

Formalizing and Verifying Semantic Type
Soundness of a Simple Compiler
(Preliminary Report)

Nick Benton
Microsoft Research, Cambridge
`nick@microsoft.com`

Uri Zarfaty
Imperial College, London
`udz@doc.ic.ac.uk`

March 2007
Technical Report
MSR-TR-2007-31

We describe a semantic type soundness result, formalized in the Coq proof assistant, for a compiler from a simple imperative language with heap-allocated data into an idealized assembly language. Types in the high-level language are interpreted as binary relations, built using both second-order quantification and a form of separation structure, over stores and code pointers in the low-level machine.

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
<http://www.research.microsoft.com>

1 Introduction

The last decade has seen an explosion of research into type systems, formal verification and certification for low-level code, ignited by the original papers on typed assembly language [34] and proof-carrying code [36], and fanned by the development of separation logic [42]. These developments have been driven partly by need (as well as the traditional arguments in favour of some level of formal verification as a way to develop software that actually works, the internet has made checkable safety of mobile code more than a purely academic problem), partly by improvements in the technology of theorem provers and model checkers, and partly by the realization that conservative techniques for verifying comparatively simple properties, such as forms of memory safety, can be much easier and more efficient to apply than complete methods for showing full functional correctness, whilst still offering useful real-world guarantees. Another driving force has been the (occasionally surprising) discovery that logical, type theoretic and semantic ideas that were originally developed in fairly abstract settings, or for very high-level programming languages, are actually applicable to realistic, low-level, ‘dirty’ languages and systems.

In the present paper, we will be concerned with *certified compilation*, proving once and for all that a compiler always produces object code that satisfies some policy, which in this case will be type safety.¹

But what do we mean by type safety? For high-level languages, there are two main approaches to formalizing type soundness properties: *syntactic* and *semantic*. The difference between the two is *not* (merely) one of proof technique; they are different kinds of result.

In the syntactic approach [46], one typically defines a small-step operational semantics for the language that gets stuck (makes no transition) in configurations that are considered to be bad. One then shows ‘preservation and progress’ – that every typable configuration is either properly terminal or makes a transition into another typeable configuration – and can then conclude by induction that well-typed programs don’t get stuck. Syntactic type soundness is often fairly straightforward to establish, but is a rather weak and fragile result. Firstly, it is closely tied to the particular set of syntactic rules that define the type system. There is no direct definition of the meaning of a type A as a property of phrases beyond ‘being assignable the type A using this particular set of rules’, so there is no real notion of what it is that the types are supposed to ensure. Secondly, the introduction of stuckness (or error states) into the operational semantics is something of a sleight of hand, changing the original problem to match the solution. For a simple type system and a high-level semantics, it seems reasonable to work with syntax for (untyped) phrases that explicitly distinguishes, say, functions from integer constants, or even booleans and integers, and which gets stuck when one tries to apply an integer or increment a boolean.

¹Certified compilation can be contrasted with *certifying compilation*, in which the compiler produces a checkable certificate for each binary, purporting to prove that the policy is satisfied in that case, and *compiler validation*, in which each output of the compiler is analysed for safety by an independent tool.

But this becomes less tenable when the type system is intended to track more interesting properties, such as the use of locks or the reading and writing of particular parts of the store. In such cases, formulating a syntactic type soundness result involves changing the operational semantics of the language one first thought of, to track extra information and possibly add new stuck states. And the more sophisticated the analysis, the more complex the instrumentation becomes. Furthermore, the notion of error is not preserved by compilation: machine code does not inherently distinguish code pointers, heap pointers, integers or booleans; no fault is raised by performing arithmetic on addresses to which one subsequently jumps or stores,² and compiled code, particularly when optimized, often depends upon such possibilities.

It is, of course, possible to mitigate the effect of bogus instrumentation by also proving an erasure theorem, showing that removing the extra information does not affect the execution of non-faulting programs, and such a result can be extended to relate evaluation in a high-level operational semantics to the behaviour of compiled code. Even then, however, the theorem about low-level code is tied to the syntactic definition of the high-level language and its type rules. This is a significant shortcoming: compiled code nearly always relies upon a runtime system and library routines that are written directly in a low-level language, and we would also like to be able to link soundly with code compiled from other high-level languages. Without an independent low-level characterization of the intended behavioural properties of code compiled from phrases with a particular high-level type, the implementer of a library function or support routine written in C or assembler does not know what specification his code should meet in order to interoperate properly with the output of the compiler.

The semantic approach to type safety, by contrast, gives a meaning to each type that is independent of any particular set of rules for assigning those types to program phrases. The meaning of a type will be (roughly) a set of values with some property; for a given language and set of types, there can be many different analyses, of varying degrees of precision, for soundly assigning types to terms.³ The meaning of a function type $A \rightarrow B$, for example, can be defined inductively to be the set of terms that when applied to a value in the meaning of type A always yield a result in the meaning of type B ; the similarity (at least at first order) of such an interpretation to the meaning Hoare triples in program logics should be apparent.

Interpretations of types as predicates over some untyped model of computation have a long history, from the early work of Scott [43] and McCracken

²This is, of course, not strictly true. Operating systems use memory management hardware to trap ‘illegal’ pointer dereferencing or jumps to addresses in pages marked as ‘no execute’, floats are passed in special registers, etc. But faults in compiled code certainly do not correspond exactly to errors in a high-level semantics, and a major goal of static verification should surely be to remove the need for such crude and expensive dynamic checks.

³This is the ‘extrinsic’ or ‘descriptive’ view of typing, traditionally associated with Curry, which one may contrast with the ‘intrinsic’, ‘prescriptive’ position of Church, which regards types as coming before terms, rather than after. Even if our source types are prescriptive, we are firmly in the Curry camp regarding their interpretations here.

[29] using retracts of universal domains through to recent research on ideals and biorthogonality [45, 30]. Particularly relevant from the point of view of the present paper is the work of Appel and his collaborators [8, 7, 9, 44, 4] on Foundational Proof Carrying Code (FPCC). The idea of FPCC is to give a semantics to high-level types as low-level specifications expressed in some sufficiently powerful program logic. This low-level logic is not tied to any particular language or type system, and proofs that a particular piece of low-level code satisfies such a specification can be generated or checked independently from any particular compilation scheme. Although the concept of FPCC is clearly parametric in just what safety property one wishes to prove and check, the only instance that has really been studied and implemented so far is *memory safety*: ensuring that ‘illegal’ accesses to memory will not occur at runtime. The rather intensional notion of which accesses are legal is formalized by changing the operational semantics of the low-level machine so that it gets stuck when certain locations are dereferenced, essentially just as in the syntactic approaches.

Whilst memory safety is undeniably important, it is not the same as type safety. Program fragments that satisfy an interpretation of a type in the style of previous work on FPCC, whilst memory safe (in the right contexts), can easily fail to have other rather basic properties one would expect, and on which security and compiler correctness can depend. Given an ML-like source language, for example, an assembly language function that simply returns its argument will be in the interpretation of the type $(\text{int} \rightarrow \text{int}) \rightarrow \text{int}$, since if one passes in the address of some closure, one will get back something that looks like an integer. But allowing the identity function to be given that type, whilst not leading to illegal memory accesses, would invalidate very basic reasoning principles for ML programs that are used by both programmers and compilers: not only are static transformations such as common subexpression elimination no longer behaviour-preserving, but the observable results of a particular compiled binary can vary according to where it is loaded in memory, the behaviour of the allocator, etc. Such possibilities violate most language-based encapsulation or security properties one can think of.

An ideal situation would be that the translation of high-level types into low-level assertions captures enough about the semantics of the source language that all typed contextual equivalences that are valid at the high level are somehow justifiable at the low-level. This would mean we had captured a sense in which our compilation scheme is *fully abstract*. Even quite small abstraction failures in compilation can easily lead to security holes [1, 26] so it seems natural to try to achieve full abstraction. One way to do that would be to take the interpretation of a type A to be (roughly) the set of code fragments that might be emitted by a fixed compiler for phrases of type A (a kind of term model), but that would clearly be silly. What we really want are *behavioural* specifications, admitting the widest possible range of well-behaved implementations. But it may be that ‘sufficient’, rather than ‘full’, abstraction is the reasonable goal. We certainly want to preserve all the reasoning principles that might be used by a compiler in optimization or on which security properties might depend, and for real languages those certainly include subtle arguments about,

for example, encapsulated state and higher-order features. On the other hand, source languages may validate very esoteric equivalences that could almost be regarded as accidental. There is some reason (e.g. some results on typed versus untyped realizability [30]) to believe that modelling all equations may lead to excessively complicated specifications and awkward restrictions on contexts, with little benefit.

What sort of interpretation of types should we use to establish sufficient abstraction of compilation? In the present paper, we work with a semantic interpretation of high-level types that uses binary relations, rather than unary predicates, over low-level code and data. One should think of these relations as carving out both a set of values *and* a type-specific notion of equality on that set of values; these are defined together because which values are judged to be in the set associated with some compound type will depend on both the sets of values and the equality relations associated with its components. The crucial case is that for functions: two values f and f' are in the relation interpreting $A \rightarrow B$ iff for any x and x' that are related by the interpretation of A , $f x$ and $f' x'$ are related by the interpretation of B . The set of values having a particular type is given by the diagonal part of the associated relation, so f has type $A \rightarrow B$ just when f is related to itself by the interpretation of $A \rightarrow B$; this is the usual notion of ‘logical’ relation [39]. Since equality should be an equivalence relation, the natural notion for interpreting a type is a set of values and an equivalence relation on that set, which is easily seen to be the same thing as a relation that is symmetric and transitive, aka a *partial equivalence relation* (PER). Interpreting types as partial equivalence relations over some untyped model of computation also has a long history, but previous work has generally taken the untyped model either to be rather high-level and abstract (e.g. a domain theoretic model of the untyped lambda calculus) or low-level but with uninteresting fine structure (e.g. some Gödel numbering of partial recursive functions). The difference here is that we work with a low-level, untyped model in whose structure we most certainly are interested, *viz.* machine code (albeit very idealized), and we work with a translation into that model that is representative of realistic compilation schemes (albeit for a rather toy language). We do not claim that the actual type soundness result we prove is, in itself, especially exciting; the interesting aspects of the work are the *form* of the result and the methodology used for showing it.

This paper describes work in progress, building on our earlier work on modular specification and verification of a simple memory allocator [13]. The results have been formalized and checked in the Coq proof assistant and most of the formal parts of the present paper are presented as extracts from the proof script, using Coq syntax.

2 Low-Level Target Machine

We work with the same completely straightforward operational semantics for an idealized assembly language that we used in our earlier work on allocation [13].

