \Phi \Rightarrow \Phi''} \text{ [R-Seq]} \qquad \vdash X := E \sim Y := E' : \Phi[E\langle 1\rangle/X\langle 1\rangle, E'\langle 2\rangle/Y\langle 2\rangle] \Rightarrow \Phi \text{ [R-Ass]}$$

$$\frac{\vdash C \sim C' : \Phi \wedge (B\langle 1\rangle \wedge B'\langle 2\rangle) \Rightarrow \Phi \wedge (B\langle 1\rangle = B'\langle 2\rangle)}{\vdash \mathtt{while}\ B\ \mathtt{do}\ C \sim \mathtt{while}\ B'\ \mathtt{do}\ C' : \Phi \wedge (B\langle 1\rangle = B'\langle 2\rangle) \Rightarrow \Phi \wedge \mathtt{not}(B\langle 1\rangle \vee B'\langle 2\rangle)} \text{ [R-Whl]}$$

$$\frac{\vdash C \sim C' : \Phi_1 \Rightarrow \Phi_2 \quad \models \Phi'_1 \leq \Phi_1 \quad \models \Phi_2 \leq \Phi'_2}{\vdash C \sim C' : \Phi'_1 \Rightarrow \Phi'_2} \text{ [R-Sub]}$$

$$\frac{\vdash C \sim C' : \Phi \Rightarrow \Phi' \quad \models PER(\Phi \Rightarrow \Phi')}{\vdash C' \sim C : \Phi \Rightarrow \Phi'} \text{ [R-Sym]} \qquad \frac{\vdash C \sim C' : \Phi \Rightarrow \Phi' \quad \vdash C' \sim C'' : \Phi \Rightarrow \Phi' \quad \models PER(\Phi \Rightarrow \Phi')}{\vdash C \sim C'' : \Phi \Rightarrow \Phi'} \text{ [R-Tr]}$$

Figure 5. Core Relational Hoare Logic

4.1 RHL Syntax and Semantics

4.1.1 Syntax

We define generalized expressions and relational assertions as follows:

$$\begin{aligned}
\mathtt{gexp} \ni GE &:= \quad \mathtt{n} \mid X\langle 1\rangle \mid X\langle 2\rangle \mid GE\ \mathtt{iop}\ GE \\
\mathtt{relexp} \ni \Phi &:= \quad \mathtt{b} \mid GE\ \mathtt{bop}\ GE \mid \mathtt{not}\,\Phi \mid \Phi\ \mathtt{lop}\ \Phi
\end{aligned}$$

We overload the notation $(\cdot)\langle 1\rangle$ and $(\cdot)\langle 2\rangle$ to stand for homomorphic embeddings $\mathtt{int\ exp} \to \mathtt{gexp}$ and $\mathtt{bool\ exp} \to \mathtt{relexp}$ in the obvious way. The basic judgement form is $\vdash C \sim C' : \Phi \Rightarrow \Phi'$ (though the use of $\sim$ for arbitrary relations is arguably bad).

4.1.2 Semantics

The semantics of generalized expressions as integer-valued functions of two states, and of relational assertions as relations on states is unsurprising:

$$\begin{aligned}
\llbracket GE \rrbracket &\in \quad \mathbb{S} \times \mathbb{S} \to \mathbb{Z} \\
\llbracket \mathtt{n} \rrbracket (S_1, S_2) &= \quad n \\
\llbracket X\langle 1\rangle \rrbracket (S_1, S_2) &= \quad S_1(X) \\
\llbracket X\langle 2\rangle \rrbracket (S_1, S_2) &= \quad S_2(X) \\
\llbracket E\ \mathtt{iop}\ F \rrbracket (S_1, S_2) &= \quad (\llbracket E \rrbracket (S_1, S_2))\ iop\ (\llbracket F \rrbracket (S_1, S_2))
\end{aligned}$$

$$\begin{aligned}
\llbracket \Phi \rrbracket &\subseteq \quad \mathbb{S} \times \mathbb{S} \\
\llbracket \mathtt{true} \rrbracket &= \quad \mathbb{S} \times \mathbb{S} \\
\llbracket \mathtt{false} \rrbracket &= \quad \emptyset \\
\llbracket E\ \mathtt{bop}\ F \rrbracket &= \quad \{(S_1, S_2) \mid (\llbracket E \rrbracket (S_1, S_2))\ bop \\
&\qquad\qquad (\llbracket F \rrbracket (S_1, S_2)) = true\} \\
\llbracket \Phi\ \mathtt{lop}\ \Phi' \rrbracket &= \quad \{(S_1, S_2) \mid (\llbracket \Phi \rrbracket (S_1, S_2))\ lop \\
&\qquad\qquad (\llbracket \Phi' \rrbracket (S_1, S_2)) = true\} \\
\llbracket \mathtt{not}\,\Phi \rrbracket &= \quad (\mathbb{S} \times \mathbb{S}) \setminus \llbracket \Phi \rrbracket
\end{aligned}$$