There is a single datatype, the natural numbers, though different instructions treat elements of that type as code pointers, heap addresses, integers, etc. The heap is simply a total function from naturals to naturals and the code heap is a total function from naturals to instructions. Computed branches and address arithmetic are perfectly allowable. There is no built-in notion of allocation and no notion of stuckness or ‘going wrong’: the only observable behaviours are termination and divergence.

The instruction set of the machine is given by Coq inductive definitions for lvalues (*dest*), rvalues (*src*) and instructions (*instruction*). Destinations are either immediate (a fixed memory location), indirect or indirect with a fixed offset. Sources are literal values, immediate (the contents of a fixed memory location), indirect or indirect with an offset.

Inductive *dest* : *Set* :=
 | *d_imm* : *nat* → *dest* | *d_ind* : *nat* → *dest* | *d_indo* : *nat* → *nat* → *dest*.

Inductive *src* : *Set* :=
 | *s_cst* : *nat* → *src* | *s_imm* : *nat* → *src* | *s_ind* : *nat* → *src*
 | *s_indo* : *nat* → *nat* → *src*.

Inductive *instruction* : *Set* :=
 | *i_halt* : *instruction*
 | *i_move* : *dest* → *src* → *instruction*
 | *i_add* : *dest* → *src* → *src* → *instruction*
 | *i_sub* : *dest* → *src* → *src* → *instruction*
 | *i_mult* : *dest* → *src* → *src* → *instruction*
 | *i_branch* : *src* → *instruction*
 | *i_brz* : *src* → *src* → *instruction*
 | *i_brnz* : *src* → *src* → *instruction*.

The mutable heap of our machine is a function from naturals to naturals, which we choose to represent using a named record type with a single field and an implicit coercion to (*nat* → *nat*):

Record *state* : *Set* := *State* { *fun_of_state* :> *nat* → *nat* }.

Definition *update* (*s*:*state*) (*n*:*nat*) (*v*:*nat*) : *state* :=
State (*fun m* ⇒ if *beq_nat n m* then *v* else *s m*).

We now give the meaning of sources, destinations, and the single-step semantics of instructions themselves. The meaning of an instruction at a particular program counter in a particular state is of an *option* type: either *None*, indicating termination, or *Some*(*s'*, *pc'*), giving a new heap and a new program counter:

Definition *sem_dest* (*de*:*dest*) (*s*:*state*) :=
 match *de* with
 | *d_imm n* ⇒ *n*
 | *d_ind n* ⇒ *s n*
 | *d_indo ofs n* ⇒ *s n + ofs*
 end.

Definition *sem_src* (*sr*:*src*) (*s*:*state*) :=

```

match sr with
| s_cst n ⇒ n
| s_imm n ⇒ s n
| s_ind n ⇒ s (s n)
| s_indo ofs n ⇒ s (s n + ofs)
end.

```

Definition *sem_instr* (*ins:instruction*) (*s:state*) (*pc:nat*) : *option (state × nat)*
:=

```

match ins with
| i_halt ⇒ None
| i_move de sr ⇒ Some (update s (sem_dest de s) (sem_src sr s), S pc)
| i_add de sr1 sr2 ⇒
  Some (update s (sem_dest de s) ((sem_src sr1 s) + (sem_src sr2 s)), S pc)
| i_sub de sr1 sr2 ⇒
  Some (update s (sem_dest de s) ((sem_src sr1 s) - (sem_src sr2 s)), S pc)
| i_mult de sr1 sr2 ⇒
  Some (update s (sem_dest de s) ((sem_src sr1 s) × (sem_src sr2 s)), S pc)
| i_branch sr ⇒ Some (s, sem_src sr s)
| i_brz srscrut srtarg ⇒
  Some (s, match sem_src srscrut s
        with 0 ⇒ sem_src srtarg s | S _ ⇒ S pc end)
| i_brnz srscrut srtarg ⇒
  Some (s, match sem_src srscrut s
        with 0 ⇒ S pc | S _ ⇒ sem_src srtarg s end)
end.

```

A program is simply a total function from labels (naturals) to instructions, whilst a program fragment is a partial function from labels to naturals:

Definition *program* : *Set := nat → instruction*.

Definition *program_fragment* : *Set := nat → option instruction*.

Definition *program_extends_fragment* (*p:program*) (*pf:program_fragment*) :=
 $\forall n, \text{match } pf \ n \text{ with } None \Rightarrow True \mid Some \ i \Rightarrow (p \ n = i) \text{ end.}$

We now define *kstepterm*, saying when a configuration comprising a program *p*, a heap *s*, and a program counter *l* terminates in *k* steps. The *terminates* predicate then holds of configurations that terminate in some number of steps:

Fixpoint *kstepterm* (*k:nat*) (*p:program*) (*s:state*) (*l:nat*) {*struct k*} : *Prop :=*
match k with
| 0 ⇒ *False*
| (S j) ⇒ *match sem_instr (p l) s l with*
| *None* ⇒ *True*
| *Some (s', l')* ⇒ *kstepterm j p s' l'*
end

end.

Definition *terminates p s l := ∃ k, kstepterm k p s l.*

Expression types

$$A ::= \text{int} \mid \text{bool} \mid A \times A'$$

Store types

$$\Gamma ::= v_1 : A_1, \dots, v_n : A_n$$

Expressions

$$\begin{array}{c} \frac{}{\Gamma \vdash \text{true} : \text{bool}} \quad \frac{}{\Gamma \vdash \text{false} : \text{bool}} \quad \frac{}{\Gamma \vdash \underline{n} : \text{int}} \quad \frac{}{\Gamma, x : A \vdash x : A} \\ \\ \frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 + e_2 : \text{int}} \quad \frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 > e_2 : \text{bool}} \\ \\ \frac{\Gamma \vdash e_1 : A_1 \quad \Gamma \vdash e_2 : A_2}{\Gamma \vdash (e_1, e_2) : A_1 \times A_2} \quad \frac{\Gamma \vdash e : A_1 \times A_2}{\Gamma \vdash \pi_1 e : A_1} \quad \frac{\Gamma \vdash e : A_1 \times A_2}{\Gamma \vdash \pi_2 e : A_2} \end{array}$$

Commands

$$\begin{array}{c} \frac{\Gamma, x : A \vdash e : B}{\Gamma, x : A \vdash (x := e) : \Gamma, x : B} \quad \frac{\Gamma \vdash C_1 : \Gamma' \quad \Gamma' \vdash C_2 : \Gamma''}{\Gamma \vdash C_1 ; C_2 : \Gamma''} \\ \\ \frac{\Gamma \vdash e : \text{bool} \quad \Gamma \vdash C_1 : \Gamma' \quad \Gamma \vdash C_2 : \Gamma'}{\Gamma \vdash \text{if } e \text{ then } C_1 \text{ else } C_2 : \Gamma'} \quad \frac{\Gamma \vdash e : \text{bool} \quad \Gamma \vdash C : \Gamma}{\Gamma \vdash \text{while } e \text{ do } C : \Gamma} \end{array}$$

Figure 1: The While Language

The major idealizations compared with a real machine are that we have arbitrary-sized natural numbers as a primitive type, rather than fixed-length words, and that we have separated code and data memory. Note also that there are no registers; we will simply adopt a convention of using some low-numbered memory locations in a register-like fashion.

3 Source Language

The source language is that of simple `while`-programs with integer- (actually natural-), boolean- and pair-valued expressions and a type system for commands that supports ‘strong updates’ to (global) variables. The syntax and type rules of the language are shown in conventional notation in Figure 1. Expressions are typed in the context of a typing Γ for the global variables. Commands, which may update variables with values of different types, are given both a pretyping and a posttyping, recording their assumptions and effects on the store.

The Coq translation of Figure 1 is fairly direct. We use natural numbers instead of names for global variables and have chosen to make elements of $EnvType$, representing store types, be *total* functions on the naturals (we will pass around a separate size as well later on, though we could have used lists or fixed-size vectors instead). Note also the use of simple ‘GADT-style’ dependent typing for expressions and commands: $Exp\ env\ t$ is the type of expressions that have type t in store environment env , and similarly for commands.

Inductive $ExpType : Set :=$
 | $TInt : ExpType$
 | $TBool : ExpType$
 | $TPair : ExpType \rightarrow ExpType \rightarrow ExpType$.

Notation " $a ** b$ " := ($TPair\ a\ b$) (at level 55).

Definition $EnvType := nat \rightarrow ExpType$.

Definition $envupdate\ (env : EnvType)\ m\ a :=$
 ($fun\ n \Rightarrow if\ beq_nat\ n\ m\ then\ a\ else\ env\ n$).

Inductive $Exp : EnvType \rightarrow ExpType \rightarrow Set :=$
 | $EInt : \forall\ env,\ nat \rightarrow Exp\ env\ TInt$
 | $EBool : \forall\ env,\ bool \rightarrow Exp\ env\ TBool$
 | $EId : \forall\ env\ m\ a\ (h : env\ m = a),\ Exp\ env\ a$
 | $EAdd : \forall\ env,\ Exp\ env\ TInt \rightarrow Exp\ env\ TInt \rightarrow Exp\ env\ TInt$
 | $EGt : \forall\ env,\ Exp\ env\ TInt \rightarrow Exp\ env\ TInt \rightarrow Exp\ env\ TBool$
 | $EPair : \forall\ env\ a\ b,\ Exp\ env\ a \rightarrow Exp\ env\ b \rightarrow Exp\ env\ (TPair\ a\ b)$
 | $EFst : \forall\ env\ a\ b,\ Exp\ env\ (TPair\ a\ b) \rightarrow Exp\ env\ a$
 | $ESnd : \forall\ env\ a\ b,\ Exp\ env\ (TPair\ a\ b) \rightarrow Exp\ env\ b$.

Inductive $Command : EnvType \rightarrow EnvType \rightarrow Set :=$
 | $CAssign : \forall\ env\ m\ a,\ Exp\ env\ a \rightarrow Command\ env\ (envupdate\ env\ m\ a)$
 | $CSeq : \forall\ env1\ env2\ env3,\ Command\ env1\ env2 \rightarrow Command\ env2\ env3 \rightarrow Command\ env1\ env3$
 | $CIf : \forall\ env1\ env2,\ Exp\ env1\ TBool \rightarrow Command\ env1\ env2 \rightarrow Command\ env1\ env2 \rightarrow Command\ env1\ env2$
 | $CWhile : \forall\ env,\ Exp\ env\ TBool \rightarrow Command\ env\ env \rightarrow Command\ env\ env$.