The intended meaning of judgements is given by

$$\begin{aligned}
&\models C \sim C' : \Phi \Rightarrow \Phi' \\
\equiv\ & \forall (S_1, S_2) \in \llbracket \Phi \rrbracket.\ (\llbracket C \rrbracket (S_1), \llbracket C' \rrbracket (S_2)) \in \llbracket \Phi' \rrbracket_\perp
\end{aligned}$$

We will also need some auxiliary semantic judgements, whose meanings are as follows:

$$\begin{aligned}
\models \Phi \leq \Phi' &\equiv \quad \llbracket \Phi \rrbracket \subseteq \llbracket \Phi' \rrbracket \\
\models PER(\Phi) &\equiv \quad (\llbracket \Phi \rrbracket \circ \llbracket \Phi \rrbracket \subseteq \llbracket \Phi \rrbracket)\ \text{and}\ (\llbracket \Phi \rrbracket^{-1} \subseteq \llbracket \Phi \rrbracket)
\end{aligned}$$

LEMMA 3.

1. *For all GE, GE', X, S, S':*

$$\begin{aligned}
&\llbracket GE[GE'/X\langle 1\rangle] \rrbracket (S, S') \\
=\ & \llbracket GE \rrbracket (S[X \mapsto \llbracket GE' \rrbracket (S, S')], S')
\end{aligned}$$

And similarly for $X\langle 2\rangle$ and S'.

2. *For all Φ, GE, X:*

$$\begin{aligned}
&\llbracket \Phi[GE/X\langle 1\rangle] \rrbracket \\
=\ & \{(S, S') \mid (S[X \mapsto \llbracket GE \rrbracket (S, S')], S') \in \llbracket \Phi \rrbracket\}
\end{aligned}$$

And similarly for $X\langle 2\rangle$ and S'.

3. *For all Φ and Φ', $\llbracket \Phi \rrbracket$ and $\llbracket \Phi \Rightarrow \Phi' \rrbracket$ are admissible relations.*

4.1.3 Inference Rules

The core rules for RHL are shown in Figure 5. Observe that, as was the case in DDCC, the basic rules ensure that the same conditional branches are taken and that loops are executed the same number of times on the two sides. Note also that one could add distinct semantic judgements for symmetry and transitivity, rather than requiring both. The assignment rule is surprisingly liberal, but there

is no reason to require the assigned variables to be the same in both commands.

4.2 Equations

As with DDCC, we will specify optimizing transformations by adding extra (sound) rules to the core. But even before we do that, RHL can justify some useful transformations. Here's an example of removing a redundant evaluation:

1. $\vdash \begin{array}{c} Z\texttt{:=}Y\texttt{+}1 \\ \sim Z\texttt{:=}X \end{array} : \begin{array}{c} X\langle 1 \rangle = X\langle 2 \rangle \wedge \\ Y\langle 1 \rangle + 1 = X\langle 2 \rangle \end{array} \Rightarrow \begin{array}{c} X\langle 1 \rangle = X\langle 2 \rangle \wedge \\ Z\langle 1 \rangle = Z\langle 2 \rangle \end{array}$ by [R-Ass]

2. $\vdash \begin{array}{c} X\texttt{:=}Y\texttt{+}1 \\ \sim X\texttt{:=}Y\texttt{+}1 \end{array} : \begin{array}{c} Y\langle 1 \rangle + 1 = Y\langle 2 \rangle + 1 \wedge \\ Y\langle 1 \rangle + 1 = Y\langle 2 \rangle + 1 \end{array} \Rightarrow \begin{array}{c} X\langle 1 \rangle = X\langle 2 \rangle \wedge \\ Y\langle 1 \rangle + 1 = X\langle 2 \rangle \end{array}$ by [R-Ass]