4 Compilation

The compiler is structured as a pair of straightforward recursive functions traversing expressions and commands in the high-level language to produce lists of low-level instructions. The correctness of the generated code relies on it being linked with a memory allocator module satisfying the specification given in our previous work [13]. We call the allocator firstly to get a statically fixed-size block of memory for storing variables and an evaluation stack and, secondly, for dynamically allocating the heap storage for values of pair types. Heap datastructures generated by programs in our language can involve non-trivial sharing,

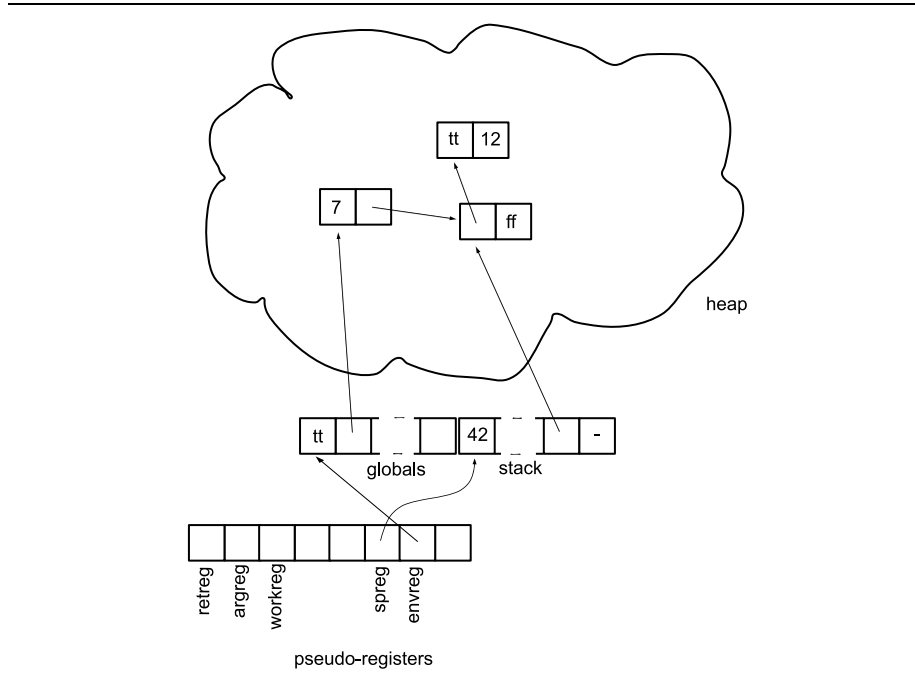


Figure 2: Memory Layout

which complicates their reclamation. We have not yet considered proving either garbage collection or any static memory management scheme, so for now just let the compiled code leak memory.

We adopt a convention of using memory locations 0 to 9 in a register-like fashion. The calling convention for the memory allocator is that a return address is passed in location 0 (*retreg*) and the size of the block requested is passed in location 1 (*argreg*); a pointer to the free block is returned in location 0. The code produced by compiling phrases of our language relies on location 6 (*envreg*) holding a pointer to the base of a contiguous block of memory, the first part of which is used to store the global variables and the remainder of which is used as a stack during the evaluation of expressions. Location 5 (*spreg*) points to the next free stack location. Figure 2 shows a view of a typical layout of the store at run-time. We have shown a situation in which the second component of the second global variable (which has type $\text{int} \times ((\text{bool} \times \text{int}) \times \text{bool})$) aliases with the value at the top of the stack. It is also worth re-emphasizing at this point that the store is *really* just a function from naturals to naturals: the intended interpretation of some of them as pointers, booleans, etc. as shown in the figure is just what we are going to formalize by giving a semantics to types.

Here is the code sequence (represented as a list of instructions, built using ‘:’ and ‘*nil*’) for pushing a particular natural number n onto the evaluation stack. We do an indirect store of the constant n to the memory location pointed to by

spreg and then increment *spreg*:

Definition *int_code* $n :=$
 $(i_move (d_ind\ spreg) (s_cst\ n)) ::$
 $(i_add (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 1)) ::$
 $nil.$

The code for boolean constants is similar, pushing 1 for true and 0 for false:

Definition *bool_code* $(b:bool) :=$
 $(i_move (d_ind\ spreg) (s_cst\ (if\ b\ then\ 1\ else\ 0))) ::$
 $(i_add (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 1)) ::$
 $nil.$

The value of a global variable n is obtained by indirection through *envreg* with an offset:

Definition *id_code* $n :=$
 $(i_move (d_ind\ spreg) (s_indo\ n\ envreg)) ::$
 $(i_add (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 1)) ::$
 $nil.$

The code for addition subtracts two from the stack pointer, adds the two values that were previously on top together, writing the result to the location pointed to by the updated stack pointer, and then increments the stack pointer again:

Definition *add_code* :=
 $(i_sub (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 2)) ::$
 $(i_add (d_ind\ spreg) (s_ind\ spreg) (s_indo\ 1\ spreg)) ::$
 $(i_add (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 1)) ::$
 $nil.$

The sequence for the greater-than test manipulates the stack in the same way, using subtraction to compare the two values. We are working with natural numbers and a subtraction operator that yields zero when the result would otherwise be negative, thus we either leave zero (representing false) or some strictly positive value, all of which we take to represent true. This encoding, or realization, of the booleans will be made more explicit when we consider the semantics of types later.

Definition *gt_code* :=
 $(i_sub (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 2)) ::$
 $(i_sub (d_ind\ spreg) (s_ind\ spreg) (s_indo\ 1\ spreg)) ::$
 $(i_add (d_imm\ spreg) (s_imm\ spreg) (s_cst\ 1)) ::$
 $nil.$

The code for creating a pair has to allocate a fresh cons cell, pop two values off the stack and write them into the fields of the new cell and finally push the address of the new cell back to the stack. It is parameterized by the starting address of the code fragment, *label*, and the entry point of the allocation routine, *alloc*.

Definition *pair_code* $label\ alloc :=$

```

(i_sub (d_imm spreg) (s_imm spreg) (s_cst 2)) ::
(i_move (d_imm argreg) (s_cst 2)) ::
(i_move (d_imm retreg) (s_cst (4 + label))) ::
(i_branch (s_cst alloc)) ::
(i_move (d_ind retreg) (s_ind spreg)) ::
(i_move (d_indo 1 retreg) (s_indo 1 spreg)) ::
(i_move (d_ind spreg) (s_imm retreg)) ::
(i_add (d_imm spreg) (s_imm spreg) (s_cst 1)) ::
nil.

```

The code for projecting from a pair decrements the stack pointer and does an indirect load via *spreg* to store the address of the relevant cons cell on the heap in *workreg*. It then pushes the value pointed to by *workreg* onto the stack:

Definition *fst_code* :=

```

(i_sub (d_imm spreg) (s_imm spreg) (s_cst 1)) ::
(i_move (d_imm workreg) (s_ind spreg)) ::
(i_move (d_ind spreg) (s_ind workreg)) ::
(i_add (d_imm spreg) (s_imm spreg) (s_cst 1)) ::
nil.

```

Definition *snd_code* := ... (similar to above) ...

We now show the function for compiling an expression *e*, which is parameterized by a starting address for the generated code, *label* and the address of the allocation routine, *alloc*.⁴ *compile_exp* returns a list of instructions and the next free code address. The *compile_expression* function wraps the compilation of a complete expression, decrementing the stack pointer at the end so that it points to the computed value, rather than the next free stack location.

Fixpoint *compile_exp* (*env*:EnvType) (*a*:ExpType) (*e*:Exp env a) (*label alloc*:nat) {*struct e*} : list instruction × nat :=

```

match e with
| EInt _ n ⇒ (int_code n, int_code_size + label)
| EBool _ b ⇒ (bool_code b, bool_code_size + label)
| EId _ n _ _ ⇒ (id_code n, id_code_size + label)
| EAdd _ e1 e2 ⇒ let (code',label') := compile_exp e1 label alloc in
                  let (code'',label'') := compile_exp e2 label' alloc in
                  (code' ++ code'' ++ add_code, add_code_size + label'')
| EGt _ e1 e2 ⇒ let (code',label') := compile_exp e1 label alloc in
                  let (code'',label'') := compile_exp e2 label' alloc in
                  (code' ++ code'' ++ gt_code, gt_code_size + label'')
| EPair _ _ _ e1 e2 ⇒ let (code',label') := compile_exp e1 label alloc in
                       let (code'',label'') := compile_exp e2 label' alloc in
                       (code' ++ code'' ++ pair_code label'' alloc,
                        pair_code_size + label'')
| EFst _ _ _ e' ⇒ let (code',label') := compile_exp e' label alloc in

```

⁴*e* is actually a typed expression in context, but the parameters *env* and *a* are Implicit, so do not appear in the recursive calls.

$$\begin{aligned}
& (code' ++ fst_code, fst_code_size + label') \\
| ESnd _ _ _ e' \Rightarrow & let (code',label') := compile_exp e' label alloc in \\
& (code' ++ snd_code, snd_code_size + label')
\end{aligned}$$

end.

Definition *compile_expression env a (e:Exp env a) label alloc :=*
let (code, label) := compile_exp e label alloc in
(code ++ ((i_sub (d_imm spreg) (s_imm spreg) (s_cst 1)) :: nil), S label).

The *compile* function compiles a command into a sequence of instructions:

Fixpoint *compile (env1 env2: EnvType) (c:Command env1 env2) (label alloc:nat) {struct c} : list instruction × nat :=*

match c with

$$\begin{aligned}
| CAssign _ m _ e \Rightarrow & let (code',label') := compile_expression e label alloc \\
& in (code' ++ (i_move (d_indo m envreg) (s_ind spreg) \\
& \quad :: nil), 1 + label')
\end{aligned}$$

$$\begin{aligned}
| CSeq _ _ _ c1 c2 \Rightarrow & let (code',label') := compile c1 label alloc in \\
& let (code'',label'') := compile c2 label' alloc in \\
& (code' ++ code'', label'')
\end{aligned}$$

| CIf _ _ b c1 c2 \Rightarrow let (code',label') := compile_expression b label alloc
in

$$\begin{aligned}
& let (code'',label'') := compile c1 (1 + label') alloc in \\
& let (code''',label''') := compile c2 (1 + label'') alloc in \\
& (code' ++ \\
& \quad (i_brz (s_ind spreg) (s_cst (1 + label'')) :: nil) \\
& \quad ++ code'' ++ (i_branch (s_cst label'') :: nil) ++ \\
& \quad code''', label''')
\end{aligned}$$

$$\begin{aligned}
| CWhile _ b c1 \Rightarrow & let (code',label') := compile_expression b label alloc in \\
& let (code'',label'') := compile c1 (1 + label') alloc in \\
& (code' ++ \\
& \quad (i_brz (s_ind spreg) (s_cst (S label'')) :: nil) \\
& \quad ++ code'' ++ (i_branch (s_cst label) :: nil), \\
& \quad 1 + label'')
\end{aligned}$$

end.

Just in case it is not immediately apparent from the above, we should explicitly remark that – even in the absence of `setcar`/`setcdr` operations – code produced by the compiler really *does* build heap datastructures that are DAGs with non-trivial sharing. For example, the program

$$X := (3, 4) ; Y := (X, X)$$

generates two `cons` cells, with both fields of the second (which is pointed to from the global variable `Y`) pointing to the first (which is also pointed to from `X`).

5 Relational Assertions

The language and compiler described in the previous sections are extremely simple. We now turn to the rather less trivial things that we would like to say about them. The next subsection gives a slightly informal account of the idea of relational specifications, which is followed by the more detailed Coq version.

5.1 Overview of relations for specification

The central idea of our approach to specifications in general, and the interpretation of types in particular, is that they are about invariance, independence, or ‘how much difference makes a difference’. With the representations we are using, there is no sensible way that even a simple statement like ‘location 74 holds a boolean’ can be interpreted as a predicate on the contents of location 74: whatever value v is stored there, it is always interpretable as either a natural number, a boolean or even a pointer. How the value is interpreted depends on how it will be used, and the difference between a piece of code that is typed assuming location 74 holds a natural and one that is typed assuming that it holds a boolean is that the latter *should only care about* whether the value is zero or not. In other words (assuming everything else is fixed), the code can have two different observable behaviours: one in the case that v is zero and the other one for *all* the non-zero values. But the notion of observable behaviour needs to be defined carefully. Consider, for example, what we might mean by saying a particular piece of code is supposed to be entered with a boolean in location 74 and will then exit with a boolean in location 74. This specification is met by a piece of code that does nothing at all, or which doubles the value in 74 (both of which implement, or realize, the identity on booleans). After the exit point however, we certainly *can* place a piece of code that behaves differently (e.g. halting or diverging) according to whether or not the initial value v was, say, 42. Clearly, we have to restrict the notion of observation (and hence observably equivalent behaviour) to take types into account. So we’ll refine our specification to say that *assuming* that the code at the exit point (the continuation) has the same behaviour for all non-zero values in location 74, *then* the code at the entry point promises to have the same behaviour whatever non-zero value is in 74 when *it* is called. We’ll write this specification something like this:

$$\text{exit} : (74 \mapsto \llbracket \text{bool} \rrbracket)^\top \vdash M \triangleright \text{entry} : (74 \mapsto \llbracket \text{bool} \rrbracket)^\top$$

Here M is the *program_fragment* that is the subject of the judgement. $\llbracket \text{bool} \rrbracket$ is a binary relation on natural numbers, defined as

$$\llbracket \text{bool} \rrbracket \stackrel{\text{def}}{=} \{(n, n') \mid (n = n' = 0) \vee (n > 0 \wedge n' > 0)\}$$

which is a (partial equivalence) relation capturing when two natural numbers are equivalent as boolean values. $(74 \mapsto \llbracket \text{bool} \rrbracket)$ is then a binary relation on *states*, relating two states whenever they hold values in location 74 that are

$\llbracket \text{bool} \rrbracket$ -related. More generally, if $r \subseteq \mathbb{N} \times \mathbb{N}$ and $x \in \mathbb{N}$,

$$(x \mapsto r) \stackrel{\text{def}}{=} \{(s, s') \mid (s x, s' x) \in r\}$$

Finally, the ‘perp’ operator, $(\cdot)^\top$, takes a binary relation on *states* to one on pairs of programs and code pointers. If $R \subseteq \text{state} \times \text{state}$, then R^\top relates two such pairs just when they behave equivalently whenever they are started in states that are R -related. The notion of equivalent behaviour we use here is *equi-termination*:

$$R^\top \stackrel{\text{def}}{=} \{(p, l), (p', l') \mid \forall (s, s') \in R, \text{terminates } p \text{ } s \text{ } l \Leftrightarrow \text{terminates } p' \text{ } s' \text{ } l'\}$$

One can think of the elements of R^\top as a collection of ‘test contexts’ for R . The meaning of the entire judgement above is then

$$\begin{aligned} & \forall p \text{ } p', \text{ program_extends_fragment } p \text{ } \mathbf{M} \\ & \Rightarrow \text{program_extends_fragment } p' \text{ } \mathbf{M} \\ & \Rightarrow ((p, \mathbf{exit}), (p', \mathbf{exit})) \in (74 \mapsto \llbracket \text{bool} \rrbracket)^\top \\ & \Rightarrow ((p, \mathbf{entry}), (p', \mathbf{entry})) \in (74 \mapsto \llbracket \text{bool} \rrbracket)^\top. \end{aligned}$$

Reducing various notions of observational equivalence to equitermination in a set of contexts is common in denotational semantics, but some readers might have expected to see something more direct, such as that two executions started at **entry** with $\llbracket \text{bool} \rrbracket$ -related values in 74 will both reach **exit** with similarly related values. Such a property is actually implied by the definition we have given, because amongst the chosen tests are ones that make all the observable distinctions one wishes to be able to make at intermediate points. For example, in $(74 \mapsto \llbracket \text{bool} \rrbracket)^\top$ there are, amongst other things, pairs of contexts that both halt starting at **exit** whenever the value in 74 represents true and diverge whenever it represents false, and which have varying termination behaviours elsewhere. By considering such contexts, one can deduce that if one jumps to **entry** with two values representing the same boolean, then either both computations diverge, both computations halt, or both computations reach **exit** with two values that represent the same boolean in location 74. For simple first-order types such as we consider here, the ‘double negation’ formulation deals neatly with possible divergence; for more complex types it has other advantages over direct-style versions, involving, for example, admissibility.

As mentioned above, one kind of relation with which we will work is relations on *stores*. To be able to reason locally and modularly about how such relations are either affected or left unchanged by the execution of pieces of code it is necessary to have some handle on which *part* of the store a given relation depends upon. For example, the store relation $(x \mapsto r)$ clearly only ‘looks at’ location x (on both sides); one consequence is that if we start with two states $(s_0, s'_0) \in (x \mapsto r)$ and make arbitrary updates to them in locations *other* than x , yielding new states s_1, s'_1 , then $(s_1, s'_1) \in (x \mapsto r)$ too. Indeed, we capture the notion of which locations a relation depends upon (its ‘support’) in terms of invariance under change. If $L \subseteq \mathbb{N}$ and s_0 and s_1 are states, then define

$$s_0 \sim_L s_1 \stackrel{\text{def}}{=} \forall x \in L, s_0 x = s_1 x$$

(In Coq, we represent subsets by maps into *Prop* and define *equpto* : (*nat* → *Prop*) → *state* → *state* → *Prop* to mean \sim .)

Now, though we shall refine this definition shortly, say that a pair of sets of locations (L, L') supports a relation $R \subseteq \textit{state} \times \textit{state}$ when

$$\forall (s_0, s'_0) \in R, \forall (s_1, s'_1), (s_0 \sim_L s_1) \wedge (s'_0 \sim_{L'} s'_1) \Rightarrow (s_1, s'_1) \in R$$

In other words, if one starts with two states in the relation then any modifications outside the support yield another pair of states in the relation. If R_1 is supported by (L_1, L'_1) and R_2 by (L_2, L'_2) , then define a form of separating conjunction [37] by

$$R_1 \otimes R_2 \stackrel{\textit{def}}{=} \begin{cases} R_1 \cap R_2 & \text{if } L_1 \cap L_2 = \emptyset \text{ and } L'_1 \cap L'_2 = \emptyset \\ \emptyset & \text{otherwise} \end{cases}$$

So two states are in $R_1 \otimes R_2$ when they are in both R_1 and R_2 and the supports are disjoint. It is easy to see that $R_1 \otimes R_2$ is supported by $(L_1 \cup L_2, L'_1 \cup L'_2)$. When working with concrete supported relations, the separating conjunction allows many specifications to be written very concisely, since it abbreviates a host of ‘absence of aliasing’ conditions that would be explicit in a traditional Hoare-logic style assertion language. The separating conjunction will prove even more useful when we reason about modules: the private invariants of modules will be captured by existentially quantifying over supported relations about which clients know nothing *except* that their support is disjoint from that of the client’s own store.

Unfortunately, the above notion of support is slightly too weak to be useful. For example, consider a relation (*List74*) expressing that two states have equal linked lists of integers in location 74. Assuming the usual representation, this will relate s and s' when *either* $s74$ and $s'74$ are both zero, *or* they are both non-zero, $s(s74) = s'(s'74)$ – capturing the equality of the first elements – and (inductively) there are equal linked lists starting at $(s74)+1$ and $(s'74)+1$. Thus the sets of locations that get looked at to test whether two stores are in (*List74*) are not constant; they depend on the contents of those stores. So we have to replace sets of locations $L \subseteq \mathbb{N}$ with functions $A : \textit{state} \rightarrow \mathbb{P}(\mathbb{N})$. We restrict attention to *accessibility maps* (first introduced in [15]), those A for which

$$\forall s s', s \sim_{A(s)} s' \Rightarrow A s = A s'$$

The accessibility map condition intuitively says that A ‘supports itself’, and has the effect of making the relation \sim_A , defined by $s \sim_A s' \Leftrightarrow s \sim_{A(s)} s'$ an equivalence relation.

We will build our specifications out of state relations supported by pairs of accessibility maps, making much use of (a suitable generalization of) the separating conjunction described above.

5.2 Relations for specifications, formally

In this section we present the formal definitions of the relations and operations on relations with which we will be working. These are slightly more complex

than those introduced in the semi-formal discussion of the previous section. The first extra complexity is that relations on *states* and naturals will both generally depend on a pair of programs, because they will involve sets of code pointers that have particular behaviours, which only makes sense relative to some program. The second additional bit of structure we shall need is a form of admissibility property to justify recursive reasoning about program fragments and recursive definitions of relations. The way we deal with this is to work with relations that are the limit of a sequence of k -indexed approximants, where the natural number k represents a number of steps in the operational semantics [9, 4, 12]. In other words, our notions of ‘equivalent’ are expressed as the limit of ‘indistinguishable for up to k steps’ as k goes to ω . As more steps allow more distinctions to be made, it is natural to work with indexed relations that are antimonotonic in k .

Here is the program- and step-indexed definition of relations on natural numbers. A *Natrel* is a record containing two fields. The first, *NRrel*, is the carrier: the relation itself. The second, *NRcond*, is a *proof* that the relation is antimonotonic in the index k :

```
Record Natrel : Type :=
  mkNR {NRrel :> program → program → nat → nat → Prop ;
        NRcond : ∀ p p' j k x x', j < k → NRrel p p' k x x' → NRrel p p' j
          x x'}.