3. $\models (Y\langle 1 \rangle = Y\langle 2 \rangle) \leq \begin{array}{c} Y\langle 1 \rangle + 1 = Y\langle 2 \rangle + 1 \wedge \\ Y\langle 1 \rangle + 1 = Y\langle 2 \rangle + 1 \end{array}$ by logic

4. $\vdash \begin{array}{c} X\texttt{:=}Y\texttt{+}1 \\ \sim X\texttt{:=}Y\texttt{+}1 \end{array} : Y\langle 1 \rangle = Y\langle 2 \rangle \Rightarrow \begin{array}{c} X\langle 1 \rangle = X\langle 2 \rangle \wedge \\ Y\langle 1 \rangle + 1 = X\langle 2 \rangle \end{array}$ by [R-Sub] applied to 2. and 3.

5. $\vdash \begin{array}{c} X\texttt{:=}Y\texttt{+}1\texttt{;}Z\texttt{:=}Y\texttt{+}1 \\ \sim X\texttt{:=}Y\texttt{+}1\texttt{;}Z\texttt{:=}X \end{array} : Y\langle 1 \rangle = Y\langle 2 \rangle \Rightarrow \begin{array}{c} X\langle 1 \rangle = X\langle 2 \rangle \wedge \\ Z\langle 1 \rangle = Z\langle 2 \rangle \end{array}$ by [R-Seq] applied to 4. and 1.

4.2.1 Basic Equations

The basic equations we presented in the context of DDCC are still valid for RHL, with the exception of self-assignment elimination.

4.2.2 Optimizing Transformations

Dead assignment:

$$\vdash X\texttt{:=}E \sim \texttt{skip} : \Phi[E\langle 1 \rangle / X\langle 1 \rangle] \Rightarrow \Phi \; \text{[R-DAssl]}$$

$$\vdash \texttt{skip} \sim X\texttt{:=}E : \Phi[E\langle 2 \rangle / X\langle 2 \rangle] \Rightarrow \Phi \; \text{[R-DAssr]}$$

These rules subsume our previous dead-assignment and self-assignment rules.

Known branch:

$$\frac{\vdash C \sim D : \Phi \Rightarrow \Phi'}{\vdash \texttt{if } B \texttt{ then } C \texttt{ else } C' \sim D : \Phi \wedge B\langle 1 \rangle \Rightarrow \Phi'} \; \text{[R-KBTl]}$$

$$\frac{\vdash C' \sim D' : \Phi \Rightarrow \Phi'}{\vdash \texttt{if } B \texttt{ then } C \texttt{ else } C' \sim D' : \Phi \wedge \neg B\langle 1 \rangle \Rightarrow \Phi'} \; \text{[R-KBFl]}$$

Plus two similar rules with the conditional on the right. (The fact that our relations are no longer assumed symmetric leads to a slightly unfortunate proliferation of variants.)

Dead while:

$$\vdash \texttt{while } B \texttt{ do } C \sim \texttt{skip} : \Phi \wedge \texttt{not} B\langle 1 \rangle \Rightarrow \Phi \wedge \texttt{not} B\langle 1 \rangle \; \text{[R-DWhll]}$$

Plus the variant with `skip` on the left.

Soundness is a simple induction, relying on Lemma 3:

THEOREM 3. *For all C, C', Φ, Φ', if $\vdash C \sim C' : \Phi \Rightarrow \Phi'$ is derivable using the rules in Figure 5 and Section 4.2 then $\models C \sim C' : \Phi \Rightarrow \Phi'$.*

4.3 Examples

With these rules, one can prove the correctness of many traditional compiler optimizations, including various forms of code motion and predicated transformation. Producing proofs in RHL is fairly straightforward, so we just give a couple of small examples of the sort of thing one can prove.

Invariant hoisting:

```
while I<N do          X := Y+1;
  X := Y+1;     ==>   while I<N do
  I := I+X;              I := I+X;
```

at type $\Phi \Rightarrow \Phi$ where Φ is $I\langle 1 \rangle = I\langle 2 \rangle \wedge N\langle 1 \rangle = N\langle 2 \rangle \wedge Y\langle 1 \rangle = Y\langle 2 \rangle$. Note that the lifting is only valid because we do not care about the final value of X. The proof makes two uses of the dead-assignment rule, which is a common pattern for performing code-motion in RHL: one effectively adds `skip`s to both sides to make them the same 'shape', shows the equivalence using the congruence rules and then removes the `skip`s.

Sophisticated dead-code:

```
if X>3 then Y=X else Y=7    ==> skip
```

at type $(X\langle 1 \rangle = X\langle 2 \rangle \wedge Y\langle 1 \rangle > 2 \wedge Y\langle 2 \rangle > 2) \Rightarrow (Y\langle 1 \rangle > 2 \wedge Y\langle 2 \rangle > 2)$. I.e. if all that matters about the value of Y in the rest of the derivation is that it is greater than 2, then the conditional has no effect. This is proved by using the conditional rule to show the equivalence of the two branches and then applying the equivalent branch rule.

The main weakness of RHL as presented here relates to its treatment of loops. Since we insist that transformed programs have the same termination behaviour as the original, but have no nontrivial termination analysis, this is hardly suprising. I believe it is possible to add sound rules which can justify some cases of loop distribution/fusion, but more ambitious loop optimizations seem to require either a language with restricted iteration constructs or a logic which can reason about termination.

4.4 Embedding Simpler Logics in RHL

RHL is powerful but hardly suitable for direct implementation in a compiler. However, it can provide a useful framework for developing sound type and transformation systems which are more specific. One would start by identifying a restricted sublanguage of relational assertions. For example, several useful analyses can be formulated using only partial equivalence relations generated from axioms such as:

$$\vdash PER(E\langle 1 \rangle = E\langle 2 \rangle) \qquad \vdash PER(B\langle 1 \rangle = B\langle 2 \rangle)$$

$$\vdash PER(B\langle 1 \rangle \wedge B\langle 2 \rangle)$$

plus rules stating that PERs are closed under conjunction, disjoint union and the arrow constructor. Our earlier DDCC system is of this form, with state relations being formed as conjunctions of primitive assertions of the forms $X\langle 1 \rangle = X\langle 2 \rangle$ and $X\langle 1 \rangle = n \wedge X\langle 2 \rangle = n$. The rules of DDCC can then be presented as derived rules in RHL.

A natural question is whether the usual Hoare logic can be embedded in RHL. One's first thought might be that a partial correctness judgement $\{P\}C\{Q\}$ would be equivalent to the 'squared' RHL judgement $\vdash C \sim C : P\langle 1 \rangle \wedge P\langle 2 \rangle \Rightarrow Q\langle 1 \rangle \wedge Q\langle 2 \rangle$, but this is not the case because C's termination behaviour might differ on two states satisfying P. Nor can one simply intersect the pre- and post-relations with the identity relation on states, since we do not have

syntax for that 'global' identity relation. If we fix C, however, we can conjoin the pre- and post-relations with $X\langle 1\rangle = X\langle 2\rangle$ for every variable X occurring in C and thus effectively recover Hoare logic.[3] Going the other way, one can soundly extend RHL with the squared versions of valid *total* correctness judgements $[P]C[Q]$.

As a simple, concrete example of the embedding approach, Figure 6 presents (a very naive version of) a type system AERC for available expression analysis and removal of redundant evaluation. State types Θ are finite sets $\{X_i = E_i \mid 1 \le i \le n\}$ of equalities between variables and expressions (in which the same variable may occur multiple times on the left) and we write $\Theta \le \Theta'$ for $\Theta \supseteq \Theta'$. The macros *kill* and *gen* are defined by

$$kill(\Theta,X) = \{(X_i = E_i) \in \Theta \mid X_i \ne X \wedge X \notin E_i\}$$

$$gen(X,E) = \begin{cases} \{X = E\} & \text{if } X \notin E \\ \{\} & \text{otherwise} \end{cases}$$

The translation of the AERC into RHL is indexed by a finite set V of variables. Define

$$\Theta_V^* = \bigwedge_{X \in V}(X\langle 1\rangle = X\langle 2\rangle) \wedge \bigwedge_{(X=E)\in\Theta}(X\langle 1\rangle = E\langle 1\rangle)$$

It is easy to see that for any Θ, $\models PER(\Theta_V^*)$ and that $\Theta \le \Theta'$ implies $\Theta_V^* \le \Theta_V'^*$. The following asserts the soundness of the translation, and hence of AERC:

THEOREM 4. *For any expressions E,F and commands C,D all of whose variables occur in V,*