```

The carrier of a *Natrel* is a relation, expressed in curried form, on pairs of programs and triples of natural numbers. The first two arguments are the left and right programs, p and p' . The third argument is the step index, $k:nat$. The fourth and fifth arguments are the natural numbers on the left and the right, x and x' . Since all values in our model are naturals, some *Natrels* we use will express properties of x and x' viewed as code pointers, whilst others will treat them as heap pointers, naturals or booleans.

There’s a natural partial order and equality relation on *Natrels*, which satisfy the usual axioms:

```
Definition Natrelleq (R1 R2 : Natrel) :=
  ∀ p p' k n n', R1 p p' k n n' → R2 p p' k n n'.

```

```
Definition Natreleq Na1 Na2 := Natrelleq Na1 Na2 ∧ Natrelleq Na2 Na1.

```

We can also lift non-indexed relations on naturals to *Natrels*:

```
Definition Natrel_lift (R : nat → nat → Prop) : Natrel.
  intro R.
  refine ( mkNR (fun p p' k => R) _ ).
  tauto.

```

Defined.

The definition of *Natrel_lift* makes use of the interactive proof language of Coq: the *refine* tactic is used to define the carrier of the lifted relation, leaving a hole (the underscore) for the *NRcond* proof component that is needed to show that the monotonicity requirement is satisfied. The proof is then filled in interactively, in this simple case just by calling the automatic tactic *tauto*. Several

of our definitions are of similar form, though generally with less trivial interactive proofs; we will henceforth usually elide the proofs, presenting just the application of the *refine* tactic.

Here is the definition of the type *Accrel* of supported, indexed relations on *states*. An element of *Accrel* is a record comprising the relation itself (*ARel*), two accessibility maps (*ARacc* and *ARacc'*), a proof (*ARcond*) that the accessibility maps are accessibility maps and do support the relation, and a proof (*ARindexed*) that the relation is antimonotonic in the step index:

```
Record Accrel : Type :=
  mkAR { ARrel :> program → program → nat → state → state → Prop;
        ARacc : state → state → nat → Prop;
        ARacc' : state → state → nat → Prop;
        ARcond : ∀ p p' k s0 s0' s1 s1',
          (ARrel p p' k s0 s0') → equpto (ARacc s0 s0') s0 s1
            → equpto (ARacc' s0 s0') s0' s1' →
          (ARrel p p' k s1 s1') ∧
          (∀ n, ARacc s0 s0' n ↔ ARacc s1 s1' n) ∧
          (∀ n, ARacc' s0 s0' n ↔ ARacc' s1 s1' n);
        ARindexed : ∀ p p' j k s s', j < k → ARrel p p' k s s' → ARrel p
  p' j s s'
  }.
```

The carrier relates two programs (*p* on the left, *p'* on the right), a step index and two states (*s* on the left, *s'* on the right). *ARacc* is the accessibility map giving the locations that are relevant on the left (i.e. in state *s*), whilst *ARacc'* is associated with the state on the right. Note that these are actually dependent on two states, rather than one as in our earlier overview; this turns out to be technically smoother, though we won't really exploit the extra generality in the present paper. The *ARcond* condition looks complex because it combines the conditions on both accessibility maps and on the relation, but if we ignore the program and index dependence then it reads as follows: if we start with two states *s0* and *s0'* in the relation, and *s1* and *s1'* are two other states, with *s1* equal to *s0* up to the left hand accessibility map (applied to the states we started with), and *s1'* equal to *s0'* up to the right hand accessibility map, then three things happen. Firstly, *s1* and *s1'* are also in the relation; this says that the accessibility maps do support the relation. Second, the left hand accessibility map yields the same set of locations when given *s0* and *s0'* as arguments as it does when given *s1* and *s1'*; this is the accessibility map condition. Finally, the same is true of the right hand accessibility map. Compared with our earlier slightly informal account, we have not only parameterized accessibility maps by both states rather than one, but we have also tied the maps and the relations closer together by only requiring the accessibility map condition to hold for states that are in the relation, which is why we do not define accessibility maps separately (as we did in our earlier work [13]).

Accrels also have a natural partial order and equality. Note that the order involves an implication between the carrier relations *and* a containment the

other way between the accessibility maps:

Definition *Accrelleq* (*Na1 Na2 : Accrel*) :=
 $\forall p p' k s s', Na1 p p' k s s' \rightarrow$
 $((Na2 p p' k s s') \wedge$
 $(\forall n, ARacc Na2 s s' n \rightarrow ARacc Na1 s s' n) \wedge$
 $(\forall n, ARacc' Na2 s s' n \rightarrow ARacc' Na1 s s' n))$.

Definition *Accreleq* *Na1 Na2* := *Accrelleq Na1 Na2* \wedge *Accrelleq Na2 Na1*.

Here are some basic *Accrels*:

Definition *Accrel_const* (*q : Prop*) (*a : nat \rightarrow Prop*) (*a' : nat \rightarrow Prop*) : *Accrel*.
intros.
refine (*mkAR* (*fun p p' k s s' \Rightarrow q*)
 $(fun s s' \Rightarrow a)$
 $(fun s s' \Rightarrow a')$
 $-$
 $-$).

...
 Defined.

Definition *Emptyrel* (*q : Prop*) := *Accrel_const q* (*fun n \Rightarrow False*) (*fun n \Rightarrow False*).

Definition *Toprel* := *Emptyrel True*.

Lemma *Accrelleq_Toprel* : $\forall r, Accrelleq r Toprel$.

We can now define some interesting constructions on *Accrels*. The first is the ordinary additive conjunction, *RelConj*, which does not require disjoint supports:

Definition *nunion* (*a1 a2 : nat \rightarrow Prop*) *n* := (*a1 n*) \vee (*a2 n*).

Definition *RelConj* (*Ar1 Ar2 : Accrel*) : *Accrel*.

intros.
refine (*mkAR* (*fun p p' k s s' \Rightarrow (Ar1 p p' k s s') \wedge (Ar2 p p' k s s')*)
 $(fun s s' \Rightarrow (nunion (ARacc Ar1 s s') (ARacc Ar2 s s')))$
 $(fun s s' \Rightarrow (nunion (ARacc' Ar1 s s') (ARacc' Ar2 s s')))$
 $-$
 $-$).

...
 Defined.

Lemma *RelConj_unit* : $\forall R, Accreleq (RelConj R Toprel) R$.

Lemma *RelConj_comm* : $\forall R1 R2, Accreleq (RelConj R1 R2) (RelConj R2 R1)$.

Lemma *RelConj_assoc* : $\forall R1 R2 R3, Accreleq (RelConj (RelConj R1 R2) R3) (RelConj R1 (RelConj R2 R3))$.

The second is the separating conjunction, *RelTensor*, introduced previously. This is also associative and commutative with *Toprel* as unit (amongst other properties whose formal statements we elide).

Definition $ndisj (a1 a2 : nat \rightarrow Prop) := \forall n, \sim(a1 n \wedge a2 n)$.

Definition $RelTensor (Na1 Na2 : Accrel) : Accrel$.

intros.

refine ($mkAR (fun p p' k s s' \Rightarrow (Na1 p p' k s s') \wedge (Na2 p p' k s s') \wedge$
 $(ndisj (ARacc Na1 s s') (ARacc Na2 s s')) \wedge$
 $(ndisj (ARacc' Na1 s s') (ARacc' Na2 s s'))$
 $(fun s s' \Rightarrow (nunion (ARacc Na1 s s') (ARacc Na2 s s'))$
 $(fun s s' \Rightarrow (nunion (ARacc' Na1 s s') (ARacc' Na2 s s'))$
 $-$
 $-)$).

...

Defined.

The $ptsto$ relation is like the ‘points to’ predicate of separation logic. It relates two states s and s' just when the values stored in location l in s and in location l' in s' are related by the $Natrel$, r . Obviously, it is supported by the accessibility map that always returns the singleton set $\{l\}$ on the left, and by the constant $\{l'\}$ map on the right:

Definition $ptsto (l l':nat) (r : Natrel) : Accrel$.

intros.

refine ($mkAR (fun p p' k s s' \Rightarrow r p p' k (s l) (s' l')$
 $(fun s s' n \Rightarrow n=l)$
 $(fun s s' n \Rightarrow n=l')$
 $- -)$).

...

Defined.

Notation " $[m, n] \mid \Rightarrow r$ " := $(ptsto m n r)$ (at level 80).

Notation " $m \mid \rightarrow r$ " := $(ptsto m m r)$ (at level 80).

The definition of the ‘perp’ operation is the place where we make careful use of the step-indexing.

Definition $Perp (R:Accrel) : Natrel$.

intro R.

refine(
 $mkNR (fun p p' k l l' \Rightarrow \forall j s s', j < k \rightarrow R p p' j s s' \rightarrow$
 $((kstepterm j p s l) \rightarrow (terminates p' s' l')) \wedge$
 $((kstepterm j p' s' l') \rightarrow (terminates p s l))))$
 $-)$.

...

Defined.

Note the way in which the indices are used: two labels l, l' are in $Perp R$ at index k just when for any strictly smaller j , and states related by R at index j , if jumping to l terminates within j steps, then jumping to l' terminates in *some* number of steps, and vice versa.⁵ The limit of $Perp R$ agrees with the definition

⁵A more naive definition might end up only relating computations that terminated in exactly the same number of steps, which is not what we want at all.

of R^\top we gave earlier, in that

$$R^\top = \{(p, l), (p', l') \mid \forall k, \text{Perp } p \ p' \ k \ l \ l'\}.$$

As one would expect, Perp is contravariant:

Lemma $\text{Accrelleq_Perp} : \forall R1 \ R2,$

$$\text{Accrelleq } R1 \ R2 \rightarrow \text{Natrelleq } (\text{Perp } R2) (\text{Perp } R1).$$

We define ‘#’ as Coq notation for RelTensor , and ‘!’ as notation for Perp .⁶

6 Specification of Allocation

This section briefly recalls the specification of a memory allocator module from our previous work [13]; see that paper for a fuller account. The module has three entry points: one for *initialization*, one to *allocate* a block and one to *deallocate* a block. We just detail the specification for allocation here, since we will not be making calls to the other routines.

We have already described the calling convention of the allocator: the return address is passed in location 0, the size of the requested block is passed in 1, and a pointer to the start of the allocated block is returned in 0. But what is the formal contract between the allocator and its clients?

After the allocator has been initialized, the heap will, like Gaul, be divided into three parts. First, the pseudo-registers 0 to 9; second, the part belonging to the allocator; and finally, the part belonging to the rest of the program. Ownership of blocks of memory is transferred between the allocator and its clients by calls to *alloc* and *dealloc*. The allocator promises not to (observably) read or write the part belonging to the clients. In return, the clients promise not to read or write the part belonging to the allocator and not to care about either the location or the initial contents of the blocks they are given.