1. *If $\vdash E \sim F : \Theta \Rightarrow \tau$ then $\models \Theta_V^* \le (E\langle i\rangle = F\langle j\rangle)$ for $i,j \in \{1,2\}$.*

2. *If $\vdash C \sim D : \Theta \Rightarrow \Theta'$ in AERC then $\vdash C \sim D : \Theta_V^* \Rightarrow \Theta_V'^*$ in RHL.*

PROOF. The interesting case is the AERC rule for assignment, which generates the following verification condition in RHL: if $\models \Theta_V^* \le (E\langle 1\rangle = E'\langle 2\rangle)$ then $\models \Theta_V^* \le (kill(\Theta,X) \cup gen(X,E) \cup gen(X,E'))_V^*[E\langle 1\rangle/X\langle 1\rangle, E'\langle 2\rangle/X\langle 2\rangle]$, which is straightforward to check. □

5 Related Work

As we said in the introduction, there has been a good deal of work on proving the correctness of optimizing transformations for functional languages, especially from the group at Northeastern [30, 28, 32, 31] but also by Amtoft on strictness analysis [5], Damiani and Giannini on dead-variables [10, 11], Kobayashi on dead-variables [17] and Benton and Kennedy on effects [8]. Damiani and Giannini explicitly use PERs in giving the semantics of their analysis system but give a more algorithmic account its use in transformation. Benton and Kennedy present optimizing transformations as equations in context, but derive those (rather clumsily) from a predicate-based semantics for the analysis.

Recently Lacey et al. [19] described how some of the classical [16, 12, 4] transformations considered here (dead code elimination, constant folding and a simple code-motion transformation) can be formulated as conditional rewrite rules on control flow graphs. The rewrites are predicated on temporal logic formulae expressing (intensionally) the contexts in which the rewrites may be applied. The

authors then use a small-step operational semantics to verify that under these conditions, their transformations preserve the observable behaviour of programs. Lacey et al. express strongly the view that more traditional semantic techniques, in particular denotational ones, are either unable to express the properties which justify optimizing transformations or can only do so at the cost of complex proofs and 'mathematically sophisticated' techniques. I believe the present paper provides some evidence to the contrary.

Lerner et al. [20] have built an implementation of a domain specific language for specifying and justifying rewrites on a simple imperative language which interfaces to a theorem prover for checking the supplied justification. This system also uses a (rather restricted) language of temporal logic formulae for specifying optimizations over flow graphs.

Kozen and Patron [18] describe an algebraic approach to proving some traditional optimizations correct. There is no mention of relations in their work, and they abstract rather severely from the actual language (there are no assignments, just unspecified atomic programs including one which makes a variable 'undefined'), but the connections between their work and ours seem worth further study.

The work that is most closely related to that presented here has been done in the contexts of *credible compilation* [24, 25] and *translation validation* [21, 36]. These both take the view that formal verification of complete optimizing compilers is impractical, but that one might realistically produce a correctness proof relating the input and output of particular compilations. Translation validation tries to do this without modifying the compiler, using an independent tool that tries to infer that the output is a correct translation of the input. Credible compilation envisages an instrumented compiler producing a putative proof that the transformations it performed in each particular case were safe; these proofs can then be examined by a comparatively simple proof-checker. The basic technical ideas used in credible and validated compilation are very close indeed to the ones presented here (developed quite independently). The main difference is that we use the language of types, denotational semantics and PERs instead of that of control-flow graphs, operational semantics and simulation relations. Inspired by Rinard's work, Yang [35] has recently used a version of relational Hoare logic in reasoning about the correctness of the Schorr-Waite graph marking algorithm.

The idea of directly axiomatising a logic of PERs [3] and more general relations was inspired by the work of Abadi et al on a formal logic for parametric polymorphism [2].

We have already mentioned some of the large amount of recent work using PERs (and domain-theoretic projections) to give semantics to analyses for non-interference, slicing, secure information flow, binding time analysis. An elegant general calculus, DCC, for such dependency-based analyses has been defined by Abadi et al. [1]. DCC seems comparable to a higher-order version of our DDCC, though it is not explicitly presented as an equational calculus and is more directly in the style of type systems for secure information flow.

The work of Hughes on type specialization [14] seems to have interesting connections with (a higher-order version of) the work presented here. Hughes has formulated a type-based analysis which essentially uses a form of singleton type, and proved the correctness of an associated transformation system which changes types. Singleton types and their PER semantics have also been studied in some depth by Aspinall [6].

---

[3]The civilised way to do this is to index all our judgements by finite sets of variable names.

$$\vdash X \sim X : \Theta \Rightarrow \texttt{int} \ [\text{A-V}] \qquad \vdash \texttt{n} \sim \texttt{n} : \Theta \Rightarrow \texttt{int} \ [\text{A-N}] \qquad \vdash \texttt{b} \sim \texttt{b} : \Theta \Rightarrow \texttt{bool} \ [\text{A-B}]$$

$$\vdash X \sim E : \Theta \cup \{X = E\} \Rightarrow \texttt{int} \ [\text{A-Red}] \qquad \vdash \texttt{skip} \sim \texttt{skip} : \Theta \Rightarrow \Theta \ [\text{A-Skp}]$$

$$\frac{\vdash E \sim E' : \Theta \Rightarrow \texttt{int} \quad \vdash F \sim F' : \Theta \Rightarrow \texttt{int}}{\vdash E \ \texttt{iop} \ F \sim E' \ \texttt{iop} \ F' : \Theta \Rightarrow \texttt{int}} [\text{A-iop}] \ (+ \text{ similar } bop \text{ and } lop) \qquad \frac{\vdash C_1 \sim C_1' : \Theta \Rightarrow \Theta' \quad \vdash C_2 \sim C_2' : \Theta' \Rightarrow \Theta''}{\vdash (C_1 ; C_2) \sim (C_1' ; C_2') : \Theta \Rightarrow \Theta''} [\text{A-Seq}]$$

$$\frac{\vdash B \sim B' : \Theta \Rightarrow \texttt{bool} \quad \vdash C \sim C' : \Theta \Rightarrow \Theta}{\vdash (\texttt{while } B \texttt{ do } C) \sim (\texttt{while } B' \texttt{ do } C') : \Theta \Rightarrow \Theta} [\text{A-Whl}] \qquad \frac{\vdash E \sim E' : \Theta \Rightarrow \texttt{int}}{\vdash X \texttt{:=} E \sim X \texttt{:=} E' : \Theta \Rightarrow (kill(\Theta, X) \cup gen(X, E) \cup gen(X, E'))} [\text{A-Ass}]$$

$$\frac{\vdash B \sim B' : \Theta \Rightarrow \texttt{bool} \quad \vdash C_1 \sim C_1' : \Theta \Rightarrow \Theta' \quad \vdash C_2 \sim C_2' : \Theta \Rightarrow \Theta'}{\vdash (\texttt{if } B \texttt{ then } C_1 \texttt{ else } C_2) \sim (\texttt{if } B' \texttt{ then } C_1' \texttt{ else } C_2') : \Theta \Rightarrow \Theta'} [\text{A-If}] \qquad \frac{\vdash C \sim C' : \Theta \Rightarrow \Theta'}{\vdash C' \sim C : \Theta \Rightarrow \Theta'} [\text{A-CSym}]$$

$$\frac{\vdash C \sim C' : \Theta_1 \Rightarrow \Theta_2 \quad \Theta_1' \le \Theta_1 \quad \Theta_2 \le \Theta_2'}{\vdash C \sim C' : \Theta_1' \Rightarrow \Theta_2'} [\text{A-CSub}] \qquad \frac{\vdash E \sim E' : \Theta \Rightarrow \tau \quad \Theta' \le \Theta}{\vdash E \sim E' : \Theta' \Rightarrow \tau} [\text{A-ESub}] \qquad \frac{\vdash F_\tau \sim F_\tau' : \Theta \Rightarrow \tau}{\vdash F_\tau' \sim F_\tau : \Theta \Rightarrow \tau} [\text{A-ESym}]$$