We capture this intent by saying that a module M_a with entry point *alloc* meets the specification of an allocator if there exists a supported relation Ra – the allocator’s private invariant – such that for all programs p, p' extending M_a , for all k , for all Rc (client invariants) and for all n (block sizes),

$$(R_al \ Ra \ n \ Rc) \ p \ p' \ k \ \text{alloc} \ \text{alloc}$$

where

Definition $R_aret \ (n:nat) : \text{Accrel}.$

intro.

$$\begin{aligned} \text{refine } (mkAR \ (&fun \ p \ p' \ k \ s \ s' \Rightarrow s \ 0 > 9 \wedge s' \ 0 > 9) \\ &(\text{fun } s \ s' \ l \Rightarrow (l = 0) \vee (l \geq s \ 0 \wedge l < n + s \ 0)) \\ &(\text{fun } s \ s' \ l \Rightarrow (l = 0) \vee (l \geq s' \ 0 \wedge l < n + s' \ 0))) \end{aligned}$$

–

⁶These are really defined as notation for operations on an inductive type of relation-denoting *expressions* that we use in doing proofs by computational reflection [25, 19], but we gloss over that technicality here.

-).

...

Defined.

and

Definition $R_al (Ra:Accrel) n (Rc:Accrel) :=$

$$\begin{aligned} & ! ((0 \mid\!-\!> !(R_aret\ n \# T_rel\ (1\ to\ 4)\ (1\ to\ 4) \# Rc \# Ra \# E)) \\ & \quad \# (1 \mid\!-\!> (Natrel_lift\ (fun\ l\ l' \Rightarrow l = n \wedge l' = n))) \\ & \quad \# T_rel\ (2\ to\ 4)\ (2\ to\ 4) \# Rc \# Ra \# E). \end{aligned}$$

This says that two calls to *alloc* are guaranteed to behave the same whenever they are started in a pair of initial states s, s' that are related by all of the following disjoint relations:

- Ra , so the allocator's invariant holds before the call;
- Rc , so the client's invariant holds before the call
- $T_rel\ (2\ to\ 4)\ (2\ to\ 4)$. This is the 'true' relation with support $\{2, 3, 4\}$ on both sides, so these locations are not looked at by the allocator;
- location 1 holds the value n on both s and s'
- The contents of location 0 on the two sides are code pointers that promise to behave the same if *they* are started in states related by
 - Ra , so the allocator invariant holds *after* the call;
 - Rc , so the client invariant holds after the call;
 - $T_rel\ (2\ to\ 4)\ (2\ to\ 4)$, so these locations are not looked at by the return addresses, i.e. they may be modified by the allocator;
 - $R_aret\ n$, which expresses that location 0 on each side points to a block of size n that doesn't overlap the pseudo-registers

In previous work, we described a very naive allocation module that satisfies this specification; we have since verified that a slightly less trivial implementation that uses a free list satisfies the same spec.

7 Formalizing and Verifying Type Soundness

This section presents the actual type soundness theorem for the simple compiler.

We start with a useful construction on *Accrels*:

Definition $pex\ (l\ l':nat)\ (h: nat \rightarrow nat \rightarrow Accrel) : Accrel.$

intros.

$$\begin{aligned} refine\ (mkAR\ (fun\ p\ p'\ k\ (s\ s':state) \Rightarrow \\ & h\ (s\ l)\ (s'\ l')\ p\ p'\ k\ s\ s' \wedge \\ & \neg\ (ARacc\ (h\ (s\ l)\ (s'\ l'))\ s\ s'\ l) \wedge \\ & \neg\ (ARacc'\ (h\ (s\ l)\ (s'\ l'))\ s\ s'\ l')) \end{aligned}$$

$$\begin{aligned} & (\text{fun } s \ s' \ n \Rightarrow (n = l) \vee (\text{ARacc } (h \ (s \ l) \ (s' \ l')) \ s \ s' \ n)) \\ & (\text{fun } s \ s' \ n \Rightarrow (n = l') \vee (\text{ARacc}' \ (h \ (s \ l) \ (s' \ l')) \ s \ s' \ n)) \\ & - \text{ -}). \end{aligned}$$

...
Defined.

Notation " $\text{peexists}' \ (x, y) \ @ \ (l, m), g$ " := $(\text{pex } l \ m \ (\text{fun } x \ y \Rightarrow g))$
(at level 200, ...).

The *peexists* operation captures a pattern of existential quantification over values in the store that is common in defining *Accrels*, viz.

$$\exists x \ x', ([l, l'] \models [x, x']) \ \# \ g(x, x')$$

i.e. s and s' are in the relation just when locations l and l' in s, s' respectively hold some values x, x' , and the relation $g(x, x')$ holds on some disjoint part of the store.⁷

Now we can inductively define the semantics of types in our source language as relations over values and stores of the low-level machine:

Fixpoint $\text{typerefssem} \ (t : \text{ExpType}) \ (l \ l' : \text{nat}) \ \{ \text{struct } t \} : \text{Accrel} :=$
match t *with*
| $\text{TInt} \Rightarrow \text{Emptyrel} \ (l = l')$
| $\text{TBool} \Rightarrow \text{Emptyrel} \ ((l = 0) \wedge (l' = 0) \vee (l \neq 0) \wedge (l' \neq 0))$
| $\text{TPair } a \ b \Rightarrow \text{peexists} \ (va, va') \ @ \ (l, l'), \ \text{peexists} \ (vb, vb') \ @ \ (S \ l, S \ l'),$
 $\text{RelConj} \ (\text{typerefssem } a \ va \ va') \ (\text{typerefssem } b \ vb \ vb')$
end.

Two states are related by $\text{typerefssem } t \ l \ l'$ just when l and l' are equal as values of type t in those states. So

- Two values are equal as natural numbers just when they are equal, independent of what the states are.
- Similarly, two values are equal as booleans just when they are in $\llbracket \text{bool} \rrbracket$, again independent of the states.
- l and l' are equal as values of type $\text{TPair } a \ b$ in states s and s' when the cons cells pointed to by l in s and by l' in s' have first components that are equal as values of type a and second components that are equal as values of type b . The use of the additive conjunction, RelConj , allows the storage used by the values pointed to in the first and second components to share with one another, but note that we have not allowed sharing with the cell itself (because of the separation built into the definition of *peexists*).

The support part of the *Accrel* returned by typerefssem will, of course, follow chains of pointers to capture what parts of the two heaps are looked at in judging relatedness; this will be a function of both the actual values in the heap and the type at which we are comparing them.

⁷The reason for this definition, suggested to us by Matthew Parkinson, is that our notion of supported relation does not (easily) allow general existential quantification. *peexists* is a form of quantification in which the witness, if any, is uniquely determined, fixing the support.




Figure 3: Typerefsem Examples

Figure 3 shows some examples. On the left is a positive example: the stores s and s' are related by

$$\text{typerefsem } (TInt^{**} TBool)^{**} (TInt^{**} TBool) \ 74 \ 56 \ s \ s'$$

because location 74 in s and location 56 in s' are both interpretable as holding the value $((42, \text{true}), (42, \text{true}))$. The right hand side of the figure shows a non-example: one might (rather deviously) think that 12 in s and 56 in s' are both interpretable as the pair of pairs of integers $((12, 12), (12, 12))$, but the punning representation on the left is actually ruled out by the use of separation in the definition of *peexists*.⁸ The essence of the positive example is that our semantics of types whose representation involves pointers makes explicit the independence of the represented value from the actual pointer values and potential sharing used in the representation.

Having defined the relational interpretation of each *ExpType*, we need to define the relational interpretation of an *EnvType*, capturing the notion of equality on the vector of globals, the evaluation stack and the heap (as was illustrated in Figure 2). This is built up by induction over the length of the environment (globals+stack), taking care to maintain the separation between individual environment entries and between the environment and the heap, whilst allowing sharing within the heap. To this end, we first define a function that builds an *Accrel* by folding *peexists* over the vectors of length n (starting at locations l and l' , respectively) in the two states, additively conjoining all the results of applying

⁸A small tweak to the definition of *typerefsem* would admit such encodings. This would actually work for our simple compiler, but would break if we later added copying garbage collection. The important thing to note is that our formalism allows precise control over such details.




Figure 4: Typesem Example

a function f to the existentially quantified values stored in the corresponding elements of the vectors.

```

Fixpoint pexconj (n m l l' : nat) (f : nat → nat → nat → Accrel) (r:Accrel)
{struct n} : Accrel :=
  match n with
  | 0 ⇒ r
  | S n' ⇒ pex l l' (fun (x1 x2 : nat) ⇒ pexconj n' (S m) (S l) (S l') f
    (RelConj (f x1 x2 m) r))
  end.

```

Using $pexconj$ to fold $typerefsem$, we can define the relational interpretation of an environment type env of length n , starting at locations $base$ and $base'$:

```

Definition typesem n (env:EnvType) base base' :=
  pexconj n 0 base base' (fun x1 x2 m ⇒ typerefsem (env m) x1 x2) Toprel.

```

Figure 4 shows an example of two states related by $(typesem\ 4\ env\ 20\ 22)$, where we assume env maps offsets to types as follows:

- 0 \mapsto $TBool$
- 1 \mapsto $(TInt**TBool)**(TInt**TBool)$
- 2 \mapsto $TInt$
- 3 \mapsto $TInt**TBool$

Having defined relations accounting for the structure of the environment and heap, we now need to define the contracts for the pieces of compiled code that come from typed expressions and commands in the source language. These will involve $Perps$, expressing that jumping to certain pairs of addresses will yield equitermination provided that the initial states are in a certain relation, which will involve a $typesem$ for the heap plus something about the pseudo-registers being suitably related.

Here is the formal definition of the prerule for commands and expressions that expect to be entered with $envsize$ global variables typed according to the

EnvType env and an empty stack (which is allowed to grow up to *maxstack* + 1 locations), assuming allocators related by *Ra* and starting addresses *envbase* and *envbase'* for the two environments. *Ro* is an arbitrary relation on the parts of memory which do *not* belong to either the allocator or the compiled code:

Definition *R_comp* (*Ra Ro : Accrel*) *envsize env envbase envbase' maxstack* :=
 let *sp* := (*envsize* + *envbase*) in
 let *sp'* := (*envsize* + *envbase'*) in
 ! (*T_rel* (0 to 4) (0 to 4)
 # (*envreg* |-> (*Natrel_lift* (*fun l l' => l = envbase* \wedge *l' = envbase'*)))
 # (*spreg* |-> (*Natrel_lift* (*fun l l' => l = sp* \wedge *l' = sp'*)))
 # *typesem envsize env envbase envbase'*
 # *T_rel* (*sp* to (*maxstack* + *sp*)) (*sp'* to (*maxstack* + *sp'*))
 # *Ra* # *Ro* # *E*).