$$\frac{\vdash C \sim C' : \Theta \Rightarrow \Theta' \quad \vdash C' \sim C'' : \Theta \Rightarrow \Theta'}{\vdash C \sim C'' : \Theta \Rightarrow \Theta'} [\text{A-CTr}] \qquad \frac{\vdash F_\tau \sim F_\tau' : \Theta \Rightarrow \tau \quad \vdash F_\tau' \sim F_\tau'' : \Theta \Rightarrow \tau}{\vdash F_\tau \sim F_\tau'' : \Theta \Rightarrow \tau} [\text{A-ETr}]$$
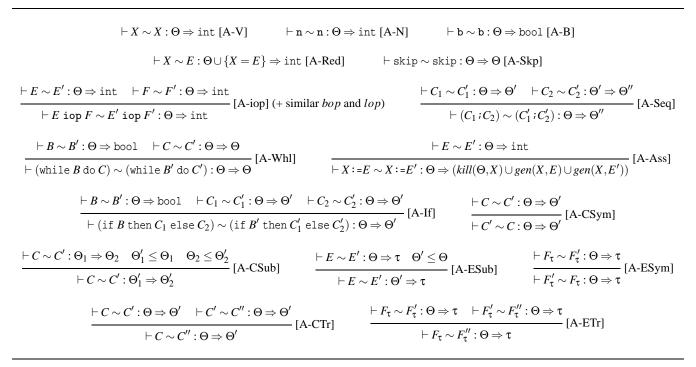
Figure 6. AERC: Available Expressions and Redundant Computation

6 Conclusions and Further Work

We have shown how some very elementary techniques can be used to prove the combined correctness of analyses and transformations for simple imperative programs. From a purely semantic perspective, there is nothing very surprising here. But that is as it should be: we have finally shown that something that appears simple actually *is* simple.

One obvious shortcoming of the present work is that it says nothing about concrete inference or transformation algorithms. Although there are benefits in factoring a correctness proof into the soundness of a declarative set of inference rules and the correctness of an inference algorithm, one does ultimately have to provide both parts. Although there seems no reason why the approach taken here should not carry over to the control-flow graphs more commonly used in optimizing compilers for imperative languages, proving directly that analyses as they are actually implemented in real compilers imply our extensional properties seems likely to be somewhat messy. Combining our extensional relational approach to correctness with a more algorithmic, but still declarative, framework for specifying transformations (such as that of Lacey et al.) seems a more reasonable next step.

Relational Hoare logic is a promising formalism which certainly merits further study. A limitation of the system presented here is that it cannot justify any transformations which remove loops, except in the special case that they can be completely unrolled at compile-time. This naturally suggests investigating a total-correctness variant of the logic, but one might also consider a version which allows termination-improving transformations. A further possibility is to axiomatise the version of relational lifting which maps Φ to $\{(d, d') \in \mathbb{S}_\perp \times \mathbb{S}_\perp \mid (d = \lceil s \rceil \wedge d' = \lceil s' \rceil) \implies (s, s') \in \Phi\}$. This is inappropriate for optimizations, but seems useful in reasoning about termination-insensitive information flow [7].

There are many natural ways to develop the ideas here, both in terms of the language features and analyses addressed (higher-types, higher-typed store, dynamic allocation) and in terms of the features of the logics (e.g. quantification, conjunctive and disjunctive types). Doing some of these would require working with relations on recursively-defined domains, for which we expect to use the techniques described by Pitts [23]. If the technique extends to imperative programs with higher-typed store, a promising idea is to look at optimizations on low level code that are justified by relational invariants passed down from a high-level compiler.

7 References

[1] M. Abadi, A. Banerjee, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *Conference Record of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 147–160. ACM Press, January 1999.

[2] M. Abadi, L. Cardelli, and P.-L. Curien. Formal parametric polymorphism. *Theoretical Computer Science*, 121, 1993.

[3] M. Abadi and G. D. Plotkin. A PER model of polymorphism and recursive types. In *Proceedings of the 5th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 355–365. IEEE Computer Society Press, June 1990.

[4] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques and Tools*. Addison Wesley, 1986.

[5] T. Amtoft. Minimal thunkification. In P. Cousot, M. Falaschi, G. Filè, and A. Rauzy, editors, *Proceedings of the 3rd International Workshop on Static Analysis, Padova, Italy*, volume

724 of *Lecture Notes in Computer Science*, pages 218–229. Springer-Verlag, September 1993.

[6] D. Aspinall. Subtyping with singleton types. In L. Pacholski and J. Tiuryn, editors, *Computer Science Logic, 8th International Workshop (CSL'94)*, number 933 in Lecture Notes in Computer Science. Springer-Verlag, 1995.

[7] A. Banerjee and D. Naumann. Secure information flow and pointer confinement in a Java-like language. In *Proceedings of the 15th IEEE Computer Security Foundations Workshop (CSFW)*, pages 253–267. IEEE Computer Society Press, June 2002.