As explained in Section 5.1, however, the entry point of the code for a command or expression will only be in the *R_comp* corresponding to the pre-type under the assumption that the code at its exit point satisfies a suitable relation for the post-type. For commands, which always expect to be entered with an empty stack, the assumption on the exit will just be another instance of *R_comp*. For expressions, however, we expect a value of a particular type to be left on the stack. That's expressed by the following variant of *R_comp* which adds the requirement that there be values related by the interpretation of type *t* on the stacks:

Definition *R_comp_exp_post* (*Ra Ro : Accrel*) *envsize (env:EnvType) t envbase envbase' maxstack* :=
 let *sp* := (*envsize* + *envbase*) in
 let *sp'* := (*envsize* + *envbase'*) in
 ! (*T_rel* (0 to 4) (0 to 4)
 # (*envreg* |-> (*Natrel_lift* (*fun l l' => l = envbase* \wedge *l' = envbase'*)))
 # (*spreg* |-> (*Natrel_lift* (*fun l l' => l = S sp* \wedge *l' = S sp'*)))
 # *typesem* (*S envsize*) (*envupdate env envsize t*) *envbase envbase'*
 # *T_rel* ((*S sp*) to (*maxstack* + *sp*)) ((*S sp'*) to (*maxstack* + *sp'*))
 # *Ra* # *Ro* # *E*).

Note the way in which the first stack location is treated as if it were the (*envsize* + 1)-th variable. At last, we can give the type soundness theorems for our compiler. There is one for expressions and one for commands. Here is the one for expressions:

Theorem *comp_expression_thm* :
 \forall (*alloc alloc' : nat*) (*Ra Ro : Accrel*) (*p p' : program*)
 (*envbase envbase' : nat*) (*envsize : nat*)
 (*env : EnvType*) (*a : ExpType*) (*e : Exp env a*),
 \forall (*h_env : env_ok_exp e envsize*)
 (*maxstack : nat*) (*h_stack : stack_ok_exp e maxstack*)
 (*k label label' : nat*)

$$\begin{aligned}
& (code:list\ instruction)\ (code':list\ instruction) \\
& (hcode : code = fst\ (compile_exp\ e\ label\ alloc)) \\
& (hcode' : code' = fst\ (compile_exp\ e\ label'\ alloc')), \\
& program_extends_fragment\ p\ (fragfromlist\ code\ label) \\
& \rightarrow program_extends_fragment\ p'\ (fragfromlist\ code'\ label') \\
& \rightarrow (\forall\ n\ Rc,\ (R_al\ Ra\ n\ Rc)\ p\ p'\ k\ alloc\ alloc') \\
& \rightarrow (R_comp_exp_post\ Ra\ Ro\ envsize\ env\ a\ envbase\ envbase'\ maxstack)\ p\ p' \\
& k\ (length\ code\ +\ label)\ (length\ code'\ +\ label') \\
& \rightarrow (R_comp\ Ra\ Ro\ envsize\ env\ envbase\ envbase'\ maxstack)\ p\ p'\ (1 + k)\ label \\
& label'.
\end{aligned}$$

And here is the theorem for commands:

Theorem *comp_thm* :

$$\begin{aligned}
& \forall\ (alloc\ alloc' : nat)\ (Ra\ Ro : Accrel)\ (p\ p' : program) \\
& \quad (envbase\ envbase' : nat)\ (envsize : nat) \\
& \quad (env1\ env2 : EnvType)\ (c : Command\ env1\ env2), \\
& \forall\ (h_env : env_ok\ c\ envsize) \\
& \quad (maxstack : nat)\ (h_stack : stack_ok\ c\ maxstack) \\
& \quad (label\ label' : nat) \\
& \quad (code : list\ instruction)\ (code' : list\ instruction) \\
& \quad (hcode : code = fst\ (compile\ c\ label\ alloc)) \\
& \quad (hcode' : code' = fst\ (compile\ c\ label'\ alloc')), \\
& program_extends_fragment\ p\ (fragfromlist\ code\ label) \\
& \rightarrow program_extends_fragment\ p'\ (fragfromlist\ code'\ label') \\
& \rightarrow \forall\ k, \\
& \quad ((\forall\ n\ Rc,\ (R_al\ Ra\ n\ Rc)\ p\ p'\ k\ alloc\ alloc') \\
& \quad \rightarrow (R_comp\ Ra\ Ro\ envsize\ env2\ envbase\ envbase'\ maxstack)\ p\ p'\ k\ (length \\
& \quad code\ +\ label)\ (length\ code'\ +\ label')) \\
& \quad \rightarrow (R_comp\ Ra\ Ro\ envsize\ env1\ envbase\ envbase'\ maxstack)\ p\ p'\ (1 + k) \\
& \quad label\ label').
\end{aligned}$$

Let's look at the theorem for commands, *comp_thm*, first to see what it says. Ignoring the checks that *maxstack* is sufficiently large and that only variables numbered less than *envsize* are used, the essence is the following:

- For any *Command*, *c*, typeable with a pretype *env1* and a posttype *env2*,
- if we compile *c* twice, once starting at *label* and once starting at *label'*, linking the first with an allocator at *alloc* and the second with an allocator at *alloc'*,
- then if we put those bits of code into contexts such that *alloc* and *alloc'* are equivalent memory allocators (according to the specification of allocation) and the exit points of the two bits of compiled code behave equivalently in all states related by the interpretation of the posttype *env2*

- then the entry points of the bits of compiled code behave equivalently in all states related by the interpretation of the pretype *env1*.

The above captures lots of information about what the behaviour of the code compiled from *c* can depend upon. For example, it is independent of where the compiled code is placed, where the allocator is, where the environment is stored, what addresses the allocator returns and what their original contents are, how booleans or pairs are represented in the initial state, and so on.

We have a strong (extensional) form of memory safety, showing that the compiled code doesn't *observably* read or write any locations that it shouldn't. The preservation of any *Ro*, for example, means that storage disjoint from both the allocator's private store and the while-program's heap neither affects the behaviour of code compiled from a command, because the poststates will be equivalent for *any* initial *Ro*, nor is affected by it, because any *Ro* (in particular extensions of singleton relations) will be preserved. Note that the notion of independence really is more liberal than a naive intensional one: code that messes with unowned memory locations but restores them before exit meets the specification, as does code that builds literally different, but equivalent according to the types, heap structures according to the contents of unowned memory. See [14] for more on how preservation of sets of relations can express not only complete independence, but also read-only and write-only effects on particular storage locations.

The theorem for expressions is similar to that for commands, except that the environments in the pre and post relations stay the same and the postrelation assumes that there is a value of type *a* on the stack.

The proofs of the above theorems are basically inductions over the source language, with each case being dealt with by forward Hoare-style reasoning, similar to that of our previous work on allocation. The indexing structure on relations is used just in the case for *CWhile*, which uses mathematical induction: we assume that the label at the entry of the loop satisfies the desired relation at index *k*, and then examine the loop body to show that the entry then satisfies the same relation to index *k* + 1.

The total size of the Coq development is around 8500 lines, which includes the low-level machine, metatheory of supported relations, the language and compiler and the actual proofs. There is scope for significant simplification here though. We are still comparatively inexperienced Coq users and were developing much of the theory *in* the prover, rather than doing post-hoc formalization of a completed paper development, so there is a lot of 'junk DNA' in those 8500 lines. We use little automation so far, but the proof scripts for particular segments of assembly code are already about an order of magnitude shorter than in our earlier efforts, averaging around 20 (instead of 200) lines of proof for each assembly language instruction. The only lemma that turned out to be tricky to prove was one used in the *CAssign* case: the intertwined mix of additive and multiplicative conjunctions in the definition of *typesem* made getting the right induction hypothesis for showing the soundness of strong updates harder than one might have expected.

8 Discussion

We have presented a semantic interpretation of the types of a high-level language as relations over configurations of a low-level machine, and used that to formulate and prove type correctness of a compiler.

One might fairly characterize this work as ‘saying complicated things about simple programs’. There seems to be more sophistication in the specifications than in the original code. This is partly because our framework is a little more general than is really needed for this simple first-order problem (the step-indexing structure and treatment of code pointers is really aimed at treating higher-order programs and recursive specifications and types). But types are, inherently, surprisingly non-trivial things. The entire point of types is compositionality: we carve out a boundary between one piece of the program and the rest and write down a contract comprising assumptions and guarantees about interactions across that boundary. Although those contracts abstract behavioural details of any particular piece of code, they also have to capture the space of all pieces of code that would do just as well, which involves making explicit many details about the interface that are implicit in the code. When one adds all the parameterization necessary to achieve compositionality, it is easy for an abstract description of what a short piece of code is supposed to do to be longer than the code itself. The trick, of course, is to choose the primitives, place the boundaries and compose systems in such a way that the interfaces of the composites don’t get still more complex (or at least, not too quickly).

A crucial feature of our approach is that the semantics of a type is a supported relation on low-level stores that makes no further reference to the source language type we started with. One might have instead defined a ‘represents’ relation between high-level values and low-level stores; two low-level stores could then be said to be equivalent at a type if there exists a high-level value such that both stores represent that value. We do not want to take such a definition as primitive (even though we used something like it in some of the intuitive explanations in Section 7) for a couple of reasons. Firstly, it does not fit with our ‘foundational’ goal of compiling different high-level type systems down to a common language-independent low-level assertion language in such a way that we can justify cross-language linking and specify run-time systems. Secondly, for languages with interesting features (such as higher-order functions and references), the question of what equal means at high-level types is about as hard as, and addressed using the same relational techniques, as what we are doing down at the low-level. Rather than construct a naive denotational semantics for the high-level language, then refine (quotient) it with a state-based logical relation (as in [15], for example) and *then* construct a relation between the refined model and the low level, we just construct low-level relations directly. The hope is that this will ultimately prove simpler and more useful, since encapsulation that is provided by language features (e.g. local references) will be treated in exactly the same way as encapsulation that is used in implementing language features (e.g. environments of closures, memory management). But that will only be tested when we consider more complex source languages.