[8] N. Benton and A. Kennedy. Monads, effects and transformations. In *Third International Workshop on Higher Order Operational Techniques in Semantics (HOOTS), Paris*, volume 26 of *Electronic Notes in Theoretical Computer Science*. Elsevier, September 1999.

[9] P. N. Benton. *Strictness Analysis of Lazy Functional Programs*. PhD thesis, Computer Laboratory, University of Cambridge, December 1992.

[10] F. Damiani. Useless-code Detection and Elimination for PCF with Algebraic Datatypes. In *4th International Conference on Typed Lambda Calculi and Applications (TLCA)*, volume 1581 of *Lecture Notes in Computer Science*, pages 83–97. Springer-Verlag, 1999.

[11] F. Damiani and P. Giannini. Automatic useless-code detection and elimination for hot functional programs. *Journal of Functional Programming*, pages 509–559, 2000.

[12] M S. Hecht. *Flow Analysis of Computer Programs*. Elsevier North-Holland, 1977.

[13] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–585, October 1969.

[14] J. Hughes. Type specialization for the lambda calculus. In *Proceedings of the Dagstuhl Seminar on Partial Evaluation*, 1996.

[15] S. Hunt and D. Sands. Binding time analysis: A new PERspective. In *Proceedings ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation (PEPM)*, June 1991.

[16] G. A. Kildall. A unified approach to global program optimization. In *Proceedings of the 1st ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)*, pages 194–206. ACM Press, 1973.

[17] N. Kobayashi. Type-based useless variable elimination. In *Proceeedings of the ACM Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, 2000.

[18] D. Kozen and M. Patron. Certification of compiler optimizations using Kleene algebra with tests. In *Proceedings of the 1st International Conference in Computational Logic*, volume 1861 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2000.

[19] D. Lacey, N. D. Jones, E. Van Wyk, and C. C. Frederiksen. Proving correctness of compiler optimizations by temporal logic. In *Proceedings of the 29th Annual ACM SIGPLAN - SIGACT Symposium on Principles of Programming Languages, Portland*, January 2002.

[20] S. Lerner, T. Millstein, and C. Chambers. Automatically proving the correctness of compiler optimizations. In *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation (PLDI)*, June 2003.

[21] G. C. Necula. Translation validation for an optimizing compiler. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 83–95, 2000.

[22] F. Nielson. Program transformations in a denotational setting. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 7(3):359–379, July 1985.

[23] A.M. Pitts. Relational properties of domains. *Information and Computation*, 127, 1996.

[24] M. Rinard. Credible compilation. Technical Report MIT-LCS-TR-776, Massachusets Institute of Technology, March 1999.

[25] M. Rinard and D. Marinov. Credible compilation with pointers. In *Proceedings of the FLoC Workshop on Run-Time Result Verification*, July 1999.

[26] A. Sabelfeld and D. Sands. A PER model of secure information flow in sequential programs. *Higher-Order and Symbolic Computation*, 14(1):59–91, March 2001.

[27] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *Conference Record of the 25th ACM Symposium on Principles of Porgramming Languages (POPL)*, January 1998.

[28] P. A. Steckler and M. Wand. Lightweight closure conversion. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, pages 48–86, January 1997. Original version appeared in Proceedings 21st ACM Symposium on Principles of Programming Languages, 1994.

[29] D. Volpano, G. Smith, and C. Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4:167–187, December 1996.

[30] M. Wand. Specifying the correctness of binding-time analysis. In *Conference Record of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 137–143. ACM, January 1993.

[31] M. Wand and W. D. Clinger. Set constraints for destructive array update optimization. *Journal of Functional Programming*, 11(3):319–346, May 2001. Preliminary version appeared in International Conference on Computer Languages, 1998.

[32] M. Wand and I. Siveroni. Constraint systems for useless variable elimination. In *Proceedings 26th ACM Symposium on Principles of Programming Languages*, pages 291–302, 1999.

[33] M. Weiser. Program slicing. *IEEE Transactions on Software Engineering*, 10(4):352–357, July 1984.

[34] G. Winskel. *The Formal Semantics of Programming Languages*. MIT Press, 1993.

[35] H. Yang. Verification of the Schorr-Waite graph marking algorithm by refinement. Slides from talk at Dagstuhl Seminar 03101, March 2003.

[36] L. Zuck, A. Pnueli, Y. Fang, and B. Goldberg. VOC: A methodology for the translation validation for optimizing compilers. *Journal of Universal Computer Science*, 9(3):223–247, 2003.