The treatment of termination in the current work could usefully be refined. The way in which we’ve used perping in the specs means that a program fragment that always diverges will satisfy any pre-post relation pair, irrespective of what effects it has on the store. This is just what one normally intends for interrelations of types in sequential languages with recursion or looping. On the other hand, the same is true of our interpretation of expression typings, whereas evaluation of expressions in this particular source language actually always terminates (unlike most real languages, of course). One could make the low-level semantics closer to fully abstract by using total correctness judgements, involving pre and postrelations, for expression evaluation in place of the ‘double negation’ judgements we’ve used here. There is a similar small weakness in the specification of allocation, which also allows for non-termination. Whilst these differences do not substantially weaken our type soundness result, they do restrict our ability to prove interesting program transformations. Type soundness is essentially about programs being in the diagonal part of the relational interpretations of their types, i.e. being related to themselves. Having a relational interpretation should also allow one to prove that two *different* pieces of machine code (possibly compiled from different high-level phrases, possibly written by hand) are equivalent modulo the contract of a particular type. But if the allocator is assumed to be able to diverge, and we make different calls to the allocator in the two programs, then such proofs don’t go through. Our more recent work uses an allocator specification that does enforce totality, so that we can reason about equations on low-level code.

There is a great deal of related work, of which we can only mention a fraction. Compiler correctness has been studied for at least four decades [28, 32] with notable early examples including the work on Piton and Micro-Gypsy using the Boyer-Moore prover [31, 48] and the manual verification of the VLISP Scheme compiler [24]. More recent work includes that of Leroy on verifying compilation of a C-like language to PowerPC assembly code [27]. Full compiler correctness is a more ambitious goal than type safety, but these projects relate high-level to low-level without making explicit the kind of language-independent low-level contracts that we are formalizing here. (Of course, one would ideally like both, and that is one subject of our current work.)

Reasoning directly about unstructured low-level code also has a long history, going right back to Floyd’s original paper [23]. Once again, the Boyer-Moore prover was used in some notable early mechanization projects [20]. The idea of developing type systems for low-level programs, and preserving typing through compilation, is more recent [34, 33] and has attracted much attention in the context of proof-carrying code [36], as well as in more traditional compiler certification. That low-level types might be given a semantic interpretation in terms of more primitive logical assertions is the key idea of *foundational* proof-carrying code, pioneered by Appel and his collaborators [7]. We have already mentioned that the step-indexing idea we use here originated in the FPCC project; the technique has recently been refined somewhat, replacing literal natural numbers with a modal operator [10], and it would be interesting to incorporate this into our formalization.

Modelling types by partial equivalence relations goes back a long way [41, 6, 22, 3, 21] and, amongst many other things, parametric logical relations have recently been used by many authors in reasoning about program equivalences in the presence of higher order functions and encapsulated dynamically allocated store [38, 40, 15]. Ahmed has used step-indexed relations to give a complete characterization of contextual equivalence in a pure language with recursive and quantified types [5]. Relational program logics have been developed by several researchers [2, 11, 47], though note that the current work does not make use of a specialized relational logic: we just work directly in CiC.

The other main influence on this work is separation logic [37, 42], though we work with relations rather than predicates, and use explicit higher-order parameterization over frames in place of the more usual ‘tight’ interpretation. Recent work on separation logic typing with higher order frame rules [17] and and extensions with quantification [16, 18] are technically very close to the present work, though working on paper and with slightly higher-level languages. Hoare type theory (HTT) is a related mixture of polymorphism, dependent type theory and separation-logic style reasoning about side effects [35].

Acknowledgements. Thanks to Josh Berdine and Andrew Kennedy for useful discussions and feedback on earlier drafts of this work.

References

- [1] M. Abadi. Protection in programming-language translations. In *ICALP '98: Proceedings of the 25th International Colloquium on Automata, Languages and Programming*, pages 868–883. Springer-Verlag, 1998.
- [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 *Proc. 5th IEEE Symposium on Logic in Computer Science (LICS)*, pages 355–365. IEEE Computer Society Press, June 1990.
- [4] A. Ahmed. *Semantics of Types for Mutable State*. PhD thesis, Princeton University, 2004.
- [5] A. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *Proc. 15th European Symposium on Programming (ESOP)*, 2006.
- [6] R. M. Amadio. Recursion over realizability structures. *Information and Computation*, 91(1), 1991.
- [7] A. Appel. Foundational proof-carrying code. In *Proc. 16th IEEE Symposium on Logic in Computer Science (LICS)*, 2001.

- [8] A. Appel and A. Felty. A semantic model of types and machine instructions for proof-carrying code. In *Proc. 27th ACM Symposium on Principles of Programming Languages (POPL)*, 2000.
- [9] A. Appel and D. McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 23(5), 2001.
- [10] A. W. Appel, P.-A. Melliès, C. D. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, January 2007.
- [11] N. Benton. Simple relational correctness proofs for static analyses and program transformations. In *Proc. 31st ACM Symposium on Principles of Programming Languages (POPL)*, January 2004. Revised version available from <http://research.microsoft.com/~nick/publications.htm>.
- [12] N. Benton. A typed, compositional logic for a stack-based abstract machine. In *Proc. 3rd Asian Symposium on Programming Languages and Systems (APLAS)*, volume 3780 of *Lecture Notes in Computer Science*, November 2005.
- [13] N. Benton. Abstracting allocation: The new new thing. In *Proc. Computer Science Logic (CSL)*, volume 4207 of *Lecture Notes in Computer Science*. Springer-Verlag, 2006.
- [14] N. Benton, A. Kennedy, M. Hofmann, and L. Beringer. Reading, writing and relations: Towards extensional semantics for effect analyses. In *Proceedings of the Fourth Asian Symposium on Programming Languages and Systems (APLAS)*, number 4279 in *Lecture Notes in Computer Science*. Springer-Verlag, November 2006.
- [15] N. Benton and B. Leperchey. Relational reasoning in a nominal semantics for storage. In *Proc. 7th International Conference on Typed Lambda Calculi and Applications (TLCA)*, volume 3461 of *Lecture Notes in Computer Science*, 2005.
- [16] B. Biering, L. Birkedal, and N. Torp-Smith. BI-hyperdoctrines and higher-order separation logic. In *European Symposium on Programming (ESOP)*, 2005.
- [17] L. Birkedal, N. Torp-Smith, and H. Yang. Semantics of separation logic typing and higher-order frame rules. In *IEEE Symposium on Logic in Computer Science (LICS)*, 2005.
- [18] L. Birkedal and H. Yang. Relational parametricity and separation logic. In *Proc. Foundations of Software Science and Computation Structures (FOSACS)*, volume to appear of *Lecture Notes in Computer Science*. Springer-Verlag, 2007.

- [19] S. Boutin. Using reflection to build efficient and certified decision procedures. In *Proceedings of the Third International Symposium on Theoretical Aspects of Computer Software (TACS)*, pages 515–529. Springer-Verlag, 1997.
- [20] R. S. Boyer and Y. Yu. Automated correctness proofs of machine code programs for a commercial microprocessor. In *Proceedings of the 11th Conference on Automated Deduction (CADE)*, volume 607 of *Lecture Notes in Computer Science*. Springer-Verlag, 1992.
- [21] L. Cardelli and G. Longo. A semantic basis for Quest. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming (LFP)*, pages 30–43, New York, NY, USA, 1990. ACM Press.
- [22] F. Cardone. Relational semantics for recursive types and bounded quantification. In *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 372 of *Lecture Notes in Computer Science*. Springer-Verlag, 1989.
- [23] R. W. Floyd. Assigning meanings to programs. In J. T. Schwartz, editor, *Proceedings of the AMS Symposium in Applied Mathematics*, volume 19, pages 19–32, 1967.
- [24] J. Guttman, J. Ramsdell, and M. Wand. VLISP: A verified implementation of scheme. *Lisp and Symbolic Computation*, 8(1/2), 1995.
- [25] J. Harrison. Metatheory and reflection in theorem proving: A survey and critique. Technical Report CRC-053, SRI Cambridge, 1995. Available on the Web as <http://www.cl.cam.ac.uk/~jrh13/papers/reflect.dvi.gz>.
- [26] A. Kennedy. Securing the .NET programming model. *Theor. Comput. Sci.*, 364(3):311–317, 2006.
- [27] X. Leroy. Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. In *33rd Symposium Principles of Programming Languages (POPL)*, pages 42–54. ACM Press, 2006.
- [28] J. McCarthy and J. Painter. Correctness of a compiler for arithmetic expressions. In *Mathematical Aspect of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*. AMS, 1967.
- [29] N. McCracken. *An investigation of a programming language with a polymorphic type structure*. PhD thesis, Syracuse University, June 1979.
- [30] P.-A. Melliès and J. Vouillon. Recursive polymorphic types and parametricity in an operational framework. In *IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, 2005.

- [31] J. Strother Moore. A mechanically verified language implementation. *J. Autom. Reason.*, 5(4), 1989.
- [32] F. Lockwood Morris. Advice on structuring compilers and proving them correct. In *Proceedings of the 1st ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)*, 1973.
- [33] G. Morrisett, K. Crary, N. Glew, and D. Walker. Stack-based typed assembly language. *Journal of Functional Programming*, 12(1), 2002.
- [34] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 21(3), 1999.
- [35] A. Nanevski, A. Ahmed, G. Morrisett, and L. Birkedal. Abstract predicates and mutable ADTs in Hoare type theory. In *European Symposium on Programming (ESOP)*, volume to appear of *Lecture Notes in Computer Science*. Springer-Verlag, 2007.
- [36] G. Necula. Proof-carrying code. In *Proc. 24th ACM Symposium on Principles of Programming Languages (POPL)*, 1997.
- [37] P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Proc. 10th Annual Conference of the European Association for Computer Science Logic (CSL)*, volume 2142 of *Lecture Notes in Computer Science*, 2001.
- [38] A. M. Pitts and I. D. B. Stark. Operational reasoning for functions with local state. In *Higher Order Operational Techniques in Semantics*. CUP, 1998.
- [39] G. D. Plotkin. Lambda definability and logical relations. Technical report, Department of AI, University of Edinburgh, 1973.
- [40] U. S. Reddy and H. Yang. Correctness of data representations involving heap data structures. *Science of Computer Programming*, 50(1–3):129–160, March 2004.
- [41] J. C. Reynolds. Types, abstraction and parametric polymorphism. *Information Processing ’83*, pages 513–523, 1983.
- [42] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. 17th IEEE Symposium on Logic in Computer Science (LICS)*, 2002.
- [43] D. Scott. Data types as lattices. *SIAM Journal of Computing*, 4, 1976.
- [44] G. Tan, A. Appel, K. Swadi, and D. Wu. Construction of a semantic model for a typed assembly language. In *5th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*, January 2004.

- [45] J. Vouillon and P.-A. Melliès. Semantic types: A fresh look at the ideal model for types. In *Proc. 31st ACM Symposium on Principles of Programming Languages (POPL)*, 2004.
- [46] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 15 November 1994.
- [47] H. Yang. Relational separation logic. *Theoretical Computer Science*, 2004. Submitted.
- [48] W. D. Young. A mechanically verified code generator. *J. Autom. Reason.*, 5(4), 1989.